



ELSEVIER

Computers and Electronics in Agriculture
30 (2001) 85–102

Computers
and electronics
in agriculture

www.elsevier.com/locate/compag

Ethics in computer software design and development

Alan J. Thomson ^{a,*}, Daniel L. Schmoltd ^b

^a *Pacific Forestry Centre, Canadian Forest Service, 506 West Burnside Road, Victoria, BC, Canada V8Z 1M5*

^b *USDA Forest Service, Southern Research Station, Department of Biological Systems Engineering, USDA Forest Service, 460 Henry Mall, Madison, WI 53706-1561, USA*

Abstract

Over the past 20 years, computer software has become integral and commonplace for operational and management tasks throughout agricultural and natural resource disciplines. During this software infusion, however, little thought has been afforded human impacts, both good and bad. This paper examines current ethical issues of software system design and development in relation to privacy, accuracy, property, accessibility, and effects on quality of life. These issues are explored in the context of simulation models, databases, geographic information systems and artificial intelligence programs, especially expert systems. New approaches to system development place a much higher emphasis on the effects of system deployment within a complex human environment. Software design decisions often depend on more than one ethical issue, possibly conflicting, where the appropriate ethical choice is not always clear cut. Professional codes of ethics do little to change peoples' behavior; rather, incentives for using an ethical approach to software development may lie in significantly increased likelihood of system success. Crown copyright © 2001 Published by Elsevier Science B.V. All rights reserved.

Keywords: Ethics; Software design and development; Traditional ecological knowledge; Indigenous knowledge; Intellectual property; Information ecologies

* Corresponding author. Tel.: + 1-250-3630632; fax: + 1-250-3630775.

E-mail address: athomson@pfc.cfs.nrcan.gc.ca (A.J. Thomson).

1. Introduction

Ethics is the study of value concepts such as ‘good,’ ‘bad,’ ‘right,’ ‘wrong,’ ‘ought’, applied to actions in relation to group norms and rules. Therefore, it deals with many issues fundamental to practical decision-making (Veatch, 1977). Computer software systems lie at the heart of modern decision making, including data/information storage and manipulation, data availability, and ‘alternatives’ formulation and selection. In fact, the very use of computer systems can often frame the types of questions that can be asked as well as their possible answers. This is particularly evident when we incorporate software systems into our knowledge management methods (Schmoldt and Rauscher, 1994), as they then play an essential role in institutional memory. The ubiquity of software systems in all aspects of public and private institutions means that the environment that they create needs to be critically examined as they are developed and deployed.

Two major ethical questions must be addressed with regard to software systems. Firstly, can these systems represent the different codes of ethics of the groups affected by software-mediated decisions? Secondly, what ethical considerations should guide the design and development of the software itself?

In regard to the first question, a range of artificial intelligence (AI) approaches has been proposed to represent different codes of ethics in environmental decision systems (Thomson, 1997). The present study addresses the second question by exploring ethical issues in design and development of systems in general, and four types of system, in particular, simulation models, databases, geographic information systems (GIS), and artificial intelligence programs.

The role of ethics in software system design has increased in importance recently. Mason (1986) gives an early perspective on ethical issues in the information age, categorizing them into privacy, accuracy, property, and accessibility concerns. More recent views add concerns about the use of knowledge in organizations (Bella, 1992) and concerns over effects on quality of life (Forester and Morrison, 1994) to the previous issues. This review will not address intentionally malicious behavior, such as computer crime, software theft, hacking, viruses, and deliberate invasions of privacy, but rather will explore the subtler, yet important, impacts that software development and deployment can have on people and their cultural, corporate and other institutions. First, some ethical implications of adopting a particular approach to system design are examined. Next, the specific ethical issues mentioned above are presented in turn. The paper concludes with a discussion of professional codes of ethics, and a rationale for adopting an ethical approach to system design and development.

2. Ethical approaches to system design

At the broadest level, ethics can be applied in the overall approach to system design; however, those approaches vary considerably in their ability to deal with ethical issues. Traditional approaches to system design and development, such as

the structured systems analysis and design method (SSADM), the most popular development methodology in the UK, focus more on technical issues than on human issues.

SSADM addresses technological aspects of system development by breaking down system development into smaller tasks. It consists of a sequence of stages, viz. feasibility, requirements analysis, business systems options, requirements specifications, technical systems options, logical design and physical design. Each stage consists of a number of steps, and each step consists of a number of tasks. At the lowest level, SSADM consists of approximately 230 tasks, not all of which may be executed in a particular implementation. Stage definitions tell the developer which preceding products feed into the next stage, and what products that stage will create. SSADM attempts to define three views of a system, how the data items in the system move through it; how the various data items are related to each other; and how each data item changes over time (Kendall and Kendall, 1999). While the 230 tasks currently composing SSADM do not specifically address ethical issues per se, it is conceivable that SSADM's stages/steps/tasks could be restructured to do so.

An alternative approach, the CATWOE model, identifies the customers, actors, transformation processes, worldview, owners, and environment. Because it explicitly includes people's roles, it is better able to support ethical concerns, especially if augmented by construction of an ethical conflict web (Wood-Harper et al., 1996). The approach attempts to identify who is doing what for whom, to whom are they answerable, what assumptions are being made, and in what environment this is happening. Drawbacks of a CATWOE approach are that unlike a SSADM approach, it does not actually tell you how to build a system. Additionally, it assumes that managers and workers openly discuss their problems.

