*Ethics and Information Technology* **4:** 1–9, 2002.


# Mapping the foundationalist debate in computer ethics

Luciano Floridi[1] and J.W. Sanders[2]
[1]*Faculty of Philosophy, Sub-factulty of Computation, and Programme in Comparative Media Law and Policy, University of Oxford*
*Address for correspondence: Wolfson College, Oxford, OX2 6UD, UK*
*E-mail: luciano.floridi@philosophy.oxford.ac.uk*
[2]*Programming Research Group, OUCL, University of Oxford*
*Address for correspondence: Wolfson Building, Parks Road, Oxford, OX1 3QD, UK*
*E-mail: jeff@comlab.ox.ac.uk*

**Abstract.** The paper provides a critical review of the debate on the foundations of Computer Ethics (CE). Starting from a discussion of Moor's classic interpretation of the need for CE caused by a policy and conceptual vacuum, five positions in the literature are identified and discussed: the "no resolution approach", according to which CE can have no foundation; the professional approach, according to which CE is solely a professional ethics; the radical approach, according to which CE deals with absolutely unique issues, in need of a unique approach; the conservative approach, according to which CE is only a particular applied ethics, discussing new species of traditional moral issues; and the innovative approach, according to which theoretical CE can expand the metaethical discourse with a substantially new perspective. In the course of the analysis, it is argued that, although CE issues are not uncontroversially unique, they are sufficiently novel to render inadequate the adoption of standard macroethics, such as Utilitarianism and Deontologism, as the foundation of CE and hence to prompt the search for a robust ethical theory. Information Ethics (IE) is proposed for that theory, as the satisfactory foundation for CE. IE is characterised as a biologically unbiased extension of environmental ethics, based on the concepts of information object/infosphere/entropy rather than life/ecosystem/pain. In light of the discussion provided in this paper, it is suggested that CE is worthy of independent study because it requires its own application-specific knowledge and is capable of supporting a methodological foundation, IE.

**Key words:** computer ethics, information and communication technology, information ethics, macroethics, metaethics, microethics, policy vacuum, uniqueness debate

**Abbreviations:** CA – conservative approach; CE – computer ethics; IA – innovative approach; ICT – information and communication technologies; IE – information ethics; NA – no resolution approach; PA – professional approach; PE – pop ethics; RA – radical approach.

## Introduction

Computer Ethics (CE) stems from practical concerns arising in connection with the impact of Information and Communication Technologies (ICT) on contemporary society. The so-called digital revolution has caused new and largely unanticipated problems, thus outpacing ethical, theoretical and legal developments (Bynum 1998; Bynum 2000; Johnson 2000 for an overview). In order to fill this policy and conceptual vacuum (Moor 1985), CE carries out an extended and intensive study of individual cases, amounting very often to real-world issues rather than mere mental experiments, usually in terms of reasoning by analogy. The result has been inconsistencies, inadequacies and an unsatisfying lack of general principles. However, CE's aim is to reach decisions based on principled choices and defensible ethical criteria, and hence to provide more generalised conclusions – in terms of conceptual evaluations, moral insights, normative guidelines, educational programs, legal advice, industrial standards and so forth – which may apply to whole classes of comparable cases. So, at least since the seventies (see Bynum 2000 for earlier works in CE), CE focus has moved from problem analysis – primarily aimed at sensitising public opinion, professionals and politicians – to tactical solutions resulting, for example, in the evolution of professional codes of conduct, technical standards, usage regulations, and new legislation. The constant risk of this bottom-up procedure has remained the spreading of ad hoc or casuistic approaches to ethical problems. Prompted partly by this difficulty, partly by a natural process of self-conscious maturation as an independent discipline, CE has further combined tactical solutions with more strategic and global analyses. The foundation-

alist debate is an essential part of this top-down development. It is characterised by a metatheoretical reflection on the nature and justification of CE and the discussion of CE's relations with the broader context of metaethical theories. Can CE amount to a coherent and cohesive discipline, rather than a more or less heterogeneous and random collection of ICT-related ethical problems, applied analyses and practical solutions? If so, what is its conceptual rationale? And how does it compare with other ethical theories? Five approaches to the foundation of CE that have emerged in the literature can be explained as resulting from different answers to those questions. Here they are referred to as the "no resolution approach" (NA), the Professional Approach (PA), the Radical Approach (RA), the Conservative Approach (CA), and the Innovative Approach (IA). The order in this list is both historical and logical. In the rest of this article it is argued that NA provides a minimalist starting point, methodologically useful, which prompts the development of the other four approaches; that PA represents a valuable professional approach to CE, which leads to the adoption of a theoretical position when metaethical issues are in question; that RA stresses the novelty of CE; that CA connects CE to other standard ethics; and that IA, relying on the previous approaches, succeeds in providing a satisfactory answer to the foundationalist question, by presenting Information Ethics (IE) as the theoretical foundation of CE.

## The "No resolution approach": CE as not a real discipline

The expression "no resolution view" (or approach) is introduced in Gotterbarn 1991:

> The "no resolution view" has been reinforced by some recent works. For example, Donn Parker 1981 uses a voting methodology to decide what is ethical in computing. . . . He says, this work was not guided by a concept of computer ethics nor was there an attempt to discover ethical principles. . . . Not only was there an absence of a concept of computer ethics, but the primary direction was an emphasis on proscribed activities. . . . Parker used the diversity of opinions expressed about these scenarios to argue that there was no such thing as computer ethics. And a fortiori, that it could not be taught in a computer science curriculum.

According to the "no resolution approach" (NA) CE problems represent unsolvable dilemma and CE is itself a pointless exercise, having no conceptual foundation. NA is convincingly criticised in Gotterbarn 1991 and 1992, which analyses Parker 1981, 1982 and 1990. Empirically, the evolution of CE has proved NA to be unnecessarily pessimistic. CE problems are successfully solved, CE-related legislation is approved and enacted, professional standards and codes have been promoted, and so forth. It is understandable, perhaps, that the view arose at a time when both public and professionals were being alerted to wide-ranging unethical uses of ICT. The reason that Parker did not infer an almost opposite conclusion (i.e., that CE is essential) from the "emphasis on proscribed activities" is presumably his voting methodology. It is dangerous to infer, from inconsistent replies to a question, that it has no answer. The same reasoning might lead one to believe, after asking a representative sample of supporters, that neither side would win the test at Lords. NA's emphasis on the wide variety of proscribed activities is characteristic. Bynum (1992) has described such an approach as "pop ethics" (PE). PE is characterised by usually unsystematic and heterogeneous collections of dramatic stories, discussed in order "to raise questions of unethicality rather than ethicality" (Parker 1981). Its goal is largely negative, "to sensitise people to the fact that computer technology has social and ethical consequences" (Bynum 1992) and it is not neutral. That is why it played a useful role at the beginning of the development of CE, at around the time when "hacker" became used disparagingly, for example. It is comparable to early work done in business ethics: it points to whatever goes wrong but fails to promote a relevant, beneficial, professional ethos. