

Neural Expert Systems in Medical Image Interpretation: Development, Use, and Ethical Issues

Athanasia Pouloudi and George D. Magoulas

Centre for Health Informatics & Computing, Department of Information Systems & Computing, Brunel University, Uxbridge UB8 3PH, UK[†]

ABSTRACT

In this paper, the appropriate level and role of neural network-based methodologies in the development and use of expert systems for medical image interpretation is investigated as technical, organizational, and social issues become intertwined. The notion of the information life cycle is applied to highlight ethical issues during the acquisition, processing and storage, dissemination and use of clinical information. These issues are further analyzed from a stakeholder perspective to accentuate the role of human agents in avoiding ethical risks. Relevant stakeholders, other than the key participants—namely system developers and medical users—are identified. The results of this analysis indicate that each stage of the development and use of a neural expert system entails ethical issues. Significantly, the responsibility for medical image interpretation is affected by contextual factors and should be shared amongst the main stakeholders. These conclusions are useful for the stakeholder groups that are conscious of their obligation to behave ethically and for researchers who wish to investigate further the ethical implications of artificial intelligence use in medicine.

[†]tel.: +44-1895-274000 ext. 3745, 2521; fax: +44-1895-251686
e-mail: {Nancy.Pouloudi, George.Magoulas}@brunel.ac.uk

KEY WORDS

ethics, stakeholders, medical image interpretation, computer-assisted medical diagnosis, neural expert systems, artificial intelligence

1. INTRODUCTION

The use of intelligent information systems in health care increasingly provides opportunities to both facilitate and enhance the work of medical experts and ultimately to improve the efficiency and quality of medical care (Aleksander & Morton 1991; Furundzic et al., 1998; Ifeakor & Rosen 1994; Lane et al., 1996). At the same time, the debate about the appropriate level and the role of information systems that assist in medical diagnosis has become more complex as technical, organizational and social issues become intertwined (Anderson, 1997; Baase, 1997; Kushniruk & Patel, 1998; Ridderikhoff & van Herk, 1999).

Advances in neural networks have opened the way for the establishment of systems that are able to learn complex associations by example (Haykin, 1994). The growing number of projects that employ neural networks in image interpretation expert systems makes it necessary to examine neural networks' development and use in this context. It is acknowledged that the appropriate use of the neural expert system methodologies in medical image interpretation could be very effective to provide rapid identification and diagnosis in real time (Ifeakor & Rosen, 1994; Hanka et al., 1996; Innocent et al., 1997; Karkanis et al., 2000; Phee et al., 1998; Zhu & Yan, 1997).

Research in neural expert systems to date remains centered on technological issues and most is application driven. Previous research and experience suggests, however, that the successful implementation of information systems (for example, Anderson, 1997; Pouloudi, 1999)—and decision support systems in particular (for example, Ridderikhoff & van Herk, 1999)—in the area of healthcare relies on the successful integration of the technology with the organizational and social context within which it is applied. In this paper, we argue that the successful implementation of neural expert systems for medical image interpretation also relies not only on their technical feasibility and

effectiveness but also on organizational and social aspects. The aim of the paper is, then, to study the ethical issues that may arise from their application in this area as clinical information is acquired, processed, used, and exchanged among professionals. Ethical issues are critical in healthcare applications because such issues ultimately reflect the quality of care that is provided.

Following an introduction to the use of neural networks in medical image interpretation in the next section, the paper focuses on the broader ethical implications of the use of neural expert systems in this area, including issues of interpretation, coordination between the technology and the human expert, validation of results, and professional responsibilities. These issues are presented for different stages of the development and use of a neural network-based expert system. The paper concludes with a more general discussion on the responsibilities of the human experts for the ethical use of intelligent systems.

2. NEURAL EXPERT SYSTEMS IN MEDICAL IMAGE INTERPRETATION

A neural network is a massively parallel distributed processor that has a natural propensity for storing experiential knowledge and making it available for use (Haykin, 1994; 2). Neural networks are modeled on the general features of biological networks and are created through hundreds or thousands of artificial neurons that are connected in a manner similar to the brain's neurons (Rumelhart & McClelland, 1986). The function of an artificial neuron is to multiply each input by a respective weight, then to sum all its inputs, and finally to determine its output according to this summation. The weights for all inputs are adjusted during the training process until the network exhibits the desired behavior. These final weights represent the knowledge in the network. The training process constitutes the main difference between neural and conventional computing, as the former uses examples to produce its own internal structure or programming in terms of neurons and interneuron connection weights, whereas the latter requires explicit instructions in the form of a computer program.

