
Ethics in Biomedical and Health Informatics: Users, Standards, and Outcomes

10

Kenneth W. Goodman, Reid Cushman,
and Randolph A. Miller

After reading this chapter, you should know the answers to these questions:

- Why is ethics important to informatics?
- What are the leading ethical issues that arise in health care informatics?
- What are examples of appropriate and inappropriate uses and users for health-related software?
- Why does the establishment of standards touch on ethical issues?
- Why does system evaluation involve ethical issues?
- What challenges does informatics pose for patient and provider confidentiality?
- How can the tension between the obligation to protect confidentiality and that to share data be minimized?
- How might computational health care alter the traditional provider–patient relationship?

- What ethical issues arise at the intersection of informatics and managed care?
- What are the leading ethical and legal issues in the debate over governmental regulation of health care computing tools?

10.1 Ethical Issues in Biomedical and Health Informatics

More and more the tendency is towards the use of mechanical aids to diagnosis; nevertheless, the five senses of the doctor do still, and must always, play the preponderating part in the examination of the sick patient. Careful observation can never be replaced by the tests of the laboratory. The good physician now or in the future will never be a diagnostic robot. – Scottish surgeon Sir William Arbuthnot-Lane (Lane 1936)

Human values should govern research and practice in the health professions. Health care informatics, like other health professions, encompasses issues of appropriate and inappropriate behavior, of honorable and disreputable actions, and of right and wrong. Students and practitioners of the health sciences, including informatics, share an important obligation to explore the moral underpinnings and ethical challenges related to their research and practice.

Although ethical questions in medicine, nursing, human subjects research, psychology, social work, and affiliated fields continue to evolve and increase in number, the key issues are generally well known. Major questions in bioethics have been addressed in numerous professional,

K.W. Goodman, PhD (✉)
University of Miami Bioethics Program, 1400 NW 10th Ave., Suite 916 (M-825), Miami 33136, FL, USA
e-mail: kgoodman@med.miami.edu

R. Cushman, PhD
Department of Medicine, University of Miami, 1400 NW 10th Ave, NW, Suite 912, Miami 33136, FL, USA
e-mail: rcushman@med.miami.edu

R.A. Miller, MD
Department of Biomedical Informatics, Vanderbilt University Medical Center, 2209 Garland Avenue, B003C Eskind Biomedical Library, Nashville 37232-8340, TN, USA
e-mail: randolph.a.miller@vanderbilt.edu

scholarly, and educational contexts. Ethical matters in health informatics are, in general, less familiar, even though certain of them have received attention for decades (Szolovits and Pauker 1979; Miller et al. 1985; de Dombal 1987). Indeed, informatics now constitutes a source of some of the most important and interesting ethical debates in all the health professions.

People often assume that the confidentiality of electronically stored patient information is the most important ethical issue in informatics. Although confidentiality and privacy are indeed of vital interest and significant concern, the field is rich with other ethical issues, including the appropriate selection and use of informatics tools in clinical settings; the determination of who should use such tools; the role of system evaluation; the obligations of system developers, maintainers, and vendors; the appropriate standards for interacting with industry; and the use of computers to track clinical outcomes to guide future practice. In addition, informatics engenders many important legal and regulatory questions.

To consider ethical issues in health care informatics is to explore a significant intersection among several professions—health care informatics per se, health care delivery and administration, applied computing and systems engineering, and ethics—each of which constitutes a vast field of inquiry. Fortunately, growing interest in bioethics and computation-related ethics has produced a starting point for such exploration. An initial ensemble of guiding principles, or ethical criteria, has emerged to orient decision making in health care informatics. These criteria are of practical utility to health informatics, and often have broader implications for all of biomedical informatics.

10.2 Health-Informatics Applications: Appropriate Use, Users, and Contexts

Application of computer-based technologies in the health professions can build on previous experience in adopting other devices, tools, and

methods. Before clinicians perform most health-related interventions (e.g., diagnostic testing, prescription of medication, surgical and other therapeutic procedures), they generally evaluate appropriate evidence, standards, available technologies, presuppositions, and values. Indeed, the very evolution of the health professions entails the evolution of evidence, of standards, of available technologies, of presuppositions, and of values.

To answer the clinical question, “What should be done in this case?” one must pay attention to a number of subsidiary questions, such as:

1. What is the problem?
2. What resources are available and what am I competent to do?
3. What will maintain or improve this patient’s care?
4. What will otherwise produce the most desirable results (e.g., in public health)?
5. How strong are my beliefs in the accuracy of my answers to questions 1 through 4, above.

Similar considerations determine the appropriate use of informatics tools.

10.2.1 The Standard View of Appropriate Use

Excitement and enthusiasm often accompany initial use of new tools in clinical settings. Negative emotions are also common (Sittig et al. 2005). Based on the uncertainties that surround any new technology, scientific evidence counsels caution and prudence. As in other clinical areas, evidence and reason determine the appropriate level of caution. For instance, there is considerable evidence that electronic laboratory information systems improve access to clinical data when compared with manual, paper-based test-result distribution methods. To the extent that such systems improve care at an acceptable cost in time and money, there is an obligation to use computers to store and retrieve clinical laboratory results. There is a small but growing body of evidence that existing **clinical expert systems** can improve patient care in a small number of practice environments at an acceptable cost in

time and money (Kuperman and Gibson 2003). Nevertheless, such systems cannot yet uniformly improve care in typical, more general practice settings, at least not without careful attention to the full range of managerial as well as technical issues affecting the particular care delivery setting in which they are used (Kaplan and Harris-Salamone 2009; Holroyd-Leduc et al. 2011; Shih et al. 2011).

Clinical expert systems (see Chap. 22) attempt to provide decision support for diagnosis, therapy, and/or prognosis in a more detailed and sophisticated manner than do simple reminder systems (Duda and Shortliffe 1983). A necessary adjunction of expert systems – creation and maintenance of their related knowledge bases – still involves leading-edge research and development. One must recognize that humans for the most part remain superior to electronic systems in understanding patients and their problems, in efficiently interacting with patients to ascertain pertinent past history and current symptoms across the spectrum of clinical practice, in the interpretation and representation of data, and in clinical synthesis. Humans might not always hold the upper hand in these tasks, and claims of their superiority must continually be tested empirically.

What has been called the “standard view” of computer-assisted clinical diagnosis (Miller 1990; cf. Friedman 2009) holds in part that human cognitive processes, being more suited to the complex task of diagnosis than machine intelligence, should not be overridden or trumped by computers. The standard view states that when adequate (and even exemplary) decision-support tools are developed, they should be viewed and used as supplementary and subservient to human clinical judgment. Quite literally: they *support* decisions; they do not make them. Progress should be measured in terms of whether clinicians using a CDS tool perform better on specific tasks than the same clinicians without the tool (Miller 1990; cf. Friedman 2009). These tools should assume subservient roles because the clinician caring for the patient knows and understands the patient’s situation and can make compassionate judgments better

than computer programs. Furthermore, clinicians, and not machine algorithms, are the entities whom the state licenses, and specialty boards accredit, to practice medicine, surgery, nursing, pharmacy, and other health-related activities.

Corollaries of the standard view are that (1) practitioners have an obligation to use any computer-based tool responsibly, through adequate user training and by developing an understanding of the system’s abilities and limitations; and (2) practitioners must not abrogate their clinical judgment reflexively when using computer-based decision aids. Because the skills required for diagnosis are in many respects different from those required for the acquisition, storage, and retrieval of laboratory data, there is no contradiction in urging extensive use of electronic health records (as became U.S. policy under the HITECH act of 2009, discussed in Chap. 27), and, for the time being, cautious deployment of expert diagnostic decision-support tools (i.e., not permitting their use in settings in which knowledgeable clinicians cannot immediately override faulty advice).

The standard view addresses one aspect of the question, “How and when should computers be used in clinical practice?” by capturing important moral intuitions about error avoidance and evolving standards. Error avoidance and the benefits that follow from it shape the obligations of practitioners. In computer-software use, as in all other areas of clinical practice, good intentions alone are insufficient to insulate recklessness from culpability. Thus, the standard view may be seen as a tool for both error avoidance and ethically optimized action.

Ethical software use should be evaluated against a broad background of evidence for actions that produce favorable outcomes. Because informatics is a science in extraordinary ferment, system improvements and evidence of such improvements are constantly emerging. Clinicians have an obligation to be familiar with this evidence after attaining minimal acceptable levels of familiarity with informatics in general and with the clinical systems they use in particular (Fig. 10.1).




Fig. 10.1 The U.S. Department of Veterans Affairs has developed “Veterans Health Information Systems and Technology Architecture” (VistA), the largest electronic health record system in the United States. This fictitious

screen shot demonstrates some of the system’s functions and utilities (Credit: Courtesy of U.S. Department of Veterans Affairs, Veterans Health Administration Office of Informatics and Analytics)

10.2.2 Appropriate Users and Educational Standards

Efficient and effective use of health care informatics systems requires prior system evaluations demonstrating utility, then education and training of new users, monitoring of experience, and appropriate, timely updating. Indeed, such requirements resemble those for other tools used in health care and in other domains. Inadequate preparation in the use of tools is an invitation to catastrophe. When the stakes are high and the domain large and complex—as is the case in the health professions—education and training take on moral significance.

