6
Ethical and Legal Issues in Decision Support

Kenneth W. Goodman

Discrete maladies or illnesses tend to produce particular signs and symptoms. This natural correlation makes possible the process of diagnosis and prognosis. In fact, so strong is our belief in the regularity of signs and symptoms that the process has long been regarded as straightforward, if not easy: ". . . there is nothing remarkable," Hippocrates suggested some 2,400 years ago, "in being right in the great majority of cases in the same district, provided the physician knows the signs and can draw the correct conclusions from them".[1]

Of course, accurate diagnosis and prognosis can be quite difficult, even given the regularity of signs and symptoms. For one thing, "knowing the signs" requires a great deal of empirical knowledge and experience. For another, there is rarely a unique and isomorphic relationship between symptom and disease. Significantly, Hippocrates smuggles into his account a presumption of the very thing being described. To say that being right is unremarkable when one can draw the "correct conclusions" is to say that it is easy to be right when you know how to be right. Or, making an accurate diagnosis or prognosis is easy if one knows how to make an accurate diagnosis or prognosis!

The need to make accurate diagnoses is not based merely on the personal satisfaction that comes from being right, as gratifying as that is. It is based on the good effects that follow more frequently from accurate diagnoses than from inaccurate diagnoses. It is also based on the bad effects that error entails.

In the context of trust and vulnerability that shape patient-physician and patient-nurse encounters, there emerges an ethical imperative: to adhere to, or surpass, educational and professional standards, to monitor changes in one's domain, to know when one is out of one's depth. Decision support systems have the potential to assist clinicians, but their use also entails a number of ethical concerns. In fact, this is evidence for the maturity of the science: new health technologies almost always raise ethical issues, and it should come as no surprise that clinical decision support would provide a number of challenges for those who use, or would use, computers to assist,

guide or test clinical decisions. Any comprehensive treatment of clinical decision support systems (CDSS) should include a review of ethical issues. In what follows, we identify a number of ethical issues and positions that emerge when intelligent machines are used to perform or support clinical decision making, and we survey key legal and regulatory issues.

Ethical Issues

Background and Current Research

It has been clear for more than a decade that health computing raises interesting and important ethical issues. In a crucial early contribution, a physician, a philosopher, and a lawyer identified a series of ethical concerns, not the least of which are several surrounding the questions of who should use a "medical computer program" and under what circumstances.[2] Another early contribution emphasized the challenges raised by threats to physician autonomy.[3]

What has emerged since has been called the "Standard View" of computational diagnosis.[4] Randolph A. Miller, M.D., a key figure both in the scientific evolution of computational decision support and in scholarship on correlate ethical issues, has argued that "Limitations in man-machine interfaces, and, more importantly, in automated systems' ability to represent the broad variety of concepts relevant to clinical medicine, will prevent 'human-assisted computer diagnosis' from being feasible for decades, if it is at all possible."[4] Another way of putting this is to say that computers cannot, either in principle or at least for the foreseeable future, supplant human decision makers. This observation entails ethical obligations, namely that computers ought not to be relied on to do what humans do best, and that a "computer diagnosis" cannot, as a matter of course or policy, be allowed to trump a human decision or diagnosis.

Happily, the Standard View has been advanced not by those hostile to the development and use of CDSS, but by leading proponents. The Standard View bespeaks a conservative and cautious approach to applications of a new technology, and, as such, captures important moral intuitions about technological change, risks, and standards.

Interest in the three-way intersection of ethics, medicine, and computing has increased significantly since initial efforts to explore these issues. On the one hand, professional societies such as the American Association for the Advancement of Science, the American College of Physicians and the American Medical Informatics Association have encouraged educational programs and other professional activities. On the other hand, the literature exploring this intersection has progressed significantly, and now includes the first book devoted to the topic.[5]

Three core areas of ethical concern have emerged in discussions of computer systems that are used to remind, consult, or advise clinicians:

(1) care standards; (2) appropriate use and users; and (3) professional relationships.[6]

Care Standards

We know a great deal about responsibility in medicine and nursing. For instance, we know that practitioners should generally not deceive their patients. We know that patients can be especially vulnerable, and that such vulnerability should be respected. And we know that physicians and nurses have a responsibility to do their best, irrespective of economic (dis)incentives, and that they should not attempt treatments that are beyond their training or expertise.

