



User Preferences for Task-specific vs. Generic Application Software

Bonnie A. Nardi

Apple Computer

1 Infinite Loop

Cupertino CA 95014 USA

+1-408-974-8708 nardi@taurus.apple.com

Jeff A. Johnson

FirstPerson

100 Hamilton Avenue

Palo Alto CA 94301 USA

+1-415-473-7230 jjohnson@firstperson.com

ABSTRACT

We conducted an ethnographic study to investigate the use of generic *vs.* task-specific application software by people who create and maintain presentation slides. Sixteen people were interviewed to determine how they prepare slides; what software they use; and how well the software supports various aspects of the task. The informants varied in how central slide preparation was to their jobs. The study was motivated by our beliefs that: 1) some software programs are task-generic, intended for use in a wide variety of tasks, while others are task-specific, intended to support very specific tasks; 2) task-specific software is preferable, but is often not used because of cost, learning effort, or lack of availability; and 3) people who infrequently perform a task tend to use generic tools, while people who frequently perform a task tend to use task-specific tools. Our findings suggest that the truth is more complex: 1) task-specificity/genericness is not a simple continuum; 2) a task cannot be looked at in isolation without reference to a higher level goal; and 3) an alternative to task-specific programs is a modular collection of independent interoperable services supporting small subtasks.

KEYWORDS:

Task-specificity, task analysis, slidemaking, end user computing, interoperability, collaborative work.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

CHI94-4/94 Boston, Massachusetts USA

© 1994 ACM 0-89791-650-6/94/0392...\$3.50

INTRODUCTION

It is widely acknowledged that personal computing has failed to achieve the ubiquity predicted in the late Seventies, at the dawn of the personal computer age. One reason is that most people are not interested in using computers *per se*, but only in performing tasks in their own familiar domains—tasks for which computers often provide poor support. There is a large gap between what users want to do to accomplish their tasks and the unfamiliar constructs presented by the computer [4, 6]. The failure of personal computers to become as ubiquitous as, say, washing machines, should thus come as no surprise. The predictions of the Seventies assumed a highly computer-centric view of the world. This view is not native to the average user who wishes to work in his or her own domain-specific idiom, rather than be forced to adopt the foreign jargon of the machine (files and directories, commands and arguments, control characters, cursors, modes, text strings, selection, etc.).

An antidote to the limitations of computer-centric software might be to provide highly task-specific applications that leverage users' familiar domain knowledge and narrow the gap between what they know and what the computer demands [1, 3, 4, 8, 10]. In this paper we try to pin down the nature of "task-specificity." What does it mean for a computer application to provide a high degree of support for a task? In what ways do applications vary in their degree of support for tasks? Under what circumstances is a high degree of task support necessary or unnecessary? Because of our interest in fostering the development of task-specific



applications as a means of promoting end user programming [5, 7], we conducted a study to gain some insight into the determinants of user preferences for task-specific vs. generic software. The task domain we studied was the creation and editing of slides for visual presentations.

In this paper, we first give some examples of software applications and discuss their varying levels of task-specificity. Next, we compare the support that various types of programs provide for slidemaking. Finally, we describe our study and findings, and offer some conclusions.

Task-Specific vs. Generic Tools

What exactly do we mean by "task-specific"? Consider the average kitchen. Most kitchens contain a large variety of tools. Some tools are used in many different tasks, e.g., knives, bowls, spoons, stoves, pots. Others are used for a relatively small set of tasks, e.g., blenders, peelers, tongs, basters, graters, cutting boards, butter knives. Still others are used for one task only, e.g., fish scalers, cheese slicers, nutmeg grinders, apple corers, cookie cutters, coffee makers.

For many tasks in which computers are used, a similarly large variety of software tools are in use, ranging from extremely generic to extremely specific. Some people use calculators for preparing income tax returns, some use spreadsheets, and some use income-tax programs designed for a particular year. Some companies do their accounting on calculators, some use spreadsheets, some use general accounting packages, some use accounting packages designed for a particular type of business (e.g., restaurants), and some use custom accounting software developed exclusively for them. For preparing organization charts, some people use painting programs, some use structured drawing programs, some use general tree-graph editors, and some use organization-chart editors.

