

A Technology for Supporting Knowledge Work: The RepTool

Brigitte Jordan

Xerox Palo Alto Research Center and Institute for Research on Learning

Ron Goldman

Bell Atlantic Science and Technology and Institute for Research on Learning

Anja Eichler

Witten-Herdecke University and Xerox Palo Alto Research Center

Chapter for Information Technology for Knowledge Management

Uwe Borghoff and Remo Pareschi, eds.

Springer Verlag, Berlin, 1998

Abstract: In this chapter we report on the RepTool, a computer based tool for the collection, analysis, and presentation of data about “worksapces”. Users of the tool collect relevant information about workplaces and workpractices in a central data structure on which the tool provides different views. Physical, social, and cognitive spaces in users’ or customers’ environments can be represented and analyzed. The RepTool supports systematic data collection and collaborative data analysis that can help people at all levels of an organization build a shared view of the formal and informal work processes they participate in. A tool for collaborative knowledge management, it nurtures intra- and inter-team conversations that can lead to an empirically grounded, shared view of current realities and potentially necessary changes.

Key Words: workscape analysis, formal and informal work practices, systematic data collection; collaborative data analysis; collective inquiry; conversational boundary objects; knowledge work; technology infrastructure; communication technologies; participatory knowledge management; team ownership; cross-organizational conversations

1 Introduction

Knowledge work characteristically takes place in multi-person work groups such as collaborative work teams, communities of practice or other kinds of social units. The activities of such groups are located in “worksapces” or “work ecologies” that include the totality of affordances and constraints provided by particular physical settings, social relationships, endogenous work practices, official work processes, and the customary artifacts and technologies that are instrumental in accomplishing the work at hand.

At present, most representations of work are prescriptive and process-oriented. In particular, work groups and their often informal work practices are not well captured by current tools. As we attempt to go beyond sequential workflow modeling and look at knowledge work with a holistic workscape perspective, we need tools that support this perspective.

Ideally such tools would show a number of specific characteristics. Minimally, they should:

- be equally useful for workers, analysts, and managers

- be usable by non-specialist participants under field conditions
- be tailorable and adaptable to the requirements of different users
- support systematic data collection
- make data available synchronously and asynchronously
- allow collaborative data analysis
- foster collective inquiry by building team-shared views of conventional or virtual workspaces
- produce representations that can act as “boundary objects” conducive to reflection among stakeholders
- support a collaborative (re)design of work processes, work places and work technologies.

When analyzing workspaces from a holistic perspective many different kinds of information might need to be considered. Workspace analysts might want to be able to investigate issues such as:

- What is the physical space like: Who inhabits it; what objects and technologies are located in it; how much physical space is available to carry out the work at hand.
- What are the social relationships between people: Who reports to whom; who works with whom; who goes bowling with whom; who asks whom the questions no M&P manual answers; who supports what projects.
- What are the relationships between the various communication and computing technologies present in the workplace: Which workstations have access to which servers; which workstations share the same software; which computers are linked to which printers; or whose desktop conferencing setup can connect to whom.
- What is the cognitive space for this work setting like: What sorts of projects, ideas, enterprises does it contain; where are they located; who works on them; what technologies do they require; what other projects or ideas are they linked to.

We have designed a representational tool, the RepTool, to help investigators explore these kinds of questions (Jordan, Goldman and Sachs 1995). As a tool for the collaborative collection, analysis and representation of information, it can aid knowledge workers and analysts of workspaces to better understand and shape particular work settings. Since the systematic representations that are produced when using it can be dynamically updated and shared among stakeholders, the tool supports reflection-in-action and facilitates conversations leading to collective interpretation and sense making.

Michael Schrage, in “Shared Minds: The New Technologies of Collaboration” predicts that collaborative tools will “converge to create new collaborative environments” and that “people will discover new modes of interaction, new styles of creativity and conversation, and a new desire to work with people to solve problems and craft innovations” (Schrage 1990:185). The RepTool is intended as a step in that direction.

2 Views provided by the RepTool

The RepTool offers a variety of different views on a tailorable, central data structure in which users/investigators collect information relevant to their concerns. What makes working with the RepTool data structure different from working with many other kinds of databases is:

- A high level of tailorability allowing users to change the structure of the information they are collecting on the fly.
- Ability to provide different views on the information collected: Printing or otherwise displaying the central data structure presents a verbal view of the data. A visual graphic view is provided by the fact that the same information can be displayed as a map that shows the objects of interest, be they artifacts or people, and optionally the relationships between these objects.
- A highlighting function that supports the combination of verbal and visual spaces: It enables the user to run queries on the data structure and display results on the map as highlighted objects or by drawing arrows and lines to indicate relationships. Thus different views can be explored that may lead to new insights.

Users typically will want to have information on the physical space, the social space, the technology space and the cognitive or idea space of the workspaces they are concerned with.

The physical space: Users of the tool might want to see the physical layout of a workplace in a map-like representation, complete with desks, doorways, windows, people, documents and relevant technologies. This would allow them, for example, to highlight the office locations of all persons who work on a particular project in order to study the effects of collocation. When trying to understand activities on a production floor, they might want to see all locations through which a product travels.

What differentiates these representations from conventional sketches and maps is that each object in such a representation is “alive”, that is to say can be associated with a record in the central data structure that contains additional information about the object.