Learning how to meet these and other responsibilities in the context of a broad variety of social problems is arguably the leading task in bioethics. We must first ask whether computing tools help or hinder attempts to meet responsibilities, and, second, whether the tools impose new or special responsibilities. The overarching question may be put thus: does the new technology improve patient care? If the answer is affirmative, we may suppose we have met an important responsibility. If the answer is negative, it seems clear we should not use the new technology. The problem is, we often do not know how to answer the question. That is, we are sometimes unsure whether care will be improved by the use of new technologies. If we want to meet the responsibility to avoid harm, for instance, we are impotent until we can determine the effects of the technology (see Chapter 7). The upshot here is that error avoidance is an ethical imperative, both to maximize positive, short-term consequences and to ensure that, in the long run, informatics is not associated with error or carelessness, or the kind of cavalier stance sometimes associated with high-tech boosterism.

The concept of error avoidance is wed to that of a standard of care. Standards evolve in the health professions because they plot the kinds of actions that are most successful in achieving certain ends. To fail to adhere to a standard is thus to increase the risk of error, at least in a mature science. Because errors or their consequences are generally regarded as harms or evils, the obligation to hew to standards is an ethical one.

But standards are empirical constructs, and so are open to revision. New evidence forces changes in standards. (This demonstrates why clinicians have an ethical obligation to monitor the scientific maturation of their disciplines by reading journals, attending conferences, etc.) To be sure, the precise content of any standard might be open to dispute. The "reasonable person" standard requires the postulation of a vague entity; this is particularly problematic when reasonable people disagree, as is often the case in medicine and nursing. A "community standard" similarly fails to identify a bright line between error and success in all circumstances in which it might be invoked. Note also that it is not always bad to forgo adherence to a

practice standard—the standard will generally be invoked in ethical and legal contexts only when there is a bad outcome, or a flagrant disregard for the risk of a bad outcome. Sometimes there are good reasons to violate a standard. This demonstrates how some clinical progress is possible: if everyone in all cases stuck to a rigid standard, there would be no internal evidence to support modifications of the standard. In other cases, standards are modified as a result of clinical trial findings, observational studies, and serendipitous discoveries.

In the case of computer-assisted diagnoses, the challenge is perhaps best put in the form of a question: does use of a decision support system increase the risk of error? Note in this regard the following three points. First, while accurate diagnosis is often linked to optimal treatment, this is not always the case: some patients are treated appropriately despite an inaccurate diagnosis, and some are treated incorrectly despite an accurate diagnosis. Second, one might still be able to provide an optimal treatment with a vague or imprecise diagnosis.[7] Third, computers can render diagnoses (or perform diagnosis-like functions) outside of clinical contexts, as, for instance, in tests for blood-borne pathogens,[8] cytology screens,[9] and the like.

To ask if a computer diagnosis increases (or decreases) the risk of diagnostic or other error is in part to ask whether it will improve patient care. If the answer is that, on balance, the tool increases (the risk of) diagnostic error, then we should say it would be unethical to use it. Significantly, though, what is sought here is an empirical finding or a reasoned judgment—where such a finding is often lacking or even methodologically hard to come by, or where such a judgment is based on inadequate epistemic support, at least according to standards otherwise demanded to justify clinical decisions.

This means that we are pressed to answer an ethical question (is it acceptable to use a decision support system?) in a context of scientific uncertainty (how accurate is the system?). Many challenges in contemporary bioethics share this feature, namely, that moral uncertainty parallels scientific or clinical ignorance.

What we generally want in such cases is a way to stimulate the appropriate use of new technologies without increasing patient risk. One approach to doing this is given the nearly oxymoronic term "progressive caution." The idea is this: "Medical informatics is, happily, here to stay, but users and society have extensive responsibilities to ensure that we use our tools appropriately. This might cause us to move more deliberately or slowly than some would like. Ethically speaking, that is just too bad."[10] Such a stance attempts the ethical optimization of decision-support use and development by encouraging expansion of the field, but with appropriate levels of scrutiny, ovesight, and, indeed, caution.

The moral imperative of error avoidance is, in other words, not antiprogressive. Rather, it is part of a large and public network of checks and balances that seeks to optimize good outcomes by regulating conflicts between

boosters and naysayers. The idea of progressive caution is an attempt to capture the core values of that regulation.

