Ethics in Health Informatics

Kenneth W. Goodman

Institute for Bioethics and Health Policy, University of Miami Miller School of Medicine, Miami, USA

Summary

Contemporary bioethics was fledged and is sustained by challenges posed by new technologies. These technologies have affected many lives. Yet health informatics affects more lives than any of them. The challenges include the development and the appropriate uses and users of machine learning software, the balancing of privacy rights against the needs of public health and clinical practice in a time of Big Data analytics, whether and how to use this technology, and the role of ethics and standards in health policy. Historical antecedents in statistics and evidence-based practice foreshadow some of the difficulties now faced, but the scope and scale of these challenges requires that ethics, too, be brought to scale in parallel, especially given the size of contemporary data sets and the processing power of new computers. Fortunately, applied ethics affords a variety of tools to help identify and rank applicable values, support best practices, and contribute to standards. The bioethics community can in partnership with the informatics community arrive at policies that promote the health sciences while reaffirming the many and varied rights that patients expect will be honored.

Keywords

Artificial Intelligence, Big Data, Bioethics, Clinical Decision Support Systems, Public Policy

Yearb Med Inform 2020;26-31

<http://dx.doi.org/10.1055/s-0040-1701966>

Introduction

Ethics is a kind of lens we use to identify issues and a lever used to formulate and motivate best practices. Applied ethics is a tool to employ widely shared if not universal values to contemporary questions and challenges in science and the professions. Bioethics is the branch of applied ethics that addresses issues in the health professions; it is often linked to other kinds of applied ethics, including business ethics, computer ethics, government ethics, and so on.

It has become commonplace to observe or argue that science usually or even always outpaces, or advances, more swiftly than applied ethics. In the case of health informatics, this is a mistake. For some four decades, albeit with some exceptions, advances in biomedical informatics have been matched step for step by scholars who have identified and addressed the ethical, legal, and social issues (ELSI) raised by the expansion of a new science. This acronym, borrowed from the Human Genome Project, has ably served the informatics community as a label or guidepost for research and pedagogy.

Though bioethics has moved forward, the same cannot be said for the law which continues to lag as a source of official governance and oversight in health informatics and other domains. It might be that the relationship between ethics and public policy represents the greatest challenge faced by health informatics and the society health informaticians seek to serve.

What follows from these two observations is this: we have an extraordinary opportunity at a crucial time to try to ensure that the insights and analyses provided by ethics continue to mature and, as important, that they are taken up and incorporated by academic and health care institutions, businesses, professional organizations, and governments.

In what follows, I expand on these points by filtering them through a number of contemporary challenges. These include artificial intelligence and machine learning; Big Data, data sharing and privacy; duties to use and manage new technology; and ethics and public policy.

1 Artificial Intelligence, Machine Learning, and Ethics

The ancient fantasy of an intelligent machine or a smart homunculus became a research project in the 17th century when Gottfried Leibniz, the philosopher and logician who, with Newton, discovered the infinitesimal calculus, suggested that human reason could be rendered in a universal language such that argumentation could be reduced to calculation. He built a primitive calculator [1], arguably the first machine to replicate an aspect of human thinking. Leibniz thought that intelligent machines would formalize reason and end disagreements. They did not. During the greatest disagreement in the history of civilization – World War II – code-breaking machines represented nontrivial instantiations of Leibniz' aspiration: "The first digital, electronic and programmable computer was developed as an instrument of warcraft. The Colossus was a room-sized collection of racks, pulleys, wires and some 2,400 bottle-sized vacuum tubes built at Britain's Bletchley Park to decipher encrypted German messages.... It became operational in 1944 and was used to prepare the D-Day invasion of Normandy. One could argue that it eventually saved more lives than most medical inventions" [2].

However, we want health informatics to save and improve lives, to reduce suffering, to help to achieve the larger goals of the

healthcare professions. If an intelligent machine can help do that, why is there a problem or any controversy?

Indeed, there are several problems, and their identification has contributed to both greater understanding and to a period of overheated handwringing.

Key and core ethical issues in the development and use of machine learning programs have already been identified repeatedly and elsewhere. They can be framed as lessons learned or as recommendations. Some are lessons for developers, some for users. Though the context here is health and health care, the lessons may be useful in many other domains.

1.1 Quality and Standards Are Ethical Issues

Good software conforms to certain standards for quality which can be assessed in terms of trustworthiness and reproducibility. The mark or measure of good software will include accuracy of documentation and transparency about the provenance or source of any code components. This is no mere courtesy to code-writing colleagues – it is an auditable track record of what code is intended to do and how it was modified along the way. This facilitates understanding, corrections, and improvements. It follows that if one is developing or modifying machine learning software, the automated learning process itself must also be monitored and documented. Relatedly, careful software version control is an essential part of high-quality programming. The values of transparency, veracity, and accountability reinforce the connection among quality, standards, and ethics [3]. This has long been true across the health professions.

1.2 Prevent and Eliminate Bias

A sure way to erode confidence in artificial intelligence is to identify ways in which a deep learning algorithm embeds racial, ethnic, gender, or other biases which shape or corrupt its results. The very nature of machine learning algorithms makes plain that one might unintentionally develop a biased

system or accept biased results. Though intent matters in ethics – many actions are praiseworthy or blameworthy precisely because of someone's intent – this problem makes equally clear that failure to attend to and correct or mitigate a problem can also be blameworthy; that failure to do so is a kind of negligence (no surgeon ever intends to cut off the wrong arm or operate on the wrong patient). If datasets used for training machine learning algorithms include or entail bias or foster biased interpretations, then careful scrutiny of such sets is a difficult and labor-intensive approach to take. Another is careful screening of output to identify and filter bias, and perhaps to identify ways to modify the algorithm to suppress it. It is now even possible to incorporate anti-bias features into the code itself [4]. This is an extraordinarily promising approach, and something like it might very well be the best way to ensure that public health surveillance and prediction, disease diagnosis and treatment, and health policy are not infected or corrupted by illicit bias.

1.3 Use Machine Learning Software for Good and not Evil

Current relativist trends and fashions notwithstanding, some actions are good and others are bad and there is little credible dispute about their moral status. Reducing suffering, eliminating disparities, and improving health are good; depriving people of rights, using people for political or economic purposes without permission given voluntarily, and harming people for profit are bad. If we start there, great progress can be made. It is the hard cases, or cases on which reasonable people disagree, that pose the most significant challenges. Is there a chance that a clinical decision support system will improve diagnosis, treatment, and prognosis and simultaneously erode confidence in the clinician-patient relationship? May a nurse with a good computer system undertake duties traditionally reserved for physicians? Will the use of intelligent machines in health care cause the erosion of healthcare practitioner clinical skills? Addressing these and other challenges successfully will require sustained research,

education, and debate. At the least, we will make great progress by acknowledging and grappling with such challenges.

1.4 Insist on and Provide Robust Education and Evaluation

No technology will ever fulfill its promise or be used in an ethically optimized manner unless we *insist on and provide robust education and evaluation*. Scientists and clinicians must not only be taught empirical methods and clinical skills, but also the appropriate uses of these methods and skills. Given our concern for appropriate uses of intelligent machines in healthcare, it follows we must also identify appropriate users, and a core criterion for demonstrating such appropriateness will be a user's fitness to use the tool. Making and acting on this assessment has long been recognized as a key ethical challenge when computers are used in healthcare [5]. We are, moreover, in urgent need of a comprehensive ethical curriculum in and for health informatics. We must also not lose sight of the importance of system (algorithm, device) evaluation in the context of its actual use; this, too, is an obligation long recognized [6]. Much more recently, system validation in real-world settings has been identified as a duty for anyone who would use an intelligent system in healthcare [7]. This may be especially true and important given the rapid progress in precision medicine.

2 Big Data, Data Sharing, and Privacy

Empirical sciences have evolved and grown by using data and information to support inferences which, if correct, come to constitute knowledge. In biology and medicine, we use that knowledge to prevent, predict, mitigate, and cure maladies that sicken, disable, or kill us.

A favorite example from the history of medicine of “calculating” to improve health care is that of bloodletting, used for millennia to treat and often unintentionally kill people

who, under the humoral theory of disease, had too much blood. For instance, George Washington one day in 1799 awoke with a sore throat, asked to be bled, lost some five pints of blood, and was dead four days later, probably because of the bloodletting. The work of Pierre Charles Alexandre Louis (1787-1872) was likely unknown to him. Louis, a French physician, analyzed cases from several hospitals using his “numerical method” and determined that bloodletting was not a cure but, rather, usually harmful. He wrote: “As to different methods of treatment, it is possible for us to assure ourselves of the superiority of one or other... by enquiring if the greater number of individuals have been cured by one means than another. Here it is necessary to count. And it is, in great part at least, because hitherto this method has not at all, or rarely been employed, that the science of therapeutics is so uncertain” [8].