Neural network-based systems exhibit the following features that make them useful in practical applications, such as medical image processing.

- They resemble the brain in that knowledge is acquired through a learning process and is stored in the interneuron connection strengths known as synaptic weights.
- They can learn from examples about the relations between input and output data, even when these relations are not explicitly prescribed. Thus, such systems are useful in situations in which associations between inputs and outputs are required or rules are unknown and difficult to explain or to formulate.
- They do not need any assumptions that are established on the statistical distributions underlying the input data and can adapt to changes in the environment through learning to improve their performance.
- They have good generalization ability—for example they can deal with previously unencountered situations
- They exhibit robustness for imprecision and uncertainty in comparison with traditional artificial intelligence methodologies and fault tolerance because of their parallel distributed processing.

Neural networks have been increasingly used in medicine and especially in the development of neural expert systems for intelligent medical image interpretation (Ifeachor & Rosen, 1994; Hanka et al., 1996; Innocent et al., 1997; Karkanis et al., 2000; Phee et al., 1998; Zhu & Yan, 1997). In most cases, the development of such systems is considered an attempt to emulate the doctor's expertise in the identification of malignant regions in minimally invasive imaging procedures (for example, computed tomography, ultrasonography, endoscopy, confocal microscopy, computed radiography, or magnetic resonance imaging). The objective is to increase the expert's ability to identify cancer regions while decreasing the need for intervention and maintaining the ability for accurate diagnosis. Furthermore, it may be possible to examine a larger area, studying living tissue in vivo—possibly at a distance (Delaney et al., 1998)—and thus minimize the shortcomings of biopsies, such as a limited number of tissue samples, a delay in diagnosis, and discomfort for the patient. The need for more effective methods of early detection—such as those that computer assisted medical diagnosis systems aim to provide—is obvious.

In technical terms, the problem in automatic image interpretation is to associate sets of pixels (structures) in an image with the unknown objects that are present in the scene from which the image has been drawn. The difficulty

increases when several objects of different kinds, related by a set of spatial-temporal relations, are present in the observed scene. In medical practice, endoscopic approaches and other minimally invasive techniques (for example, computed tomography and magnetic resonance imaging) are now permitting visualization of previously inaccessible regions of the body. In diagnostic endoscopy, the medical expert, based on a distributed percept of local changes, interprets the physical surface properties of the tissue—such as the roughness or the smoothness, the regularity, and the shape—to detect abnormalities. Adjacent surfaces showing different surface properties are distinguished on the basis of the texture differences. It is important to note, however, the vast difficulties in physical attributes of the organs. For example, in colonoscopy, no two colons are alike. Even within the same colon, one section may have very different characteristics from another. This fact introduces severe limitations in the use of computer-assisted endoscopy for interpreting colonoscopic images (Kwoh, 1995). Given a medical image, the ‘true’ features associated with the physical surface properties of the tissue are not exactly known to the image-interpretation system developer. Usually, one or more feature-extraction models (Looney, 1997) are used to provide values for each feature’s parameters. The findings are then used to infer the correct interpretation. On this same task of interpretation on the basis of local changes on the properties of the tissue under examination, the performance of human perception is considered outstanding. Furthermore, medical experts have the ability to either add or remove components from an image to give meaning to what they see. Medical experts can also adapt to changes to the extent that even a distorted image can be recognized.

Neural network methodologies present some human-like qualities, such as learning from experience, generalization, and handling uncertainty and ambiguity in distorted or noisy images. Thus, such methods provide human experts with significant assistance in medical diagnosis (Innocent et al., 1997; Karkanis et al., 2000; Phee et al., 1998; Zhu & Yan, 1997). From an ethical perspective, however, the interaction of the medical expert with the medical image interpretation neural expert system presents a number of challenges that may lead to ethical dilemmas for both the clinician and the neural network-based system developer. These are presented in detail in the next section.