Who should use a health care-related computer application? Consider expert decision-support systems as an example. An early paper

on ethical issues in informatics noted that potential users of such systems include physicians, nurses, physicians’ assistants, paramedical personnel, students of the health sciences, patients, and insurance and government evaluators (Miller et al. 1985). Are members of all these groups appropriate users? One cannot answer the question until one precisely specifies the intended use for the system (i.e., the exact clinical questions the system will address). The appropriate level of training must be correlated with the question at hand. At one end of an appropriate-use spectrum, we can posit that medical and nursing students should employ decision-support systems for educational purposes; this assertion is relatively free of controversy once it has been verified that such tools convey accurately a sufficient quantity and quality of educational content. But

it is less clear that patients, administrators, or managed-care gatekeepers, for example, should use expert decision-support systems for assistance in making diagnoses, in selecting therapies, or in evaluating the appropriateness of health professionals' actions. To the extent that some systems present general medical advice in generally understandable but sufficiently nuanced formats, such as was once the case with Dr. Benjamin Spock's 1950s era print-based child-care primer, one might condone system use by laypersons. There are additional legal concerns related to negligence and product liability, however, when health-related products are sold directly to patients rather than to licensed practitioners, and when such products give patient-specific counsel rather than general clinical advice (Miller et al. 1985).

Suitable use of a software program that helps a user to suggest diagnoses, to select therapies, or to render prognoses must be carry of goals and best practices for achieving those goals, including consideration of the characteristics and requirements of individual patients. For example, the multiple, interconnected inferential strategies required for arriving at an accurate diagnosis depend on knowledge of facts; experience with procedures; and familiarity with human behavior, motivation, and values. **Diagnosis** is a process rather than an event (Miller 1990), so even well-validated diagnostic systems must be used appropriately in the overall context of patient care.

To use a **diagnostic decision-support system** (Chap. 22), a clinician must be able to recognize when the computer program has erred, and, when it is accurate, what the output means and how it should be interpreted. This ability requires knowledge of both the diagnostic sciences and the software applications, and the strengths and limitations of each. After assigning a diagnostic label, the clinician must communicate the diagnosis, prognosis, and implications to a patient, and must do so in ways both appropriate to the patient's educational background and conducive to future treatment goals. It is not enough to be able to tell patients that they have cancer, human immunodeficiency virus (HIV), diabetes, or heart disease and then simply hand over a prescription.

The care provider must also offer context when available, comfort when needed, and hope as appropriate. The reason many jurisdictions have required pretest and posttest HIV counseling, for instance, is not to vex busy health professionals but rather to ensure that comprehensive, high-quality care—rather than mere diagnostic labeling—has been delivered.

This discussion points to the following set of ethical principles for appropriate use of decision-support systems:

1. A computer program should be used in clinical practice only after appropriate evaluation of its efficacy and the documentation that it performs its intended task at an acceptable cost in time and money.
2. Users of most clinical systems should be health professionals who are qualified to address the question at hand on the basis of their licensure, clinical training, and experience. Software systems should be used to augment or supplement, rather than to replace or supplant, such individuals' decision making.
3. All uses of informatics tools, especially in patient care, should be preceded by adequate training and instruction, which should include review of applicable product evaluations.

Such principles and claims should be viewed as analogous to other standards or rules in clinical medicine and nursing.

10.2.3 Obligations and Standards for System Developers and Maintainers

Users of clinical programs must rely on the work of other people who are often far removed from the context of use. As with all complex technologies, users depend on the developers and maintainers of a system and must trust evaluators who have validated a system for clinical use. Health care software applications are among the most complex tools in the technological armamentarium. Although this complexity imposes certain obligations on end users, it also commits a system's developers, designers, and maintainers to

adhere to reasonable standards and, indeed, to acknowledge their moral responsibility for doing so.

10.2.3.1 Ethics, Standards, and Scientific Progress

The very idea of a **standard of care** embodies a number of complex assumptions linking ethics, evidence, outcomes, and professional training. To say that nurses or physicians must adhere to a standard is to say, in part, that they ought not stray from procedures previously shown or generally believed to work better than alternatives. The difficulty lies in how to determine if a procedure or device “works better” than another. Such determinations in the health sciences constitute progress, and provide evidence that we now know more. Criteria for weighing such evidence, albeit short of proof in most cases, are applied. For example, evidence from well-designed randomized controlled trials merits greater trust than evidence derived from uncontrolled retrospective studies (see Chap. 11). Typically, verification by independent investigators must occur before placing the most recent study results into common practice.

People who develop, maintain, and sell health care computing systems and their components have obligations that parallel those of system users. These obligations include holding patient care as the foremost value. The duty to limit or prevent harm to patients applies to system developers as well as to practitioners. Although this principle is easy to suggest and, generally, to defend, it invites subtle, and sometimes overt, resistance from people for whom profit or fame are primary motivators. (This is of course also true for other medical devices, processes and industries.) To be sure, quests for fame and fortune often produce good outcomes and improved care, at least eventually. Even so, some approaches fail to take into account the role of intention as a moral criterion.

In medicine, nursing, and psychology, a number of models of the **professional–patient relationship** place trust and advocacy at the apex of a hierarchy of values. Such a stance cannot be maintained if goals and intentions other than patient well-being are (generally) assigned

primacy. The same principles apply to those who produce and attend to health care information systems. Because these systems are health care systems—and are not devices for accounting, entertainment, real estate, and so on—and because system under performance can cause pain, disability, illness, and death, it is essential that the threads of trust run throughout the fabric of clinical system design and maintenance.

System purchasers, users, and patients must rely upon developers and maintainers to recognize the potentially grave consequences of errors or carelessness, trust them to care about the uses to which the systems will be put, and rely upon them to value the reduced suffering of other people at least as much as they value their own personal gain. This reliance emphatically does not entail that system designers and maintainers are blameworthy or unethical if they hope and strive to profit from their diligence, creativity, and effort. Rather, it implies that no amount of financial benefit for a designer can counterbalance bad outcomes or ill consequences that result from recklessness, avarice, or inattention to the needs of clinicians and their patients. Purchasers and users should require demonstrations that systems are worthy of such trust and reliance before placing patients at risk, and that safeguards (human and mechanical) are in place to detect, alert, and rectify situations in which systems underperform.

Quality standards should stimulate scientific progress and innovation while safeguarding against system error and abuse. These goals might seem incompatible, but they are not. Let us postulate a standard that requires timely updating and testing of knowledge bases that are used by decision-support systems. To the extent that database accuracy is needed to maximize the accuracy of inferential engines, it is trivially clear how such a standard will help to prevent or reduce decision-support mistakes. Furthermore, the standard should be seen to foster progress and innovation in the same way that any insistence on best possible accuracy helps to protect scientists and clinicians from pursuing false leads, or wasting time in testing poorly wrought hypotheses. It will not do for database maintainers to insist that they are busy doing the more productive or

scientifically stimulating work of improving knowledge representation, say, or database design. Although such tasks are important, they do not supplant the tasks of updating and testing tools in their current configuration or structure. Put differently, scientific and technical standards are perfectly able to stimulate progress while taking a cautious or even conservative stance toward permissible risk in patient care.

This approach has been described as **progressive caution**. “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.” (Goodman 1998b).

A more recent concern, with both ethical and legal implications, is the responsibility of software developers to design and implement software programs that cannot easily be hacked by malicious code writers. This concern goes beyond privacy and confidentiality issues (discussed below), and includes the possibility that medical devices with embedded software might be nefariously “reprogrammed” in a manner that might cause harm to patients. (See, for example, Robertson 2011). A more detailed discussion of this topic appears under the Sect. 10.5 below.

10.2.3.2 System Evaluation as an Ethical Imperative

Any move toward “best practices” in biomedical informatics is shallow and feckless if it does not include a way to measure whether a system performs as intended. This and related measurements provide the ground for quality control and, as such, are the obligations of system developers, maintainers, users, administrators, and perhaps other players (see Chap. 11).

Medical computing is not merely about medicine or computing. It is about the introduction of new tools into environments with established social norms and practices. The effects of computing systems in health care are subject to analysis not only of accuracy and performance but of acceptance by users, of consequences for social and professional interaction, and of the context of use. We suggest that system evaluation can illuminate social and ethical issues in medical computing, and in so

doing improve patient care. That being the case, there is an ethical imperative for such evaluation (Anderson and Aydin 1998).