The concept of 'information ecologies' (Davenport and Prusak, 1997) is a new approach to system development that views technology in the context of the people, organization and practices that relate to system use. This approach is based on analogies with ecology such as information diversity (species diversity), information environment evolution (biological evolution), emphasis on observation and description, and a focus on people and information behavior. This focus on people and the way in which they use information should be an improvement over traditional approaches, although it lacks the specific focus on ethics discussed by Wood-Harper et al. (1996). In any system design methodology, however, there should not only be a focus on how people use information, but also on how people misuse information or systems, explicitly indicating how a system should not be used.

3. Privacy

Improper access to personal information is the issue that 'privacy' usually brings to mind. In this section, privacy is examined from the standpoint of data fusion, location privacy, public information, and Internet technologies. Each creates unique problems for software design, development, and deployment.

3.1. Data fusion

Any unauthorized access to information can be an invasion of privacy. However, even authorized access may lead to privacy concerns, when access to separate data sources is used to combine information (Mason, 1986). For example, one institution may record an individual's name and employee number, while another may be authorized to store employee number and health insurance claims, but the combination of the individual's name and the health insurance claims may be an invasion of privacy. This is not yet a major issue in system development for natural resources and agriculture. However, as environmental databases increase in size, complexity, and connectivity, software projects that involve combining data or knowledge sources must consider the ethical implications of those activities.

3.2. Location privacy

In recent years, a new privacy issue has arisen in the field of GIS, related to location protection. For example, the Forest Practices Code of British Columbia mandates protection of First Nations' (indigenous peoples') cultural sites, yet entering site locations in a GIS may disclose locations for unethical use. In this case, a polygon can be defined to contain a site or group of sites, without disclosing exact point of locations. A similar situation exists in relation to biodiversity and rare species protection. Innovative approaches are required to facilitate resource monitoring and protection while simultaneously ensuring there is no loss of privacy resulting from location disclosure.

3.3. Public information

Privacy issues arose during development of the Canadian Forestry Researchers Directory (CFRD)¹. The system was designed to permit scientists to enter information about their research programs. However, to ensure full coverage of research programs, it was also technically feasible to automate production of entries for in-house staff based on other on-line documents (all of which were in the public domain). It was finally decided that privacy issues related to involuntary participation should be the guiding criteria, rather than the corporate benefits resulting from full and inclusive coverage of programs. Similar concerns related to re-packaging of information may arise in the US, where public funding of government research places researcher and research information in the public domain. It is conflicting interests such as these that makes ethical conduct a gray area and makes engineers and computer scientists uncomfortable.

¹<http://www.pfc.cfs.nrcan.gc.ca/cfrd/>.

3.4. Internet technologies

The use of autonomous software agents that roam the Internet raises a range of new ethical issues. “An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future” (Franklin and Graesser, 1996). While such agents currently perform useful tasks such as indexing the World Wide Web, they can also lead to privacy problems if they are not specifically excluded from Web sites. Danielson (1992) has explored how agents should interact among themselves. He argues that, for machines or agents that face unstable social co-operation (i.e. some agents adhering to rules and others ignoring those rules), well-designed principles of moral constraint are rational. In particular, it is rational for machines to be able to keep promises, e.g. to co-operate with promise-keeping agents. Just as there are hardware, software, and data protocols that constrain computer system interactions, we can eventually expect, and will need, similar standards for inter-agent behavior. Many of these inter-agent protocols will likely be couched in social/ethical terms.

Another side to the privacy issue stems from ‘push’ technology, which in turn may also depend on agents. Push technology (Aragon, 1997) delivers, or ‘pushes’ specific content automatically to your computer. Ethical concerns relate to whether recipients actually want what is pushed, and if they have the appropriate bandwidth and computer capability to receive it. As many Internet companies rely on commercial advertising to support their Web site businesses, we are constantly bombarded with graphical images that may not be desired, that can significantly affect browsing speed, and can consume general Internet bandwidth. Push technology is currently under investigation as an extension to a virtual ‘adaptive environmental management system’ (Thomson, 2000b). In that system, participants (stakeholders) can define their value systems by weighting socio-economic and biological indicators (Fig. 1). The eventual goal is to provide customized feedback and reporting (‘push’) to an individual or group, based on a participant’s value rankings. In this way, users can specify the type and amount of information they receive.

4. Accuracy

Accuracy is a broad topic and so has many associated ethical issues. System inputs, internal processing, and system outputs can all affect accuracy, and at each level there are several important ethical problems. The following sections enumerate some of these.

4.1. Software complexity and accuracy

A system analyst’s ability to know and predict all states (especially error states) is low for complex systems. This leads to several ethical issues related to software

accuracy. Documentation of assumptions (if we are aware of them), development of appropriate test conditions and performing thorough system validation and verification are well-known approaches to such issues. At first sight, it would appear that a system developer would be ethically bound to correct all system errors. However, dealing with errors can raise ethical dilemmas — 15–20% of attempts to remove program errors introduce one or more new errors. For programs with more than 100 000 lines of code, the chance of introducing a severe error when correcting an original error is so large that it may be better to retain and work around the original error rather than try to correct it (Forester and Morrison, 1994)! The frequency of disclaimers, software updates and patches, as well as the lack of substance to software warranties, result from software developers' recognition of this problem (Forester and Morrison, 1994). The ultimate effect is larger and more complex software, whose size is less related to functional capability than it is related to software age and the battery of 'fixes' that it has received over time.