It has been clear since the first efforts to address ethical issues in medical informatics that as computers help the sciences of medicine and nursing to progress, they will also contribute to changes in the standard of patient care. When that happens, however, it increases the likelihood that computer use will come to be required of clinicians. Put differently, in a comparatively short time, there has been a major shift in the availability and use of informatics tools. To the degree that informatics can improve the practice of the health professions, there is a requirement that its tools be used.

This point is often the most disturbing for practitioners. It is troublesome that one might have an obligation to use a tool that has been presented as controversial and in need of further validation. But there is no contradiction here. In fact, it appears that the rise of medical informatics parallels the emergence of other exciting and controversial tools, ranging from organ transplantation techniques and advanced life support to laparoscopic surgical procedures and genetic testing and therapy. It is often the case in history that progress involves this tension. What is wanted is evidence that people of good will can both advance science and safeguard against abuses. Research studies that examine not just the accuracy of the systems, but how they are used, are crucial to collecting that evidence.

Appropriate Use and Users

One way to abuse a tool is to use it for purposes for which it is not intended. Another is to use a tool without adequate training. A third way is to use a tool incorrectly (carelessly, sloppily, etc.) independently of other shortcomings.

There are a number of reasons why one should not use tools in unintended contexts. First, a tool designed for one purpose has a greater likelihood of not working, or not working well, for other purposes. To be sure, one might successfully perform an appendectomy with a kitchen knife, or dice vegetables with a scalpel, but it is bizarre to suggest that one should try either, except in an emergency. A medical computer system may be used inappropriately if, for instance, it was designed for educational purposes but relied on for clinical decision support; or developed for modest decision support (identifying a number of differential diagnoses) but used in such a way as to cause a practitioner to abandon a diagnosis arrived at by sound clinical methods.

In ethically optimizing the use of CDSS, it is perhaps reassuring to know that we have many models and precedents. From advanced life support and organ transplantation to developments in pharmacotherapy and genetics, society regularly has had to cope with technological change in the health

sciences. Managing change requires that new tools are used appropriately and by adequately qualified practitioners. Education is at the core of such management.

Identifying qualifications and providing training must be key components of any movement to expand the use of decision support software. Ethical concerns arise when we are unsure of the appropriate or adequate qualifications and levels of training.[6]

The fear is that: (1) a healthcare novice, or (2) a healthcare professional ignorant of a system's design or capacity will use a decision support system in patient care. The reason the former is worthy of concern is that, as above, the practice of medicine and nursing remain human activities. A nonphysician or nonnurse cannot practice medicine or nursing, no matter how much computational support is available. This is also a concern in the context of consumer health informatics, or the widespread availability of online health advice to the untrained (see Chapter 11). What this means is that the novice might not know when the system is in error or producing flawed output, when it is operating on insufficient information, when it is being used in a domain for which it was not designed, and so on.

There are several reasons we must also focus ethical attention on the use of decision support software by computationally naive health professionals. Such professionals might not use such software to good effect (either by over- or underestimating its abilities), might not be using it properly, or, like the novice, might not know when the system is being used in inappropriate contexts.

Such fears can be addressed by requirements that users of CDSS have appropriate qualifications and be adequately trained in the use of the systems. Unfortunately, it is not yet clear what those qualifications should be or how extensive a training program would be adequate. It is clear, however, that the use of diagnostic software cannot, in the long run, advance ethically without a better sense of where to establish guideposts for qualifications and training. This will be an increasingly important area of research in coming years.

A further ethical concern about appropriate use and users emerges from the potential to deploy decision support systems in contexts of practice evaluation, quality assessment, reimbursement for professional services, and the like. One can imagine an insurance company or managed care organization using decision support to evaluate, or even challenge, clinical decisions. What makes such use problematic is precisely the same ensemble of concerns that led us to disdain applications in other contexts: the primacy of human cognitive expertise, uncertainty about adequate qualifications, and doubt about the consequences for improved patient care. This is not to say that a machine cannot give a correct answer in a particular case but, rather, that there are inadequate grounds to prefer machine decisions as a matter of general policy.

Professional Relationships

Many patients believe, mistakenly, that their physicians are omniscient. Many physicians believe, mistakenly, that their patients are ignoramuses. Recognition of these mistakes has led, in recent years, to the development of the idea of "shared decision making," namely, that patients and providers are most productively seen as partners.[11] If this is so, and there is much to recommend it in many (though not all) instances, then we need to assess the effect of a third partner—the computer.