CREATING PRESENTATION SLIDES

Slide preparation is a task for which a wide variety of computer-based tools are used. People prepare slides using text editors, desktop publishing systems, painting programs, drawing programs, spreadsheets, statistical-analysis programs, business-graphics programs, animation programs, and, recently, presentation-making programs.

Programs such as Charisma, Persuasion, and PowerPoint¹ are highly task-specific programs

intended only for slidemaking. Slides comprising a presentation are contained in one file. The format and common content of slides are specified once for the presentation, rather than separately for each slide. Typical slide editing actions such as removing or adding a level of detail are explicitly supported.

Despite the existence of these programs, many people use structured drawing programs such as MacDraw for making presentation slides. With drawing programs, users place manipulable graphical objects on a canvas. Text and graphics, once placed, can be edited, moved, or copied. However, drawing programs' degree of support for slidemaking is limited. They offer, for example, no notion of a presentation *set* of slides. Users can compensate by putting slides into separate files and grouping them in folders or directories, or by placing all the drawings of a presentation together on one canvas, but such workarounds are inconvenient and inefficient. Drawing programs provide no support for consistency of format, font, and layout in a presentation. For example, to have the same margins on every slide, users must painstakingly arrange things that way, separately for each slide. If the formatting requirements change, users must change every slide. Standard slide content, such as logos, headers, and footers, must be explicitly placed on every slide, and changes require editing every slide. Drawing programs provide little help in changing the structure of a presentation's content; e.g., splitting one slide into two requires much explicit copying, moving, and deleting of graphic objects. Drawing programs have no provision for task-specific actions such as attaching speaker's notes to a slide, as is possible with task-specific slidemaking software.

In short, drawing programs lack the concepts of *slides*, *relations between slides*, and a set of slides comprising a *presentation*. Most of the other types of software used for preparing slides also lack support for the process of creating and editing presentation slides. Nonetheless, these programs are commonly used for slidemaking. We will try to explain why as we proceed.

A True Story

In the early 1980's, at one company where the second author worked, employees made presentation slides using a text editor in conjunction with a slide-formatting program that had a conventional text-based user interface. To make a set of slides, users created a text file containing the textual slide content and embedded formatting commands. The text file was then "compiled" using the slide formatter, producing a set of graphics files containing images of the slides. Users had many complaints about this process. The formatting commands were hard to learn; it was difficult to tell how a particular slide

¹ Trademarks: Charisma, MicroGrafx; Persuasion, Aldus Corporation; PowerPoint, Microsoft.



would look from its source file; one had to "compile" the entire set to check how a single slide looked.

When interactive painting and drawing programs became available, users switched to them immediately. The new programs were easier to learn and provided much better feedback than the text editor/formatter combination.

However, it soon became apparent that with the new programs it was hard to obtain consistent formatting, hard to manage sets of slides, and hard to edit slides. After a short honeymoon, most users switched back to the old programs, occasionally using the new ones to enhance slides generated by the tried-and-true slide formatter. At the time, this mass retreat to the old tools was difficult to understand. Users had, unaccountably, eschewed the advantages of direct manipulation interfaces and WYSIWYG interaction to return to an old-fashioned text-based command language and tedious compile-and-debug cycle.

MOTIVATION FOR THE STUDY

As we began our research on end user programming in the late 1980's, we believed we understood the reason why the users in our story had abandoned direct manipulation software for making their slides: *it was not task-specific*. The new programs were not *slidemaking* programs, but rather generic painting and drawing programs. We reasoned that though the new programs made some *slidemaking* tasks easier, they made other, more crucial tasks harder. The aspects of *slidemaking* that the new tools made more difficult were the deeper, more task-related aspects. While users could, with effort and talent, make nicer looking presentations with the new tools, they could produce *acceptable* presentations much more quickly with the old tools.

As we were to discover in the present study, this analysis was too simple. The goal of this paper is to explain what we have learned about task-specificity and how it affects users' preferences.

Our initial analysis led us to a set of beliefs, the gist of which is captured in the previous story, that can be stated as follows:

1. Software applications occupy a position along a continuum, from completely generic to completely task-specific.
2. Task-specific applications are preferable. It is better to have a tool designed specifically for the task one is performing. For example, for working with schematic drawings, a schematic editor would be preferable to a drawing editor.
3. People use generic applications when more task-specific applications are either unavailable, cost

too much, or require too much learning effort. Users "get by" with generic tools because their level of need does not justify the cost of obtaining and learning to use task-specific tools.

4. People prefer task-specific tools when they perform a task frequently. Their level of use justifies the overhead of acquiring and learning to use the tool.

Based on these beliefs, we (with other colleagues) set about developing an Application Construction Environment (ACE) designed to facilitate the development of task-specific software applications [5, 7]. One of our goals was to change the economics of software development to make application development easy and cheap enough so that it would be cost-effective to develop highly task-specific applications for small, specialized markets. A major goal of ACE was to move the development process closer to users, who best understand their goals and tasks.

As we developed ACE, we were aware that we should regard our beliefs as working assumptions to be empirically tested. Over time, we began to have reason to suspect that our initial assumptions weren't quite right. For example, a graphic artist friend who had done a lot of work producing presentation slides for others indicated that important parts of what we believed were false. She claimed that experienced professional *slidemakers* prefer generic drawing and painting software for creating slides because it doesn't restrict them from doing what they want, while dedicated *slidemaking* programs often impose over-simplified views of the task and restrict the results that can be produced. Based on this counterclaim and on our own further analysis of the nature of task-specificity, we decided to conduct an empirical study as a first step in evaluating and correcting our assumptions. We chose the domain of *slidemaking* because of the large variety of software tools used for the task, the accessibility of users as informants, and, not least, the challenge posed by our friend.

METHOD

We conducted an ethnographic study [2] to examine how people create presentation slides. The informants² were sixteen people whose jobs involved creating, editing, and maintaining slide presentations. All were college educated with several years experience making slides. They worked for a variety

² In an ethnographic study, participants are called *informants* as their role is to actively inform the investigator. This is in contrast to the use of the term *subjects* for experimental studies, where participants are subjected to experimental conditions and observed.



of companies, ranging from single-person independent consultancies to large multinational corporations in the San Francisco Bay area (most outside of Silicon Valley). Six of the sixteen informants worked in research or marketing, and made slides for their own use in presentations, with slidemaking being only one of many of their job responsibilities. The other ten can be considered professional slidemakers; they had as a significant (for some, dominant) part of their job the creation of presentation slides for others, in a variety of business areas: legal, advertising, research, and general business.

We developed a set of questions (available upon request from the first author) that we asked each informant. Interesting conversational threads were opportunistically pursued as they arose. The interviews were audiotaped at the informant's workplace, often with a computer slidemaking system ready-at-hand so that we could see the user's work on-line.

During the interviews, we began by explaining that the purpose of the study was to learn what is involved in making slide presentations, what sorts of software people use for the task, and what people's reasons are for using or not using various software tools. We asked the informant to describe the entire slidemaking process, from start to finish. We allowed the conversation to flow naturally rather than strictly following the list of questions, but made sure that answers to each of the predetermined questions were captured on tape. We did not explain the distinction between task-specific vs. generic software, or our initial working assumptions. Interviews ranged from 1–3 hours per informant.

We found that our informants were quite happy to talk about their slidemaking software. Several warned when making interview appointments that their busy schedules could accommodate only a brief interview, but once the interview was underway they seemed willing to talk for as long as we would listen. People have strong opinions, both positive and negative, about the software they use.

The taped interviews were transcribed, resulting in about 250 pages of text.

Data Analysis

We read the transcripts of each interview, in some cases referring to the audiotape to clarify transcription problems or informant intent. A summary was made of each interview that included: the informant's job role and involvement in slidemaking, the context in which slides were being produced, a summary of the slide-production process as described by the informant, the software the

informant uses or has used for slidemaking, the informant's reasons for using it, software features that the informant considered useful or a hindrance in slidemaking, and informant comments (if any) that seemed especially germane to the study.

RESULTS AND DISCUSSION

In a nutshell, we found that our original assumptions were right in some respects and wrong in others, and that the truth is more complex than either we or our graphic artist friend understood. As we predicted, people who infrequently make slides—that is, those for whom slidemaking is peripheral to their main job—typically use whatever general purpose software they use for other work, e.g., word-processing. One person, for example, used a combination of EMACS and LaTeX to produce slides. In terms of learning and installation time, and software purchase price, it just doesn't pay such users to bother with task-specific slidemaking programs.

However, contrary to our assumptions, but in agreement with the claims of our graphic artist friend, we found some users who are professional slidemakers—i.e., those who spend most or all of their time making slides—using generic software such as drawing, painting, and word-processing programs extensively, rather than sticking to task-specific slidemaking programs. We found other professional slidemakers using task-specific slidemaking programs, in tune with our original assumptions. In the following discussion, we focus on the professional slidemakers in our study and describe two of the factors affecting their differing software choices. (A longer paper describing other factors of interest is in preparation.)

Factors in Choosing Task-specific vs. Generic Software

Factor #1: Presentation Quality

One factor that influences the choice of software for professional slidemakers is desired presentation quality. Some presentations are considered "ordinary" while others are "fancy" or "very important." Most presentations are relatively mundane; others have millions of dollars riding on the impression they make. We found that in cases where slidemakers want to produce simple presentations quickly, they use task-specific slidemaking programs (though factors such as familiarity and availability sometimes limit this tendency). This is what we had predicted. But we were surprised to learn that for fancy presentations, which are usually planned and prepared long in advance, slidemakers—generally highly skilled graphic artists in this case—usually use generic



drawing, painting, desktop publishing, or animation software. Why?

Expert slidemakers use generic programs for fancy presentations because the illustration and/or text formatting capabilities of slidemaking programs are insufficient and limiting from their point of view. For example, one user said,

"I think the main point about why we use MacDraw is because, yeah...Persuasion would be better for a lot of word slides...But nobody's willing to simplify their graphs that much. You know? It's like they would have to...work at such a simple level to make a presentation, that...the [clients] can't cut down on the complexity of their slides, to be able to fit in with the limitations of a program like that."

What this tells us then, is that when we speak of "task-specificity" we must be precise about *which aspects of a task we are talking about*. For slidemaking, is it illustration or text formatting or managing a set of slides in terms of common formats and file manipulation? To fully understand which aspects of the task are most important in predicting tool use we must back up one level further, and look at the *goal behind the task*. The goal of producing a fancy presentation leads to a different set of choices than the goal of producing an ordinary presentation. Each kind of presentation entails illustration, text formatting, and slide organization subtasks, but the differing goals mandate optimizing different aspects of the overall process.

The professional slidemakers doing fancy presentations provide a mirror image example to our story in which users abandoned generic illustration software in favor of their old dedicated slide formatter. The users in that story did not want to optimize illustration; for them the goal was to get out decent slides quickly, rather than to produce beautifully illustrated slides.

Our original assumption that software should always be as "task-specific" as possible thus comes into question when we consider the variability in something seemingly as simple as slidemaking. In terms of software choice, "slidemaking" has fractured into "making ordinary presentations" vs. "making fancy presentations." What we see here is that *any goal may involve the execution of subtasks that are common to many goals* (e.g., subtasks of drawing, illustration) *as well as subtasks that are more narrowly conceived* (e.g., containing and managing a set of slides). When we consider the differing software choices dictated by ordinary vs. fancy presentations, what at first appeared to be "a task"—i.e., slidemaking—has proven not to be a conceptually clean cut at the problem. It is more helpful to begin the analysis with the user's *goal*,

and then to consider the subtasks necessary to fulfill the goal. The subtasks involved in slidemaking vary crucially depending on the quality and turnaround time desired for a presentation. Different tools are entailed by the differing demands of a given presentation.

Perhaps, though, we have erred in framing "slidemaking" as the task, instead of declaring the task to be "making a fancy presentation" or "making an ordinary presentation." Such an analysis seems incorrect to us in that so many different factors may be relevant to a given situation; with slidemaking, for example, in predicting software choice we have also to consider the frequency of use of a program, as we discussed earlier. Many variable factors can affect software choice. Without a clear sense of the user's goal and some context—e.g., a user who infrequently makes slide presentations wants to make an ordinary presentation quickly—it is not meaningful to speak of "a task" as an isolated disembodied entity.

What a preliminary analysis of our data might seem to suggest is that slidemaking software should provide high quality illustration/text formatting capabilities *and* slide manipulation and format capabilities. This would certainly seem best if software cost, in terms of both development time for the vendor and dollar cost to the user, were not at issue. But cost *is* an issue, which may explain why programs tend to optimize some but not all subtasks relevant to the achievement of users' goals.³

How then, can we support users' varying goals and the differing contexts of software use? One answer is suggested by the practice of some of the professional slidemakers in our study who create fancy presentations. To mitigate the limitations of the different software programs that can be used to make fancy presentations, they use *both* generic and dedicated slidemaking software: for example, illustrations are produced with drawing tools, and then the files are exported to a slidemaking program such as Persuasion which is used to contain and organize them. (Slidemaking programs that do not easily accept text and graphics files from other programs, are, needless to say, not popular with these slidemakers.)

Thus we see slidemakers devising a strategy of the use of multiple, interoperable services to achieve their goals. This example suggests that perhaps we should be evaluating whether it is preferable to

³ Many programs seem to provide good support for one subtask while being surprisingly deficient in others. For example, Instant Update (ON Technology, Inc.) is very good at update management but has an unimpressive text editor. For collaborative writing, both are critical.



design task-specific programs that attempt to cover every aspect of "a task," or whether we should be re-thinking the whole issue, considering a software architecture that provides modular, integrated services that users can pick and choose as needed, services that play together easily.

Factor #2: Support for Teamwork

Another factor in software choice for slidemaking that we had not even considered in our original working assumptions is the degree of teamwork supported or allowed by the software. Most slidemaking programs are designed to support an individual who produces slides alone. However, in our study we found that in many settings, people work in teams to produce and maintain slides. Existing slidemaking programs make it difficult to do this. According to one informant:

"I looked at Persuasion, because everybody was saying that Persuasion was great. And I think a bunch of the secretaries...used it as well. And I think the programs that are that specific are very well designed for a person who is going to sit down and think up a presentation and create the presentation right there. But the way we work is that...there are dozens of people out there thinking up things, and we integrate presentations for all of them. And so for us to be able to *distribute that work amongst enough people to get it done, we need to break it down into smaller units*..For each job here, if we used Persuasion, each job...would have its own...file with all of its slides in it. But slides get used from one job to another...And so I think because of that, [Persuasion] wouldn't work. The outlining...you know, is wonderful. But [Persuasion is] really designed for a different type of work atmosphere. It's designed for the guy who's sitting down and going to do his own presentation" (emphasis added).

Though more generic tools don't provide real support for team production, they at least don't interfere with it in the sense that they impose little structure on the process at all. While task-specific programs certainly could, in principle, support teamwork, today's task-specific slidemaking systems do not encourage sharing of slides and files; thus we find another reason why some of our informants chose not to use task-specific programs.

The need to support teamwork argues for preferring a set of modular interoperable services over task-specific systems. In many settings, because "slides get used from one job to another," as our informant noted, interoperability is again a key feature needed by users. An architecture supporting the interplay of small modular services to be applied to the slides for varied purposes by different users would enable collaborative development distributed across a set of cooperating users.

CONCLUSIONS

Before our study, we assumed that task-specific software is generally a good thing. The primary hindrance to its use, we thought, was acquisition and learning costs: how much does it cost to get the software, and how much of an incremental learning "hump" does using the tool require? Companies don't like to spend money. Users avoid learning new things. Thus, tools users already have and know have a strong advantage.

But we found that cost, in terms of money and learning time, isn't the whole story (see [9] on identifying costs). Several additional criteria are important in deciding which tools will be used to make presentations:

- **Power:** Can the required slides be produced with the task-specific tool? For example, if beautiful illustrations are needed, can they be created?
- **Support for teamwork:** Does the tool support people working together on a presentation, if that is how presentations are produced at the worksite in question? Today's slidemaking tools assume a single user working alone, but in fact, many presentations are created by teams, with different people contributing different parts.
- **Interoperability:** Can the tool easily take input from other tools such as might be needed; e.g., can a drawing made with one program be put into a slide or presentation made using another? Programs that provide good interoperability are preferred.

We went into our study thinking that software applications vary along a continuum of task-specificity and found that it isn't that simple. There are several aspects to task-specificity, including a user's goal in performing a task and the necessary subtasks that enable the fulfillment of the goal. What we call "task-specific" software programs today really only support some subtasks relevant to a given goal. Users' goals vary widely enough that it is difficult to create the many individual off-the-shelf programs that would be necessary to support these goals. As one user observed,

"Persuasion and PowerPoint are sort of integrated programs and they're good for someone who isn't a power word processor, who isn't a power graphic artist, where they basically want to type in their own headers and dot points, and it's great for that...But if you have to go beyond that where you're...doing real serious word processing, or doing some real elaborate graphics, it just doesn't cut it either way...*There hasn't been any software that does everything well*" (emphasis added).



We found that professional slidemakers creating fancy presentations prefer using collections of interoperable tools. Their usage pattern suggests a modular set of interoperable software services as an alternative to individual highly task-specific programs, which often fall short on some crucial subtask necessary for getting a job done. The subtask of outlining, for example, mentioned in a previous quotation, is surely a service useful for many kinds of work. Outlining could be selected from a set of services and applied to the task of slidemaking when appropriate.

Given a set of interoperable services, users might come to bundle individual services together into personalized packets to capture regularities in their daily work, depending on the kinds of goals they pursued. Local developers ("office gurus") might also be enlisted to bundle services for a group of users, or help them do it (see [7]). They could develop shareable templates as happens in many domains [7] and users themselves could share individually developed templates.

In terms of documentation, instead of a single program documented by a fat manual with a great deal of daunting, irrelevant detail, users would work with "recipe books" suggesting sets of ingredients that could be combined to get started on a task. Heavily annotated examples and rich indexes would also support the learning task.⁴ Users would not be faced with trying to use a software program with a non-optimal bundle of services (as many of the slidemaking systems have, as we have tried to illustrate); rather they would construct their own environments from small, modular, interoperable services. This is essentially what we observed the professional slidemakers doing, and without benefit of programs explicitly designed to be used together.

Of course, the generality of the findings reported here should be validated through comparable studies in other task domains. We believe the present study argues strongly for in-depth ethnographic research on how users work: our original assumptions, which we felt to be well-reasoned, were actually rather shallow and wrong in important ways. A focused look at the work of real users in a simple domain of application development led us to rethink our key assumptions and to consider an alternative way of supporting end user application development.

⁴ Examples and indexes would be useful in any case, but since we are trying to suggest a new paradigm here, we'd like to underscore the importance, from the beginning, of taking the time to properly document software in ways not typically done today.

ACKNOWLEDGEMENTS: The research described in this paper was conducted while the authors were employed at Hewlett-Packard Labs, Palo Alto, CA. The authors thank Michelle Gantt, the student intern and graphic artist who served as the interviewer for the study and also—by questioning our initial assumption about who uses task-specific vs. generic software—as the primary catalyst for our conducting it. Many thanks to Jim Miller, Vicki O'Day, Dan Russell, and Craig Zarmer for helpful comments on earlier drafts of this paper. Last but not least, thanks to our informants for taking time out of their busy schedules to educate us about their work practices.

REFERENCES

1. Casner, S. (1991). A task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics* 10: 111–151.
2. Glaser, B. and Strauss, A. (1967). *The Discovery of Grounded Theory*. New York: Aldine Publishing.
3. Gould, J., Boies, S. and Lewis, C. (1991). Making usable, useful, productivity-enhancing computer applications. *CACM* 34: 75–85.
4. Hutchins, E., Hollan, J. and Norman, D. (1986). Direct manipulation interfaces. In D. Norman and S. Draper, eds., *User Centered System Design*. Hillsdale, NJ: Erlbaum Publishers.
5. Johnson, J., Nardi, B., Zarmer, C. and Miller, J. (1993). ACE: building interactive graphical applications. *CACM* 36: 40–55.
6. Lewis, C. and Olson, G. (1987). Can principles of cognition lower the barriers to programming? In *Empirical Studies of Programmers: Second Workshop*. Norwood, NJ: Ablex. Pp. 248–263.
7. Nardi, B. (1993). *A Small Matter of Programming: Perspectives on End User Computing*. Cambridge, Mass: MIT Press.
8. Olsen, D., McNeill, T. and Michell, D. (1992). Workspaces: An architecture for editing collections of objects. *Proceedings CHI'92*. Monterey, CA. 3–7 May. Pp. 267–272.
9. Russell, D., Stefik, M., Pirolli, P. and Card, S. (1993). The cost structure of sensemaking. *Proceedings INTERCHI'93*. 24–29 April. Amsterdam. Pp. 269–276.
10. Vlissides, J. and Linton, M. (1990). Unidraw: A framework for building domain-specific graphical editors. *ACM Transactions on Information Systems* 8: 237–268.