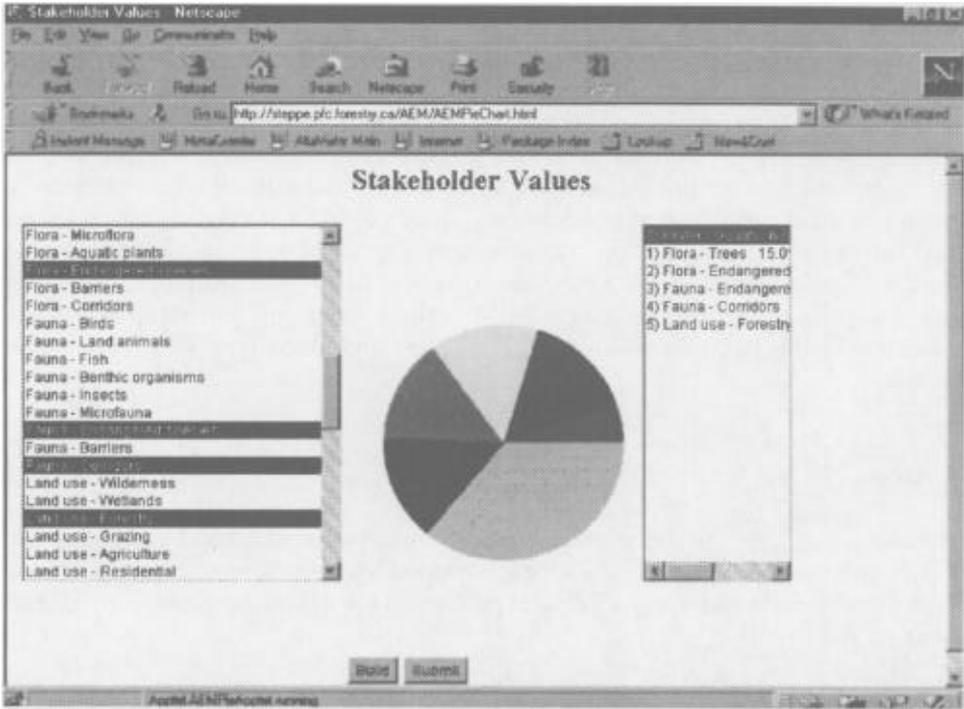


Fig. 1. Elicitation of stakeholder values. First, the indicators of interest are selected, then the relative values are established through manipulation of the pie chart (from Thomson, 2000b). This expression of interest is currently being used as the basis of an information 'push technology'.

4.2. Indicators and models

Another problem related to accuracy is determining which specific information to use. For example, it is often difficult to select appropriate socio-economic or biological indicators or to choose among predictive models. An *indicator* is something that points to an outcome or condition, and shows how well a system is working in relation to that outcome or condition. For example, in a forest simulation model, tree diameter at breast height (dbh) is a key indicator of treatment effects. However, there may be a range of potential equations available to predict dbh. One equation may simply predict dbh from tree height, while another equation may predict it from both height and crown width. The equation selected will have different consequences with regard to accuracy, precision, data costs, and suitability for extrapolation. This choice relates, in turn, to precision and bias in the estimators used. Requirements of the intended user and usage should guide the choice.

Goodenough et al. (1994) address this issue in a case-based reasoning system that attempts to answer queries that combine GIS and remote sensing. Their system includes knowledge about a range of processors, GIS platforms, and a wide range of aircraft and space-borne optical and radar sensors. The system, developed in Prolog, is based on a generic ‘solve algorithm’ for multiple goals,

```
solve((Goal and Goals), Solution) if
  solve_one_goal(Goal, Sol1) and
  solve(Goals, Sol2) and
  merge( Sol1, Sol2, Solution).
```

Accuracy and other information about images and GIS layers are maintained as part of the metadata as they pass through a sequence of transformations. Accuracy issues can then be part of the *solve_one_goal* clause.

Accuracy may also be influenced by the sequence in which operations are applied. In theory, error limits of predictions should be supplied; however, while error limits of individual equations may be known, it is rare that models actually compute the consequences of combining multiple equations. Mowrer (2000) examines error propagation in simulation models and presents several approaches (Monte Carlo simulation and Taylor series expansion) to project errors. This has become an active research topic recently (q.v. Mowrer, 2000), as often many models are used in combination to predict future conditions.

Even when accuracy figures are provided, they generally refer to accuracy over the whole range of potential applications. If a system is only used for difficult cases, the accuracy achieved may be much less than the quoted accuracy (Saveland et al., 1988). This implies that consideration should be given to provide case-specific error estimates that reflect the uncertainty and complexity of the case at hand.

When a social or economic indicator is being used, ethical considerations are even more significant. If the indicator misrepresents a value set, then it cannot be considered accurate. Indicators have long been used in predictive systems (Holling, 1978); such indicators must be relevant, understandable, reliable, and timely. In natural resource disciplines, with their current emphasis on sustainability, indicators

must have additional characteristics. Sustainability indicators must include community carrying capacity; they must highlight the links between economic, social and environmental well-being; they must be usable by the people in the community; they must focus on a long range view; and they must measure local sustainability that is not at the expense of global sustainability (Hart, 1999). Scale is a key determinant of indicator usefulness; some indicators that are useful at the household or community level are difficult to measure at the regional level, and some regional indicators may have little meaning at the community or household level. Because indicators compress so much ecological, economic, or social information into a single variable or set of variables, it is especially crucial that they are chosen, measured, and interpreted carefully.