To give a flavor of how a comprehensive evaluation program can ethically optimize implementation and use of an informatics system, consider these ten criteria for system scrutiny (Anderson and Aydin 1994):

1. Does the system work as designed?
2. Is it used as anticipated?
3. Does it produce the desired results?
4. Does it work better than the procedures it replaced?
5. Is it cost effective?
6. How well have individuals been trained to use it?
7. What are the anticipated long-term effects on how organizational units interact?
8. What are the long-term effects on the delivery of medical care?
9. Will the system have an impact on control in the organization?
10. To what extent do effects depend on practice setting?

Another way to make this important point is by emphasizing that *people* use computer systems. Even the finest system might be misused, misunderstood, or mistakenly allowed to alter or erode previously productive human relationships. Evaluation of health information systems in their contexts of use should be taken as a moral imperative. Such evaluations require consideration of a broader conceptualization of “what works best” and must look toward improving the overall health care delivery system rather than only that system’s technologically based components. These higher goals entail the creation of a corresponding mechanism for ensuring institutional oversight and responsibility (Miller and Gardner 1997a, b).

10.3 Privacy, Confidentiality, and Data Sharing

Some of the greatest challenges of the Information Age arise from placing computer applications in health care settings while upholding traditional

principles and values. One challenge involves balancing two competing values: (1) free access to information, and (2) protection of patients' privacy and confidentiality.

Only computers can efficiently manage the now-vast amount of information generated during clinical encounters and other health care transactions (see Chap. 2); at least in principle, such information should be easily available to health professionals and others involved in the administration of the care-delivery system, so that they can provide effective, efficient care for patients. Yet, making this information readily available creates greater opportunities for inappropriate access. Such access may be available to curious health care workers who do not need the information to fulfill job-related responsibilities, and, even more worrisome, to other people who might use the information to harm patients physically, emotionally, or financially. Clinical system administrators must balance the goals of protecting confidentiality by restricting use of computer systems and improving care by assuring the integrity and availability of data. These objectives are not incompatible, but there are trade-offs that cannot be avoided.

10.3.1 Foundations of Health Privacy and Confidentiality

Privacy and confidentiality are necessary for people to evolve and mature as individuals, to form relationships, and to serve as functioning members of society. Imagine what would happen if the local newspaper or gossip blog produced a daily report detailing everyone's actions, meetings, and conversations. It is not that most people have terrible secrets to hide but rather that the concepts of solitude, intimacy, and the desire to be left alone make no sense without the expectation that at least some of our actions and utterances will be kept private or held in confidence among a limited set of persons.

The "average" sentiment about the appropriate sphere of private vs. public may vary considerably from culture to culture, and even from generation to generation within any particular

culture; and it may differ widely among persons within a culture or generation, and evolve for any particular person over a lifetime. Even the "born digital" generation, for which Facebook and other social networking is a fact of everyday life, has its boundaries (Palfrey and Gasser 2010).

The terms privacy and confidentiality are not synonymous. As commonly used, "privacy" generally applies to people, including their desire not to suffer eavesdropping, whereas "confidentiality" is best applied to information. One way to think of the difference is as follows. If someone follows you and spies on you entering an AIDS (acquired immunodeficiency syndrome) clinic, your privacy is violated; if someone sneaks into the clinic without observing you in person and looks at your health care record, your record's confidentiality is breached. In discussions of the electronic health record, the term privacy may also refer to individuals' desire to restrict the disclosure of personal data (National Research Council 1997).

There are several important reasons to protect privacy and confidentiality. One is that privacy and confidentiality are widely regarded as rights of all people, and such protections help to accord them respect. On this account, people do not need to provide a justification for limiting access to their identifiable health data; privacy and confidentiality are entitlements that a person does not need to earn, to argue for, or to defend. Another reason is more practical: protecting privacy and confidentiality benefits both individuals and society. Patients who know that their identifiable health care data will not be shared inappropriately are more comfortable disclosing those data to clinicians. This trust is vital for the successful physician-patient, nurse-patient, or psychologist-patient relationship, and it helps practitioners to do their jobs.

Privacy and confidentiality protections also benefit public health. People who fear disclosure of personal information are less likely to seek out professional assistance, increasing the risks that contagion will be spread and maladies will go untreated. In addition, and sadly, people still suffer discrimination, bias, and stigma when certain health data do fall into the wrong hands. Financial

harm may occur if insurers are given unlimited access to family members' records, or access to patient data, because some insurers might be tempted to increase the price of insurance for individuals at higher risk of illness or discriminate in other ways if such price differentiation were forbidden by law.

The ancient idea that physicians should hold health care information in confidence is therefore applicable whether the data are written on paper or embedded in silicon. The obligations to protect privacy and to keep confidences fall to system designers and maintainers, to administrators, and, ultimately, to the physicians, nurses, and others who elicit the information in the first place. The upshot for all of them is this: protection of privacy and confidentiality is not an option, a favor, or a helping hand offered to patients with embarrassing health problems; it is a duty, regardless of the malady or the medium in which information about it is stored.

Some sound clinical practice and public health traditions run counter to the idea of absolute confidentiality. When a patient is hospitalized, it is expected that all appropriate (and no inappropriate) employees or affiliates of the institution—primary physicians, consultants, nurses, therapists, and technicians—will have access to the patient's medical records, when it is in the interest of the patient's care to do so. In most communities of the United States, the contacts of patients who have active tuberculosis or certain sexually transmitted diseases are routinely identified and contacted by public health officials so that the contacts may receive proper medical attention. Such disclosures serve the public interest and are and should be legal because they decrease the likelihood that more widespread harm to other individuals might occur through transmission of an infection unknowingly.

A separate but important public health consideration (discussed in more detail below) involves the ability of health care researchers to anonymously pool data (i.e., pool by removing individual persons' identifying information) from patient cases that meet specified conditions to determine the natural history of the disease and the effects of various treatments. Examples of

benefits from such pooled data analyses range from the ongoing results generated by regional collaborative chemotherapy trials to the discovery, nearly four decades ago, of the appropriateness of shorter lengths of stay for patients with myocardial infarction (McNeer et al. 1975). Most recently, the need for robust **syndromic surveillance** has been asserted as necessary for adequate bioterrorism preparedness, as well as for earlier detection of naturally occurring disease outbreaks (see Chap. 16).

10.3.2 Electronic Clinical and Research Data

Access to electronic patient records holds extraordinary promise for clinicians and for other people who need timely, accurate patient data (see Chap. 12). Institutions that do not deploy electronic health record systems are falling behind, a position that may soon become blameworthy. Failure to use such systems may also disqualify institutions for reimbursements from public and private insurance, making it effectively an organizational death sentence. Conversely, systems that make it easy for clinicians to access data also make it easier for people in general to access the data, and electronic systems generally magnify number of persons whose information becomes available when a system security breach occurs. Some would consider failure to prevent inappropriate access as at least as blameworthy as failure to provide adequate and appropriate access.

Nonetheless, there is no contradiction between the obligation to maintain a certain standard of care (in this case, regarding minimal levels of computer use) and ensuring that such a technical standard does not imperil the rights of patients. Threats to confidentiality and privacy are fairly well known. They include economic abuses, or discrimination by third-party payers, employers, and others who take advantage of the burgeoning market in health data; insider abuse, or record snooping by hospital or clinic workers who are not directly involved in a patient's care but examine a record out of curiosity, for blackmail, and so

on; and malevolent hackers, or people who, via networks or other means, copy, delete, or alter confidential information (National Research Council 1997). Identity theft to commit insurance or other financial fraud could now be added to the list. Moreover, widespread dissemination of information throughout the health care system often occurs without explicit patient consent. Health care providers, third-party payers, managers of pharmaceutical benefits programs, equipment suppliers, and oversight organizations collect large amounts of patient-identifiable health information for use in managing care, conducting quality and utilization reviews, processing claims, combating fraud, and analyzing markets for health products and services (National Research Council 1997).

The proper approach to such challenges is one that will ensure both that appropriate clinicians and other people have rapid, easy access to patient records and that others do not have access. Is that a contradictory burden? No. Is it easy to achieve both? No. There are many ways to restrict inappropriate access to electronic records, but all come with a cost. Sometimes the cost is explicit, as when it comes in the form of additional computer software and hardware; sometimes it is implicit, as when procedures are required that increase the time commitment by system users.

One way to view the landscape of protective measures is to divide it into technological methods and institutional or policy approaches (Alpert 1998):

10.3.2.1 Technological Methods

Computers can provide the means for maximizing their own security, including authenticating system users with passwords, tokens or biometrics, to make sure that they are who they say they are; using access controls to prohibit people without a professional need from accessing particular health information within a system; and using audit trails, or logs, of people who do inspect confidential records so that automated security auditors, authorized facility administrators, as well as patients can review who accessed what. Encryption can protect data in transit and at

rest (in storage). These technical means are complemented by protecting the elements of the electronic infrastructure with physical barriers when operations allow it. Auditing works best when appropriately severe punishments are widely known to be policy, and when policy breaches are uniformly punished in a semi-public manner.

10.3.2.2 Policy Approaches

In its landmark report, the National Research Council (1997) recommended that hospitals and other health care organizations create security and confidentiality committees and establish education and training programs. These recommendations parallel an approach that had worked well elsewhere in hospitals for matters ranging from infection control to bioethics. The Health Insurance Portability and Accountability Act (HIPAA) requires the appointment of privacy and security officials, special policies, and the training of health care workforce members who have access to health information systems.

Such measures are all the more important when health data are accessible through networks. The rapid growth of **integrated delivery networks (IDNs)** (see Chap. 14) and Health Information Exchanges, for example, illustrate the need not to view health data as a well into which one drops a bucket but rather as an irrigation system that makes its contents available over a broad—sometimes an extremely broad—area. It is not yet clear whether privacy and confidentiality protections that are appropriate in hospitals will be fully effective in a ubiquitously networked environment, but it is a start. System developers, users, and administrators are obliged to identify appropriate measures in light of the particular risks associated with a given implementation. There is no excuse for failing to make this a top priority throughout the data storage and sharing environment.

10.3.2.3 Electronic Data and Human Subjects Research

The use of patient information for **clinical research** and for quality assessment raises interesting ethical challenges. The presumption of a right to confidentiality seems to include the idea

that patient records are inextricably linked to patient names or to other identifying data. In an optimal environment, then, patients can monitor who is looking at their records. But if all unique identifiers have been stripped from the records, is there any sense in talking about confidentiality?

The benefits to **public health** loom large in considering record-based research (Chap. 16). A valuable benefit of the electronic health record is the ability to access vast numbers of patient records to estimate the incidence and prevalence of various maladies, to track the efficacy of clinical interventions, and to plan efficient resource allocation (see Chap. 16). Such research and planning would, however, impose onerous or intractable burdens if informed, or valid consent had to be obtained from every patient whose record was represented in the sample. Using confidentiality to impede or forbid such research fails to benefit patients at the same time it sacrifices potentially beneficial scientific investigations.

A more practical course is to establish safeguards that better balance the ethical obligations to privacy and confidentiality against the social goals of public health and systemic efficiency. This balancing can be pursued via a number of paths. The first is to establish mechanisms to **anonymize** the information in individual records or to decouple the data contained in the records from any unique patient identifier. This task is not always straightforward; it can be remarkably difficult to anonymize data such that, when coupled with other data sets, the individuals are not at risk of re-identification. A relatively rare disease diagnosis coupled with demographic data such as age and gender, or geographic data such as a postal code, may act as a surrogate unique identifier; that is, detailed information can in combination serve as a data fingerprint that picks out an individual patient even though the patient's name, Social Security number, or other (official) unique identifiers have been removed from the record. Challenges and opportunities related to de-identifying and re-identifying data are among the most interesting, difficult and important in all health computing (Atreya et al. 2013; Benitez and Malin 2010; Malin and Sweeney 2004; Malin et al. 2011; Sweeney 1997; Tamersoy et al. 2012).

Such challenges point to a second means of balancing ethical goals in the context of database research: the use of institutional panels, such as **medical record committees** or **institutional review boards**. Submission of database research to appropriate institutional scrutiny is one way to make the best use of more or less anonymous electronic patient data. Competent panel members should be educated in the research potential of electronic health records, as well as in ethical issues in epidemiology and public health. Scrutiny by such committees can also give appropriate weight to competing ethical concerns in the context of internal research for quality control, outcomes monitoring, and so on (Goodman 1998b; Miller and Gardner 1997a, b).

10.3.2.4 Challenges in Bioinformatics

Safeguards are increasingly likely to be challenged as genetic information makes its way into the health care record (see Chaps. 24 and 25). The risks of bias, discrimination, and social stigma increase dramatically as **genetic data** become available to clinicians and investigators. Indeed, genetic information “goes beyond the ordinary varieties of medical information in its predictive value” (Macklin 1992). Genetic data also may be valuable to people predicting outcomes, allocating resources, and the like. In addition, genetic data are rarely associated with only a single person; they may provide information about relatives, including relatives who do not want to know about their genetic risk factors or potential maladies, as well as relatives who would love dearly to know more about their kin's genome. There is still much work to be done in sorting out and addressing the ethical issues related to electronic storage, sharing, and retrieval of genetic data (Goodman 1996).

Bioinformatics or **computational biology** provides an exciting ensemble of new tools to increase our knowledge of genetics, genetic diseases, and public health. Use of these tools is accompanied by responsibilities to attend to the ethical issues raised by new methods, applications, and consequences (Goodman and Cava 2008). Identifying and analyzing these issues are among the key tasks of those who work at the

intersection of ethics and health information technology. The future of genetics and genomics is utterly computational, with data storage and analysis posing some of greatest financial and scientific challenges. For instance:

- How, to what extent, and by whom should genomic databases be used for clinical or public health decision support?
- Are special rules needed to govern the study of information in digital genetic repositories (or are current human subjects research protection rules adequate)?
- Does data mining software present new challenges when applied to human genetic information?
- What policies are required to guide and inform the communication of patient-specific and incidental findings?
- Are special protections and precautions needed to address and transmit findings about population subgroups?

It might be that the tools and uses of computational biology will eventually offer ethical challenges—and opportunities—as important, interesting and compelling as any technology in the history of the health sciences. Significantly, this underscores the importance of arguments to the effect that attention to ethics must accompany attention to science. Victories of health science research and development will be undermined by any failures to address corresponding ethical challenges. We must strive to identify, analyze, and resolve or mitigate important ethical issues.

10.4 Social Challenges and Ethical Obligations

The expansion of **evidence-based medicine** and, in the United States, of managed care (now sometimes called **accountable care** since the passage of health reform legislation in 2010; see Chap. 27) places a high premium on the tools of health informatics. The need for data on clinical outcomes is driven by a number of important social and scientific factors. Perhaps the most important among these factors is the increasing unwillingness

of governments and insurers to pay for interventions and therapies that do not work or that do not work well enough to justify their cost.

Health informatics helps clinicians, administrators, third-party payers, governments, researchers, and other parties to collect, store, retrieve, analyze, and scrutinize vast amounts of data—though the task of documenting this is itself a matter of research on what has come to be called “meaningful use.” The functions of health informatics might be undertaken not for the sake of any individual patient but rather for cost analysis and review, quality assessment, scientific research, and so forth. These functions are critical, and if computers can improve their quality or accuracy, then so much the better.

Challenges arise when intelligent applications are mistaken for decision-making surrogates or when institutional or public policy recommends or favors computer output over human cognition. This may be seen as a question or issue arising under the rubric of “appropriate uses and users.” That is, by whom, when, and under what constraints may we elicit and invoke computational analysis in shaping or applying public policy? The question whether an individual physician or multispecialty group, say, should be hired or retained or reimbursed or rewarded is information-intensive. The question that follows, however, is the key one: How should the decision-making skills of human and machine be used, and balanced (cf. Glaser 2010)?

10.4.1 Vendor Interactions

Motivated if not inspired by both technological necessity and financial opportunity, humble private practices and sprawling medical centers have—or should have—begun the transition from a paper patient record to an electronic one. The need to make such a transition is not in dispute: paper (and handwriting) are hard to store, find, read and analyze. **Electronic Health Records** (EHR) are not, or should not be. While there are important debates about the speed of the transition and regarding software quality, usability and ability to protect patient safety, it is

widely agreed that the recording and storage of health information must be electronic.

Public policy has attempted to overcome some of the reluctance to make the change because of financial concerns. Notably, the U.S. Health Information Technology for Economic and Clinical Health (HITECH) Act, a part of the American Recovery and Reinvestment Act of 2009 (Blumenthal 2010), authorizes some \$27 billion in incentives for EHR adoption. These incentives help address but do not eliminate financial concerns in that they offset only some of the cost of converting to an e-system. Still, while a number of companies had previously found opportunity in developing hospital and other clinical information systems, HITECH accelerated the pace (see Chap. 27).

The firms that make and sell EHRs are not regulated in the same way as those that manufacture pharmaceutical products or medical devices (see Sect. 10.5.3). In an increasingly competitive environment, this has led to controversy about the nature of vendor interactions with the institutions that buy their products. An EHR system for a mid-sized hospital can cost upwards of \$100 million over time, including consulting services, hardware and training. It follows that it is reasonable to ask what values should guide such vendor interactions with clients, and whether they should be similar to or different from values that govern other free-market dealings.

While many or most contracts between vendors and hospitals are confidential, it has been reported that some HIT vendors require contract language that indemnifies system developers for personal injury claims or malpractice, even if the vendor is at fault; some vendors require system purchasers to agree not to disclose system errors except to the vendor (Koppel and Kreda 2009). Such provisions elicit concern to the extent they place or appear to place corporate interests ahead of patient safety and welfare. In this case, a working group chartered by AMIA, the society for informatics professionals (see Chap. 1), issued a report that provided guidance on a number of vendor interaction issues (Goodman et al. 2010). Importantly, the working group comprised industry representatives as well as scientists and other

academics. The group's recommendations included these:

- Contracts should not contain language that prevents system users, including clinicians and others, from using their best judgment about what actions are necessary to protect patient safety. This includes freedom to disclose system errors or flaws, whether introduced or caused by the vendor, the client, or a third party. Disclosures made in good faith should not constitute violations of HIT contracts. This recommendation neither entails nor requires the disclosure of trade secrets or of intellectual property....
- Because vendors and their customers share responsibility for patient safety, contract provisions should not attempt to circumvent fault and should recognize that both vendors and purchasers share responsibility for successful implementation. For example, vendors should not be absolved from harm resulting from system defects, poor design or usability, or hard-to-detect errors. Similarly, purchasers should not be absolved from harm resulting from inadequate training and education, inadequate resourcing, customization, or inappropriate use.

While some of the debates that led to those conclusions were about political economy (regulation vs. free enterprise) as much as ethics (right vs. wrong), the opportunity for rapprochement in the service of a patient-centered approach may be seen as an affirmation of the utility of an applied ethics process in the evolution of health information technology.

10.4.2 Informatics and Managed Care

Consider the utility of **prognostic scoring systems** that use physiologic and mortality data to compare new critical-care patients with thousands of previous patients (Knaus et al. 1991). Such systems allow hospitals to track the performance of their critical-care units by, say, comparing the previous year's outcomes to this year's or by comparing one hospital to another. If, for instance,




Fig. 10.2 “Severity adjusted daily data” in fictitious APACHE® Outcomes screen shot. Using prognostic scoring systems, clinicians in critical-care units can monitor events and interventions and administrators can manage staffing based on patient acuity. Clinicians can also use such systems to predict mortality, raising a number of

ethical issues. This image shows 10 CCU patients. For the second one in the leftmost column, for instance, the “acute physiology score” is 128; the risk of hospital mortality is 96 % and the risk of ICU mortality is 92 % (Credit: Courtesy of Cerner Corporation, with permission)

patients with a particular profile tend to survive longer than their predecessors, then it might be inferred that **critical care** has improved. Such scoring systems can be useful for internal research and for quality management (Fig. 10.2).

Now suppose that most previous patients with a particular physiologic profile have died in critical-care units; this information might be used to identify ways to improve care of such patients—or it might be used in support of arguments to contain costs by denying care to subsequent patients fitting the profile (since they are likely to die anyway).

An argument in support of such an application might be that decisions to withdraw or withhold care are often and customarily made on the basis of subjective and fragmented evidence; so it is preferable to make such decisions on the basis of objective data of the sort that otherwise underlie sound clinical practice. Such **outcomes data** are precisely what fuels the engines of managed care, wherein health professionals and institutions compete on the basis of cost and outcomes. Why should society, or a managed-care organization,

or an insurance company pay for critical care when seemingly objective evidence exists that such care will not be efficacious? Contrarily, consider the effect on future scientific insights of denying care to such patients. Scientific progress is often made by noticing that certain patients do better under certain circumstances, and investigation of such phenomena leads to better treatments. If all patients meeting certain criteria were denied therapy on the basis of a predictive tool, it would become a self-fulfilling prophecy for a much longer time that all such patients would not do well (Miller 1997).

Now consider use of a decision-support system to evaluate, review, or challenge decisions by human clinicians; indeed, imagine an insurance company using a diagnostic expert system to determine whether a physician should be reimbursed for a particular procedure. If the expert system has a track record for accuracy and reliability, and if the system “disagrees” with the human’s diagnosis or treatment plan, then the insurance company can contend that reimbursement for the procedure would be a mistake. Why

pay for a procedure that is not indicated, at least according to a computational analysis?

In the two examples just offered (a prognostic scoring system is used to justify termination of treatment to conserve resources, and a diagnostic expert system is used to deny a physician reimbursement for procedures deemed inappropriate), there seems to be justification for adhering to the computer output. There are, however, three reasons why it is problematic to rely exclusively on clinical computer programs to guide policy or practice in these ways:

1. As we argued earlier with the standard view of computational diagnosis (and, by easy extension, prognosis), human cognition is, at least for a while longer, still superior to machine intelligence. Moreover, the act of rendering a diagnosis or prognosis is not merely a statistical or computational operation performed on un-interpreted data. Rather, identifying a malady and predicting its course requires understanding a complex ensemble of causal relations, interactions among a large number of variables, and having a store of salient background knowledge—considerations that have thus far failed to be grasped, assessed, and effectively blended into decisions made by computer programs.
2. Decisions about whether to treat a given patient are often value laden and must be made relative to treatment goals. In other words, it might be that a treatment will improve the quality of life but not extend life, or vice versa (Youngner 1988). Whether such treatment is appropriate cannot be determined scientifically or statistically (Brody 1989). The decisions ultimately depend on human preferences—those of the provider or, even more importantly, the patient.
3. Applying computational operations on aggregate data to individual patients runs the risk of including individuals in groups they resemble but to which they do not actually belong. Of course, human clinicians run this risk all the time—the challenge of inferring correctly that an individual is a member of a set, group, or class is one of the oldest problems in logic and in the philosophy of science. The point is that

computers have not solved this problem, yet, and allowing policy to be guided by simple or unanalyzed correlations constitutes a conceptual error.

The idea is not that diagnostic or prognostic computers are always wrong—we know that they are not—but rather there are numerous instances in which we do not know whether they are right. It is one thing to allow aggregate data to guide policy; doing so is just using scientific evidence to maximize good outcomes. But it is altogether different to require that a policy disallow individual **clinical judgment** and expertise.

Informatics can contribute in many ways to health care reform. Indeed, computer-based tools can help to illuminate ways to reduce costs, to optimize clinical outcomes, and to improve care. Scientific research, quality assessment, and the like are, for the most part, no longer possible without computers. But it does not follow that the insights from such research apply in all instances to the myriad variety of actual clinical cases at which competent human clinicians excel.

10.4.3 Effects of Informatics on Traditional Relationships

Ill patients are often scared and vulnerable. Treating illness, easing fear, and respecting vulnerability are among the core obligations of physicians, nurses, and other clinicians. Health informatics has the potential at some future time to complement these traditional duties and the relationships that they entail. We have pointed out that medical decisions are shaped by nonscientific considerations. This point is important when we assess the effects of informatics on human relationships. Thus:

The practice of medicine or nursing is not exclusively and clearly scientific, statistical, or procedural, and hence is not, so far, computationally tractable. This is not to make a hoary appeal to the “art and science” of medicine; it is to say that the science is in many contexts inadequate or inapplicable: Many clinical decisions are not exclusively medical—they have social, personal, ethical, psychological, financial, familial, legal, and other components; even art might play a role. (Miller and Goodman 1998)

10.4.3.1 Professional–Patient Relationships

If computers, databases, and networks can improve physician–patient or nurse–patient relationships, perhaps by improving communication, then we shall have achieved a happy result. If reliance on computers impedes the abilities of health professionals to establish trust and to communicate compassionately, however, or further contributes to the dehumanization of patients (Shortliffe 1993, 1994), then we may have paid too dearly for our use of these machines.

Suppose that a physician uses a decision-support system to test a diagnostic hypothesis or to generate differential diagnoses, and suppose further that a decision to order a particular test or treatment is based on that system’s output. A physician who is not able to articulate the proper role of computational support in his decision to treat or test will risk alienating those patients who, for one reason or another, will be disappointed, angered, or confused by the use of computers in their care. To be sure, the physician might just withhold this information from patients, but such deception carries its own threats to trust in the relationship.

Patients are not completely ignorant about the processes that constitute human decision making. What they do understand, however, may be subverted when their doctors and nurses use machines to assist delicate cognitive functions. We must ask whether patients should be told the accuracy rate of decision support automata—when they have yet to be given comparable data for humans. Would such knowledge improve the informed-consent process, or would it “constitute another befuddling ratio that inspires doubt more than it informs rationality?” (Miller and Goodman 1998).

To raise such questions is consistent with promoting the responsible use of computers in clinical practice. The question whether computer use will alienate patients is an empirical one; it is a question for which, despite many initial studies, we lack conclusive data to answer. (For example, we cannot yet state definitively whether all categories of patients will respond well to all specific types of e-mail messages from their doctors. Nevertheless, as a moral principle discussed above, one should not convey a new diagnosis of a malignancy via

email.) To address the question now anticipates potential future problems. We must ensure that the exciting potential of health informatics is not subverted by our forgetting that the practice of medicine, nursing, and allied professions is deeply human and fundamentally intimate and personal.

10.4.3.2 Consumer Health Informatics

The growth of the World Wide Web and the commensurate evolution of clinical and health resources on the Internet also raise issues for professional–patient relationships. **Consumer health informatics**—technologies focused on patients as the primary users—makes vast amounts of information available to patients (see Chap. 17). There is also, however, misinformation—even outright falsehoods and quackery—posted on some sites. If physicians and nurses have not established relationships based on trust, the erosive potential of apparently authoritative Internet resources can be great. Physicians once accustomed to newspaper-inspired patient requests for drugs and treatments now face ever increasing demands that are informed by Web browsing. Consequently, the following issues will gain in ethical importance with each passing year:

- Peer review: How and by whom is the quality of a Web site to be evaluated? Who is responsible for the accuracy of information communicated to patients?
- Online consultations: There is no standard of care yet for online medical consultations. What risks do physicians and nurses run by giving advice to patients whom they have not met or examined in person? This question is especially important in the context of **telemedicine** or **remote-presence health care**, the use of video conferencing, image transmission, and other technologies that allow clinicians to evaluate and treat patients in other than face-to-face situations (see Chap. 18).
- Support groups: Internet support groups can provide succor and advice to the sick, but there is a chance that someone who might benefit from seeing a physician will not do so because of anecdotes and information otherwise attained. How should this problem be addressed?

That a resource is touted as worthwhile does not mean that it is. We lack evidence to illuminate the utility of consumer health informatics and its effects on professional–patient relationships. Such resources cannot be ignored given their ubiquity, and they often are useful for improving health. But we insist that here—as with decision support, appropriate use and users, evaluation, and privacy and confidentiality—there is an ethical imperative to proceed with caution. Informatics, like other health technologies, will thrive if our enthusiasm is open to greater evidence and is wed to deep reflection on human values.

10.4.3.3 Personal Health Records

At the same time as institutions have moved to computer-based health records systems, the tools available to individuals to keep their own health records have been making a similar transition. Electronic **personal health record** (PHR) systems, whether designed for use on a decoupled storage device or accessible over the Web, are now available from a rapidly expanding set of organizations (see Chap. 17) (Figs. 10.3 and 10.4).

PHRs provide a storage base for data once kept on paper (or in the patient’s head) and repeatedly extracted with each institutional encounter for inclusion in that entity’s records system, typically:

- Allergies, current medications
- Current health status and major health issues (if any)
- Major past health episodes and the condition of oneself and (sometimes) relatives
- Vaccinations, surgeries and other treatments

All these data can be kept on something simple (and un-networked) like a flash drive. It is becoming more common to store the data on a Web site, where PHR data can also be linked to other health information relevant to the person. The PHR data can also be linked to a health care provider institution’s records, to allow updating in both directions, or be free of any such tie. A flash drive can be forgotten or lost, whereas a Web site can be centrally updated and uniformly available via any properly authenticated device on the Internet.


Traditional insurers and health care providers are duty-bound by privacy laws and regulations to protect the information under their control. PHRs have a somewhat shakier set of protections



Fig. 10.3 Project HealthDesign is a program sponsored by the Robert Wood Johnson Foundation’s Pioneer Portfolio and intended to foster development of personal health records. Here is a barcode scanner that recognizes medication labels. Designed by researchers at the University of

Colorado at Denver, the “Colorado Care Tablet” allows elderly users to track prescriptions with such scanners and portable touch-screen tablets (Credit: Courtesy of Project HealthDesign (<http://projecthealthdesign.org>); Creative Commons Attribution 3.0 Unported License)

Fig. 10.4 A portable blood glucose communicator is part of the personal health record system developed by the T.R.U.E. Research Foundation of Washington, DC. The diabetes management application analyzes, summarizes, displays and makes individualized recommendations on nutritional data, physical activity data, prescribed medications, continuous blood glucose data, and self-reported emotional state (Credit: Courtesy of Project HealthDesign (<http://projecthealthdesign.org>); Creative Commons Attribution 3.0 Unported License)



given their relatively short history. The legal obligations of institutions that provide PHRs, but do not fully manage the content of those records nor their use, as well as the obligations (if any) of the individuals who “manage” their own health records, remain to be sorted out (Cushman et al. 2010).

PHRs are now commonly linked to so-called “personal health applications” (PHAs) which provide ways of moving beyond simple static storage of one’s medical history. Most provide some sort of primitive decision support, if only in linking to additional information about a particular disease or condition. Others include more ambitious decision-support functionality. All the concerns about the accuracy of Web-based information recur in this context, with concerns about the reliability of decision support added to that. Compounding concerns about accuracy are the inherent limitations of the “owner-operator”: If it can be difficult for trained health care providers to evaluate the quality of advice rendered by a decision support system, the challenges for patients will be commensurately greater.

Traditional health care institutions may see the PHR as a device for patient empowerment because it adds a way for persons to keep track of their own data; but they can also be used as a way

of preserving “loyalty” to a particular institution in the health care system. It has been proposed that PHRs be subject to standards allowing “interoperability”—in this case, easy movement from one type of PHR to another—to prevent leveraging it as an impediment to patients’ movements when they wish to change providers or other preferences change. Whether such standards will evolve enough to make it easy to move from one PHR to another remains to be seen, given economic incentives to impede patient movement (provided the patient is economically desirable in terms of insurance status or personal wealth).

It is also unclear whether PHRs will reach the majority of patients. PHRs and their associated applications may be compelling for persons who must manage for themselves or a dependent a chronic health condition with complex treatment regimes. Persons who deal with less complex current conditions or histories may prefer to leave records management to their providers. There is also a concern that PHRs may replicate the “digital divide” in the context of health care, exacerbating rather than reducing health disparities. That is, PHRs are likely to be differentially beneficial to persons with the income and education to make full use of them.

10.5 Legal and Regulatory Matters

The use of clinical computing systems in health care raises a number of interesting and important legal and regulatory questions.

10.5.1 Difference Between Law and Ethics

Ethical and legal issues often overlap. Ethical considerations apply in attempts to determine what is good or meritorious and which behaviors are desirable or correct in accordance with higher principles. Legal principles are generally derived from ethical ones but deal with the practical regulation of morality or behaviors and activities. Many legal principles deal with the inadequacies and imperfections in human nature and the less-than-ideal behaviors of individuals or groups. Ethics offers conceptual tools to evaluate and guide moral decision making. Laws directly tell us how to behave (or not to behave) under various specific circumstances and prescribe remedies or punishments for individuals who do not comply with the law. Historical precedent, matters of definition, issues related to detectability and enforceability, and evolution of new circumstances affect legal practices more than they influence ethical requirements.

10.5.2 Legal Issues in Biomedical Informatics

Prominent legal issues related to the use of software applications in clinical practice and in biomedical research include liability under tort law; potential use of computer applications as expert witnesses in the courtroom; legislation governing privacy and confidentiality; and copyrights, patents, and intellectual property issues.

10.5.2.1 Liability Under Tort Law

In the United States and in many other nations, principles of tort law govern situations in which harm or injuries result from the manufacture and

sale of goods and services (Miller et al. 1985). Because there are few, if any, U.S. legal precedents directly involving harm or injury to patients resulting from use of clinical software applications (as opposed to a small number of well-documented instances where software associated with medical devices has caused harm), the following discussion is hypothetical. The principles involved are, however, well established with voluminous legal precedents outside the realm of clinical software.

A key legal distinction is the difference between products and services. **Products** are physical objects, such as stethoscopes, that go through the processes of design, manufacture, distribution, sale, and subsequent use by purchasers. **Services** are intangible activities provided to consumers at a price by (presumably) qualified individuals.

The practice of clinical medicine has been deemed a service through well-established legal precedents. On the other hand, clinical software applications can be viewed as either goods (“products”) (software programs designed, tested, debugged, placed on DVDs or other media, and distributed physically to purchasers) or services (applications that present data or provide advice to practitioners engaged in a service such as delivering health care). There are few legal precedents to determine unequivocally how software will be viewed by the courts, and it is possible that clinical software programs will be treated as goods under some circumstances and as services under others. It might be the case that that software purchased and running in a private office to handle patient records or billing would be deemed a product, but the same software mounted on shared, centralized computers and accessed over the Internet (and billed on a monthly basis) would be offering a service.

Three ideas from tort law potentially apply to the clinical use of software systems:

(1) **Harm by intention**—when a person injures another using a product or service to cause the damage, (2) the **negligence theory**, and (3) **strict product liability** (Miller et al. 1985). Providers of goods and services are expected to uphold the standards of the

community in producing goods and delivering services. When individuals suffer harm due to substandard goods or services, they may sue the service providers or goods manufacturers to recover damages. **Malpractice** litigation in health care is based on negligence theory.

