Journal of Universal Computer Science, vol. 1, no. 6 (1995), 410-422
submitted: 20/3/95, accepted: 19/6/95, appeared: 28/6/95

From Personal Computer to Personal Assistant

Jennifer Lennon
(Department of Computer Science, University of Auckland, New Zealand)

Arnould Vermeer
(Institute for Applied Computer Science I, University of Hagen, Germany)

Abstract: Much of the confusion that surrounds electronic personal assistants arises from the open-ended complexity of their development. In this paper we categorise some of their more common uses before suggesting several thought-provoking extensions.
Key Words: Electronic assistants, electronic agents, message pads, searches, prediction, programming by example, voice commands.
Category: H.m

1 Introduction

Evolution in human-computer interface design has brought us a long way from the days when computers were all but dictators, when users obeyed enigmatic "rules" or risked ignominious "crashes"! Years of research and many trials involving users (willingly and unwillingly) have steadily improved computer interfaces to the point where today's operators are much more in control. However we believe that the next few years will see another quiet revolution as the computer becomes more of a versatile supportive ally - a true Personal Assistant (PA for short).

The idea of an electronic assistant has already been promoted with much hype by major computer companies. But the impression often given is that computers will eventually dominate our lives, relaying frantic messages between offices and homes, frustrating even our rest and relaxation. This is certainly not our vision: we look to the computer to help us, not discombobulate us.

A PA, as we define it, is not to be confused with the devices that several companies are developing as "personal digital assistants" or "message pads" (see Section 12). Those are still extremely limited in the applications they support. We believe that a PA system should - and can - do a great deal more.

Because application programs are providing constantly better interfaces, the exact nature of an electronic personal assistant is becoming blurred. Some of its functions may be controlled best by a continuous background processor. Others may be delegated to various applications. However, just as there is a move towards "universal documents" that allow text, graphics, spreadsheet and database editing all in the same document, we should also expect a PA to help in activities ranging from simple wordprocessing tasks to scientific research. A PA should support our work in ALL areas of computing.

In this paper we consider the assistance we might expect from a PA in a few everyday activities.

2 Writing Documents

There are probably countless ways in which a PA can aid us in the exacting labour of writing documents. Certainly, as we shall see in Section 7, PAs will

relieve us of a large amount of repetitive editing by making predictions based on our personal work patterns. They will do a great deal more than this, however.

2.1 Making Literature Searches

Making a literature search is a frequent necessity, a time-consuming, exacting chore, yet vital whether our need is to find information or to check our work. Additional electronic assistance with literature searches would be a great relief. In particular, we should be able to vet our work much more efficiently than is currently possible. At the very least, a PA will help us set up and manage a filter system so that our searches are as effective as possible. This is particularly important when searching in hypermedia1 systems (see Section 6.3). The PA will "learn" what filters we prefer to have set under various conditions and apply them, [KM94]. It will provide intelligent cross-links so that better use can be made of huge databases. Obviously everything we write should be open to verification. One such system, "Ways 2", produced by the Swiss researcher Keller, has been marketed by Vobis in Germany with considerable success. In a PA system, data in wordprocessing text will be checked against data in our spreadsheets and databases. Running as a background processor, it can make continuous searches for relevant references, aiding us in the verification of both our own work and the reference material we are reading (see Section 6). Having such checking done automatically, or semi-automatically, is possibly the only way of ensuring that the increasing amount of archived work taken to be authoritative is at least relatively free of errors.

2.2 Improving Spelling Checkers and Grammar Checkers

Spelling and grammar checkers are often inefficient and cumbrous. Most spelling checkers are not even context driven; grammar checkers are clumsy, throwing up too many red herrings. Users should be able to tailor grammar checkers to their personal preferences. An intelligent editor should also be able to "learn" rules of style appropriate to the particular type of document being edited - such as appropriate levels of word difficulty, or suitable sentence length. The system should be flexible enough to allow the user a choice between simultaneous checking and checking only on a File Save or Quit command.

Programs such as these need to be able to "learn" from their "mistakes". For example, in a maths paper the second occurrence of the word "group" probably should not be replaced by "collection"! If a user has to override the "second occurrence", an intelligent editor will query whether it is a technical term and act accordingly.