Necessary to count? Numerical method? Is it – can it really be – that simple?

Yes, sometimes. Observational science and basic arithmetic were, at first, the best tools we had to support efforts to learn how the world worked. Observe and count (or measure), find the mean, note changes over time, and then make inferences about physical reality. An early form of statistics, frequency analysis, originated in the Middle Ages with the Arab philosopher Al-Kindi as a code-breaking tool more than a millennium before Colossus. His method also required counting – of letters. With an adequate “plaintext... long enough to fill one sheet or so,” one could determine the frequency of letters and combinations of letters as a baseline and use it to determine the values of letters in a coded document [9]. It was a very early anticipation or foreshadowing of the data sets used to train artificial intelligence programs.

What became clear in the 17th and 18th centuries was the utility of using data about observations to make predictions. In the 1660s, John Graunt of London’s “Natural and Political Observations Made upon the Bills of Mortality” constituted the first actuarial tables and spurred efforts to predict and control the causes of death. It is also the basis for modern insurance, including the business of using statistics to sell health “coverage” to people in need of health care in countries with dysfunctional health care systems.

The preceding part of this section embeds several important ethical issues. Chief among them are these:

- Are data reliable and who is responsible for ensuring reliability?
- How exactly does the calculation work?
- How should it be determined who should use these calculators, and for what purposes?

These, of course, parallel the “lessons learned” we earlier identified for appropriate use of machine learning or any other medical software.

The datasets created by Al-Kindi, Louis, and Graunt were actually tiny; we might even call this “small data.” Yet their work suggests the first grains or nuclei of a vastly larger project. Framed very broadly, one might argue that the work of these three data collector-analyzers raised ethical issues that parallel those we find when collecting and analyzing some of the more than 2,000 exabytes of medical data in the world today (an exabyte, recall, is 1,000 petabytes, with one petabyte being 1,000 terabytes, and one terabyte being 1,000 gigabytes, etc.). But their work did not cause any ethical consternation at the time; it was just science advancing.

Though we ought to regard privacy and confidentiality as universal rights, they are neither absolute nor have they ever been regarded as absolute. There are tradeoffs to be made: public health surveillance gives us better public health, hospital patient monitoring gives us better hospital care, the sharing of research results and other data gives us better medical treatments. Put differently: if a physician learns something about effective practice from Patient Epsilon’s response to a treatment, does the physician violate Epsilon’s privacy or confidentiality six months later by using what she learned to help Patient Omega [10]? Indeed, such a *reductio ad absurdum* argument makes clear that not only does the physician not violate confidentiality but also that she would be irresponsible not to use the information from one patient to help another.

With very large datasets, however, our challenge is magnified. Until recently, the ability to reconnect a datum with a person,

or “re-identify” the datum, was either impossible or difficult. Now it is comparatively easy, and this entails that the tools of applied ethics must in a world of Big Data rise to the occasion. One way to put this, according to my colleague Richard Bookman in a personal communication, is that “Ethical boundaries move with scale. As little data goes to Big Data, onetime practices at small scale may be rendered offensive and unacceptable at broader scale.”

The same could be said for data sharing. Few would argue that science progresses when investigators communicate openly and generously with each other. It also progresses when scientists are motivated by curiosity and service instead of profit and priority – at least historically. When the financial stakes were low, history shows discovery science was fertile and productive. That slowed when fame attached to empirical success and reputations were made by being the first to make a discovery. We then commodified science so that not only would a good scientist learn something, he could also make a lot of money. In 1955, Jonas Salk famously disdained patenting his polio vaccine (“Could you patent the sun?”); it was the same year the Dartmouth Workshop, at which Artificial Intelligence as a field was fledged, was proposed. In 2020, one would be regarded as a fool for not leveraging a discovery or method for profit. Big Data means big profit, and the ethical boundaries have correspondingly shifted.