3. EMERGENT ISSUES IN THE DEVELOPMENT AND USE OF MEDICAL IMAGE INTERPRETATION SYSTEMS

The development and use of neural expert system-based medical image interpretation applications, similar to other information technologies, can be considered to follow a series of stages that correspond to an *information life cycle* presented in Mason et al. (1995) as a sequence of functions through which information is handled: *acquisition, processing, storage, dissemination, and use*. This life cycle is important to our discussion of ethical issues because each stage may lead stakeholders to “an ethical crossroads” (ibid., p. 7). The following discussion highlights the challenges that are presented to both developers and users of medical image interpretation systems as clinical information is processed or exchanged at each stage and explains how such challenges can result in important ethical issues.

3.1 Acquiring Information

Knowledge acquisition is one of the key problem areas in artificial intelligence applications because of the difficulties in expert selection and the elicitation of the expert’s knowledge. These difficulties are as follows:

- difficulty in finding an expert who is willing or available to participate in the knowledge acquisition process;
- difficulty in finding the relevant parts for a particular problem solving activity from the wealth of expert knowledge available; and
- problems of expert bias, and so on (Bolger, 1995; Gaines & Boose, 1988; Firlej & Hellens, 1991; Hart, 1986; Whitley et al., 1992).

Neural expert systems have relative advantages in relation to other intelligent systems at this stage of the information life cycle because they can handle incomplete information and provide reasoning under uncertainty (approximate reasoning) (Kasabov, 1996). Also, they do not require the formulation of expertise in a rule form as in rule-based systems (Dreyfus, 1987). Nevertheless, neural expert systems still present several challenges for knowledge acquisition, as outlined in the following paragraphs. Challenges have an impact on the effectiveness and reliability of neural expert systems; thus,

challenges signify ethical implications for any decisions that are made at this stage using such systems.

The selection of experts remains problematic in neural network applications. This problem relates to the difficulty of finding available experts that are recognized in the field as knowledgeable. Often intelligent systems' developers assume human expertise to be 'the ideal' and do not question the suggestions of the experts and, similarly, clinicians also seem to have high confidence in their decisions (Carroll, 1994; Ridderikhoff & van Herk, 1999). In practice, however, there may be alternative views on expertise—a result of expertise being a socially constructed and historically situated variable (Agnew et al., 1995; Stein, 1992; Whitley, 1995). Also, experts often behave as 'stakeholders'—that is as parties whose behavior may be shaped by their interests and political roles in the knowledge acquisition process (Pouloudi, 1997). The choice of influential experts may therefore be difficult but is likely to have a positive influence in the final acceptance of the system as it increases the confidence of other users in it.

At the same time, elicitation and articulation of the expert's knowledge creates further challenges for the developer as well as for the expert. First, it is important that the developer understands the application context (for example, in the case of medical imaging, the features of an image that are relevant for diagnosis). In neural expert systems imaging applications, the main challenge of knowledge acquisition is the selection of appropriate examples. Examples are critical for accurate and suitable representation of knowledge, particularly because the expert is called to select representative images or regions in images, based on previous experience. The main danger here lies in missing important information about conditions that may have to be recorded. Another serious problem is the inability of the clinician to express competence in a metric or language that is meaningful to the researcher (Bolger, 1995), namely, in a form that can be represented and processed by a neural network (Kasabov, 1996). Overconfidence in the decision parameters as defined by previous cases may be misleading if such parameters prevent retraining the network when necessary. The participants in the process and future users will have to be aware of these limitations. The problem is that the impact of selecting the appropriate decision parameters (or their absence) can rarely be anticipated *a priori*.

Further problems may be encountered in cases for which no previous experience in the interpretation of medical images has been accrued. Such limitations, inherent in the historical context, may be due to a lack of previously interpreted images—as in the case of the use of new imaging technologies—or in cases, especially for cancer diagnosis, in which no histological confirmation is available for the images of, possibly, malignant areas used for training the network (for example Hanka, 1996). Such problems ultimately reflect on the effective use of the medical image interpretation system and hence signify ethical dilemmas about its adoption for its users.