We need to be able to refer to a whole range of dictionaries - maybe at just the click of a button. Using techniques such as those described in Section 6, we look forward to much better success rates in our searches.

An interesting application applies to natural language translation. Obviously we are not referring to full natural language translation - a complex and subtle specialty still struggling for success. But word-for-word translations are practical and this opens up several rather thought-provoking possibilities. For example, a researcher trying to read foreign journal articles with only a smattering of the language would be greatly helped by a handy translation option. Similarly, if

correspondents restrict their sentences to grammatically simple ones, translation programs will have a better success rate. The qualified success of existing style-checking programs suggests that certain types of documents may be amenable to automatic rewriting in simplified words and phrases - a facility of great use to new immigrants, for example.

3 Reading Documents

Browsing electronic text consumes so much of our time that all personalised assistance with it will be valuable. As software systems become daily more complex, the problems associated with training operators in their efficient use multiply exponentially. Since users have neither time nor motivation to wade through extensive manuals, many firms are applying the principle of "just-in-time learning". This will become more effective when employees have their own PA to control the Help systems. The PA will be able to take into account the user's previous experience - or lack of it!

Many people find reading from the screen more tiring than reading print, particularly when windows are densely packed with text. The displayed text is either too small for comfort or too clumsily large. One solution may be to indicate the particular line we are reading by moving the cursor up and down manually (or by methods mentioned in Section 8) so that the editor can enlarge it. It may also be desirable that a line or two above or below the current line be enlarged to some degree as well. This will obviously be dependent on personal preferences, but each user should be able to individually "train" their text windows.

Since reading is so much faster than writing, any reference system that is provided needs to be very efficient for us not to feel unnecessarily held up by it. Hence, though facilities for referencing and checking are just as vital for reading as for writing, they will have to be uninhibiting to be really useful. Perhaps in "Browse mode" a single click on a word could initiate a dictionary look-up or a reference check. Each person will decide which dictionaries or databases are to be filtered and used.

Yet another necessary feature for most applications is the presence of date computation routines. For example, if I read that a colleague is arriving on 18th May, my immediate question will probably be "What day of the week is that?" I should be able to click on the "18" and "May" and have a formula (such as Snell's) calculate the day - given any particular year. Or maybe I need to know what date Easter falls on. Even more important, when making travel arrangements I must know which days are public holidays in the countries I am visiting, so that I can either avoid them or make sure I see any special festivities.

4 E-mail

We all agree that there must be better ways of managing e-mail than keeping our names off mailing lists [Den82]. M.I.T. has made considerable progress with its Information Lens and Object Lens projects [MGT+87] [Rob91]. Certainly we should be able to set up our own preferential filtering systems so that a PA, perhaps working as a background processor, can search all incoming mail for names and keywords. It will search the body of each item, as well as the

subject line, and classify the mail according to our own prescriptions, [KAF+94]. Furthermore, we do need additional options such as ranking and grouping by author. Answering mail from your boss may be more urgent than attending to personal matters; and if he has mailed more than one item, it is important that they appear together so that we do not waste time answering the first item only to find that it was superseded by a later item!

As more and more correspondence is sent by e-mail, we are looking forward to additional electronic assistance with all the chores associated with daily avalanches of mail. A personal assistant will serve us by keeping address lists up to date. It will assist us in writing form-letters (see Section 8) and help with addresses. It can manage certain housekeeping tasks such as archiving necessary information before deleting messages. Notices in template format can enable the system to use the advertised dates to generate expiry dates - these may, of course, be overridden by the user. Ideally the user should also be able to define what is to be done with the notice once the expiry date has passed: delete it, or place it in a Past Events list to be maintained by the system. This list will be invaluable for compiling such things as end-of-year reports, although decisions will have to be made on how to maintain it. E-mail must also support structured discussions and systematic collaboration between two or more participants, i.e., conferencing. Furthermore, we should have the option of attaching notes to individual letters in group dispatches. In fact, as information proliferates in all systems, the ability to annotate at all levels of the system becomes increasingly important.

It is intriguing to ponder the consequences of having e-mail answered automatically. For example, we may wish to generate replies that indicate we are unavailable. Of course we would avoid being naive and publicising to all and sundry that a house or office is unattended and hence is a good burglary risk! We could, however, have different responses to different people. Acquaintances could be informed that I am at a meeting or that I am very busy - whereas close friends are told that I am really away at the beach!

5 Intelligent Calendars

Electronic diaries will become much more useful when they are more widely linked into electronic mail systems. Calendars should also be capable of generating mail for us - either mail to ourselves or mail to others. On a personal level, they should be able to send us reminder notices that contain much more than just bald statements of events and dates. Perhaps, when we are prompted that next week is an important anniversary, we also need to be reminded that last year's "bright idea" for a gift wasn't a success!

On a more workaday note, an intelligent PA will be able to set up appointments for us. Incoming e-mail about meetings will be searched, our calendar checked and updated, and outgoing e-mail messages generated. The PA system should be intelligent enough to be capable of adaptive behaviour. Conceivably, it could even "learn" when one person in a group is being unreasonably difficult about the times they are available and it could take a "tough" negotiating line! On the other hand, although I may wish my system to negotiate for me I may also not want all my timetable to be known!

The logical extension of this is, of course, that we shall eventually have networks of PAs communicating with each other and the user will, at best, only be

required to approve the final result.

6 Information Retrieval

Electronic journals such as J.UCS (Journal of Universal Computer Science) have significant advantages over traditional and existing electronic journals [CMS94], [JUC94], [MS95]. Since it is based on Hyper-G, [KAF+94], its structure lends itself to user-controlled filtering and automatic notification of new material and updates - all of which are necessary components of any PA.

6.1 Searches

To make the best use of resources such as JUCS [CMS94], [JUC94], researchers need more assistance in searching and checking documents - their own and other people's. This is becoming critically important as daily so much highly questionable data is being quoted as fact.

There are several important requirements to be met by search programs, and we list just a few:

- Appropriate use must be made of filters (see Section 2.1).
- The system should help the user by prompting for alternatives. For example, a search for "kiwi" could well generate the prompt, "Bird, person or fruit?"
- We look to better hit rates in searches. An intelligent search program will tailor its algorithms to fit the type of document being searched, deducing probable error types. For example, when text is entered from the keyboard it will look for transposed characters. Scanned documents on the other hand commonly produce different types of errors. Since scanning produces only a simple bit map, an Optical Character Recognition (OCR) program must be used to convert the bit map into text. Unfortunately, even the best OCR programs introduce new errors into the scanned text such as "g"s being confused with "q"s [COD90]. Here again we look to increasingly intelligent search algorithms to improve translations. We also look forward to a much better hit rate for words with multiple errors by using improvements that fuzzy1 search techniques seem likely to achieve [BYG92], [WM92].
- The system should be able to make fuzzy searches based on semantic nets instead of just lists of synonyms. A search for "air" could lead on to "gas", "oxygen", etc.
- If a search is not successful then the system should be capable of making alternative suggestions - it is frustrating when unsuccessful searches leave us no better off.
- Certainly we would like to have phonetic searches, as well as the option of searching dictionaries that include inflected forms of words, abbreviations, accents - and entire foreign languages.
- Obviously intelligent fuzzy searches should not be restricted to just text files. We should be able to efficiently search file names, directory lists, program names, and so on.

  Above all, such a system must be easy to use. To take a trivial example, if I read that the Sun is 20% helium and 80% hydrogen I will check the figures if and only if it can be quickly done - perhaps with PA assistance, on just the click of a mouse.

6.2 Retrieving Information from Libraries

A big step forward for all researchers will be a PA giving us easier access to library information. We believe that in the foreseeable future it will become standard practice for librarians to insert the table of contents and/or a brief summary into the hypermedia database so that researchers can rapidly access the new information. This much at least can be done without infringing copyright laws. With a system such as this, a PA can constantly check all new books and journals, and using our own categorisation system, notify us of those that are relevant.

The application PC-Bibliothek (PC-Library) is a recent development from Graz University of Technology for PCs running Windows [MMS94]. PC- Bibliothek's attractive and easy-to-use graphical interfaces, coupled with powerful search algorithms, highlight the benefits to be enjoyed once complete reference libraries are literally at our fingertips.