It would perhaps not be so difficult if the accumulation of wealth from others’ information could rely on their concurrence and support. Worse, the erosive use of Big Data for profit in business and industry has polluted public perceptions of all uses of Big Data. We learned earlier that applied ethics can contribute to the identification of appropriate uses of information technology. Here we have an opportunity to protect fundamental privacy rights by making a distinction between using social media, for instance, to track people’s behavior to target them for political purposes, and using electronic health records to analyze their data to improve clinical outcomes. There is indeed evidence that ordinary people are prepared to accept this distinction, at least if the user is trustworthy [3].

We have several ways to meet the privacy challenges posed by Big Data for health care. All of them should be seen as trust-enhancing.

- *The first is security backed up by law.* Make it difficult for unauthorized – inappropriate – parties to access or use the data, and punish them severely if they do. Constant and evolving research improves privacy-protecting software. The 2018 *IMIA Yearbook of Medical Informatics*, for instance, included a special section (“Between Access and Privacy: Challenges in Sharing Health Data”) that featured a suite of articles describing research devoted to improving digital privacy protection [11]. It is clear that the good guys write better software than the bad guys. Use it.
- *Next, insist on accountable oversight.* From bioinformatics laboratories to hospitals, those who collect, store, and analyze health information are stewards of a precious resource. They must, absolutely must, be held to trust-enhancing standards. The “trusted broker” model of scrutiny and oversight can ensure both innovative research and the protections that are demanded by the sources of health data and information, i.e., usually, patients [3]. Institutions and governments need to require and help implement mechanisms to review, scrutinize, and approve collections and analyses of health data. This builds trust and makes clear that those who use health data are motivated by values different than those who scrape online clickstreams to beguile shoppers, those who vacuum social media postings to monetize the intimate behavior of strangers, and those who weaponize the Web to undermine democracies. They wanted, rather, to heal the sick.
- *Finally, educate lay people to improve digital health literacy, and teach scientists and clinicians to learn to appreciate that the opportunity to study others’ information is a privilege.* If the former are willing to share data for better health and the latter are motivated by higher values than those that govern the marketplace, then we reinforce the foundation of a collective, international edifice that both

protects legitimate privacy interests while simultaneously ensuring enjoyment of the right to benefit from science [12].

It might be that the greatest utility of applied ethics is in the balancing of rights, in part by the identification of responsibilities.

3 Duties to Use New Technology

Information technology has transformed daily life, business, entertainment, and health care. The list could, of course, go on, and on. The proper criterion to use in determining if any new tool should be used is, generally, whether its advantages outweigh its disadvantages and, if so, whether any antecedent rights are being violated. We learned more than three decades ago that if health information technology could improve health, then, *ceteris paribus*, this entailed a duty to use it for that purpose [5].

Such an insight does not of itself begin to answer the many questions that follow – that is the job of scholars and analysts who scrutinize appropriate uses, privacy implications, and governance. It nevertheless signals a powerful moral commitment to reject basic Luddism, the movement that began in the 19th century with British workers smashing mechanical looms because the machines imperiled their jobs. The best approach is to control machines, not destroy them.

One of the most promising and systematic approaches to achieve that control at scale is the Learning Health System (LHS) paradigm [13], an evidence-building, safety-promoting system “capable of continuous self-study and improvement” [14]. Such a system is a data-driven creature; it learns from every patient encounter, every lab result, every outcome tracked and measured; it learns from public health prevention and intervention. It captures Louis’ fabulous and simple understanding that counting can count for a lot [8]. Moreover, LHS embodies core values of evidence-based practice, most especially that it is neglectful and blameworthy to lose information that could have been used to improve the wellbeing of individuals and the health of populations.

The LHS also provides an excellent example of and justifies what has been called the “secondary use” of health information. In many respects, the term “secondary use” is a misnomer: it implies that health information collected in a clinic, for instance, is and ought to be collected for a primary purpose, and that any other purpose enjoys lesser warrant. Similarly, it would be peculiar to collect public health information and then regard that information’s use in a clinic as secondary and requiring additional permission. Rather, it would be irresponsible to acquire information in a clinic, research study or public health surveillance effort and not use it to support one of the others [15–17]. Widespread views and focused regulations regarding how best to protect these three kinds of data and information are largely historical artifacts. Clinic data is protected by privacy laws that evolved to protect, well, clinical data. Research data use is governed by rules aimed at protecting human subjects, including their privacy. Public health findings are, and always have been, acquired with the presumed or tacit consent of those who benefit from – and trust – competent public health authorities.

