ADAPTING KNOWLEDGE MANAGEMENT SYSTEMS FOR COLLABORATIVE TECHNICAL PROBLEM SOLVING

DINESH RATHI

Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign, 501 E. Daniel, Champaign, IL 61820, USA
E-mail: drathi@uiuc.edu

MICHAEL TWIDALE† and VANDANA SINGH‡

Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign, 501 E. Daniel, Champaign, IL 61820, USA
†E-mail: twidale@uiuc.edu, ‡E-mail:vsingh1@uiuc.edu

The paper considers the potential of a representation to facilitate the sharing of technical how-to help in workplace environments. The method uses screenshots of trails through a computer application made by others and sharable through a Knowledge Management System. A pilot study was conducted comparing this type of help both with conventional built-in application help files, and with the ad hoc individualized face-to-face help provided by colleagues in informal workplace knowledge sharing interactions. The findings indicate the potential of the method as an additional knowledge sharing resource, both as a standalone method and as a way to reduce the interaction time during informal face-to-face help interactions. The study also reveals the power of a lightweight pilot study for uncovering issues for future systems redesign and more detailed evaluations.

1. Introduction

Although being powerful tools for supporting work processes, the complexity of computer applications can be a serious barrier to their effective adoption, and innovative adaptive use within an organization to continually increase productivity and address changing business needs. Furthermore the continual arrival of new technologies, applications and new versions of existing software (not to mention new business processes and new employees) imposes a severe burden just to learn how to use the technologies, let alone to use them to do the work at hand in an effective and efficient manner. Thus how to use applications in order to do desired work tasks is an important element of knowledge sharing and should be a part of any Knowledge Management System (KMS), including knowledge about how best to use the KMS itself.

Many researchers have observed that people who are having difficulty accomplishing a work task with a piece of software look around their physical space for a colleague and ask them for help (see Twidale, 2005 for a review). In the case of informal workplace how-to help about the use of some aspect of an application, this has been termed ‘Over the Shoulder Learning’ (OTSL) (Twidale, 2005). One challenge is to be able to provide similar types of help in the situation when a colleague’s help is not immediately or easily available. What alternative help can be provided to people? How should this help be different from the online help already available in the software? Can the results of successful informal ad hoc help-giving sessions be saved and re-used in the future by others, perhaps in the form of an FAQ? What are the alternatives to conventional software help that can aid distributed employees? Finally, how can a corporate KMS used for sharing many different kinds of knowledge be used to share this potentially rather fiddly and confusing how-to knowledge?

This study considers how we might extend the use of an existing KMS as a channel of help giving. The idea is that the KMS stores the complete interactions of a help giving session that can be searched and retrieved at any time by others whenever they need to solve a related problem. In order to do this, we need to consider the way in which the help giving session might be stored and represented. Davenport and Prusak (1998) define knowledge as “a fluid mix of framed experience, values, contextual information, and expert insight that provides a framework for evaluating and incorporating new experiences and information”. The knowledge in an organization is not static but dynamic (Bannon and Kutti, 1996). In organizations, it often becomes embedded not only in documents or repositories but also in organizational routines, process, practices, and norms (Davenport and Prusak, 1998). As Eales (2004) notes: “the development of computer-related skills is an increasingly important area for organizational productivity”.

The next section provides an overview the knowledge management problem of sharing how-to expertise. In subsequent sections, we describe the study conducted to explore the idea. We then briefly describe the key observations taken during the study. The paper ends with a discussion and future research directions.

2. Collaborative Problem Solving

In earlier OTSL workplace studies (Twidale and Ruhleder, 2004) (Twidale, 2005) it was found that usually people at first try to find a solution to their problem on their own. It is only when they are stuck that they call in a colleague for help. However, what if neither colleague nor remote help is available at that instance? We need to move from managing knowledge to ‘*sharing expertise*’ (Ackerman, Pipek and Wulf, 2003). The natural choice is a KMS acting as a repository of documented help giving or problem solving sessions that took place among employees. For example, when an employee learns from one of her colleagues how to add columns and rows in MS Excel, the system could record all the steps of the solution. Stored help giving sessions have the potential to exploit much richer local contextual information than conventional built-in application help. They can also provide more visual clues and show more details of the process of solving the problem. The OTSL workplace studies revealed that a large number of real work tasks are resolved by compositions of multiple applications, often by users copying and pasting elements between applications to create sophisticated and adaptable workflows. Inevitably these in-house solutions, invented and tailored to the unique needs of the organization by its members, are unlikely to be found within the help system of any single application. By contrast a KMS help trail across several applications means that such end user innovations (von Hippel, 2005) can be shared around more effectively.