3.2 Processing and Storing Information

The development and training of the neural expert system involves

- the selection of appropriate network architectures, training methods, and features;
- the development of appropriate representations for the network input and output information; and
- the preparation of the training, testing, and validation data.

The latter rely on the human experts (both the developer and the user), who are responsible for the quality of the data, to avoid data inconsistencies and to cover all the possible cases.

Algorithms have been developed to increase the speed of the training process by computing the best direction in which to proceed to reach an error minimum in the weight space—that is, increasing the effectiveness of the neural system (Magoulas et al., 1997; 1999). The nature of the data may indicate the type of the training that is likely to be most appropriate for the given problem. When the outputs associated with given inputs are well known, supervised learning may be used. In this case, a set of inputs with their corresponding correct outputs, which are known from past experience, is required for teaching the network how to respond. As mentioned in the previous stage, however, the historical context may also be limiting. Uncertain classifications of inputs, in contrast, usually require unsupervised learning to discover correlations in the data. Clearly, a domain expert must examine and interpret the categories—and perhaps recommend training the network again—to obtain a different, more suitable, number of categories.

Similarly, the resulting outputs will have to be inspected by experts to provide interpretations of their meanings. The importance of selecting appropriate experts becomes evident once more, as the expert will influence the reliability of the recommendations for retraining the neural system and consequently improving its effectiveness.

The key problematic issue in the use of neural system methods in a medical image interpretation system, however, is that how the results are obtained is unclear. Most neural systems suffer from the opaqueness of their learned associations. In medical applications, this black box nature (Kasabov, 1996) may make clinicians reluctant to use a neural expert system, no matter how great the claims made for its performance, creating ethical dilemmas for users during this stage too. Thus, there is a need to enhance neural network-based systems with explanation capabilities. At the same time, there are also certain potential disadvantages of automatic explanation generation. That a medical expert be able to check through the generated explanations is essential, as these should not vary in unpredictable ways and should not be seen to take up the role of the clinician—for example by adding information to the patient's record—so that the human expert maintains control over the system's use.

3.3 Disseminating Information

In disseminating information, it is important that developers are aware of the limitations of the system *as well as* of their importance in the context of use. Thus in image interpretation, the overall accuracy of the neural expert system is not a sufficient measure of its performance (Karkanis et al., 2000). For example, when discriminating amongst normal and cancer regions in endoscopic images, misinterpreting a cancer region as normal (false negative) is more critical than misinterpreting a normal region as cancerous (false positive). Developers should be able to understand such limitations and to convey them accurately to the users.

The importance of user involvement becomes evident once more, primarily because it facilitates the communication between the 'experts of development' and the 'experts of the context' and can help deal with resistance to change phenomena (see next section). Developers often fulfill their obligation to take

into account the experts' perspectives not only in acquiring knowledge but also by inviting them to participate in pilot implementations. These, however, may be less effective if participants have different values and norms (Wagner, 1994), different assumptions about the purposes and intended consequences of the systems (Yuthas & Dillard, 1999), or if different user groups present different agendas once the system is in place (Pouloudi, 1999).

Given the importance of medical image interpretation applications, the developer has an obligation to make the results of innovative research known. Nevertheless, on the one hand, ensuring that the information reaches all relevant stakeholders is difficult. On the other hand, neural network experts often rely on the use of heuristic techniques that identify their work, are part of their expertise, and which, consequently, the experts are unwilling to make public.

3.4 Using Information

The ethical implications stemming from the challenges that were described in the previous stages become evident at the stage at which the output—namely, the result of the medical image interpretation system—is available. Most problems at this stage relate to human-system interaction, but are rooted in the confidence of the users and developers in dealing with the problems that were experienced during the previous stages and presented in the previous sections.

An important issue in the implementation and hence the use of the information capabilities of the systems is resistance to change—one of the most widely referenced 'problems' in the implementation of computer applications since Markus' (1983) seminal paper. The use of information technologies depends on political, economic, and social factors (Bloomfield, 1997; Mumford, 1987; Walsham, 1993). Resistance is important in this respect because it is the defense mechanism of expert users against systems that may be unreliable. In such cases, users have a moral obligation to identify unreliable systems as such and to refrain from relying on them for their judgment. In fact, many clinical experts distrust or disregard the output of intelligent technologies that seem to challenge human thinking and the human

need to be in control of the technology (Beerel, 1989; Ridderikhoff & van Herk, 1999).