Because the law views delivery of health care as a service (provided by clinicians), it is clear that negligence theory will provide the minimum legal standard for clinicians who use software during the delivery of care. Patients who are harmed by clinical practices based on imperfect software applications may sue the health care providers for negligence or malpractice, just as patients may sue attending physicians who rely on the imperfect advice of a human consultant (Miller et al. 1985). Similarly, a patient might sue a practitioner who has not used a decision-support system when it can be shown that use of the decision-support system is part of the current standard of care, and that use of the program might have prevented the clinical error that occurred (Miller 1989). It is not clear whether the patients in such circumstances could also successfully sue the software manufacturers, as it is the responsibility of the licensed practitioner, and not of the software vendor, to uphold the standard of care in the community through exercising sound clinical judgment. Based on a successful malpractice suit against a clinician who used a clinical software system, it might be possible for the practitioner to sue the manufacturer or vendor for negligence in manufacturing a defective clinical software product, but cases of this sort have not yet been filed. If there were such suits, it might be difficult for a court to discriminate between instances of improper use of a blameless system and proper use of a less-than-perfect system.

In contrast to negligence, strict product liability applies only to harm caused by defective products and is not applicable to services. The primary purpose of strict product liability is to compensate the injured parties rather than to deter or punish negligent individuals (Miller et al. 1985). For strict product liability to apply, three conditions must be met:

1. The product must be purchased and used by an individual.
2. The purchaser must suffer physical harm as a result of a design or manufacturing defect in the product.
3. The product must be shown in court to be “unreasonably dangerous” in a manner that is the demonstrable cause of the purchaser’s injury.

Note that negligence theory allows for adverse outcomes. Even when care is delivered in a competent, caring, and compassionate manner, some patients with some illnesses will not do well. Negligence theory protects providers from being held responsible for all individuals who suffer bad outcomes. As long as the quality of care has met the prevailing standards, a practitioner should not be found liable in a malpractice case (Miller et al. 1985). Strict product liability, on the other hand, is not as forgiving or understanding.

No matter how good or exemplary a manufacturer’s designs and manufacturing processes might be, if even one in ten million products is defective, and that one product defect is the cause of a purchaser’s injury, then the purchaser may collect damages (Miller et al. 1985). The plaintiff needs to show only that the product was unreasonably dangerous and that its defect led to harm. In that sense, the standard of care for strict product liability is 100-percent perfection. To some extent, appropriate product labeling (e.g., “Do not use this metal ladder near electrical wiring”) may protect manufacturers in certain strict product liability suits in that clear, visible labeling may educate the purchaser to avoid “unreasonably dangerous” circumstances. Appropriate labeling standards may similarly benefit users and manufacturers of clinical expert systems (Geissbuhler and Miller 1997).

Health care software programs sold to clinicians who use them as decision-support tools in their practices are likely to be treated under negligence theory as services. When advice-giving clinical programs are sold directly to patients, however, and there is less opportunity for intervention by a licensed practitioner, it is more likely that the courts will treat them as products, using strict product liability, because the purchaser of the program is more likely to be the

individual who is injured if the product is defective. (As personal health records become more common, this legal theory may well be tested.)

A growing number of software “bugs” in medical devices have been reported to cause injury to patients (Majchrowski 2010). The U.S. Food and Drug Administration (FDA) views software embedded within medical devices, such as cardiac pacemakers and implantable insulin pumps, as part of the physical device, and so regulates such software as part of the device (FDA 2011). The courts are likely to view such software using principles of strict product liability (Miller and Miller 2007)

Corresponding to potential strict product liability for faulty software embedded in medical devices is potential negligence liability if such software can easily be “hacked” (Robertson 2011). Malicious code writers might mimic external software-based “radio” controllers for pacemakers and insulin pumps and reprogram them to cause harm to patients. While such “hackers” should face criminal prosecution if they cause harm by intention, the device manufacturers have a responsibility to make it difficult to change the software code embedded in devices without proper authorization.

10.5.2.2 Privacy and Confidentiality

The ethical basis for privacy and confidentiality in health care is discussed in Sect. 10.3.1. For a long time, the legal state of affairs for privacy and confidentiality of electronic health records was chaotic (as it remains for written records, to some extent). This state of affairs in the U.S. had not significantly changed in the three decades since it was described in a classic *New England Journal of Medicine* article (Curran et al. 1969).

However, a key U.S. law, the Health Insurance Portability and Accountability Act (HIPAA), has prompted significant change. HIPAA’s privacy standards became effective in 2003 for most health care entities, and its security standards followed 2 years later. A major impetus for the law was that the process of “administrative simplification” via electronic recordkeeping, prized for its potential to increase efficiency and reduce costs, would also pose threats to patient privacy and confidentiality. Coming against a backdrop of a

variety of noteworthy cases in which patient data were improperly—and often embarrassingly—disclosed, the law was also seen as a badly needed tool to restore confidence in the ability of health professionals to protect confidentiality. While the law has been accompanied by debate both on the adequacy of its measures and the question whether compliance was unnecessarily burdensome, it nevertheless established the first nationwide health privacy protections. At its core, HIPAA embodies the idea that individuals should have access to their own health data, and more control over uses and disclosures of that health data by others. Among its provisions, the law requires that patients be informed about their privacy rights, including a right of access; that uses and disclosures of “protected health information” generally be limited to exchanges of the “minimum necessary”; that uses and disclosures for other than treatment, payment and health care operations be subject to patient authorization; and that all employees in “covered entities” (institutions that HIPAA legally affects) be educated about privacy and information security.

As noted above, the HITECH Act provided substantial encouragement for Electronic Health Record (EHR) development, particularly the encouragement of billions of dollars in federal subsidies for “meaningful use” of EHRs. However HITECH also contained many changes to HIPAA privacy and security requirements, strengthening the regulations that affect the collection, use and disclosure of health information not only by covered entities, but also the “business associates” (contractors) of those covered entities, and other types of organizations engaged in health information exchange.

The Office of Civil Rights in the U.S. Department of Health and Human Services remains the entity primarily charged with HIPAA enforcement, but there is now a role for states’ attorneys general as well as other agencies such as the Federal Trade Commission. HITECH increases penalty levels under HIPAA and includes a mandate for investigations and periodic audits, shifting the enforcement balance away from voluntary compliance and remediation plans.

HITECH's changes to HIPAA, those from other federal laws such as the Genetic Information Nondiscrimination Act of 2008 (GINA) and the Patient Safety and Quality Improvement Act of 2005, and the new attention to information privacy and security in most states' laws, comprise significant changes to the legal-regulatory landscape for health information.

10.5.2.3 Copyright, Patents, and Intellectual Property

Intellectual property protection afforded to developers of software programs, biomedical knowledge bases, and World Wide Web pages remains an underdeveloped area of law. Although there are long traditions of copyright and patent protections for non-electronic media, their applicability to computer-based resources is not clear. **Copyright law** protects intellectual property from being copied verbatim, and **patents** protect specific methods of implementing or instantiating ideas. The number of lawsuits in which one company claimed that another copied the functionality of its copyrighted program (i.e., its "look and feel") has grown, however, and it is clear that copyright law does not protect the "look and feel" of a program beyond certain limits. Consider, for example, the unsuccessful suit in the 1980s by Apple Computer, Inc., against Microsoft, Inc., over the "look and feel" of Microsoft Windows as compared with the Apple Macintosh interface (which itself resembled the earlier Xerox Alto interface).

It is not straightforward to obtain copyright protection for a list that is a compilation of existing names, data, facts, or objects (e.g., the telephone directory of a city), unless you can argue that the result of compiling the compendium creates a unique object (e.g., a new organizational scheme for the information) (Tysyer 1997). Even when the compilation is unique and copyrightable, the individual components, such as facts in a database, might not be copyrightable. That they are not copyrightable has implications for the ability of creators of biomedical databases to protect database content as intellectual property. How many individual, unprotected facts can someone copy from a copyright-protected

database before legal protections prevent additional copying?

A related concern is the intellectual-property rights of the developers of materials made available through the World Wide Web. Usually, information made accessible to the public that does not contain copyright annotations is considered to be in the public domain. It is tempting to build from the work of other people in placing material on the Web, but copyright protections must be respected. Similarly, if you develop potentially copyrightable material, the act of placing it on the Web, in the public domain, would allow other people to treat your material as not protected by copyright. Resolution of this and related questions may await workable commercial models for electronic publication on the World Wide Web, whereby authors could be compensated fairly when other people use or access their materials. Electronic commerce might eventually provide copyright protection (and perhaps revenue) similar to age-old models that now apply to paper-based print media; for instance, to use printed books and journals, you must generally borrow them from a library, purchase them or access them under Creative Commons or similar open-access platforms.

10.5.3 Regulation and Monitoring of Computer Applications in Health Care

In the mid-1990s, the U.S. Food and Drug Administration (FDA) held public meetings to discuss new methods and approaches to regulating clinical software systems as medical devices. In response, a consortium of professional organizations related to health care information (AMIA, the Center for Health Care Information Management, the Computer-Based Patient Record Institute, the American Health Information Management Association, the Medical Library Association, the Association of Academic Health Science Libraries, and the American Nurses Association) drafted a position paper published in both summary format and as

a longer discussion with detailed background and explanation (Miller and Gardner 1997a, b). The position paper was subsequently endorsed by the boards of directors of all the organizations (except the Center for Health Care Information Management) and by the American College of Physicians Board of Regents.