The social space: At the same time, an analyst might want to generate a representation of the work-relevant social relationships that exist in the workplace, be they friendship networks, mentoring relationships, kin networks, trade relations, co-residence, bowling team membership, alma maters, or current and previous work associations. The RepTool can represent reciprocal (e.g. friend/friend) relationships as well as complementary ones (e.g. supervisor/subordinate). In addition to formal reporting relationships expressed in an organization chart, analysts might want to understand such things as who gets advice from whom; who says they work with whom; who mentors whom; who are the informal leaders that people go to when they are stuck with a difficult problem, and the like. This kind of analysis can be used to examine an organization from many angles in order to identify key players, bottlenecks, communities of practice, communications between teams, and social networks (Stamps 1997).

Simultaneously representing the social and the physical space allows investigators to map the movement of various kinds of objects through the company. One could, for example, look at the distribution of documents in relation to the people who need access to them, at the movement of hoarded spare parts through the social channels of a service organization, or the informal exchange of favors, factors that lead to a better

understanding of how an "informal economy" is constituted in the workplace. By examining what is normally invisible in the organization, it becomes more possible to see how the organization really works and to show these findings to others in a convincing way -- as printouts on paper, as overheads or in live electronic form.

The technology space: a representation of the infrastructure consisting of computing and communication technologies such as fax machines, scanners, computers, speaker phones, computerized whiteboards, video conferencing setups and the like allows investigating which machines are connected to which other machines, which computers run what software, or who is connected to whom through what kinds of communication devices. This might be useful for such things as tracing breaks and impasses in information flow in the company.

The cognitive space: Non-physical spaces can be represented in the RepTool as well. For example, managers might want to map out a company's "project space" in order to identify project success factors or cross-division connections over projects; corporate lawyers might want to represent a company's patents together with known infringements in a "patent space"; and researchers might want to model the emergence of particular ideas across a research lab.

In sum, since the RepTool allows us to map out all manner of physical, social, organizational and idea spaces it should prove useful in all cases where investigators need to understand what kind of formal and informal situated knowledge work goes on in their own workspaces and those of their customers.

3 Technical Details

The RepTool consists of an object-oriented database connected to a drawing tool. The user gathers information about a workplace, stores it in the database, and then creates multiple graphic views, or maps, to represent different aspects of the work. Values and relationships in the database are then used to highlight the objects displayed in these graphic representations and to draw lines between them.

The RepTool is designed to be easy to learn and to be easy for the user to tailor to his or her needs. We have been using an iterative, participatory design process for its development. After identifying user needs, new features are designed and implemented. Every four to six weeks a new prototype is released for more testing and the process is repeated. The RepTool was developed in a collaboration between IRL and Bell Atlantic (formerly NYNEX Science and Technology).

3.1 The RepTool Database

The RepTool uses an object-oriented database with no inheritance. The database consists of one or more record types, each of which is made up of a list of data fields and a number of individual records of that type. For example a database may have records for people, computer systems and projects as shown in figure 1. The record for a person might contain their name, position, phone number and pointers to all of the projects they are involved with.

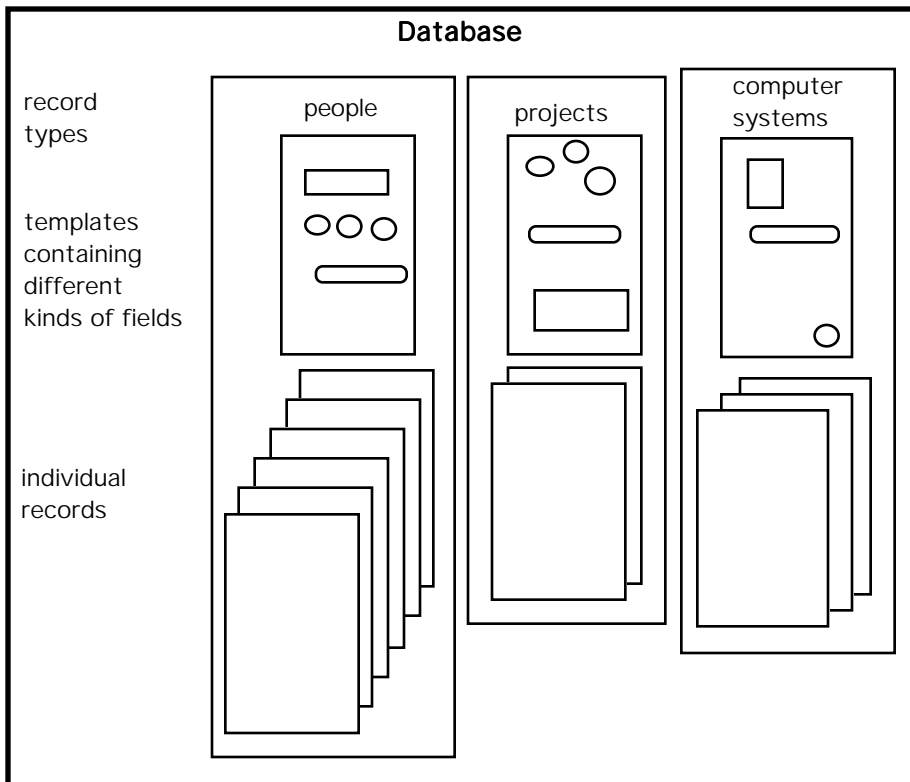


Figure 1: Database structure

The fields in each database record can contain one of the following types of data:

- text this can be a short string, like a name, a multi-line comment, or a longer text note that will appear in a separate window
- number this can be a simple integer or one with decimal digits (e.g. 3.14)
- date/time this consists of the date (i.e. day, month, year) and/or time of day (i.e. hour, minute, second)
- Boolean this is for yes/no or true/false values
- multiple choice this can be either one or several choices from a set
- pointer this is a pointer to another record in the database; a pointer field can point to one or several other records

Figure 2 is an example of a record for a fictional employee showing how the fields appear when displayed.