Ackerman’s Answer garden (Ackerman and Malone, 1990) was an early system that was effective both in introducing people to colleagues with appropriate expertise whom they might ask for help, and as a mechanism for recording, sharing and reusing the information obtained for a particular need. It mostly focused on text based discussions and solutions. Text has the great merit of ease of creation and consumption, ease of computation over it for indexing and ease of storing, sharing and processing. It seems to work very well for many kinds of abstract problems about workplace knowledge, norms, practice and processes that can be expressed verbally, particularly in cases where participants are members of the same community of practice (Wenger, 1998) and so can use specialist terminology or jargon. The power of verbal or textual storytelling as a knowledge sharing mechanism is well known (Denning, 2000). Words also work well for more complex how-to technical problems – provided that the participants are relatively expert and in a community of practice with a shared technical vocabulary (Orr, 1996). In such cases it is possible to talk effectively and intelligibly about how to do something complex with a computer application. In other cases it is not. A text-only description away from the computer may be incomprehensible to the learner who would prefer to point and gesture at the screen alongside the helper, using the interface as a common point of reference to help ensure that each knows which part of the problem the other is talking about. Consequently the question arises of how to represent prior help giving interactions in the KMS, given that text alone is probably insufficient. One option is as a movie of the whole interaction, as is possible with existing software such as Camtasia Studio (produced by TechSmith). However this creates some problems for indexing the help, and also for certain kinds of use. A movie is temporal so that all the successive stages of a complex interaction can only be seen one after the other and must be remembered by the user in order to get an overview of the problem as a whole. Comparing temporally non-adjacent steps, moving back to earlier steps and keeping a sense of the elements of the overall activity are possible but rather laborious.

Therefore we chose to investigate the much simpler approach of a list of multiple screenshots. Figure 1 illustrates the idea, showing two successive screenshots from a rather complex 33 step activity in Microsoft Excel used in the pilot study. This kind of visualization is strongly influenced by Carroll’s work on minimal manuals (Carroll, 1998) and also the work of Van der Meij and Gellevij (1998) on the importance of screen captures in documentation. Typically software help-giving interactions involve a demonstration and explanation accompanied by pointing, gestures and other kinds of annotations in addition to the action sequence of actually doing the task (Twidale, 2005). However for this study we chose as a design starting point to use an extremely minimalist approach of only using screenshots of an imaginary prior help-giving session with no accompanying text or explanation. The questions we wanted to ask were; are screenshot-only trails good enough for people to cope with? How might they be improved? And if needed, how much extra text or other kinds of contextual information is required? The study would help us explore further points in the design space.

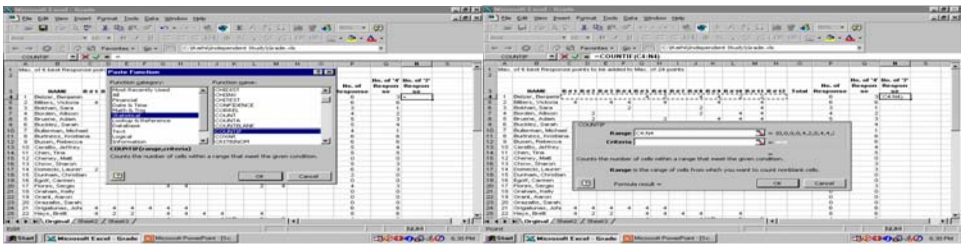

Figure 1: Examples of KM-Help: Screenshots 9 & 10 of how to calculate a grade.

In particular we wanted to know how this new kind of help, which we call KM-Help, compared against two others: the traditional online help accompanying most applications, and the personalized human-human help of OTSL. The aim was not to do a summative evaluation of KM-Help since it is just an interim design idea and can and should be improved in the light of ongoing formative evaluation. But we did want to know if it at least had sufficient potential to warrant subsequent investigation. Even in its minimalist form, KM-Help benefits from much more context than traditional application help, but has much less than OTSL. It would face a somewhat daunting comparative challenge since for all the problems of traditional textual help, users are at least somewhat familiar with it, and we consider the one-to-one interactions of OTSL to be a gold standard of idealized effective help, although not always available and always involving a significant workplace opportunity cost.