6.3 Hypermedia Document Searches

When hypermedia systems were just growing like Topsy, it was not surprising there were problems. In large hypermedia systems users still lose their sense of orientation within the environment. Recent research has seen interesting developments, particularly in the design of hypermedia structures [MKSS93] [MSS93] with links, and even without links [MPS94]. At the very least we should be to be able to signal "Help!" and have the system help us backtrack. However a good PA system will help us avoid disorientation in the first place by indicating our current position on a graphical map of our environment. Furthermore the system will remember from day to day, and week to week, which paths we have used, and how often, so that by making predictions it can help us navigate much more efficiently than we can now.

6.4 Webs and Guided Tours

There is a tremendous amount of work now being done on electronic guided tours such as those in the Hyper-G system [KMS93], [FM94]. The successes and shortcomings of "second generation" hypermedia systems receive a good introductory survey in the paper "Reflections on Notecards: Seven Issues for the Next Generation of Hypermedia Systems" [Hal88]. Interfaces such as those described in "Intermedia: The Concept and the Construction of a Seamless Information Environment" [YHMD88], and "IRIS Hypermedia Services" [HKR+92] help the user navigate hypermedia webs. Intermedia, for example, shows dynamic "tracking maps" that display the user's current position in relation to its predecessor and successor links. In Intermedia, operations "behave identically across all applications" (just one feature of beautiful design). Features such as these will be utilised by an all-embracing electronic assistant.

In an interesting program under development at the University of Auckland [MS94] viewers of a university information system will be given both two and three dimensional guided tours of the campus and buildings. It is hoped that much of the research work will be general enough to apply to other information systems.

Another move in the right direction is the increasing use of the guided task paradigm [TO90] , where users are introduced to new applications by interactively participating in guided demonstrations.

6.5 Sending Electronic Agents Through the Networks

Programs that search for information in wide-area networks are proliferating. Although the trend is being fuelled by commercial interests such as those mentioned briefly in Section 9, obviously all types of archived information are amenable to computer controlled searches. One well known system, WAIS [Ste91], automatically updates "dynamic folders" with relevant information from selected servers.

A PA could be trained to control electronic "agents" sent through Internet, [Com94], so that information relevant to the users' current needs is gathered - and summarised into their own personal newsletter!

A detailed description of collaborative agents in computer conferencing environments is given in "A Framework for Controlling Cooperative Agents" [LMS93].

7 Observing and Predicting: Electronic Assistance

During the past twenty years many papers have been written describing specific applications that incorporate forms of intelligent electronic assistance under various titles: intelligent editors, cooperative agents, programming by example, programming by rehearsal, to list just a few. An exceptionally well presented overview is given in the book *Watch What I Do: Programming by Demonstration* [Cyp93]. In 1985 Zissos and Witten described a "computer coach" that helps users avoid repetitive formatting tasks in a wordprocessing environment [ZW85]. It "unobtrusively monitors interaction with a system and offers individualised advice". The work has been greatly extended and several such systems have now been developed [MW92] [WM93]. A more generalised personal assistant has been proposed by the second author of the present report [Mau93].

Several major computer companies are committed to developing what they are terming "electronic agents". Apple Macintosh recently demonstrated the prototype version of an electronic agent that can learn from simple, mouse-controlled, repetitive actions. For example, as the user steps through each stage of a process, each selection and menu choice may be highlighted with a colour. This indicates to the user that his actions are being shadowed by the agent. If all actions are shadowed correctly, then the user has the option of letting the electronic agent help from then on: for example, selected pieces of text may be copied from one document and tabulated into another automatically.

Producing animated graphical sequences is a very repetitive task and this must surely be one area that will greatly benefit from having an electronic personal assistant or agent to help. In fact some of the earliest work done in the field of programming by example was done using graphics programs. In the paper "Metamouse: Specifying Graphical Procedures by Example" [MWK89] the authors describe a system that "induces picture-editing procedures from execution traces". It incorporates a very likeable icon called Basil - a turtle in the best LOGO tradition [Pap80]. The work has been significantly extended [MW93] [MWKF92].