The recommendations from the consortium include these:

- Recognition of four categories of clinical system risks and four classes of monitoring and regulatory actions that can be applied based on the level of risk in a given setting.
- Local oversight of clinical software systems, whenever possible, through the creation of autonomous **software oversight committees**, in a manner partially analogous to the institutional review boards that are federally mandated to oversee protection of human subjects in biomedical research. Experience with prototypical software-oversight committees at pilot sites should be gained before any national dissemination.
- Adoption by health care-information system developers of a code of good business practices.
- Recognition that budgetary, logistic, and other constraints limit the type and number of systems that the FDA can regulate effectively.
- Concentration of FDA regulation on those systems posing highest clinical risk, with limited opportunities for competent human intervention, and FDA exemption of most other clinical software systems.

The recommendations for combined local and FDA monitoring are summarized in Table 10.1. Even as the question of regulation continues to challenge the health information technology community, there has been a noteworthy move to attempt to certify and accredit software. Whether such certification efforts will have a meaningful impact on health care outcomes mediated by clinical systems has yet to be determined. Similarly, we do not yet know whether improved outcomes would occur if vendors were to give qualified (i.e., informatics-capable) institutional purchasers greater local control over system functionality.

10.5.4 Software Certification and Accreditation

If, as above, (1) there is an ethical obligation to evaluate health information systems in the contexts in which they are being used, and if, as we just saw, (2) there are good reasons to consider the adoption of software oversight committees or something similar, then it is worthwhile to consider the ethical utility of efforts to review and endorse medical software and systems.

Established in 2004, the Certification Commission for Health Information Technology, in collaboration with the Office of the National Coordinator for health information technology, assesses electronic health records according to an array of criteria, in part to determine their success in contributing to “meaningful use.” These criteria address matters ranging from electronic provider order entry and electronic problem lists to decision support and access control (cf. Classen et al. 2007; Wright et al. 2009). The criteria, tests and test methods are developed in concert with the National Institute of Standards and Technology. Practices and institutions that want to receive government incentive payments must adopt certified electronic health record technologies.

Conceived under the American Recovery and Reinvestment Act, these processes aim to improve outcomes, safety and privacy. Whether they can accomplish this—as opposed to celebrate technology for its own sake—is an excellent source of debate (Hartzband and Groopman 2008). What should be uncontroversial is that any system of regulation, review or certification must be based on and, as a matter of process emphasize, certain values. These might include, among others, patient-centeredness, ethically optimized data management practices, and what we have here commended as the “standard view,” that is, human beings and not machines practice medicine, nursing and psychology.

The move to certification has unfortunately engendered precious little in the way of ethical analysis, however. To make any system of regulation, review or certification ethically credible, government and industry leaders must eventually make explicit that attention to ethics is a core component of their efforts.

Table 10.1 Consortium recommendations for monitoring and regulating clinical software systems^a

Variable	Regulatory class			
	A	B	C	D
Supervision by FDA	Exempt from regulation	Excluded from regulation	Simple registration and postmarket surveillance required	Premarket approval and postmarket surveillance required
Local software oversight committee	Optional	Mandatory	Mandatory	Mandatory
Role of software oversight committee	Monitor locally	Monitor locally instead of monitoring by FDA	Monitor locally and report problems to FDA as appropriate	Assure adequate local monitoring without replicating FDA activity
Software risk category				
0: Informational or generic systems ^b	All software in category	—	—	—
1: Patient-specific systems that provide low-risk assistance with clinical problems ^c	—	All software in category	—	—
2: Patient-specific systems that provide intermediate-risk support on clinical problems ^d	—	Locally developed or locally modified systems	Commercially developed systems that are not modified locally	—
3: High-risk, patient-specific systems ^e	—	Locally developed, non commercial systems	—	Commercial systems

Source: Miller and Gardner (1997a)

^aFDA Food and Drug Administration

^bIncludes systems that provide factual content or simple, generic advice (such as “give flu vaccine to eligible patients in mid-autumn”) and generic programs, such as spreadsheets and databases

^cSystems that give simple advice (such as suggesting alternative diagnoses or therapies without stating preferences) and give ample opportunity for users to ignore or override suggestions

^dSystems that have higher clinical risk (such as those that generate diagnoses or therapies ranked by score) but allow users to ignore or override suggestions easily; net risk is therefore intermediate

^eSystems that have great clinical risk and give users little or no opportunity to intervene (such as a closed-loop system that automatically regulates ventilator settings)

10.6 Summary and Conclusions

Ethical issues are important in biomedical informatics, and especially so in the clinical arena. An initial ensemble of guiding principles, or ethical criteria, has emerged to orient decision making:

1. Specially trained human beings remain, so far, best able to provide health care for other human beings. Hence, computer software should not be allowed to overrule a human decision.
2. Practitioners who use informatics tools should be clinically qualified and adequately trained in using the software products.
3. The tools themselves should be carefully evaluated and validated.
4. Health informatics tools and applications should be evaluated not only in terms of performance, including efficacy, but also in terms of their influences on institutions, institutional cultures, and workplace social forces.
5. Ethical obligations should extend to system developers, maintainers, and supervisors as well as to clinician users.
6. Education programs and security measures should be considered essential for protecting confidentiality and privacy while improving appropriate access to personal patient information.
7. Adequate oversight should be maintained to optimize ethical use of electronic patient information for scientific and institutional research.

New sciences and technologies always raise interesting and important ethical issues. Much the same is true for legal issues, although in the absence of precedent or legislation any legal analysis will remain vague. Similarly important challenges confront people who are trying to determine the appropriate role for government in regulating health care software. The lack of clear public policy for such software underscores the importance of ethical insight and education as the exciting new tools of biomedical and health informatics become more common.

Suggested Readings

- Cushman, R., Froomkin, M. A., Cava, A., Abril, P., & Goodman, K. W. (2010). Ethical, legal and social issues for personal health records and applications. *Journal of Biomedical Informatics*, 43, S51–S55. This article provides an overview of ethical issues related to the use of personal health records. Topics include privacy and confidentiality, decision support and the clinician-patient relationship.
- Goodman, K. W. (Ed.). (1998). *Ethics, computing, and medicine: Informatics and the transformation of health care*. Cambridge: Cambridge University Press. (2nd edn in press.) This volume, the first devoted to the intersection of ethics and informatics, contains chapters on informatics and human values, responsibility for computer-based decisions, evaluation of medical information systems, confidentiality and privacy, decision support, outcomes research and prognostic scoring systems, and meta-analysis.
- Goodman, K. W., Berner, E. S., Dente, M. A., Kaplan, B., Koppel, R., Rucker, D., Sands, D. Z., & Winkelstein, P. (2011). Challenges in ethics, safety, best practices, and oversight regarding HIT vendors, their customers, and patients: a report of an AMIA special task force. *Journal of the American Medical Informatics Association*, 18(1), 77–81. This document is one of the first to examine issues related to the manufacturers and vendors of electronic health records, including their relationships with hospitals.
- Miller, R. A. (1990). Why the standard view is standard: People, not machines, understand patients' problems. *Journal of Medicine and Philosophy*, 15, 581–591. This contribution lays out the standard view of health informatics. This view holds, in part, that because only humans have the diverse skills necessary to practice medicine or nursing, machine intelligence should never override human clinicians.
- Miller, R. A., Schaffner, K. F., & Meisel, A. (1985). Ethical and legal issues related to the use of com-

puter programs in clinical medicine. *Annals of Internal Medicine*, 102, 529–536. This article constitutes a major early effort to identify and address ethical issues in informatics. By emphasizing the questions of appropriate use, confidentiality, and validation, among others, it sets the stage for all subsequent work.

National Research Council. (1997). *For the record: Protecting electronic health information*. Washington, DC: National Academy Press. A major policy report, this document outlines leading challenges for privacy and confidentiality in medical information systems and makes several important recommendations for institutions and policymakers.

Questions for Discussion

1. What is meant by the “standard view” of appropriate use of medical information systems? Identify three key criteria for determining whether a particular use or user is appropriate.
2. Can quality standards for system developers and maintainers simultaneously safeguard against error and abuse and stimulate scientific progress? Explain your answers. Why is there an ethical obligation to adhere to a standard of care?
3. Identify (a) two major threats to patient data confidentiality, and (b) policies or strategies that you propose for protecting confidentiality against these threats.
4. Many prognoses by human beings are subjective and are based on faulty memory or incomplete knowledge of previous cases. What are the two drawbacks to using objective prognostic scoring systems to determine whether to allocate care to individual patients?
5. People who are educated about their illnesses tend to understand and to follow instructions, to ask insightful questions, and so on. How can the World Wide Web improve patient education? How, on the other hand, might Web access hurt traditional physician–patient and nurse–patient relationships?