Partly to even the imbalance somewhat, and chiefly to simplify the study, we chose to ignore certain key KM problems: creating the data, processing it for shared use, indexing it and accessing it for future use. All of these are very important and merit their own study. Our design choice of only using screenshots clearly minimizes the effort required to create use-trails for the KMS, both programmatically and in terms of effort of the participants. It is entirely feasible to imagine automating this process, and would not require the initial help-giver or recipient to do extra work to make their interaction into a form acceptable for subsequent sharing and re-use – a common barrier to the acceptability of many KM applications. Some additional contextual information would need to be captured manually, automatically or by a combination of the two. For example who was involved, when and where the interaction occurred and something about the goal of the work task would be helpful and hopefully not too onerous. The problem of indexing and search is also significant. Assuming that ideally the KMS has many use-trails, how does a user navigate it to determine if it contains a trail that might be of use? We do not explore that in depth here, only noting the possibility of using recommender-like systems to give a potential user clues about trails that cross through elements of application interfaces that she is using to complete her work task. Of course such a user would have to be extremely lucky to find a pre-existing trail that did exactly what they wanted to do, and so in our examples we provided trails that were similar to but not identical to the user's task. By just assuming that these prerequisite challenges can be successfully addressed we can devote ourselves to the central question, once such a trail has been successfully located, what must be in it to be useful?

Finally we should note that the following is only a small scale pilot study, undertaken to provide qualitative data for ongoing refinement and iterative design of the KM-Help concept. We want to explore how such low-cost high-speed methods can be used both in iterative systems development and also by practitioners who under workplace constraints may lack the resources or time to undertake large scale, rigorous, detailed and systematic experiments. It is thus closely related to our work on extreme evaluation (Marty and Twidale, 2005a) (Marty and Twidale, 2005b) that investigates minimalist methods for usability evaluation that can better meet the practical workplace needs of systems development cycles of daily builds than traditional usability experiments.

3. The Pilot Study

The study consisted of two tasks: 1) Calculation of Net Present Value and 2) Calculation of a Grade. In both cases the participants had one or more types of help: Online Help, Human Help, and KM-Help. There were three participants (S1, S2, S3) each performing both tasks. So in total there are six sets of observations, two for each participant. In setting I, each participant was given one type of help and in setting II, each participant was given two or more types of help.

Participants had similar educational background, and similar levels of understanding of the concepts to be used in the tasks, but had no prior knowledge of either task. None of them had used Excel before, although they were numerate and familiar with many other applications. The tasks were designed to be rather challenging, so as to be likely to need help. For both tasks the details provided in the MS Excel worksheet were sufficient to solve the problem. Figures 2 & 3 show the screenshot tasks as given to the participants.

3.1. *Task 1: Net Present Value (NPV) Calculation*

The first task was to use MS Excel to calculate the Net Present Value using the Goal Seeking method. This task was designed to use two formulae, first to calculate the NPV for each year for 'n' years, second to find the Internal Rate of Return (IRR).

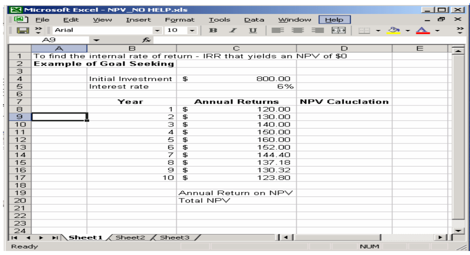

Fig. 2. Screenshot of task given to calculate NPV.

Fig.3. Screenshot of task given for Grade Computation.

3.2. *Task 2: Student Grade Computation*

This task was to calculate the grades of 62 students in a class. The students could have submitted any number of responses up to 12, each of which earned a score of 1 to 5. Any score of 4 or 5 was to be treated as a 4. The best six scores were to be chosen for grading purposes. Participants were required to count the total number of responses submitted and also count the number of responses with 2 and 4 points (including 5 points).

a. Types of Help Available to Participants

Assignment to conditions was random. Participants were told what help they could get if they needed it to solve the given problem. The study involved three types of help:

1. <u>Online Help</u>: This is the built-in help of Excel. Users could access it just by clicking on the Help Icon. Users were familiar with the concept from other applications.
2. <u>Human Help</u>: The administrator of the study offered to give help to the participant if he/she wanted. An OTSL scenario was created where the participant was asked to imagine that the administrator was a colleague, available to help if needed. The help

provided involved a combination of verbal directions to the help seeker and the help giver demonstrating solving the task on the computer.