At the same time, particularly in medical applications, resistance phenomena can create ethical problems if the use of intelligent technology is likely to improve health care provision. For example, doctors are still reported to be hesitant in using intelligent systems (Essex, 1994). There are several reasons: doctors' professional pride, overconfidence in their own clinical judgment (Ridderikhoff & van Herk, 1999), a lack of familiarity with the system (cf. Kushniruk & Patel, 1998), and a high sense of responsibility and accountability in case of error.

Indeed, human experts are expected to "compensate for the deficiencies of the artifacts in such a way that the social group continues to function as before" (Collins, 1990) and, therefore, should not become over reliant on technology (Dreyfus & Dreyfus, 1986). There are also organizational reasons for the doctors' resistance to technology, such as the interference of the systems with the way that physicians organize their work, interact with colleagues, and approach patient care (Anderson, 1997).

Finally, with the system in use, it is the responsibility of the user to evaluate the interpretations of the system. Users (medical experts) must be aware of any limitations to reaching appropriate decisions. Thus, users in this respect have a right *and* a responsibility to collaborate with developers in evaluating and understanding the limitations of the system. Furthermore, users should be responsible for conveying any limitations in diagnosis accurately to their 'customers'—in this case the patients. Depending on the social and cultural context, doctors either may or may not have an obligation to inform the patient, directly or indirectly, of the diagnosis and the potential error rate using this system.

4. ETHICAL ISSUES AS STAKEHOLDER ISSUES

Medical information is vital for the diagnosis and treatment of patients, and therefore the ethical issues presented during its life cycle are critical. The previous section illustrated some of the issues that developers and medical users face during the development and use of neural expert systems for

medical image interpretation. As such technologies become pervasive, understanding these issues becomes imperative (Herman et al., 1998). Some of these issues are system-centered—namely, related to the inherent problems of neural expert systems technology. It is humans, not systems, however, who can act as moral agents (Friedman & Kahn, 1997), meaning that only humans can identify and deal with ethical issues. The aim of this section is, therefore, to revisit the challenges and the ethical issues, presented in the previous sections, from a human-centered perspective by considering the motivations and the ethical dilemmas of stakeholders in medical image interpretation.

The primary motivation for the use of neural expert systems in medical image interpretation is to contribute to an improvement of the techniques for medical diagnosis on the basis of their advantages in relation to more traditional artificial intelligence techniques. Such an improvement is driven by the desire for more effective (desirable by both patients and medical professionals) and more cost-effective health care provision (desirable by health care organizations). Nevertheless, the interaction of the users with the technology entails ethical risks if primarily based on curiosity and fascination with the capabilities of the intelligent technology (Summers, 1997), to the extent that the role of human expertise is underestimated or an unreliable system is adopted. Similarly, unjustified trust and over reliance on the system can be problematic because of a risk to make inappropriate decisions. Medical professionals have, therefore, a moral obligation to challenge the technology and to use their expert judgment in making clinical decisions.

As discussed in the previous section, clinicians *are*, in practice, skeptical about medical diagnosis systems and have difficulty integrating them into their work practice. This reluctance, however, may also be a result of the characteristics of the system. For example, the speed of response of a medical image interpretation system, as this functions within a consultation context—for example, within a hospital information system environment—is critical for its adoption by clinicians who expect quick response rates. Also, the system may not be user friendly or cannot readily be integrated with existing work practices or with other systems. For example, research evidence concerning the organizational and contextual issues related to the integration of new technologies into hospital information systems has been published (Anderson, 1997; Kurihara et al., 1999; Pouloudi, 1999). The successful integration of

new technologies with existing systems in a particular hospital context relies on the successful collaboration of developers, users, and the organizations where they work.

Developers should be particularly sensitive to the context in which the technology will be used, the culture, and the work practices. In medicine, it is particularly important that developers take into account the need of clinicians to work with patients and to communicate decisions to them. Patients have, in turn, their own views of technology and expectations from a medical professional. Both patients and clinicians may feel uncomfortable if technology dominates the diagnosis procedure and seems to take over the professional's role or seems to interfere with the consultation.