Bernadas (1991), discussed in Chambers (1992) describes a case where an indicator, 'soil fertility,' was derived from a questionnaire based on researchers' preconceived view of priorities. However, 2 years of research based on the survey did not match farmers' needs and circumstances. Subsequently, researchers established 'duration of fallow' as a more relevant indicator, as it was directly affected by weeds of concern to the farmers. This more successful indicator resulted from a participatory process that involved informal dialogues and open-ended interviews, rather than the structured questionnaire used previously. This example illustrates that appropriate selection processes are required to ensure adoption of meaningful indicators.

4.3. Subjective judgment

When an indicator is based on subjective probabilities, accuracy can be affected by a whole array of biases (Poulton, 1994) that require a system developer to make ethical decisions. In the case of wilderness fire management, Saveland et al. (1988) discuss examples of some of these biases, such as hindsight bias, preference for certainty, and loss versus benefit. In hindsight bias, people consistently exaggerate what could have been anticipated in foresight. Also, people tend to favor certainty over uncertainty, so that if a fire is controlled immediately, there is no need for a manager to worry about its future development. Finally, potential losses have a greater influence on fire decisions than potential benefits (Saveland et al., 1988) due to the risk-averse behavior of fire management personnel. People possess 'bounded rationality' (Simon, 1979) so their decision making and subjective judgments are influenced by these biases, and many others.

4.4. Language and culture

Language and terminology used to frame a question can significantly influence the accuracy of the information elicited. This is true for any system in which the system user is forced to converse with software using concepts unfamiliar to them. This cultural mismatch is of special significance in studies of 'traditional ecological knowledge', where the interview subject may have sets of concepts and values very different from those of the questioner. For example, the term 'forest' is a key

concept for resource management but certain native peoples in Canada have no concept for forest in their culture or any word in their native tongue. Instead, they have a more holistic view of the land that includes trees, plants, animals and people (Thomson, 2000a)². Once such basic cultural differences are identified, the important challenge becomes one of understanding the ramifications of those differences, which may include a complete re-examination of all other ecological knowledge that borrows from those basic concepts and tenets.

4.5. *Software system output*

Ethics also apply to the accuracy with which results are portrayed by software systems. Many of the biases that affect information collection can also affect the interpretation of results (Poulton, 1994). While many of the issues related to knowledge presentation arise inadvertently, presentation style can be deliberately selected to influence perception, as illustrated in the book ‘How to lie with statistics’ (Huff, 1954).

As described in the book ‘How to lie with maps’ (Monmonier, 1991), when system results appear as maps, any deliberate influence on perception is ethically less clear cut. ‘As a scale model, the map must use symbols that almost always are proportionally bigger or thicker than the features they represent. To avoid hiding critical information in a fog of detail, the map must offer a selective, incomplete view of reality.’ The process becomes unethical when there is a deliberate intent to mislead. When features are generated automatically in a GIS, the potential for inadvertent misleading is high. For example, map generalization (Buttenfield and McMaster, 1991) is a common GIS function, but key features of interest may disappear when map scale changes. Monmonier (1991) discusses the role of this process in attempts to hide adverse results in environmental impact assessments. As with statistical measures, one should not provide a single value (or map) as completely representative of the current situation. A variety of maps at different scales should be provided to obtain a more complete picture of reality.

4.6. *Information filtering*

Graphs and maps are forms of images. When an image is presented in conjunction with accompanying audio, as in video presentations, the audio content can modify the perception of the image significantly (Thomson and Sivertson, 1994). Statistics, images, graphs, and maps are all methods of summarizing or filtering information. Push technology, described above, is another approach to filtering knowledge. Ethical decisions behind the selection and transformation of the material to ‘push’ can significantly affect the accuracy with which the recipient may perceive a situation.

²See <http://www.pfc.forestry.ca/main/programs/fnfp/tek/index.html>.

The ability to influence perception by selectively filtering information as it is passed from one level to another within an organization is often the source of erroneous decisions (Bella, 1992). Pauline Comeau, associate editor of *Canadian Geographic*, describes the case of the Atlantic cod fisheries in Canada, in which reports were allegedly delayed internally within the department responsible for fisheries management'. The result was that overfishing was allowed to continue to the extent that the fishery eventually collapsed.

Management information systems (MIS) can provide direct access by upper management to lower level data and information summaries. This helps bypass intervening distortions, resulting in more accurate perceptions. Greater accuracy is dependent, however, on an MIS that has itself been developed with appropriate ethical considerations. Furthermore, to achieve distortion-free information sharing, higher level managers must be willing and able to use the MIS.

5. Property

Property includes (1) the knowledge possessed by individuals and organizations and used in the software development process and (2) the communication medium that delivers the software to users (bandwidth). Participants in the software development process contribute knowledge and skills, which can lead to problems of intellectual property rights. The recent growth and development of Internet technologies means that accessibility is also an entity that can be controlled and owned. For both types of property, disagreements can occur among individuals and also between individuals and organizations.