Developers of Learning Health Systems (or “the” system if there is one and it is global) are wise and quick to identify the values that reinforce it: person-focused, privacy, inclusiveness, transparency, accessibility, adaptability, governance, cooperative and participatory leadership, scientific integrity and value [14]. Indeed, one would be forgiven for suggesting that those values should underlie any health system.

How such values are embraced will continue to challenge our best inclinations. The effort to make the most of health data and information, however, is a collective recognition of the overarching duty to use tools that improve health. Perhaps the most difficult challenge is shaped by the values of inclusiveness and accessibility; that is, inclusiveness and accessibility for whom?

Any duty to use a tool is cheapened or diluted if the tool benefits only those in some jurisdictions or some socio-economic classes. If we intend a health system to serve all, and if we know that social determinants are key etiologic sources of illness, then per-

haps a Learning Health System can be used to reduce disparities. Universalizing access to health care, which has so far eluded most countries, would require far more than data analysis – but such analysis in conjunction with tailored social commitment might at the least effect some improvement. To be sure, such a course would require the inclusion of social determinants, including poverty, in electronic health records or other repositories of patient information. This in turn could cause or magnify stigma, and it might be the case that some people would disdain being thus identified.¹

A robust attention to ethics and informatics cannot provide a script or formula for righting wrongs. It can however usefully identify issues to be addressed, scrutinize approaches to confronting those issues, and help point the way to the best possible solutions. There is, as well, an element of professionalism in this regard, and major informatics organizations have developed codes of ethics to guide those working in the field to follow such precepts as help to secure the integrity of the profession [18, 19].

4 Ethics, Standards, and Public Policy

Once we figure out how to get something right, it can be irrational to insist it should be ignored. The evolution of standards in health care has improved quality, increased safety, and conserved resources. This is especially true for health informatics. What follows should be obvious: if continuously refined and improved standards can achieve so much, and if those achievements improve human health and welfare, then there is an ethical imperative to develop and improve them.

Any flaws in the discovery and communication process – there are a lot of standards, some conflict, and others are not very good – are remediable. As was once suggested, “What is wanted here, in informatics, is in

an important if vague respect what is wanted everywhere in the world of standards: we want to be as good as we can be, along with our efforts and tools, and to improve. That part of ‘standard’ that implies regularity or consistency captures the intuitions that order is more successful than chaos, and that clarity and order produce goods and services that are easier to explain and share. We have something very important to accomplish in health information technology, and good standards are necessary for the process” [3].

Standards are also an important component of public policy and one of the ways by which applied ethics informs and shapes public governance.

From the foundations of software engineering to the design of electronic health records to embedded privacy protections to evaluation and interoperability, ethical principles and standards serve as both guiderails and signposts. Is there a problem with biased algorithms? Adopt standards for better software and testing. Do members of communities distrust those who collect and analyze personal information? Follow standards for trust enhancement. Confused about whether to adopt a new technology? Then turn to ethical standards for harm reduction and rights protection.

As elsewhere in life, we sometimes in health informatics are over-eager to make problems seem intractable or lose sight of the fact that we have available many tried-and-true conceptual devices and publicly accountable processes for solving these problems. To say, as some are wont, for instance, “privacy is dead; get over it,” is a sad surrender to an ancient challenge tricked out with very large data bases and powerful computing tools. Such surrender is neither necessary nor apt.

If we care enough about these issues, and we should, then applied ethics emerges as the silent partner or loyal opposition to speak truth to a power whose benefits we deserve and whose harms we must protest. When appropriate, in civil society, the law itself can be a partner in ensuring that standards capture best practices. Ethics should always come first, both to guide us and to reduce the likelihood of health-system laws that permit or even encourage unethical conduct.

Conclusion

Ordinary language loves the locution “ethical dilemma.” But to be on the horns of an ethical or moral dilemma is to be in a position such that no matter what one does, one does something wrong. Most ethical problems or challenges are not dilemmas, and even the dilemmas, if they be such, can, as in logic, be escaped. The project here has been to argue and so to make clear that the tools of applied ethics are adequate to the task of guiding developers, users, and institutions as they adopt and try to make the most of health information technology. This has been framed as statements of issues and problems (machine learning, privacy, and duty to use the new tools) and as suggestions for addressing them. As elsewhere in the health professions, the ethical issues raised by a new technology are sources of interesting and significant challenges. These challenges are difficult, but not intractable. Finding their solutions presents opportunities to use critical thinking in the service of shared values. Chief among these values is that of health itself. Producing better and ethically optimized tools for healers is a contribution requiring the collaboration of the informatics and ethics communities. Such collaboration affords a rare and unparalleled opportunity.