3. KM-Help: A scenario of seeking re-used help from the proposed knowledge management system was created by providing screenshots of the solution to a similar (but not identical) problem. For this study, similar problems (both NPV calculation and grade computation) with numerical value variation were solved in advance and screenshots of each step of how the solution was achieved were saved as a file of Power Point slides. The file was kept open and the window was minimized.

Table 1 gives an overview of the study with type of task and help given to the three participants.

Table 1. Summary task done by participant and type of help given.

Participant	Task	Type of Help
S1		KM-Help
S2	*Calculation of NPV*	Human Help
S3		Online Help
S1		KM-Help and Online Help
S2	*Grade Calculation*	Human Help and Online Help
S3		All Three Kinds of Help

4. Observations from the Study

Although only a pilot study with an insignificant sample size, we believe that the observations and analyses made and outlined below are useful in their own right for informing both future studies and the ongoing iterative design of better help systems. Although we can make no claims for the relative frequency of the findings to be expected in a larger population, the pilot study does give us substantial insights into the *range* of activities that occur and issues that arise in the use of help. It is worth noting how rich the qualitative data is that can be obtained form relatively tiny studies, so long as care is taken with the limitations of small samples (Marty and Twidale, 2005a), (Marty and Twidale, 2005b). In particular, small scale formative evaluations need to give greater attention to use problems that arise than they do to use successes. The problems are real (although we do not know how widespread) and so are candidates for inspiring redesign fixes. The successes may be just the happenstance of a particular individual failing to encounter a lurking problem.

4.1. *S1 doing NPV Calculation with KM-Help*

Participant S1 began by writing the formulae in each cell and calculated the NPV for each year. Then he searched for some functions in the menu by reading the drop down menu, but did not use any of them. This shows how the GUI itself can help learning by presenting the available options. Unfortunately in this instance that kind of help was insufficient. S1 continued by dragging down the column and the values of NPV got copied into all the rows in that column. He wrote formulae in each row manually. S1 was not able to solve the task, so tried KM-Help. He first looked at all the screen shots in one go and then looked at each step and worked on the problem using same steps as in screen shots. He tried to copy the solution in the help directly into his task. That is, in the task the participant was asked to give internal rate of return to have an NPV of $100, but instead created an answer for $0 (as shown in the help). S1 returned a value as $0 but realized that the desired output was different than as shown in KM-Help. Thereafter he corrected the mistake and returned the correct value.

This is a classic transfer error. S1 did exactly what was done in the KM-Help example rather than using it as a template to adapt to the slightly different needs of his own unique task. It reveals the problem with using concrete examples that although easier to obtain from real life use, and able to exploit valuable contextual information about workplace needs and practices, still may not precisely match a subsequent user's needs. The KM-Help trail may give a clue to solving a related problem, but the user has to do additional work in order to do the correct mapping.

4.2. *S2 doing NPV Calculation with Human Help*

S2 started by typing in the values and manually trying to calculate the values for each cell but he could not get started as he was missing the basic syntax of MS Excel: he was not typing '=' before a formula. He hovered over the menu bar and looked at the options but was not able to figure out anything. He then inserted the '=' sign and now could do the calculation. S2 asked the administrator (here acting as the OTSL colleague-helper) if he could change the interest rate. The administrator said he could, so he changed the interest, but still could not complete the task despite trying several different options. Then he turned to the administrator again. The administrator began by asking him what the problem was. He described the point where he was stuck and not the whole problem. The administrator asked, "If I show you the calculation for one cell, can you do it for the rest?" Then the administrator demonstrated the calculation, and S1 started doing it. He could do that particular calculation, then he turned back to the administrator and asked "Do I have to do all of these by trial and error"? The administrator then told him how to use the goal seeking function. The administrator also told him the other things that could go wrong and showed him by doing it on the worksheet. S2 kept on doing the further steps as the administrator was telling him and he kept on showing the administrator and confirming is he was doing it correct and finally reached the solution.

This example illustrates both the power and the cost of human help. It enables the provision of detailed tailored incremental help, but can be very laborious. We treat human help as our gold standard not in order to assess whether other methods are inferior (we would be amazed if they weren't), but the extent to which they can be used as much cheaper but second best alternatives that are good enough in some circumstances.