5.1. Authorship

There are many ways to acknowledge intellectual property rights (Posey and Dutfield, 1996). However, there is a general perception among computer professionals that 'the knowledge base in an expert system is owned by the organization that developed it' (Belohlav et al., 1997). This point of view is driven by a traditional interpretation of compensation, wherein an organization's property rights are based on a fee-for-service agreement between that organization and an individual. When the individual is employed by the organization, ownership and responsibility issues are fairly straightforward. This may not be the case when the individual is outside the organization. Intellectual property rights in expert system development would appear to be analogous to copyright ownership in traditional print media, where an author owns the ideas (knowledge), and a publisher can own its printed form (or 'expert system'). Future revisions then rely on the publisher (system developer) and author (knowledge source) working together.

³<http://www.cangeo.ca/cod.html>.

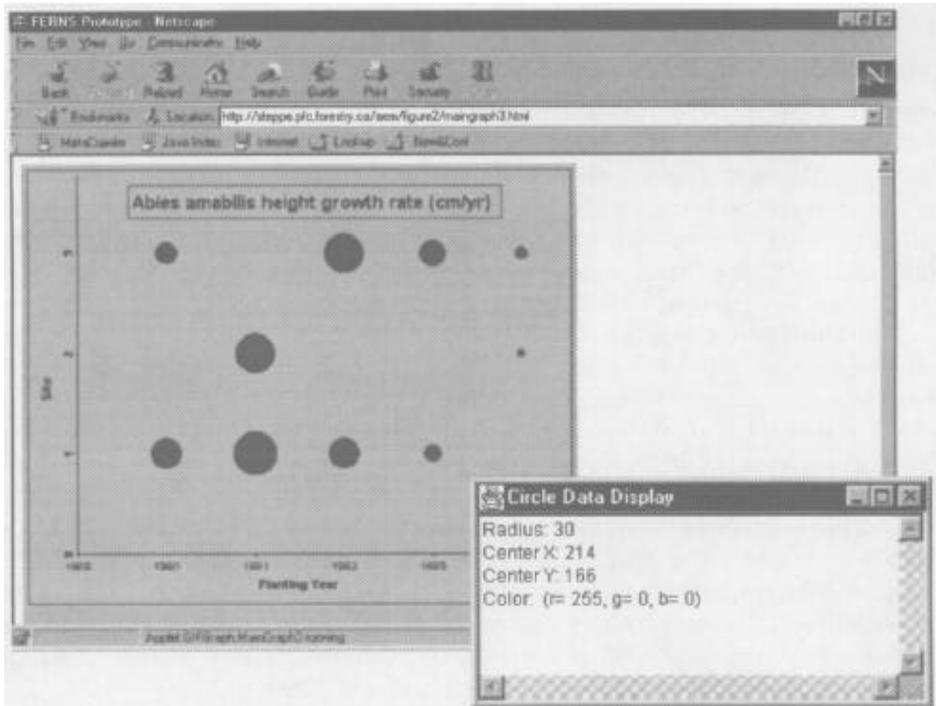


Fig. 2. An applet for collaborative construction of graphical relationships. The individual points are objects that can be manipulated, and have associated metadata that can be retrieved by clicking on them (from Thomson and Mitchell, 1999). The metadata can include ownership of the data point, source organization or other information as required.

In a computer system, metadata may be used to acknowledge contributions of information and knowledge. When knowledge elicitation software is used, inclusion of knowledge authorship in the metadata can be automated (Thomson, 2000b; Thomson and Mitchell, 1999) (Fig. 2). While it is possible to automate acknowledgment of the information provider in this way, it may not be desirable if the information contributor wishes to remain anonymous. A contributor of information may wish to remain anonymous (i.e. not have their intellectual property rights acknowledged) if adverse repercussions are anticipated from providing that knowledge. For example, Bella (1992) suggests that sources of information unfavorable to an organization tend to be isolated or reorganized out of existence. As system developers may need to know authorship in order to revisit the author at some later time, some form of security for metadata access may be necessary, analogous to the GIS example (protection of cultural sites) discussed earlier.

5.2. *Bandwidth*

An aspect of property that is often overlooked is the issue of bandwidth ownership (Mason, 1986). This issue is of special interest to Internet-based systems. Because the general bandwidth of the Internet provides a conduit for information, it may be regarded as behaving as a ‘commons’ (Mason, 1986). Problems emerging in a ‘commons’ environment always pose great difficulty owing to the large number of participants affected. When specific individuals or target systems can be identified, however, the ethical issues become clearer. In those cases, dominance by a single user, for example, can be readily identified and corrected.

Bandwidth limitation ethics apply not only to recipients of push technology as discussed above, but also to data providers. For example, assume organization A has an on-line gazetteer listing places within a given distance of a specified location. Organization B has an on-line database listing forest sector unemployment statistics by location. It is possible to write a program that will interrogate A’s database to retrieve a list of locations, then iterate through that list to interrogate B’s database to obtain regional forest sector unemployment statistics. B’s database may have been set up under the premise that it would be used for single queries entered manually. By programming queries into a loop, a user could monopolize or swamp B’s bandwidth. Because Internet technology is still in its infancy, readily detecting such bandwidth violations would currently require extra effort by software developers in organization B.

6. **Accessibility**

6.1. *Physical access*

Appropriate access to software systems has both technical and intellectual components. To use a system, a person must have access to the required hardware and software technology, must be able to provide any required input, and must be able to comprehend the information presented. For example, for a Web-based system, the user must have a reliable connection to the Internet, as well as the appropriate connection speed (especially for graphical content). The end-user must also have a browser compatible with the material sent to it (including such things as the appropriate Java classes for use with applets) and any helper applications or browser plug-ins for viewing and hearing content. If the audience is in a developing country, or in a remote area (such as the Canadian North), such technological issues may be critical. For this reason, when a system is developed, its implementation should be part of an integrated process that includes alternative knowledge elicitation and knowledge delivery methods that are appropriate to the full range of affected individuals (Thomson, 2000b). This may include specifying duties for a range of ‘actors’ such as technology transfer officers or field personnel.

6.2. *Language access*

The role of language and culture has already been discussed in relation to accurate system use, where concepts unfamiliar to end-users can be a problem. The language selected for user interaction, such as French or English, is another issue related to accessibility of systems. This is especially so in countries such as Canada where there is more than one official language. All federal government systems in Canada must be delivered in both English and French, which requires bilingual reasoning for expert systems (Thomson et al., 1998) and bilingual data structures in databases (Thomson, 2000b). Our tests of bilingual expert systems were performed only in the language in which the source material was provided, as all internal reasoning was performed using generic symbols that could be translated into the appropriate language for display. We relied here on common concepts and meanings in both languages within the limited application domain. Where concepts and meanings vary with culture and language, there would be significant extra effort required for system testing and validation.

6.3. *Skill level*

Accessibility to a system is also limited if results are presented using language and concepts beyond the end-user's understanding. Expert systems, for example, can be specifically designed to present results using appropriate concepts and language for the intended user, including appropriate filtering of knowledge to prevent information overload (Thomson and Taylor, 1990). It is also important to provide sufficient information for domain experts to validate a system, even though typical software output may be geared toward a particular target audience. This is a separate issue from developing a system for multiple target audiences, where it may be more appropriate to develop separate systems geared to the requirements of each target group rather than try to develop a single, generic system for all users.

6.4. *Decision -making environments*

A potential drawback exists when providing access to sophisticated software. Such technology may increase considerably the power of users to make or influence decisions that were formerly beyond the limits of their knowledge and experience (Belohlav et al., 1997). Decision tools should not be used blindly; thus, system developers must fully consider the interactions of all people using and/or affected by a system and ensure appropriate training to avoid misuse (Collins et al., 1994). This situation has been a bane of statisticians for years. Very powerful software packages allow users to perform all manner of inappropriate statistical tests on data without full knowledge of what they are doing. While current statistical software manuals contain a great deal of information regarding model specification and assumptions, they cannot replace a well-founded understanding of basic statistics by the experimenter.

Increasing users' decision-making power beyond their former knowledge and experience can also have positive impacts. In particular, when systems are largely based on existing publications or manuals, ethical issues typically involve knowledge delivery questions alone. Expert systems (Thomson et al., 1993, 1998) illustrate this effect. For example, the nursery diagnostic system described by Thomson et al. (1993)⁴ was developed at a time when forest nurseries had been privatized in British Columbia. Because any record of a disease problem could affect a nursery's competitive business position, use of diagnostic services declined. The use of in-house diagnoses by nonprofessionals increased, often with erroneous results. The diagnostic expert system was deployed in those nurseries and improved the accuracy of in-house diagnoses, while preserving confidentiality in a competitive situation.

7. Quality of life

Computer systems are generally intended to improve the quality of life — initially in the workplace, but now in peoples' personal lives, also. However, systems may actually degrade the quality of working life through (1) deskilling the workforce, which reduces control, responsibility and job satisfaction; (2) increasing stress, depersonalization, fatigue and boredom; and (3) health and safety concerns such as eyestrain, headaches and repetitive strain injury (Forester and Morrison, 1994). While workplace productivity typically increases with upgraded software systems, there is often a time lag before organizations are able to fully assimilate new technology with their existing work culture. It is during this interim period when quality of life impacts are often realized because old software systems can be quickly scrapped for new ones, but the human element takes longer to adjust.

Personal computers have caused many changes in work processes. A major change has arisen from widespread availability of statistical software packages, which put sophisticated analytical power in the hands of individuals with little statistical knowledge (as noted above). Also, the widespread use of word processing systems has resulted in the authors of documents being responsible for producing final copies themselves, rather than having them typed by secretarial staff. These types of changes afford workers greater control over their work and work products, but also require greater skill. If workers are not retrained or cannot otherwise upgrade their skill dossier, productivity, performance, and worker confidence/satisfaction will actually decline.

In evaluating the effects of a system, Collins et al. (1994) advocate consideration of the 'software penumbra', i.e. the people, other than providers, buyers or users, who can be affected by the software. In particular, a system should not increase the harm to the least advantaged, or risk increasing harm in already risky environments. This is certainly an ambitious ethics test for software. Given the financial,

⁴Now available at <http://www.pfc.cfs.nrcan.gc.ca/nursery/>.

intellectual, economic, political, and social reach that software can have, this idea has operational limitations, such as how far, or in what ways, to extend the penumbra to include affected parties.

8. Discussion

Computer system involvement in decision making leads to new versions of old moral issues — right and wrong, honesty, reliability, loyalty, responsibility, confidentiality, trust, accountability and fairness (Forester and Morrison, 1994). Without an ethical approach, systems may be put into operational use in spite of faulty trials and persistent errors, and without appropriate consideration of the people affected. ‘Given the incidence of faulty software, and of system failure in general, it is not surprising that software developers rarely provide their clients or purchasers with warranties of any substance’ (Forester and Morrison, 1994). Because guidelines and standards to direct software developers in ethical conduct are nonexistent, it is up to the end-user in most cases to provide critical feedback. The availability of software Web sites and e-mail contact addresses provides avenues for end-user responses to unethical practices.

Privacy, accuracy, property, accessibility, and effects on quality of life, are all issues that must be considered in developing and delivering computer software systems. Choosing a particular approach to system development can either hinder or facilitate addressing these issues in an ethical manner. Past approaches to system development were oriented towards technical aspects of system development, whereas current approaches emphasize the need to consider all aspects of the human environment in which the systems are being developed and used. Because recent approaches focus strongly on ethics, they are less adept at technical design issues. A hybrid methodology, that combines both ethical and technical needs, might adequately address both concerns simultaneously.

Privacy has long been considered an inherent right of individuals in a ‘free’ society. Initially, this involved protection of the individual from unwanted or unwarranted invasion of their physical space. More recently, privacy has been extended into an individual’s information space, as well. For software systems currently under development in natural resource and agricultural domains, more real threats are likely to arise from unintentional and unforeseen information breaches than from any intentional conspiracy. These occur when information sources are combined or used in unintended ways or when overzealous knowledge-solicitors inundate individuals with information requests that they may neither desire nor be able to eliminate from their lives. As long as information about individuals exists and is accessible by others, individual privacy can potentially be compromised. During the design of software, developers need to be cognizant of users, co-developers, publics, cultures, special interest groups, commercial enterprises, governments, and other groups that might be affected directly or indirectly by their software. Designers must also consider the information their software uses or generates, and the decision-making landscape that it affects or creates.

Use of appropriate language is at the heart of many accuracy issues, both in the knowledge elicitation phase of system development and in the presentation of system outputs. Even if a system does not estimate accuracy of results explicitly, it is important to make end-users aware of the variability in potential outcomes, and the assumptions and trade-offs that have contributed to it. Graphical knowledge-elicitation tools and graphical rendering of system outputs should also be designed to address accuracy concerns in the flow of knowledge through the system. It is also essential to address the way in which knowledge flows through organizational hierarchies, and to ensure its appropriate use at different organizational levels.

Intellectual property issues surface in many situations. Although some issues can be addressed through ethical software development, many exogenous issues exist and are driven by rapidly changing legislative, legal, and policy edicts. For example, in a recent news item in *Science* titled 'Scientists Decry Antipiracy Bill', Kaiser (1999) describes a bill under consideration by the US Congress that could 'severely hinder how everyone from molecular biologists to environmental scientists use electronic databases'. The troubling aspect of the bill was that companies, which repackage data freely available from the government, could claim ownership of the raw information. Given the wealth of information now available electronically to just about anyone, such ownership debates will become more commonplace.

As with accuracy issues, language is at the heart of many accessibility issues. Knowledge delivery must be geared to concepts appropriate to the intended audience, and information overload avoided, as knowledge can be inaccessible if the recipient is swamped with information. Limitations of technical accessibility by some groups may require developing an integrated range of systems and processes to ensure access by all stakeholders in a decision environment.

Organizations such as the Association for Computing Machinery (Anderson et al., 1993), the Institute for Electrical and Electronics Engineers, the British Computer Society, and the International Federation for Information Processing have developed codes of ethics for computing professionals. However, these codes contain few sanctions, are never used, and the language is never interpreted (Forester and Morrison, 1994). The codes also hold individuals at fault, not whole organizations, and do not make people more ethical. The organization, Computer Professionals for Social Responsibility⁵, does not advocate a formal code of ethics, but rather addresses a wide range of ethical issues.

While there will always be some ethical culpability on the individual's part, much responsibility still rests with organizations to institute standards of ethical conduct that create an atmosphere of social morality for their employees and members. Especially in the case of commercial enterprises, moral responsibility should also lead to legal responsibility. The lack of sanctions in these codes suggests that there are no incentives to consider ethical issues in software system design and development. However, real incentives may lie in the form of increased success in system delivery. In the commercial arena, this will translate into buyer preference for

⁵<http://www.cpsr.org/>.

ethical software (greater sales and increased popularity), and for noncommercial software, into broader audience use and satisfaction. The failure rate of software development projects has been estimated to run as high as 70% (Slofstra, 1999). Failure in information systems often results from an inability or unwillingness to understand the human context (Wood-Harper et al., 1996). This implies that greater attention to human perceptions of ‘good’ and ‘bad’ (ethics) will result in greater software success. Through an ethical approach to system design and development, addressing the issues discussed above, a higher likelihood of software success can be expected.

Acknowledgements

The authors would like to thank Steve Glover of the Pacific Forestry Center for many helpful comments and suggestions.