The screenshot shows a window titled "Record: Brown, Bill" with a "Cancel" button and a "Done" button. The form contains the following fields and controls:

- First Name:
- Last Name:
- Age:
- Sex: Male Female Don't Know
- Department: (dropdown menu)
- Committees: Change Management ISO 9000 Entertainment
 Long Range Planning New Products
- Hire Date:
- Reports To:
- Works With:
- Work Start Time:
- Bowling Team Member: True False Don't Know
- Comments:
- Field Notes:

Figure 2: Example database record

The database is designed so that it is easy for a user to add new fields or edit old ones at any time. All of the records in the database are immediately updated with the new fields.

The RepTool also allows the user to describe relationships between records in the database. These relationships are shown graphically as lines connecting the objects associated with

the records. This is useful for showing relationships such as who reports to whom, which databases does a computer system have access to, where does a technician go for help on questions, etc. The information connecting the records is stored in pointer fields in one or both of the records.

Currently the database is memory resident and only supports one user or team at a time. It is possible though to merge records from one database into another. Eventually we intend the RepTool to support simultaneous use of the shared database by multiple users or teams.

3.2 Graphic Views and Maps

To create a graphic representation the user can select an object and place it on a map. Objects can be moved, rotated, resized, duplicated, etc. Text labels can be added as can notes, which appear like a yellow post-it note and when clicked on open a separate window for longer comments. The display can be zoomed in or out.

Figure 3 is an example of a simple office map created with the RepTool. The circles represent people. Double clicking on the note labeled "Fieldnotes 9/22/97" would open a window containing the text of the fieldnotes.

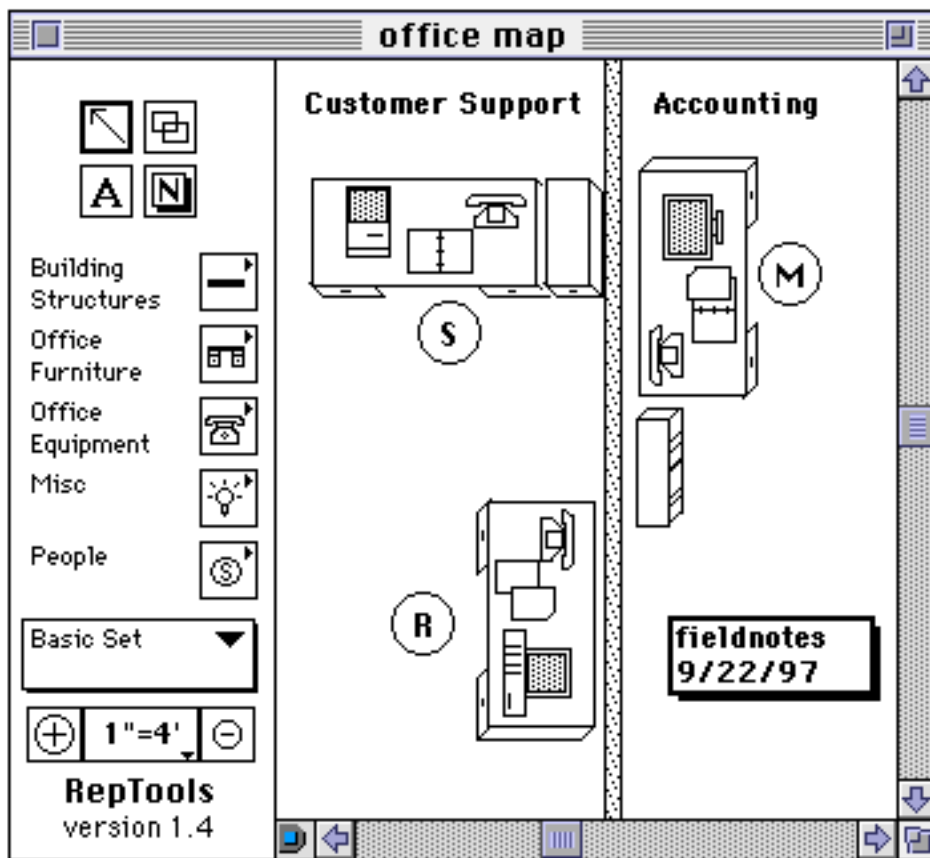


Figure 3: Sample map

The RepTool currently comes with about three dozen predefined objects such as walls, desks, tables, telephones, computer terminals and employees. When additional objects are needed they can easily be defined by the user. Objects are made up of simple drawing primitives such as lines, rectangles, ovals, arcs, polygons and text. For those objects that are associated with a record in the database, the text displayed can be the value of one of the record's fields.

Each map is stored in a separate file and consists of a list of the objects in it and information on how these objects are to be highlighted (described below).

Slide shows can be constructed specifying a series of maps to display. Each "slide" specifies what part of the map to display, what scale to display it at and what highlights to use when drawing the objects and relationships between them.

3.3 Connection between the database and graphical representations

After creating a database and a map the user can associate objects on the map with records in the database as shown in figure 4.