4.3. *S3 doing NPV Calculation with Online Help*

S3 started by copying the values of one column into another column. He looked for the Goal Seeking Function from the menu bar, found it and clicked on it but then did not know what to do with it. He wanted to ask the administrator though he knew that his only available help type was online help. The administrator did not help him. He started doing calculations on the sheet manually. He also had the same error as S2 in that he was not typing '=', but he eventually figured that out on his own. He got some negative values, which he assumed meant that he was doing something wrong, so he deleted everything. Now the participant tried to look for the online help. He got the pop up menu (paper clip), typed in “goal seeking”, and got no results. He tried other keywords and again did not get an answer but did get some hyperlinks in the help. On clicking on these hyperlinks he reached the main help pages and was able to read through them and understand the operation of the Goal Seeking Function. He minimized the help window and started working on the worksheet. He got an error three times and then revisited the help. He went deeper into help by clicking on the hyperlinks and reading them. Finally he understood the function and got the correct answer.

In this case online help did eventually enable a determined and intelligent user to complete the task, but only after considerable effort. It illustrates the all-too common frustrations people have with locating the information they need in conventional help.

4.4. *S1 doing Grade Computation with KM-Help and Online Help*

S1 started by looking at the KM-Help. He read through a couple of slides and then went to his worksheet. He clicked on the function button and got the options but did not understand what to do, so he went back to the help. This time he focused more on the slides and spent a little more time examining each slide carefully. When he was reading it on the screen he pointed on to the various points on the screen while trying to understand. Then he did the same thing in the worksheet and got answers to that stage in the task. He then started looking at the different menu options again to decide what to do next. He did not do anything on the worksheet and opened the minimized help which was at the point where he left it, exactly at the step where he was now stuck. He looked through the following couple of slides and did the same calculation but now using multiple columns. After doing this he confirmed the calculations by doing them separately on paper. He then repeated this pattern of following each step from the KM-Help by reading it thoroughly and implementing it in his own modified way on the worksheet. He did this back and forth three more times. Each time he finished one step of calculation and confirmed it by doing it on paper. In the end he got the correct result. However he did not use the KM-Help to copy the solution but rather for understanding the concept of using formulae and then did it in his own (longer) way.

In this case KM-Help gives the user some clues about how to generate his own solution, but it requires considerable effort. The lack of explicit support in terms of proving an overall conceptual overview (other than being able to glance over all the screenshots) and the lack of a mechanism for tailoring meant that the solution took far more time and effort than might be expected with a typical human help interaction.

4.5. *S2 doing Grade Computation with Human Help and Online Help*

S2 began by selecting the first row, did not know what to do next, and so immediately went to Online Help. He typed in a query, got back some text, read it and went back to the work sheet. He still did not do anything on the sheet, went back to the Help twice more, and tried different queries. He then wrote a formula and calculated the value for one cell. He tried to copy the same formulae to other cells by dragging and dropping, twice getting it wrong before succeeding. He referred to Help again and realized that he needed to use a different formula for ‘SUM’. He tried to calculate this but did not get an answer. He kept on referring to Help about ten times, trying with different keywords but did not get any satisfactory answer. Finally after 30 minutes of struggle he asked the administrator for help saying that he did not want to take any human help before because he wanted to solve it on his own. The administrator showed him how to solve the problem explaining each step. S2 kept saying that there was no help available for this kind of activity in the Online Help.

In this case human help acts as a last resort for a user who has tried nobly on his own, a situation often seen in workplace studies where a request for help is often prefaced by a litany of all the things the help-seeker has tried before bothering the colleague.

4.6. *S3 doing Grade Computation with Online help, KM-Help and Human Help*

S3 started by looking at the task. He did not know how to proceed, so went to Online Help. He typed in queries such as “count number only between specific values”, “countif” reading through the resultant help pages. With the Help window open on one side, he explored the drop down menu options in the main application. He switched back and forth between the Help window and the main application, unable to find what he wanted or to make sense of what he found. He again opened the Help window on one side of screen, scrolled through the drop down menu of the main application and clicked on one of the icons, opening a function window. He placed this window and the Help window adjacent to each other. Then he minimized the Help window and highlighted the row. He was still unable to understand what to do and so closed the function window. He opened the help window yet again, typed ‘count’ and found a help page. He minimized the window and clicked on the drop down menu and looked for the ‘count’ function but was unable to find ‘count’ in the drop down menu. He wrote the formula in the cell himself. He first checked the result of the formula by doing it in one cell and then dragged & dropped to copy the formula to the other cells in the column.