There are two overriding areas of ethical concern here. The first is that the computer will create conceptual or interpersonal distance between provider and patient. Communicating about uncertainty, especially when the stakes are high, has long been a challenge for clinicians. That a computer might be used to (help) render a diagnosis causes us to run the risk of what we will call the "computational fallacy." This is the view that what comes out of a computer is somehow more valid, accurate, or reliable than human output. Providers and patients who take such a view introduce a potentially erosive, if not destructive, element into shared decision-making contexts. Anything that increases the likelihood that a patient decision or choice will be perceived as misguided or stupid adds to the problem that shared decision making was supposed to solve.

Now, it might be supposed that the physician or nurse can eliminate at least some of this tension by not disclosing to a patient that decision support software was used in his or her case. But this introduces our second area of ethical concern, namely, the question whether patients should be given this information. The answer to this question must be determined against a background shaped by: (1) patient sophistication and understanding of medical and statistical information, and (2) clinician sophistication and understanding of communication approaches and strategies. In any case, it is inappropriate to use computer data or inferences to trump hesitant patients, or bully them into agreeing with a health professional.[12]

This point has been made most clearly in the discussion of prognostic scoring systems, or software used in critical care medicine in part to predict patient mortality. On the one hand, patients with poor prognoses might still benefit from extensive interventions, and these benefits might be important enough for the patient and/or family to seek them; on the other hand, patients with good survival odds might judge the prolongation of life to be of little value when weighed against the difficulty or burden of extensive interventions.[13]

A related issue is likely to arise with increased frequency as patients gain access to decision support software and use it to make demands on physicians, or at least to challenge or second-guess them. The difficulties raised by these demands and challenges will multiply as these systems improve. As discussed in Chapter 11, there is a sense in which one might regard such access as an important tool in the process of shared decision

making: it will not do to expect patients to become involved in their own care and simultaneously constrain their sources of information. Contrarily, a patient might constitute a paradigm case of an inappropriate decision support system user, especially in those cases in which the system causes someone to forgo appropriate medical care.

We might compare patient use of clinical decision support systems to patient use of medical texts and journals. In years past, there was an inclination to regard such access as risky and hence inappropriate. While a little knowledge can be dangerous, a position that does not go beyond such a view seems to miss an opportunity to educate patients about their illnesses and the relation between medical literature on the one hand, and medical knowledge and practice on the other. Much the same point can be made about patient use of diagnostic tools: a physician should respond to such use by making clear that computers are not surrogates for health professionals and that the practice of medicine or nursing entails far more than statistical induction from signs, symptoms, and lab values. To be sure, it would be well if actual practice embodied this insight.

As long as the healing professions are practiced in a matrix of scientific uncertainty and patient values, we err if we appoint computational decision support as a surrogate for compassionate communication, shared decisions, and quality care by competent humans.

Legal and Regulatory Issues

Computers and software raise conceptually fascinating and important practical questions about responsibility and liability. Further, the question of whether a decision-support system is a medical device needing governmental regulation is a source of tension and debate. In both domains, scientists, clinicians, philosophers, lawyers, and government and policy officials must grapple with a variety of knotty problems.

The intersection of medicine, computational decision support, and law, has been addressed mostly in speculative terms. The use of CDSS is not widespread enough to have stimulated legislation or illuminating precedent. Moreover, medicine and computing share little in the way of a common legal history. The following observation is as apt today as it was more than twenty years ago:

> "The introduction of computerized decision making will require the merger of computer science and medical care; two areas with fundamentally different legal traditions. The legal differences between the computer field and medicine are striking. Medicine is tightly regulated at all levels. Most health-care providers are licensed, and a rigid hierarchical system is the norm. Yet computer systems and companies are created in a totally unregulated competitive environment in which "software piracy" is common, standardization is in its infancy, licensing is a method of transferring trade secret software, and companies begin in garages."[14]

Liability and Decision Support

The overriding legal issue related to computational decision support is liability for use, misuse, or even lack of use of a computer to make or assist in rendering medical decisions.[15–18] In the United States, tort law holds providers of goods and services accountable for injuries sustained by users. Because of legal and regulatory variation, there are similarities and differences in other countries.[19–21] Such accountability is addressed by either the negligence standard or the strict liability standard.