A unique visual programming environment is described in the paper "Programming by Rehearsal" [FG84], where a "stage" is set and peopled with "performers". In the book Creating User Interfaces by Demonstration [Mye88] the author describes PERIOT (Programming by Example for Real-time Interface

Design Obviating Typing). Graphical menus and windows are created by example. MARQUISE is another interesting example of an interactive tool that creates graphical user interfaces by demonstration [MMK93].

Considered all together, applications such as these suggest that there is a wealth of experience waiting to be pooled into an all-embracing and indispensable interface [PY93] .

8 Making Use of the Computer's Ears, Eyes and Voice

Now that microphones and video cameras are almost standard computer attachments, we shall undoubtably see dramatic innovations stemming from their use. For example, when a significant number of PCs are connected to law enforcement networks then automatic reporting of break-ins will surely result in a proportional reduction in the number of unsolved burglaries!

With the addition of "eyes" and "ears", our PC will be capable of a wide variety of new functions. It will be able to inform us if someone is at the door when our back is turned. As we work it will adjust the screen brightness when the room brightness changes, or boost the volume of our headsets when background noises increase. We shall be able to dictate our e- mail, and have voice prompt and help systems. It is no longer inconceivable that an intelligent PA could even note our mood when we first enter the office in the morning and give appropriate responses - perhaps we might appreciate an occasional joke or even some artificial intelligence doctoring in the style of the Eliza program [Wei66]. Certainly a very practical welcome would be for the machine to open the document we last worked on and have the cursor sitting ready where we left off - surely not too much to ask, considering that there is at least one lap-top already on the market that leaves files open when re-booted.

The ability of the machine to recognise simple gestures such as a nod or shake of the head opens up many possibilities. Work on eye tracking, such as that done by M.I.T.'s Media Lab [SB90] and the Washington Naval Research Laboratory [Jac90], suggests that in the future we shall be able to look at a section of the screen, and, perhaps with only a simple nod, have that particular section enlarge automatically. This will help us in many types of searching tasks - imagine, for example, searching screenfuls of small windows as in applications that use miniaturised photos from CDs or thumbnail pictures as in Hypercard [Hyp89].

If we would like to have lines of text enlarged for working on in a word-processing task, eye movements may prove to give us better control than the cursor method explained in Section 3. Even if eye tracking proves to be too imprecise, then perhaps finger tracking may still be a better alternative. Text scroll bars as well as computer movie options could also be controlled by eye movements. And when presented with choices in a dialogue box we may prefer to focus on an option and then simply nod.

Computers that can carry out a whole range of voice commands are no longer fiction - a fact that has been widely demonstrated recently:

- "Computer open Word file" - the word processing package is opened.
- "Computer write letter to John Jones" - the letter is headed with the sender's and receiver's addresses, plus opening and closing sentences.

– "Computer include thank you message" - a whole paragraph is generated....

And of course this leads us to speculate how widespread computer-generated speech will become and how long it will be before we can actually converse with our personal assistant.

9 Help with Transaction Processing

Electronic shopping is here to stay, and the next question we must address is how we are going to make the best use of yet another barrage of information. Newsgroups have already advertised software that can send electronic agents through the networks to find "best buys" as well as making and changing hotel, restaurant or airline reservations for us. Of course, as with any new technology, we can expect problems both small and large - cases of electronic agents running "amok" in the networks have already been reported too.

In many European countries "telebanking" has become a routine aspect of life due to the continuing spread of Videotex [MS82]. Ideally we could do all our banking and paying of bills via our computer. Gone would be paper invoices, checks and forms. The conversation with the computer might go something like this: Computer please pay 400 dollars to Mike. Mike who? Computer, to Mike Melon of course. PIN no? (I type in my PIN number.) Transaction number? (I type in the correct transaction number from the current block.)

Of course, some security mechanism will be necessary to guard against impersonators. The computer will encrypt messages for security and use a different transaction number with each message to ensure that any electronic eavesdropper cannot simply copy previous messages to repeat transactions.

Although jokes are still made about the "paperless office", with the advent of electronic funds transferral we have made significant steps in that direction and indications are that the current expansion of the home computing market will lead to further advances.