To be able to evaluate and to understand the characteristics of the context, developers should work closely with clinicians. The importance of the participation of users in systems development is well documented (for example, Avgerou & Cornford, 1993), but in medical image interpretation the participation, involvement, and close collaboration of developers and users is critical for the effectiveness, acceptance, and reliability of the system. Both parties have, therefore, an ethical obligation to collaborate at all the stages of the development and use of the system, as explained in the previous section, communicating not only their expert knowledge but also any limitations that are relevant to the system. Professionals are usually uncomfortable about communicating any limitations of their expert knowledge to others, yet they have a moral obligation to do so. Indeed, it is already evident from our discussion so far that dealing with ethical issues relies on the professionalism of the experts, both the neural expert system developers (who have technical expertise) and medical users (who have medical expertise). Both stakeholders groups have an obligation to behave as professionals for achieving the most reliable diagnosis result possible.

This obligation is not only relevant for the use of the system soon after its development but also complicated by the characteristics of neural network methods. Neural expert systems differ from more traditional medical information systems with respect to the dynamic nature resulting from their autonomous learning and generalization capabilities. In practice, this capability means that neural expert systems can respond to or learn to respond to new, unknown cases—for example, to attributes of cancerous tissue

regions that medical experts may not have previously experienced or described during the training of the network. The ethical consequence for the developers is that to ensure that appropriate learning strategies are followed and that correct generalizations are made, they must monitor the development and evolution of the system that they have developed and continue to collaborate with the medical experts well after the system has been adopted. Similarly, increased confidence in the system's performance and trust in the system's recommendation should be challenged by the medical expert, regardless of the system's success. In other words, human expertise should continue to be valued above technological achievements. Medical experts should not be reluctant to use intelligent systems because they feel that such systems challenge their professional judgment or their 'sense of self'. On the contrary, as systems become more effective, the use of human expertise becomes more critical for avoiding problems of over reliance on technology and the resulting critical errors (cf. Baase, 1997).

Both professional responsibility and the importance of human expertise are intrinsically related to accountability. In medicine, this concept is particularly important as the use of clinical information may make the difference between life and death. Information ethics is replete with problems of attributing moral responsibility for outcomes (Mason, 1995). Such problems are indeed evident in the use of neural expert systems in medical image interpretation and are particularly complicated by the need of experts from two domains to communicate and to collaborate closely, as discussed previously. Significantly, the effectiveness of the suggestions of one professional depends on the professionalism of the other. As a result of this close collaboration and interdependence, tracing errors to either professional may be difficult. This ethical issue is further complicated by the inherent weakness of neural networks in providing explanations for their reasoning process. In this respect, the scientific community engaged in neural network research has an obligation to promote research that can support the explanation process. The consideration of the responsibilities of such broader scientific communities is important.

So far, our discussion of ethical issues has concentrated mainly on the role and moral obligations of neural network developers and medical users. These professionals are indeed the two key stakeholder groups that are most

actively involved in the life cycle of clinical information. Nonetheless, both professional groups are members of and work in organizations. Within healthcare organizations in particular, several other professional groups—such as nurses or radiography technicians—interact with the technology and the clinicians. Clearly, many of the ethical issues that clinicians face concerning the use of intelligent systems also affect the other groups. Their involvement also signifies that the pattern of stakeholders becomes more complex as further groups react to technology, become engaged in the organizational discourse, or are anxious to promote their own interests.

At the same time, the ethical obligations of medical organizations, as well as those of organizations employing intelligent systems developers, should not be underestimated—particularly as such organizations have their own agendas and also need to respond to a number of pressures from their environment. Neural expert systems developers work in organizations that are anxious to publish and undertake novel research, to obtain funding, or to gain profits from selling advanced technology. Similarly, clinicians work in organizations that must compete successfully with other hospitals and healthcare provision organizations, not only on the basis of cost savings but also on public image and reputation. Thus, such organizations must be seen as using advanced technology systems and employing experts that work toward their development.