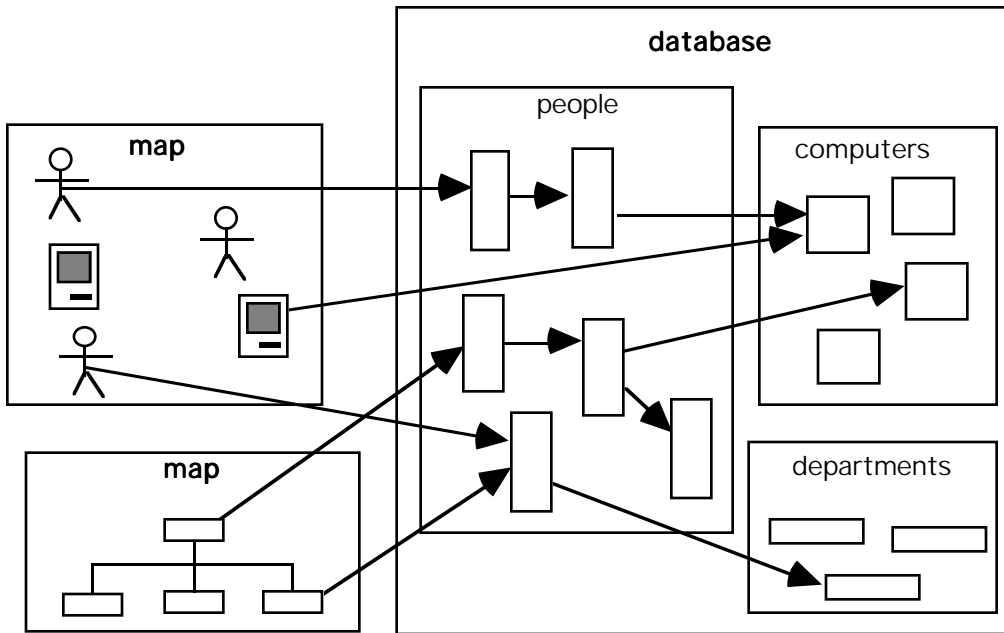


Figure 4: Connecting maps and the database

These objects can then be connected by lines showing the relationships between the records they are associated with. For example, in figure 5, the lines show the flow of orders through four legacy computer systems. The single circle “R” shows a customer representative who must take orders from the SOP system and manually enter them into the WFA-D system -- a potential bottleneck where system throughput might be improved.

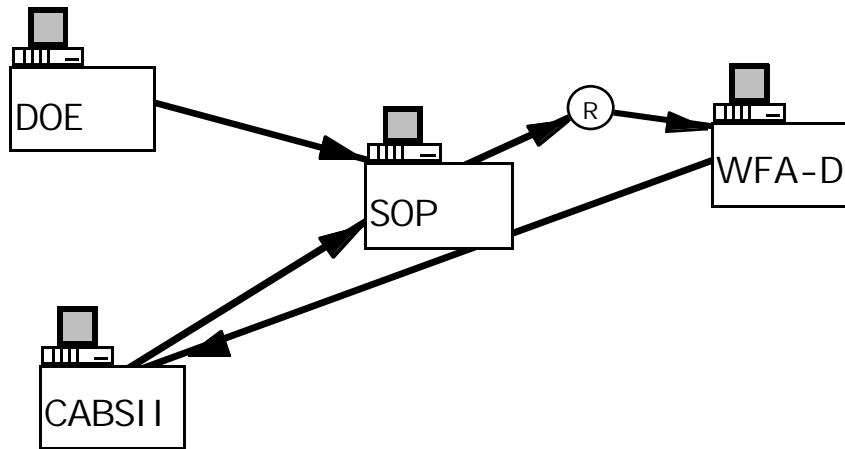


Figure 5: Displaying relationships between objects

The RepTool can selectively highlight objects based on the value of fields in the associated database record. This is done using a special highlight editor to specify the color to use and the record fields to examine. Using this editor both simple and complex queries into the database can be created. A simple example is shown in figure 6.




Figure 6: Simple database query

For example, we might highlight in red all those employees who have worked more than seven years for their company and highlight in green all those who have worked there less than three years. Relationships can also be selectively displayed. For example, we might ask the tool to draw lines between each technician and the people they go to for advice on technical problems; and specify to draw the line in red if the other person is also a technician and draw it in blue if the person is an engineer.

It is this ability to easily alter the graphic display based on the underlying database that distinguishes the RepTool from other similar applications. Conventional drawing/presentation programs, such as PowerPoint or Visio, provide greater graphic possibilities, but the resulting representation is static (what you see is all you get). Geographical information system (GIS) software has an underlying database and can highlight objects, but cannot display or represent relationships between the objects. The RepTool can use the information in the database to both highlight objects and show relationships between them. As we continue to develop the RepTool we will be building on

this strength to allow users to create new types of graphic representations based on information stored in the database.

3.4 System Requirements

The RepTool currently runs on Macintosh computers. We are looking into porting it to the PC. It is written in C++ using Symantec's Think Class Library. It requires at least 2 Mbytes of memory.

4 Working with the RepTool

The following section describes how the RepTool could be employed in an advertising agency. In this particular company it is notable, first of all, that employees consciously use the RepTool for defining their relationship to their work community. The central data structure has been designed by them collaboratively and represents categories of information about anything they think might be relevant to their working together. This includes such things as their skills, likes and dislikes, past projects, education and professional background, hobbies and anything else they want their colleagues and management to know about them. Employees update their entries whenever changing interests and developing expertise make that advisable. They are free to fill in fields in templates or leave them empty. On occasion, there is a proposal to rename, delete or add new record types and those decisions are always made collectively. The RepTool and its associated views of the company have proven especially useful for newcomers who have profited immensely from being able to familiarize themselves with their new environment and from the challenge to represent themselves.

Beyond this mirror on the company, teams within the company also use the RepTool very effectively as a tool for generating shared views and organizing their thinking about particular issues as they come up. In the following sections we are going to look at the use of the RepTool in detail by exploring two different settings: The first shows potential RepTool use in a client-focused campaign in the agency; the second revolves around the ways in which the tool can support the planning and decision-making activities involved in the agency's day-to-day operations.