S3 went back and forth in the help window, minimized it, worked on the worksheet and then opened the help window again. This process continued many times. Finally he thought he got the calculation right and said, “yes” in excitement. After working on one column, he proceeded to do the calculation into another column and the same process of shifting between help window and worksheet continued, but again he was unable to understand what to do and how to proceed. S3 called for administrator intervention and KM-Help was opened. He was told that if he cannot solve the problem, he can ask administrator. S3 looked through most of the screenshots. He called the administrator and said, “I have done the same thing as shown in the KM based help”. The administrator gave him a hint about which formula to use to solve the problem. S3 looked through the KM-Help screenshots and stopped at some for a long time. He returned to his worksheet and called the administrator again for some more clarification. He looked at the worksheet and then opened KM-Help. After looking at the KM-Help, he opened online help, copied something from the Online help and redid the formula in one of the column. Then he made changes in the formula of the other cells too and finally said “I am done”.

In this complex case S3 thought he had followed the instructions in KM-Help, but actually had not managed to do the mapping successfully. Human help could efficiently bridge this gap, and also use the KM-Help to enable a very short, efficient human-human interaction compared for example to the case above of S2 doing NPV.

5. DISCUSSION

Based on the above results we can begin to identify some common themes for each type of help. Of course all these findings are tentative, but can be used to inform the design of future larger studies, and the redesign of KM-Help.

Online Help: This is neither easy to understand nor comprehensive and its interface leaves much to be desired. It is organized in multiple levels and navigation via hyperlinks can be confusing. Access is hard as a user may need to type in multiple guesses for the name of the needed feature(s) or indexed keywords to get to what is needed. Despite its complexity, online help does not cover everything that can be done with the software.

KM-Help: KM-Help was relatively easy to follow. The time that the administrator spent in helping activities was much lower when also using KM-Help. It was used as a guiding tool to understand the details of which icons to click, which formulae to use to solve the problem, etc. KM-Help needs to operate in different scenarios. First, the easy case where the problem that the person wants to solve is exactly same as what she finds in the KMS. Second, KM-Help may be useful to explore the concepts required to solve the problem and the example can be modified to suite the problem at hand. KM-Help may not be helpful if there is more than one way to solve a problem, where the user may take a quite different solution approach to tackle the problem. There is a risk that KM-Help might lead to careless errors as people might put too much faith in it and just copy the solutions, and may not realize that their work task need requires them to use KM-Help as inspiration to invent a new tailored solution, and not just exactly reapply someone else's solution to a different problem. KM-Help was in the form of screenshots in PowerPoint. The help had a good visualization of the solution but was discrete in steps and missed certain dynamic aspects like how to click the mouse to drag & drop the formula across the solution cell. This is a weakness of the decision to not use a movie-based solution to recording the problem solving interaction.

Human Help: Participants were more comfortable in approaching the OTSL helper-administrator for help and seemed to mostly prefer this human help mode. It was interesting that the participants only requested assistance from time to time and only for the piece where they were stuck. For example they asked particular questions like "Can I change the interest rate?" "How do I define a variable?" and did not ask for the complete solution of the given task. They used trial and error approach to find the solution on their own. These findings are confirmed by the OTSL workplace studies. When combined with KM-Help, the human help required less effort on the part of the administrator. The help given was in the form of hints such as to use 'IF' function and the participant then referred to KM-Help to learn more about it. The administrator could easily (even literally) point the participant in the right direction because of the presence of KM-Help.

All Kinds of Help: Participants were inclined to seek help from the administrator (human help) even though they knew when the option was not available. For example the KM-Help visualization was very comprehensive but still the participant wanted help from administrator. On the other hand, users often try to figure out things on their own by reading through the menus and options when stuck. Indeed one of the participants did not want to take help from the administrator. We suspect that this was because he wanted to show that the task can be solved with minimum assistance from administrator. In both online help and KM-Help all the participants would go back and reconfirm each time they did anything. Sometimes they had to refer to them multiple times. Participants used KM-Help extensively and went back and forth between help and the worksheet. KM-Help was used for understanding what to do to solve the problem. It was also used for reconfirmation of the problem solving process. Users may need to revisit help items several times as they may not understood or remember it all in one instance of looking at it, a finding confirmed by the OTSL workplace studies (Twidale and Ruhleder, 2004), (Twidale, 2005). Users can modify their solutions and do not necessarily need an exact solution. Sometimes additional tools are required such as pencil and paper even with KM-Help. Some users are more comfortable in doing their calculations on paper in order to understand what they are doing.