References

- Anderson, R.E., Johnson, D.G., Gotterbarn, D., Perolle, J., 1993. Using the new ACM code of ethics in decision making. *Commun. ACM* 36 (2), 98–107.
- Aragon, L., 1997. When push comes to shove. *PC Week* February 10 On Line: http://www.pcweek.com/business/02_10/10push.html.
- Bella, D.A., 1992. Ethics and the credibility of applied science. In: Reeves, G.H., Bottom, D.L., Brookes, M.H. (Technical coordinators), *Ethical Questions for Resource Managers*. USDA Forest Service General Technical Report PNW-GTR-288, pp. 19–32.
- Belohlav, J., Drehmer, D., Raho, L., 1997. Ethical issues of expert systems. *Online Journal of Ethics* 1(1). <http://www.depaul.edu/ethics/expert.html>.
- Bernadas, C.N., 1991. Lesson in upland farmer participation: the case of enriched fallow technology in Jaro, Leyte, Phillipines. *For. Trees People Newslett.* 14, 10–11.
- Buttenfield, B.P., McMaster, R.B., 1991. *Map Generalization: Making Rules for Knowledge Representation*. Longman, Essex, UK.
- Chambers, R., 1992. *Rural Appraisal: Rapid, Relaxed and Participatory*. Discussion Paper 311. Institute of Development Studies, University of Sussex, UK.
- Collins, W.R., Miller, K.W., Spielman, B.J., Wherry, P., 1994. How good is good enough? An ethical analysis of software construction and use. *Commun. ACM* 37 (1) 81–91.
- Danielson, P., 1992. *Artificial Morality: Virtuous Robots for Virtual Games*. Routledge, London.
- Davenport, T.H., Prusak, L., 1997. *Information Ecology: Mastering the Information and Knowledge Environment*. Oxford University Press, Oxford.
- Forester, T., Morrison, P., 1994. *Computer Ethics*. MIT Press, Cambridge.
- Franklin, S., Graesser, A., 1996. Is it an agent, or just a program? A taxonomy for autonomous agents. In: *Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages*. Springer, Berlin.
- Goodenough, D.G., Charlebois, D., Matwin, S., MacDonald, D., Thomson, A.J., 1994. Queries and Their Application to Reasoning with Remote Sensing and GIS. *IEEE Press*, New York, pp. 1199–1203.
- Hart, M., 1999. *Guide to Sustainable Community Indicators*. Hart Environmental Data. North Andover, MA.
- Helling, C.S., 1978. *Adaptive Environmental Assessment and Monitoring*. Wiley, Chichester.
- Huff, D., 1954. *How To Lie With Statistics*. W.W. Norton and Company, New York.

- Kaiser, J., 1999. Scientists decry antipiracy bill. *Science* 286, 1658.
- Kendall, K.E., Kendall, J.E., 1999. *Systems Analysis and Design*. Prentice-Hall, New Jersey
- Mason, R.O., 1986. Four ethical issues of the information age. *MIS Q.* 10 (1), 5–12.
- Monmonier, M., 1991. *How to Lie with Maps*. University of Chicago Press, Chicago.
- Mowrer, T., 2000. Uncertainty in natural resource decision support systems: sources, interpretation, and importance. *Comput. Electron. Agric.*, 27, 139–154.
- Posey, D.A., Duffield, G., 1996. *Beyond Intellectual Property: Toward Traditional Resource Rights for Indigenous Peoples and Local Communities*. International Development Research Centre, Ottawa.
- Poulton, E.C., 1994. *Behavioral Decision Theory: A New Approach*. Cambridge University Press, Cambridge.
- Saveland, J.M., Stock, M., Cleaves, D.A., 1988. Decision-making bias in wilderness fire management: implications for expert system development. *AI Appl.* 2 (1) 17–30.
- Schmoldt, D.L., Rauscher, H.M., 1994. Knowledge management and six supporting technologies. *Comput. Electron. Agric.* 10 (1), 11–30.
- Simon, H.A., 1979. *Models of Thought*. Yale University Press, New Haven, CT.
- Slofstra, M., 1999. Do projects fail as often as that. *Infosyst. Execut.* 4 (6) 7–8.
- Thomson, A.J., 1997. Artificial Intelligence and environmental ethics. *AI Appl.* 11 (1), 69–73.
- Thomson, A.J., 2000a. Elicitation and representation of Traditional Ecological Knowledge, for use in forest management, *Comput. Electron. Agric.*, 27, 155–165.
- Thomson, A.J., 2000b. Knowledge elicitation tools for use in a virtual adaptive environmental management workshop. *Comput. Electron. Agric.*, 27, 57–70.
- Thomson, A.J., Mitchell, A., 1999. Collaborative knowledge management for long-term research sites. *For. Chron.* 75 (3) 491–496.
- Thomson, A.J., Sivertson, C., 1994. Video knowledge bases and computer-assisted video production: a forestry example. *AI Appl.* 8 (1), 29–39.
- Thomson, A.J., Sutherland, J.R., Carpenter, C., 1993. Computer-assisted diagnosis using expert system-guided hypermedia. *AI Applications* 7(1), 17–27 (with disk).
- Thomson, A.J., Taylor, C.M.A., 1990. An expert system for diagnosis and treatment of nutrient deficiencies of Sitka spruce in Great Britain. *AI Appl.* 4 (1), 44–52.
- Thomson, A.J., Allen, E., Morrison, D., 1998. Forest tree diagnosis over the World Wide Web. *Comput. Electron. Agric.* 21, 19–31.
- Veatch, R.M., 1977. *Case Studies in Medical Ethics*. Harvard University Press, Cambridge.
- Wood-Harper, A.T., Corder, S., Wood, J.R.G., Watson, H., 1996. How we profess: the ethical systems analyst. *Commun. ACM* 39 (3), 69–77.