4.1 Supporting teams in a client focused campaign

A team of five people plans to work on an advertising campaign for winning a new account. For the first meeting the following people are present: Andy Nord, head of the graphic design department, Loraine Martins, graphic designer, Marco Terrotti, the account manager, his assistant Marc O'Brian and Lindy McShirly, an intern. The purpose of this meeting is to share what information they have individually, such as which people they know in the target organization, the work of the agency that has been in charge of the client's product campaigns so far, what kinds of campaigns the client prefers, and the like. It turns out that collectively they already hold a lot of information but what they have resembles a patch work quilt rather than a systematic data collection effort. In order to build a winning strategy it will be necessary to get a shared view on the situation. But which is the relevant data? The team engages in a lively discussion about what they need to know that they don't know yet and explore what skills and other resources they might need.

Working on the data already assembled, the group first tries to figure out which are the major issues they need to be concerned with. The first record type they make up in the RepTool's data structure concerns their potential client. They start to define several fields

which are going to capture the client's profile: The name and address of the company, the type of business in which it is engaged and other items are entered into the data structure. After this the team pulls down an organizational chart of the client's company to check if they already have contacts to anybody in important positions. It turns out there are two persons who they know already.

Marc O'Brian suggests to create a new record type named "contact persons" to gather information on this issue. But Andy Nord is opposed: He wants to be able to have a look at these data in close connection to the information concerning the client. Marco Terrotti proposes to connect the two records through a pointer field. A double click on this field would switch between the two records. Andy Nord agrees. The team goes on to discuss what kinds of information about contact persons might be relevant to their campaign and how to collect it.

After they have done this, they take a few minutes to look back on what they have accomplished so far and what they need to do next. Lindy McShirley asks "What about the client's former advertising? Who is our competition and what have they been doing?" Marco Terrotti proposes they analyze the three latest campaigns which everybody thinks is a good idea. For a few moments several people talk at the same time until Marco Terrotti intervenes with the suggestion to make up a new record type named "Former campaigns" which contains records for each campaign. Thus they would be able to gather information more systematically.

During their conversation, Loraine Martins has already scribbled down some ideas for the new campaign which she now presents to the group. A lively discussion starts. After a few minutes of brainstorming the group has developed four competing ideas on how they might design their campaign. One member of the group proposes to make up a new record type named "envisioned campaign" in the central data structure to capture their ideas. This record type will be the template for the four proposals on what the envisioned campaign might look like. Again the team has to decide what information might be relevant to capture their emerging ideas. They decide to give a short description of the idea and to gather some data about the resources they would need for carrying it out. Marco Terrotti and Loraine Martins start to work on budget estimates for the different proposals. Finally, the team adds a comment field where they can collect further ideas as they emerge before their next face-to-face meeting.

The group is now about to finish their first session. Marc O'Brian raises the question how they can make their strategy clear to people who are going to join the group in the future. After some silence Loraine Martins suggests to make up a RepTool map to explain how they proceeded in building up their strategy. She sketches her idea on a piece of paper as she talks about it. The other team members think her sketch represents their ideas well and she maps it in the RepTool. Figure 7 shows the result.

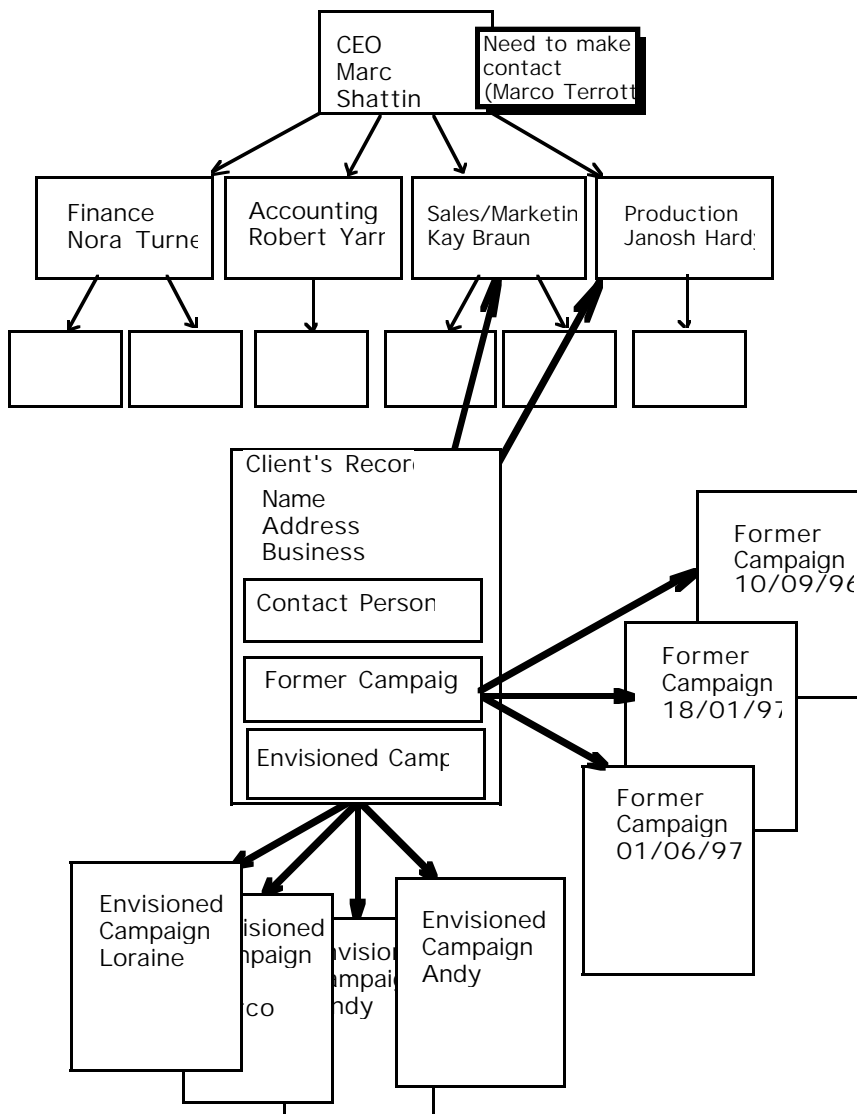


Figure 7: Building an information structure with the RepTool

The RepTool supports the process of systematic, collaborative data collection as it structures the course of discussion between team members. The team needs to discuss what information is important for their particular purposes in order to configure the tool's central data structure appropriately. When in former meetings the exchange of opinions and views had sometimes been chaotic as everybody had contributed what came into his or her mind, team members now try to fill in the fields they have jointly defined as relevant. The RepTool supports the group in brainstorming ideas, organizing ideas, evaluating ideas and, as a result of these activities, generating an outline for the envisioned campaign. It thus helps the group shift the character of their discussion from merely exchanging information to building up a shared view. Furthermore, as all records of a certain type share the same template, the RepTool supports the group in collecting data systematically. This allows the

team to gain new views and insights as it does not only benefit from the analyses of each individual campaign but from a comparison between them.

In the future course of the project the RepTool will make it possible for them to share their data with people who could not attend the first meeting but who will carry out important tasks. Moreover, the RepTool makes the gathering of additional data more efficient and more systematic. If team members go out to visit the client, they can use RepTool-generated representations to check their perceptions of client needs with the client. In such sessions, the templates in the records will remind team members of the data they still need to collect. Finally the resulting information collected by different team members, possibly at different times and in different locations, will be directly comparable. In fact it will all go into a single database which will make it easy to share field data among team members. In the next session, team members' attached notes will be a good starting point for generating a joint discussion.

In the following section we will see how the RepTool supports people in running the advertising agency in its daily operations.

4.2. Supporting operational work

A few steps away from the conference room where our team is working on the outline of their emerging campaign, Robert Lurry, in charge of facilities, is sitting in front of his computer. He is studying space utilization on the second floor. The group of Andy Nord, the head of the graphic design department, has acquired several new members and thus needs more space. In the past, problems had occurred each time people and groups had to move. Sometimes they found that after the move they didn't have enough space to spread out their working tools in their new environment or additional cables had to be pulled through the whole building to give people access to their network. And often people complained after a move that they now had to cover long distances to maintain face-to-face contact with other teams they were collaborating with.

This time Robert is going to find out how Andy's group uses their current workspace and with whom they are collaborating before he is going to make any plans for moving them. The RepTool supports him in doing spot observations on the second floor and mapping them to see how the various work spaces are used and by whom. He also queries the central data structure to see which other groups Andy Nord's team is collaborating with. On the screen, he provisionally tries to move the graphic design department to the third floor and has to notice that this does not seem to be a good idea. As he had mapped all facilities to scale he can see that the cubicles on the third floor are too small to accommodate the technical equipment of the group. Furthermore Andy Nord's group collaborates with a team of texters on the second floor, the people who create the text for his advertisements, so that makes a third floor location even more undesirable.

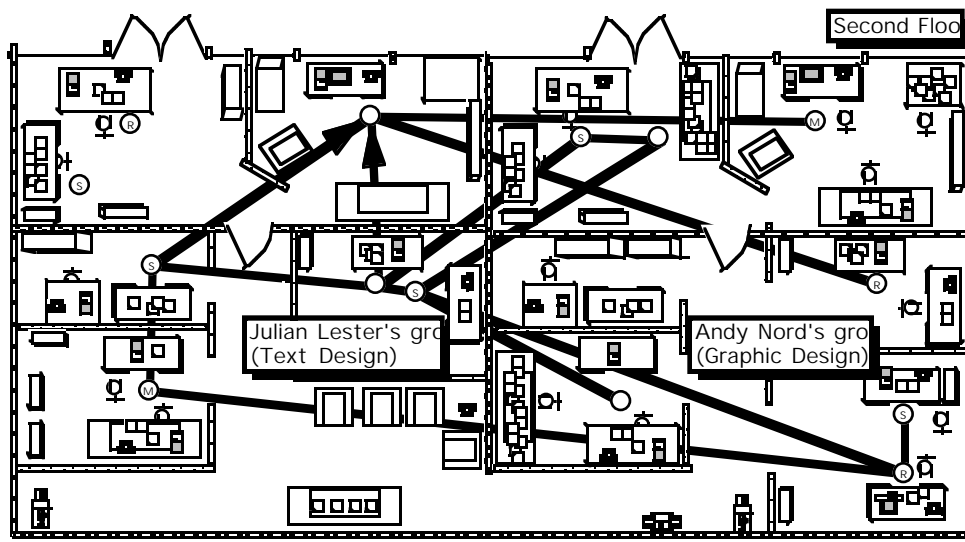


Figure 8: Representing working relationships between people with the RepTool

Robert looks for other possibilities and finally decides to make a proposal for moving the accounting department from the second to the third floor. He quickly sends email messages to several people: Andy Nord, Sandy Simmon, the head of the accounting department, and Alice Turner, the computing infrastructure analyst, asking them to pull up the proposed solution in a RepTool map and comment on it.