It is therefore necessary that a counterbalance exists for such pressures, and that stakeholders who are engaged directly or indirectly in the development and use of intelligent technologies behave ethically. In this respect, regulators and professional bodies have a responsibility to safeguard the interests of both the community and the patients by directing stakeholders to take into account the ethical issues. Existing codes of ethics for the medical and computing professions are already in place (for example, Code of Medical Ethics of the American Medical Association or the ACM Code of Ethics and Professional Conduct) and should be reinforced and applied in the case of intelligent systems for medical applications. Although intelligent systems developers do not face the decisions about treatment and practice that clinicians face, unprofessional conduct that results in a system hindering medical practice should be as severely disciplined as unprofessional clinician behavior is. As information technology—and intelligent technologies in

particular—become pervasive in the medical context, considering a dedicated code of ethics in medical informatics may be appropriate, incorporating elements from both the computing and the medical codes to which both developers and clinicians are bound.

Furthermore, the two professional groups should realize the significance of acting professionally and collaborating and communicating as needed with other professionals. The organizations in which they work should not only be sensitive to the ethical dimensions of their work but also encourage and promote it. This attitude is also in the interest of healthcare organizations—for example, recognizing the ethical risks of underestimating or being unaware of technological limitations. Such risks may lead to the misuse or inappropriate use of intelligent technologies, which ultimately reflects on the quality of care. For healthcare organizations, this misuse may also reflect on their relations with their customers and viability.

Our discussion in this section indicates that the ethical risks in the use of neural expert systems in medical image interpretation are predominantly the result of the interests, behavior, and interactions of human agents. In other words, such risks are the result of issues that stem from the interaction of stakeholders with the technology or with other stakeholders for the effective deployment of this technology, rather from than the technology itself. Thus, the ethical issues in the development and use of neural expert systems for medical image interpretation, and for medical applications more generally, can be better understood and managed if they are considered stakeholder issues.

5. CONCLUSIONS

In this paper, we have concentrated on the development and use of neural expert systems for medical image interpretation. The characteristics and the context of use make these systems prone to a number of ethical issues, some of which result from the nature of neural expert systems, while others are relevant to a broader spectrum of artificial intelligence and information technology applications in health care. Each stage of the development and use of a neural expert system—and by extension, of an intelligent system—entails

ethical issues. The responsibility for image interpretation is affected by contextual factors and should be shared amongst developers and medical users. Although users participate in medical image interpretation projects and assume professional responsibility, to what extent they understand the capabilities and shortcomings of the technology is not clear. The stakeholder perspective adopted in this paper will be beneficial for the stakeholder groups that understand their responsibility to behave ethically. Future research in information ethics in this application context should empirically investigate how developers and researchers of neural expert systems can help users to understand and use ethically intelligent systems. Future research can thus expand our analysis of ethical issues with further examples that consider dilemmas for the stakeholders of artificial intelligence applications in medicine, as these can be critical to the promotion of health care.

ACKNOWLEDGMENT

The authors would like to thank the reviewers for their constructive comments.

REFERENCES

- Aleksander, I. and Morton, H. 1991. *An introduction to neural computing*, Chapman & Hall, London, UK.
- Anderson, J.G. 1997. Clearing the way for physician's use of clinical information systems, *Communications of the ACM*, **40**, 83–90.
- Agnew, N.M., Ford, K.M. and Hayes, P.J. 1994. Expertise in context: personally constructed, socially selected, and reality-relevant? *International Journal of Expert Systems*, **7**, 65–88.
- Avgerou, C. and Cornford, T. 1993. *Developing information systems*, Macmillan, London, UK.
- Baase, S. 1997. *A gift of fire: Social, legal, and ethical issues in computing*, Prentice Hall, New Jersey, USA.