10 The Individualisation Process

There is a rather interesting inference to be made from the current vogue of having personalised emblems put on everything from coffee mugs to designer jeans. With PA assistance in mastering intricacies of design, we shall be design our own emblems, embellishments, and even, if we are so inspired, our own works of art. We can then let our PA scan the networks to find the "best" supplier and have our specifications e-mailed directly to them!

11 Secretary? Tutor? or Even Parent

From all that we have discussed, it is obvious that when we use the term Personal Assistant the emphasis remains on the word "Personal". A good PA system will be "trainable" so that after a few weeks my PA will act in a very different way from yours - even if we purchased identical software to begin with.

Ideally, it will also be desirable to have alternative sets of characteristics built into a PA that modify it for specific tasks. At the office I shall expect my PA to

act as my own personal secretary - it will help me with all my correspondence, arrange meetings (see Sectionn 5), help with time management, and so on. Of significant assistance will be a "My archive" program that helps classify and archive all types of information: text, graphics, video and audio clips, important documents and interesting little snippets. All archived material will be amenable to retrieval by powerful search techniques such as those described in Section 7.

At university or school, the PA will become my personal tutor, helping me organise work and make the most of opportunities that are offered. In the paper "Lecturing Technology: A Future With Hypermedia" [LM94], we show what tremendous potential lies in the integration of multimedia lectures with a good database when all students have their own portable PCs. Excellent computer-aided instruction material will be generated. Since students will be able to replay the lecture away from the classroom, important distance learning opportunities will be provided.

However, when I lend my PC to my six-year-old child I hope my PA will not only protect all my work from sabotage but will transact on quite a different level. It could act as a "parent" - even to pulling the plug on games and bringing up some homework!

Thus when designing a PA we shall need to ask ourselves what the characteristics are of a manager's best secretary, a student's favourite lecturer, and even more elusively a "good" parent, to see if we can simulate at least some of their attributes. A challenge indeed!

12 Looking To the Future

As small gets smaller, and more powerful still, we can expect that computers the size of notebooks will support more PA functions than we have described. The initial interest created by pen-based electronic notebooks such as Apple Macintosh's Newton or Casio's Zoomer suggests that users enjoy that environment, limited as it still is. Users particularly enjoy sketching in environments that help them work more efficiently by neatening up their work, letting them work with constraints, and supporting incompleteness [Kim89] [Zha93]. There is even a certain perverse satisfaction in erasing errors by scribbling over them! All this is certainly an indication of things to come.

In the article "To Forecast Information Technology is Impossible Yet Necessary" [ML94] , we argue that future advances in information technology are quite unpredictable. However we go on to surmise: "In ten or fifteen years from now everyone will carry small but powerful Notebook computers around with them. The much heralded Newton is certainly a first step in this direction! You will be able to talk into your notebook and have more commands, programs, and facilities available than we can imagine. For example, if you go to a foreign country and talk into your notebook in English out will come Greek or French. A global positioning system will display maps for you and show you at any time exactly where you are located on the surface of the earth. And of course a mobile telephone will be integrated into your notebook, giving access to all the databases of the world - so you can look up theatre programmes and bus and train connections. It will be your digital photo camera, and it will replace your wallet and credit cards. It will be indispensable."

13 Conclusion

Our definition of a true personal assistant obviously bears little relation to any currently available commercial product. It is not just a glorified Newton. It is part background processor (continuously scanning the networks), part consistent graphical interface (across all applications), and part special routines integrated into application programs. It thus supports the user at all levels of activity. By making predictions from repetitive tasks, it saves us both time and frustration. It manages our e-mail, classifies and archives our work, and employs powerful fuzzy search algorithms to retrieve documents from complex hypermedia systems. It is much more than a generalised help system. It can be a model secretary, tutor and baby-sitter, a police officer who patrols our surroundings while simultaneously ensuring that we do not inadvertently break copyright laws or lose ourselves in hyperspace. It will be our augmented eyes and ears, an alter ego we create for ourselves.

14 Acknowledgement

The authors wish to thank Professor Hermann Maurer for support, encouragement, and many enjoyable discussions during the writing of this paper. Many of the ideas were suggested by him.