6. Conclusion

Despite its minimalist form of only screenshots and no accompanying text, the study gives us at least a clue that KM-Help has potential to support learning from reuse of previous human help giving interactions. It can also be refined in the light of the weaknesses identified. Even with a tiny number of participants there was considerable variety in the way they addressed the tasks, the difficulties they had and the approaches they used. This serves to remind us that design to support knowledge work has to deal with a wide variety of users, uses and contexts. A single optimized solution is unlikely to be effective or even possible. Instead what is needed is a range of resources that people can choose to use in ways useful to them. An unexpected finding was how the KM-Help and human help combination may lead to solving very complex problems with relative ease and requiring less interaction time as the human helper provided only hints to the solution. A KM-Help representation can serve as a route map to guide interaction but also discussion of the use of a complex interface.

Future work includes exploring alternate versions of the design KM-Help that address the problems identified in the study, but still meet the perhaps conflicting constraints of ease of generation and minimizing the effort required of the initial help-giving interactions that generated them. This can be followed by another iteration of a small scale study and system redesign, followed by running of a larger scale study. This larger study should be designed in the light of the issues uncovered in the pilot studies, particularly those that merit more detailed attention, such as the potential efficiency gains of combining KM-Help and human help. We must also not forget to address the simplifying assumptions made for this study and examine ways to capture, annotate, index and search large sets of KM-Help trails stored in a KMS.

Finally, our use of a small scale pilot study has revealed how much valuable qualitative data can be derived cost effectively and in a relatively short time. Even from this pilot study, we have gained a richer understanding of the diversity of issues that need to be addressed in analyzing and developing better help-giving systems.

Acknowledgement

We would like to thank our volunteers for their kind support and valuable feedback. This research was supported by the National Science Foundation under Grant No. 0081112.

References

Ackerman, M.S. and Malone, T.W. (1990) "Answer Garden: A Tool for Growing Organizational Memory", *Proceedings of ACM Conference on Office Information Systems:* 31-39, New York: ACM Press.

Ackerman, M.S., Pipek, V. and Wulf, V. (2003) *Beyond Knowledge Management: Sharing Expertise*. Cambridge, MA: MIT Press.

Bannon, L.J. and Kutti, K. (1996) *Shifting Perspectives on Organizational Memory: From Storage to Active Remembering*, Proceedings of the Twenty-ninth Hawaii Conference on System Sciences, vol. 4, 156-167.

Carroll, J.M. (1998) *Minimalism Beyond the Nurnberg Funnel*, Cambridge, MA: MIT Press.

Davenport, T.H., and Prusak, L (1998) *Working Knowledge: How Organizations Manage What They Know*. Cambridge, MA: Harvard Business School Press.

Denning, S. (2000) *The Springboard: How Storytelling Ignites Action in Knowledge-Era Organizations*. Boston, MA: Butterworth-Heinemann.

Eales, R.T.J. (2004) "A Knowledge Management Approach to User Support", *5th Australasian User Interface Conference* (AUIC 2004), vol. 28, Dunedin.

Marty, P.F. and Twidale, M.B. (2005a) "Usability@90mph: Presenting and Evaluating a New, High-Speed Method for Demonstrating User Testing in Front of an Audience", *First Monday*, 10(7).

Marty, P.F. and Twidale, M.B. (2005b) "Extreme Discount Usability Engineering", *Technical Report ISRN UIUCLIS--2005/1+CSC*W.

Orr, J.E. (1996) *Talking about machines: an ethnography of a modern job.* Ithaca, NY: Cornell University Press.

Twidale, M.B. and Ruhleder, K. (2004) "Where am I and Who am I? Issues in collaborative technical help", *Proceedings of CSCW04*. 378-387, New York: ACM Press.

Twidale, M.B. (2005) "Over the shoulder learning: supporting brief informal learning", Submitted to *Computer Supported Cooperative Work: The Journal of Collaborative Computing*

Van der Meij, H. & Gellevij, M. (1998) "Screen captures in software documentation", *Technical Communication*, 45(4): 529-543.

von Hippel, E. (2005) *Democratizing Innovation.* Cambridge, MA: MIT Press.

Wenger, E. (1998) *Communities of Practice: Learning, Meaning, and Identity.* Cambridge, MA: Cambridge University Press.