The negligence standard applies to services, and strict liability applies to goods or products, although negligence can sometimes also apply to goods, as in cases of negligent product design. There is no consensus about whether decision-support systems are services or products, in part because these systems have properties that resemble both services and products.[2,14–15,22–23] For instance, a physician's diagnosis is clearly a service, and any liability for erroneous diagnoses is judged by the negligence standard. If a human diagnosis is considered a service, then, it is argued, a computer diagnosis (or the task of writing the computer code that rendered the diagnosis) should have the same status. Contrarily, commercial CDSS are manufactured, mass-marketed, and sold like entities uncontroversially regarded to be products.

An additional complication is that these systems are sold to hospitals, physicians, patients, and others, and, indeed, are now available on the World Wide Web. If a patient is injured by a defective system, it remains to be determined who used the system (the physician? the patient?) and whether it was misused. Also, it can be exquisitely difficult to identify the defect in a computer program,[15] as well as to answer the important question as to whether a physician could have intervened and prevented the application of mistaken advice.[2]

Neither is there a clear standard of care for use of decision-support software by clinicians. Physicians or nurses might someday be found negligent either for accepting a mistaken computer diagnosis or, having erred in diagnosis themselves, for failing to have used a decision-support system that might have proved corrective. In either case, the determination of negligence will have to be weighed against prevailing community or reasonable-person standards. As with other areas of practice, errors will increase liability accordingly as the practitioner is seen to have fallen behind, or moved too far ahead of, such standards.

There is a clear need for additional conceptual analysis to assist the law in sorting out these puzzles. Local trial courts and juries will often be out of their depth if called on to adjudicate liability claims that challenge fundamental conceptions of responsibility, accountability, and blame. Similar difficulties arise in other areas, such as in the intellectual property arena, when there is a need to determine whether computer software is an invention or a work of art. In one interesting approach to these questions, Prof. John Snapper attempts an account of responsibility that will not impede the

future—and presumably salutary—development of mechanical decision support. On this account, the attribution of responsibility and duty to computers for certain actions will maximize the good that will result from increased use of improved decision-support systems.[24] The idea is that use of conceptually inadequate legal tools to punish system designers, owners, and users, might have a chilling effect on the evolution of decision-support technology. Spreading responsibility around, and including computers as agents to which responsibility may be assigned, is said to offer the potential of stimulating system design and the benefits this would entail.

This much is clear: physicians and nurses who revile and disdain computers will be ignorant of machines that can, in principle, improve their practice and, hence, patient care. Zealots who take computers to constitute adequate or even superior human surrogates will have lost touch with the human foundations of their profession. At either extreme, the risk is high of falling outside emerging standards. This is a mistake—in ethics and at law.

Regulation of Decision-Support Software

While the history of governmental regulation of healthcare products is traceable to the Pure Food and Drug Acts of 1906, the regulation of medical devices was not formalized until the Federal Food, Drug, and Cosmetic Act of 1938. There, medical devices were defined as "instruments, apparatus, and contrivances, including their components, parts, and accessories intended: (1) for use in diagnosis, cure, mitigation, treatment, or prevention of diseases in man or other animals; or (2) to affect the structure or any function of the body of man or other animals."[25–26] In 1976, motivated by the increased complexity of devices and by reports of some devices' shortcomings and failures, Congress approved comprehensive Medical Device Amendments to the 1938 regulations; the amendments were to "ensure that new devices were safe and effective before they were marketed."[27–28] In 1990, a new regulation replaced that emphasis on premarket approvals with an emphasis on postmarket surveillance.[29] Proposals to regulate diagnostic software have been evaluated against the 1976 and 1990 laws and a broad array of draft policies and statements.

The U.S. Food and Drug Administration (FDA) unequivocally regards medical software as a device. The FDA identifies four types of devices:

1. **Educational and Bibliographic Software**
   Federal authorities regard the following as exempt from, or not falling under, existing regulation:
   - Software intended only for use in performing traditional "library" functions, such as storage, retrieval, and dissemination of medical information (i.e., functions that are traditionally carried out using medical textbooks and journals).