References

[BYG92] R. Baeza-Yates and G Gonnet. A new approach to text searching. *Communications of the ACM*, 35(10):74–82, 1992.

[CMS94] C. Calude, H. Maurer, and A. Salomaa. JUCS: The Journal for Universal Computer Science and its applications to science and engineering teaching. Computer Science Report Report 91, University of Auckland, Auckland, NZ, 1994.

[COD90] W. Cushman, P. Ojha, and C. Daniels. Usable OCR: What are the minimum performance requirements? In *Proc. CHI'90: Empowering People*, pages 145–151. ACM, 1990.

[Com94] Special issue on intelligent agents. Communications of the ACM, 37, 7, 1994.

[Cyp93] A. Cypher, editor. *Watch What I Do: Programming by Demonstration.* MIT Press, Cambridge, MA, 1993.

[Den82] P Denning. Electronic junk. *Communications of the ACM*, 25(3):163–165, 1982.

[FG84] W. Finzer and L. Gould. Programming by rehearsal. *Byte*, 9(6):187–210, 1984.

[FM94] B. Fenn and H. Maurer. Harmony on an expanding net. *ACM Interactions*, 1(4):26–38, 1994.

[Hal88] F. Halasz. Reflections on notecards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, 31(7):836–852, 1988.

[HKR+92] B. J. Haan, P. Kahn, V. A. Riley, J. H. Coombs, and N. K. Meyrowitz. IRIS hypermedia services. *Communications of the ACM*, 35(1):36–51, 1992.

[Hyp89] Apple Computer Inc., Cupertino, CA. *Hypercard Reference Manual*, 1989.

[Jac90] R. Jacob. What you look at is what you get: Eye movement-based interaction techniques. In *Proc. CHI'90: Empowering People*, pages 11–18. ACM, 1990.

[JUC94] About J.UCS. http://www.iicm.tu-graz.ac.at/Cabout_JUCS, 1994.

[KAF+94] F. Kappe, K. Andrews, J. Faschingbauer, M. Gaisbauer, H. Maurer, M. Pichler, and J. Schipflinger. Hyper-G: A new tool for distributed hypermedia. In *Proc. Distributed Multimedia Systems and Applications Conference*, pages 209–214. ISMM-ACTA Press, 1994.

[Kim89] T. D. Kimura. Pen-based user interface (Panel Session). In *Proc. 1989 IEEE Workshop on Visual Languages*, pages 168–173. IEEE Computer Society Press, 1989.

[KM94] F. Kappe and H. Maurer. From hypertext to active communication/ information systems. *Journal of Micro Computer Applications*, 17:333–344, 1994.

[KMS93] F. Kappe, H. Maurer, and N. Sherbakov. Hyper-G - a universal hypermedia system. *Journal of Educational Multimedia and Hypermedia*, 2(1):39–66, 1993.

[LM94] J. Lennon and H. Maurer. Lecturing technology: A future with hypermedia. *Educational Technology*, 34:5–14, 1994.

[LMS93] K. Lee, W. Mansfield, and A. Sheth. A framework for controlling cooperative agents. *Computer*, 26(7):8–16, 1993.

[Mau93] H. Maurer. *Spekulationen Über die Multimediale Zukunft*. CAP, Debis, 1993.

[MGT+87] T. W. Malone, K. R. Grant, F. A. Turbak, S. A. Brobst, and M. D. Cohen. Intelligent information-sharing systems. *Communications of the ACM*, 30(5):309–402, 1987.

[MKSS93] H. Maurer, F. Kappe, N. Scherbakov, and P. Srinivasan. Structured browsing of hypermedia databases. In *Proc. HCI'93*, pages 51–62, Vienna, 1993. Springer.

[ML94] H. Maurer and J. Lennon. Forecasting: An impossible necessity. *NZ Science Monthly*, 5(1):12–13, 1994.

[MMK93] B. A. Myers, R. G. McDaniel, and D. S. Kosbie. Marquise: Creating complete user interfaces by demonstration. In *Proc. INTERCHI'93: Human Factors in Computing Systems*, pages 293–300, Amsterdam, 1993. ACM.