References

1. Morar F-S. Reinventing machines: the transmission history of the Leibniz Calculator. *Brit J Hist Sci* 2015;48(1):123–46.
2. Goodman KW. Health information technology and globalization. In: Chadwick R, ten Have H, Meslin EM, editors. *Health Care Ethics: Core and Emerging Issues*. Los Angeles: Sage; 2011. p. 117–25. See also references therein.
3. Goodman KW. *Ethics, Medicine and Information Technology: Intelligent Machines and the Transformation of Health Care*. Cambridge: Cambridge University Press; 2016.
4. Thomas PS, Castro da Silva B, Barto AG, Giguere S, Brun Y, Brunskill E. Preventing undesirable behavior of intelligent machines. *Science* 2019;366:999–1004.
5. Miller RA, Schaffner KF, Meisel A. Ethical and legal issues related to the use of computer programs in clinical medicine. *Ann Int Med* 1985;102:529–36.
6. Anderson JG, Aydin CE. Evaluating medical information systems: social contexts and ethical

¹ I am here again indebted to my colleague Richard Bookman for thoughts about the intersection of social determinants and health information technology.

- challenges. In: Goodman KW, editor. *Ethics, Computing, and Medicine: Informatics and the Transformation of Health Care*. Cambridge: Cambridge University Press 1998. p. 57-74.
7. Liu Y, Cameron Chen P-H, Krause J, Peng L. How to read articles that use machine learning. *JAMA* 2019;322(18):1806-16.
 8. Louis PCA. *Essay on Clinical Instruction*. P. Martin, trans., London: S. Highley; 1834. (Cited in: Huth EJ, Murray TJ, editors. *Medicine in Quotations: Views of Health and Disease Through the Ages*. Philadelphia: American College of Physicians; 2006.)
 9. Cozzens M, Miller SJ. *The Mathematics of Encryption: An Elementary Introduction*. Providence, Rhode Island: American Mathematical Society; 2013.
 10. Goodman KW. Ethics, information technology and public health: new challenges for the clinician-patient relationship. *J Law Med Ethics* 2010;38(1):58-63.
 11. Malin B, Goodman KW. Between access and privacy: challenges in sharing health data. *Yearb Med Inform* 2018;55-9.
 12. UNESCO. *The Right to enjoy the benefits of scientific progress and its applications*. Paris: UNESCO, 2009. Available from: <https://unesdoc.unesco.org/ark:/48223/pf0000185558>
 13. Institute of Medicine. *The Learning Healthcare System: Workshop Summary (IOM Roundtable on Evidence-Based Medicine)*. Washington, DC: The National Academies Press; 2007.
 14. Friedman CP, Rubin JC, Sullivan KJ. Toward an information infrastructure for global health development. *Yearb Med Inform* 2017;16-23.
 15. Lee LM, Gostin LO. Ethical collection, storage, and use of public health data: A proposal for a national privacy protection. *JAMA* 2009;302(1):82-4.
 16. Faden RR, Kass NE, Goodman SN, Pronovost P, Tunis S, Beauchamp TL. An ethics framework for a learning health care system: a departure from traditional research ethics and clinical ethics. *Hastings Cent Rep* 2013;43(S1):S16-S27.
 17. Salerno J, Knoppers BM, Lee LM, Hlaing WM, Goodman KW. Ethics, big data and computing in epidemiology and public health. *Ann Epidemiol* 2017;27:297-301.
 18. Petersen C, Berner ES, Embi PJ, Fultz Hollis K, Goodman KW, Koppel R, Lehmann CU, et al. AMIA's code of professional and ethical conduct. *J Am Med Inform Assoc* 2018;25(11):1579-82.
 19. Kluge E-H. *The IMIA Code of Ethics: Revision 2016*. International Medical Informatics Association 2016. Available from: <https://imia-medinfo.org/wp/wp-content/uploads/2015/07/IMIA-Code-of-Ethics-Short.pdf>

Correspondence to:

Kenneth W. Goodman
University of Miami
POB 016960 (M-825)
Miami, Florida, 33101
USA
E-mail: kgoodman@miami.edu