- Software intended only for use as general accounting or communications functions.
- Software solely intended for educational purposes, rather than to diagnose or treat patients.[30]

2. **Software Components**
   Some software is incorporated into medical devices and is actively regulated. Examples include the software in:
   - infusion pumps;
   - pacemakers;
   - ventilators;
   - magnetic resonance imaging devices;
   - diagnostic X-ray systems;
   - clinical laboratory instruments;
   - blood grouping instruments.[30]
3. **Software Accessories**
   Software accessories are attached to, or used with, other devices, and as such are also actively regulated. These include software for:
   - radiation treatment planning;
   - conversion of pacemaker telemetry data;
   - conversion, transmission or storage of medical images;
   - off-line analysis of EEG data;
   - digital analysis and graphical presentation of EEG data;
   - calculation of rate response for a cardiac pacemaker;
   - perfusion calculations for cardiopulmonary bypass;
   - calculation of bone fracture risk from bone densitometry data;
   - statistical analysis of pulse oximetry data;
   - calculation of refractive power of intraocular lenses.[30]
4. **Stand-Alone Software**
   The most controversial class, stand-alone software, includes CDSS and other decision support systems. Whether or how such systems should be regulated is a matter of continuing debate. Examples include:
   - blood bank software systems which control donor deferrals and release of blood products;
   - software designed to assist a healthcare practitioner in arriving at a diagnosis of a particular patient;
   - software which analyzes for potential therapeutic interventions for a particular patient;
   - software which records medical information for later recall, analysis, or action by a healthcare practitioner (e.g., hospital information systems, prescription ordering, drug interaction information systems, emergency room triage software, and various calculators which automate calculations of drug doses).[30]

In 1989, an FDA draft policy proposed regulatory exemption for "Previously unclassified information management products . . . such as expert or

knowledge-based systems, artificial intelligence, and other types of decision-support systems intended to involve competent human intervention before any impact on human health occurs."[31] The question then became whether CDSS were intended to involve competent human intervention. This remains an interesting and important policy—and conceptual—issue. In Chapter 5, Miller and Geissbuhler examine some of the issues connected with FDA regulation.

While the FDA regards software as a device, there are a number of reasons why it might be best if medical decision-support software were not subjected to thorough federal regulation. The most common arguments against regulation include the following:

- Software is most accurately regarded as a mental construct or abstract entity, i.e., the sort of thing not customarily falling within the FDA's regulatory purview.
- Practitioners—not software—have traditionally been subjected to licensing requirements.
- Software evolves rapidly and locally, and any sort of national software monitoring is likely to be ineffective or impossible.
- Software is imperfect, and so improvement and refinement—not perfection—must be the standard to be striven for and met. Yet at law, strict liability standards (usually applied to devices or goods but not services) require perfection.

Several of these points could be in line with an influential stance held by a former commissioner of the agency, namely that the FDA should "apply the least regulation allowed to remain consistent with the requirements of public health and safety."[32]

The debate over medical software regulation represents one of the most important controversies of the Computer Age. The balancing of risks and benefits, as well as public safety and technological progress, means that scientists, clinicians, and policy makers have one of civilization's most interesting—and challenging—tasks.

Conclusion and Future Directions

Clinicians, philosophers, lawyers, and policy makers have grappled for more than a decade with social, ethical, and legal issues raised by the growth of health informatics, perhaps especially by progress in development of tools for clinical decision support. What has emerged is a recognition that future scientific growth must be guided by corresponding attention to ethical issues. These issues address the role of: error avoidance and standards; appropriate use and users; and professional relationships. Scientific programs and publications may be regarded as duty-bound to foster environments in which further attention to ethical, legal, and social issues is

encouraged. Indeed, to the extent that morality guides the law, vigorous programs to identify and debate ethical issues will be of no small service to society as legislatures, courts, and government regulators and policy makers attempt to apply the insights of ethics to practical problems in health informatics.

More research on ethical issues involved in use of CDSS is essential for this process. We have, for instance, only begun to address issues that arise when diagnostic tools are made available on the World Wide Web. We are in no way clear about the level of ethics education that is appropriate for students in health informatics, and there is much work to be done at the intersections of ethics and system evaluation, and of ethics and standards of care.

Elsewhere in the history of science and technology, such challenges are often taken to constitute evidence of the growth and maturation of an applied science. This is no less true for clinical decision support systems and, indeed, for all of health informatics.