[MMS94] H. Maurer, H. Muelner, and A. Schneider. An electronic library and its ramifications. Technical Report 382, IIG, Graz, Austria, 1994.

[MPS94] H. Maurer, A. Philpott, and N. Scherbakov. Hypermedia systems without links. *Journal of Microcomputer Applications*, 17:321–332, 1994.

[MS82] H. Maurer and I. Sebestyen. 'Unorthodox' videotex applications: Teleplaying, telegambling, telesoftware and telecomputing. *Information Services & Use*, 2:19–34, 1982.

[MS94] H. Maurer and A. Schneider. New aspects of a hypermedia university representation. In T. Ottmann and I. Tomek, editors, *Proc. ED-MEDIA 94*, pages 504–509. AACE, 1994.

[MS95] H. Maurer and K. Schmaranz. J.UCS and extensions as paradigm for electronic publishing. DAGS'95 Electronic Publishing and the Information Superhighway, 1995.

[MSS93] H. Maurer, N. Scherbakov, and P. Srinivasan. A new hypermedia data model. In *Proc. DEXA'93*, pages 685–696, Prague, Czech Republic, 1993. Springer, LNCS 720.

[MW92] D. Mo and I. Witten. Learning text editing tasks from examples: a procedural approach. *Behaviour and Information Technology*, 11(1):32–45, 1992.

[MW93] D. Maulsby and I. Witten. Metamouse: An instructible agent for programming by demonstration. In A. Cypher, editor, *Watch What I Do: Pro-*

gramming by Demonstration, pages 155–181. MIT Press, Cambridge, MA, 1993.

[MWK89] D. Maulsby, I. Witten, and K. Kittlitz. Metamouse: Specifying graphical procedures by example. *Computer Graphics*, 23(3):127–136, 1989.

[MWKF92] D. L. Maulsby, I. H. Witten, K. A. Kittlitz, and V. G. Franceschin. Inferring graphical procedures: The compleat metamouse. *Human-Computer Interaction*, 7(1):47–89, 1992.

[Mye88] B. Myers. *Creating User Interfaces By Demonstration.* Academic Press, Boston, 1988.

[Pap80] S. Papert. *Mindstorms: Children, Computers, and Powerful Ideas.* Harvester Press, Brighton, 1980.

[PY93] P. P. Piernot and M. P. Yvon. The AIDE project: An application-independent demonstrational environment. In A. Cypher, editor, *Watch What I Do: Programming by Demonstration*, pages 383–401. MIT Press, Cambridge, MA, 1993.

[Rob91] M. Robinson. Through a lens smartly. *Byte*, 16(5):177–187, 1991.

[SB90] I. Starker and R. A. Bolt. A gaze-responsive self-disclosing display. In *Proc. CHI'90: Empowering People*, pages 3–9. ACM, 1990.

[Ste91] R. Stein. Browsing through terabytes - wide-area information servers open a new frontier in personal and corporate information services. *Byte*, 16(5):157–164, 1991.

[TO90] R. Tuck and D. R. Olsen. Help by guided tasks; utilizing UIMS knowledge. In *Proc. CHI'90: Empowering People*, pages 71–78. ACM, 1990.

[Wei66] J. Weizenbaum. Eliza: A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, 1966.

[WM92] S. Wu and U. Manber. Fast text searching allowing errors. *Communications of the ACM*, 35(10):83–91, 1992.

[WM93] I. H. Witten and D. Mo. TELS: Learning text editing tasks from examples. In A. Cypher, editor, *Watch What I Do: Programming by Demonstration*, pages 183–203. MIT Press, Cambridge, MA, 1993.

[YHMD88] N. Yankelovich, B. J. Haan, N. K. Meyrowitz, and S. M. Drucker. Intermedia: The concept and the construction of a seamless information environment. *IEEE Computer*, 21(1):81–96, 1988.

[Zha93] R. Zhao. Incremental recognition in gesture-based and syntax-directed diagram editors. In *Proc. INTERCHI'93: Human Factors in Computing Systems*, pages 95–100, Amsterdam, 1993. ACM.

[ZW85] A. Y. Zissos and I. H. Witten. User modelling for a computer coach: A case study. *International Journal of Man-Machine Studies*, 23:729–750, 1985.