Alice Turner, the computing infrastructure analyst, routinely uses RepTool maps for resource analysis. She has just got Robert's mail and checks out if she has enough workstations for the new members of Andy Nord's team. She uses a map on which she can display all employees and all computer terminals and data structures. With the help of the map she wants to look at which computer systems Andy Nord's team uses. She starts a query with RepTool's highlight function. She first highlights all terminals on the second floor that are actually used by someone. She finds out that in one of the rooms there is a workstation nobody seems to be working on. In a second step she highlights on her map which databases can be accessed by this computer terminal and queries which data bases the members of Andy Nord's team use. It turns out that the computer system is not connected to the data bases Andy Nord's team uses. Alice attaches a note to the map referring to this and sends an email message to Robert to ask him about solutions. Maybe it would be easily possible to connect the free terminal with the server for the needed data bases.

In the meantime Robert Lurry, the facilities man, got feedback on his proposal for the move. Andy Nord agrees, but Sandy Simmon, the head of accounting, is pretty upset about the idea of moving to the third floor. She claims moving would cost her department many days of work. Furthermore she states that the art directors' clique is always given priority space whereas the work of her department is constantly underappreciated. She proposes that Robert Lurry himself with his staff should move to the third floor. Sandy also used RepTool maps to support her argument regarding the best solution to the space problem. Robert decides to arrange a meeting between himself, Andy Nord, Sandy Simmon and CEO Leigh McMiller. In this meeting he wants to discuss the reasons why he thinks that moving the accounting group might be the best solution to solve their space problem. To support his argument he will also show other possibilities for moving people around the building, which seem to him to be even more undesirable than moving the accounting

department. Robert prepares further maps and compiles them in a slide show. He hopes that it will be possible to eliminate some of the rancor that has been creeping into their exchanges by inviting everybody to a face-to-face meeting where they could refer to the objects on the map and the information associated with them. During their conversation, it will be possible to show different views and to highlight objects and show relationships according to the individual questions of each. Thus the RepTool might help to keep these discussion on a factual level and prevent them from getting personal and emotional.

As the RepTool provides its users with multiple views on the data they have collected, it merges the processes of data collection, data representation, and data analysis to expose patterns and relationships. RepTool's central data structure is designed to be shared amongst team members so that it also provides them with the ability to reflect on different data sets collected from similar environments, possibly by multiple investigators at different times, thus facilitating agreement on important patterns and relationships. It supports building an empirically grounded, shared view of the work practices, physical spaces and social relationships that define the workscape. Thus it becomes a tool for learning in organizational team settings.

Even though our example dealt with an advertising agency, the RepTool is of course equally useful for other organizations that need to go beyond the official documentation of work processes to understand how work actually happens in real work situations. The following sections suggest other work environments in which the tool might be of use.

5. Using the RepTool

We believe that the RepTool would be extremely useful in all situations where a quick, principled, assessment of the important features of a work site is necessary. In particular, the RepTool allows representing the often neglected but nevertheless crucial informal aspects of work, such as who seeks advice from whom; who is collocated with whom; who went to the same school, and the like. Since the database is built up by the user with locally relevant information, it supports asking locally relevant questions. This gets us beyond the generic "plain vanilla" approach of most workflow tools. We anticipate that the RepTool could substantially improve the productivity of individuals and teams who need to make decisions for which a detailed knowledge of particular worksapes is essential.

Imagine the following situations...

- a sales team strategizing about a new account
- a group of software designers worried about an application's fit into users' ongoing work
- a facilities manager having to relocate several teams
- a sales team preparing for a new account presentation
- a department manager trying to figure out where to locate a new employee
- a reengineering team consolidating different functions in a central call center
- professional ethnographers collecting data in the field

- customer analysts reviewing customer needs
- management consultants preparing feedback to their clients
- real estate agents answering potential customers' questions about properties
- and many others

6. RepToolian visions: the RepTool in the world

6.1. Views of knowledge and knowledge management

There appear to be two paradigms for knowledge management. The one that has the greatest currency at this point is concerned with managing information contained in some kind of depository or other, be that a set of electronic data bases, a professor's collected lecture notes, or a company's archives going back 37 years stored in a horse barn that nobody dares dispose of because of fear of a law suit. This archival paradigm is well represented in corporate processes, management consultants' expertise and a variety of tools. The second paradigm for knowledge management has to do less with the administration of existing information but rather with the creation of new knowledge and knowing practices and the construction of workspaces where knowledge is generated. This is management *for knowing* rather than management *of knowledge*. This kind of knowledge management to date is poorly understood and ill-supported by methods and tools.

Another productive way of thinking about the role of tools in knowledge management is to consider what sorts of knowledge and knowing practices are readily visible in an organization, are easily identifiable, and draw in general the greatest effort. In figure 9, they appear in quadrant one. They are concerned with all those things that information systems have dealt with since time immemorial, from accountants ledgers to university libraries to training modules stored in three ring binders and CBT modules. This quadrant is well supported by current information technologies, search engines, access control algorithms and the like.

The second quadrant is concerned with what an organization knows it needs to know but doesn't currently, the knowing practices that deal with updates on the current state of the organization and it's environment, market intelligence, information about suppliers, and so on. Here, too, we have well fairly well established methods and tools.

The third quadrant becomes more problematic. It deals with what is there but often invisible to management and outsiders, the tacit knowledge present in communities of practice where the real work gets done. Given that the very existence of this kind of knowledge is overlooked in most organizations, there is little support for it except by default.

Finally, the fourth quadrant concerns the sphere of the unimaginable, the conditions and states an organization doesn't even know it might want to know about. Beyond futures scenarios as a way to open up this dimension, there is no support there at all.