References

1. Hippocrates. Prognosis. In Lloyd GER, ed. Hippocratic writings, translated by Chadwick J, Mann WN. London: Penguin Books; 1983:170–185.
2. Miller RA, Schaffner KF, Meisel A. Ethical and legal issues related to the use of computer programs in clinical medicine. Ann Intern Med 1985;102:529–536.
3. de Dombal FT. Ethical considerations concerning computers in medicine in the 1980s. J Med Ethics 1987;13:179–184.
4. Miller RA. Why the standard view is standard: people, not machines, understand patients' problems. J Med Philos 1990;15:581–591.
5. Goodman KW, ed. Ethics, computing and medicine: informatics and the transformation of health care. Cambridge and New York: Cambridge University Press; 1997.
6. Miller RA, Goodman KW. Ethical challenges in the use of decision-support software in clinical practice. In: Goodman KW, ed. Ethics, computing and medicine: informatics and the transformation of health care. Cambridge and New York: Cambridge University Press; 1997:102–115.
7. Berner ES, Webster GD, Shugerman AA, et al. Performance of four computer-based diagnostic systems. N Engl J Med 1994;330:1792–1796.
8. Sorace JM, Berman JJ, Carnahan GE, Moore GW. PRELOG: precedence logic inference software for blood donor deferral. Proc Annu Symp Comput Appl Med Care 1991:976–977.
9. Boon ME, Kok LP. Neural network processing can provide means to catch errors that slip through human screening of Pap smears. Diag Cytopathol 1993;9:411–416.
10. Goodman KW. Bioethics and health informatics: an introduction. In: Goodman KW, ed. Ethics, computing and medicine: informatics and the transformation of health care. Cambridge and New York: Cambridge University Press; 1997:1–31.
11. Forrow L, Wartman SA, Brock DW. Science, ethics, and the making of clinical decisions. JAMA 1988;259:3161–3167.

12. Goodman KW. Outcomes, futility, and health policy research. In: Goodman KW, ed. Ethics, computing and medicine: informatics and the transformation of health care. Cambridge and New York: Cambridge University Press; 1997:116–138.
13. Brody BA. The ethics of using ICU scoring systems in individual patient management. Prob Crit Care 1989;3:662–670.
14. Brannigan VM, Dayhoff RE. Medical informatics: the revolution in law, technology, and medicine. J Leg Med 1986;7:1–53.
15. Miller RA. Legal issues related to medical decision-support systems. Int J Clin Monit Comput 1989;6:75–80.
16. Mortimer H. Computer-aided medicine: present and future issues of liability. Computer Law J 1989;9:177–203.
17. Turley TM. Expert software systems: the legal implications. Computer Law J 1988;8:455–477.
18. Hafner AW, Filipowicz AB, Whitely WP. Computers in medicine: liability issues for physicians. Int J Clin Monit Comput 1989;6:185–194.
19. Beier B. Liability and responsibility for clinical software in the Federal Republic of Germany. Comput Methods Programs Biomed 1987;25:237–242.
20. Brahams D, Wyatt J. Decision aids and the law. Lancet 1989;ii:632–634.
21. Allaërt FA, Dussere L. Decision support system and medical liability. Proc Annu Symp Comput Appl Med Care 1992:750–753.
22. Birnbaum LN. Strict products liability and computer software. Computer Law J 1988;8:135–156.
23. Gill CJ. Medical expert systems: grappling with the issues of liability. High Technol Law J 1987;1:483–520.
24. Snapper JW. Responsibility for computer-based decisions in health care. In: Goodman KW, ed. Ethics, computing and medicine: informatics and the transformation of health care. Cambridge and New York: Cambridge University Press: 1997:43–56.
25. Munsey RR. Trends and events in FDA regulation of medical devices over the last fifty years. Food and Drug Law J 1995;50:163–177.
26. Public Law No. 75-717, 52 Stat. 1040 (1938), as amended 21 U.S.C. Sections 301 et seq (1988).
27. Kessler DA, Pape SM, Sundwall DN. The federal regulation of medical devices. N Engl J Med 1987;317:357–366.
28. Public Law No. 94-295, 90 Stat. 539 (1976), codified at 21 U.S.C. Sections 360c et seq (1982).
29. Brannigan VM. Software quality regulation under the Safe Medical Devices Act of 1990: hospitals are now the canaries in the software mine. Proc Annu Symp Comput Appl Med Care 1991:238–242.
30. Food and Drug Administration. FDA regulation of medical device software. (Document prepared for an FDA Software Policy Workshop, Sept. 3–4, 1996, National Institutes of Health, Bethesda, Md.). http://www.fda.gov//cdrh/ost/points.html.
31. Food and Drug Administration, Center for Devices and Radiological Health. Policy for the regulation of computer products, draft, 13 November 1989. Rockville, Maryland: FDA, CDRH; 1989.
32. Young FE. Validation of medical software: present policy of the Food and Drug Administration. Ann Intern Med 1987;106:628–629.