- Beerel, A. 1989. *Expert systems: Strategic implications and applications*, Ellis Horwood, Chichester, UK.
- Bloomfield, B.P., Coombs, R., Knights, D. and Littler, D. 1997. *Information technology and organizations*, Oxford University Press, Oxford, UK.
- Bolger, F. 1995. Cognitive expertise research and knowledge engineering, *The Knowledge Engineering Review*, **10**, 3–19.
- Carroll, B. 1994. Artificial intelligence: Expert systems for clinical diagnosis: Are they worth the effort? in: *Social Issues in Computing*, edited by Juff, C. and Finholt, T., McGraw-Hill, New York, USA, 247–272.
- Collins, H. 1990. *Artificial experts: Social knowledge and intelligent machines*, MIT Press, Cambridge, Massachusetts, UK.
- Delaney, P.M., Papworth, G.D. and King, R.G. 1998. Fibre optic confocal imaging (FOCI) for in vivo subsurface microscopy of the colon, in: *Methods in disease: Investigating the gastrointestinal tract*, edited by Preedy V.R. and Watson R.R., Greenwich Medical Media, London, UK.
- Dreyfus, H. 1987. Misrepresenting human intelligence, in: *Artificial intelligence: the case against*, edited by Born, R., Croom Helm, London, UK.
- Dreyfus, H. and Dreyfus, S. 1986. *Mind over machine: The power of human intuition and expertise in the era of the computer*, Basil Blackwell, Oxford, UK.
- Essex, B. 1994. *Doctors, dilemmas, decisions*, BMJ Publication Group, London, UK.
- Furundzic D., Djordjevic M. and Jovicevic Bekic, A. 1998. Neural networks approach to early breast-cancer detection, *Journal of Systems Architecture*, **44**, 617–633.
- Firlej, M. and Hellens, D. 1991. *Knowledge elicitation: A practical handbook*, Prentice Hall, New York, USA.
- Friedman, B. and Kahn, P.H. Jr. 1997. People are responsible, computers are not, in: *Computers, ethics, and society*, edited by Ermann, M.D, Williams, M.B. and Shauf, M.S., Oxford University Press, Oxford, UK.
- Gaines B. and Boose J. 1988. *Knowledge acquisition for knowledge based systems*, Academic Press, London, UK.
- Hanka, R., Harte, T.P., Dixon, A.K., Lomas, D.J. and Britton, P.D. 1996. Neural networks in the interpretation of contrast-enhanced magnetic

- resonance images of the breast, in: *Proceedings of Healthcare Computing 1996*, Harrogate, UK, 275–283.
- Hart, A. 1986. *Knowledge acquisition for expert systems*, Kogan Page, London, UK.
- Haykin, S. 1994. *Neural networks: A comprehensive foundation*, Macmillan, New York, NY, USA.
- Herman, W.A., Marlowe, D.E. and Rudolph, H. 1998. *Future trends in medical device technology: Results of an expert survey*, Center for Devices and Radiological Health, <http://www.fda.gov/cdrh/ost/trends/TOC.html>, (accessed on October 2nd, 2000).
- Ifeachor, E.C. and Rosen, K.G. 1994. *Proceedings of the international conference on neural networks and expert systems in medicine and healthcare*, Plymouth, UK.
- Innocent, P.R., Barnes, M. and John, R. 1997. Application of the fuzzy ART/MAP and MinMax/MAP neural network models to radiographic image classification, *Artificial Intelligence in Medicine*, **11**, 241–263.
- Karkanis, S., Magoulas, G.D. and Theofanous N. 2000. Image recognition and neuronal networks: Intelligent systems for the improvement of imaging information, *Minimal Invasive Therapy & Allied Technologies*, **9**, 225–230.
- Kasabov, N. 1996. *Foundations of neural networks, fuzzy systems and knowledge engineering*, MIT Press, Cambridge, Massachusetts, USA.
- Kurihara, Y., Okuhara, Y., Narita, Y., Kitazoe, Y., Sawada, A. and Yoshida, S. 1999. Integration of medical images into the total hospital information system—experiences at Kochi Medical School, in: *Proceedings of the thirtieth-second hawaii international conference on systems sciences (HICSS-32)*, Maui, Hawaii.
- Kushniruk, A.W. and Patel, V.L. 1998. Cognitive evaluation of decision making processes and assessment of information technology in medicine, *International Journal of Medical Informatics*, **51**, 83–90.
- Kwoh, C.K. 1995. *Probabilistic reasoning from correlated objective data*, Ph.D. Thesis, Imperial College, London, UK.
- Lane, V.P., Lane, D. and Littlejohns, P. 1996. Neural networks for decision making related to asthma diagnosis and other respiratory disorders, in: *Proceedings of healthcare computing 1996*, Harrogate, UK, 85–93.