	WHAT WE KNOW	WHAT WE DON'T KNOW
WE KNOW	<p>Answered Questions: what we know that we know (expertise in hand)</p> <ul style="list-style-type: none"> • <i>depositories</i> • <i>archives</i> • <i>data banks</i> 	<p>Unanswered Questions: what we know that we don't know ... but have ways of finding out about; expertise we know we have to get</p> <ul style="list-style-type: none"> • <i>market intelligence</i> • <i>state of the organization</i>
WE DON'T KNOW	<p>Unquestioned Answers: what we don't know that we know</p> <ul style="list-style-type: none"> • <i>local knowledge</i> • <i>tacit knowledge and practice</i> • <i>competence in invisible communities of practice</i> 	<p>Unquestioned Questions: what we don't know that we don't know</p> <ul style="list-style-type: none"> • <i>the unimaginable</i> • <i>the questions we never thought of asking</i> • <i>an evolutionary quantum jump</i>

Figure 9: Knowing and not knowing

The RepTool has a number of characteristics that allow it to address at least three out of the four quadrants and maybe even provide the foundation for productively thinking about the fourth. Let us summarize then, some of the ways in which the RepTool can support knowledge management.

6.2 RepToolian data collection: systematic and relevant

It is difficult to manage patchy, irrelevant data, though that is what most organizations have to deal with much of the time. The RepTool is designed in such a way that the categories for data collection have to be determined in each case by the analyst (or team of analysts) as they decide what information is important for their particular purposes. Since they themselves decide how to configure the tool appropriately (by setting up records and making templates), only data relevant to the group will be assembled. Given that the tool makes it easy to display which fields have no entries, it also becomes possible to see where data is lacking and to remedy that situation. The tool also eliminates duplication of effort since individual investigators' data structures can be merged so that all team members have access to the collective effort. Organizational theorists have long argued that practically all corporate data collection efforts suffer from too much data -- data that is, to boot, often irrelevant and unlikely to get acted on (Feldman and March 1981). The RepTool enforces thinking through what might be required in a particular situation and then keeps track of the extent to which that goal has actually been achieved.

6.3 RepToolian conversations

A central feature of the RepTool is that it easily produces electronic and paper representations that can act as conversational "boundary objects" (Star and Griesemer 1989); that is to say, they can support a variety of conversations among users, stakeholders and outsiders who have an interest in the particular workscape and its issues. One great advantage is that the e-representations, viewed either on a screen or, for a larger audience, projected onto a wall, are "alive"; that is to say, if questions arise, it is easy to click on an

object and pull up relevant contextual information from the data structure or to switch to a different view altogether. For presentations, a slide show can be assembled which has the same advantages.

Clearly, e-representations have the greater functionality, but we have found that even paper printouts of RepTool-generated views can be a powerful conversation stimulator that brings agreement to the surface and makes differing perceptions of reality visible for resolution. That is as true for internal discussions among analysts as well as for checking the quality of the information collected with inhabitants of the workscape

The RepTool also facilitates participation of workers in the design process. On-site workers can be actively drawn into the data collection and analysis process through the representations created.

6.4 Support for participatory knowledge management

Field conversations facilitated by RepTool maps provide important validity checks but also establish a relationship between analysts and the people in the workplace that gives workers a stake. In some cases they may become involved in data collection themselves, suggesting new records and fields and even taking over data entry. This gives the analyst a more valid, co-constructed picture while workers gain a comprehensive picture of their own worksite and begin to own the representations of their work as they are drawn into the data collection and analysis process. Workers thus acquire the necessary skills for documenting the realities of their own work and that of their peers, to represent *their* view of work practices, social relationships and productivity impasses to relevant others such as managers, system designers, technologists and decision makers.

Using the RepTool to its greatest benefit in organizational settings probably means giving employees continuous access to what information the central data structure contains about them, so they can update it with information they want their peers, managers, and the organization as a whole to have about them and the workscapes they are a part of. The RepTool thus can play an important role in generating a common, shared, and in some sense more “accurate” view of the distribution of resources and skills. It can generate shared knowledge about the state of the organization that is “owned” by the community as a whole, thereby building employee morale and involvement. Such a tool can also stimulate and enable conversations between different levels of the organization, from front-line workers to senior management, thereby nurturing cross-organizational conversations that build a shared view of current realities and necessary changes. For example, a printout of an organization’s formal reporting structure juxtaposed to a worker-assembled view of informal leadership in the workplace may generate fruitful discussions about how these two kinds of structures can profitably interact.

As Senge et al. have proposed for dialogue (1994), we would suggest that the RepTool can serve as a tool for improving organizations, enhancing communication, building consensus and solving problems. As people work with the tool, they learn how to think together, not just in the sense of analyzing a shared problem, but also in the sense of occupying a collective sensibility in which the thoughts, emotions and resulting actions belong not to one individual but to all of them together.

7. REFERENCES

- Feldman, Martha S. and James G. March
1981 Information in Organizations as Signal and Symbol. *Administrative Science Quarterly* 26:171-186.
- Jordan, Brigitte, Ron Goldman and Patricia Sachs
1995 Tools for the Workplace. Sidebar. *Communications of the ACM*, 38:9:42 (Sept.)
- Schrage, Michael
1990 *Shared Minds: The New Technologies of Collaboration*. New York: Random House
- Senge, Peter, Charlotte Roberts, Richard Ross, Bryan Smith and Art Kleiner
1994 *The Fifth Discipline Fieldbook: Strategies and Tools For Building A Learning Organization*. New York: Doubleday
- Stamps, David
1997 Off the Charts. *Training Magazine*, October 1997:77-83.
- Star, Susan Leigh and James R. Griesemer
1989 Institutional Ecology, 'Translations' and Boundary objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-1939. *Social Studies of Science* 19:387-420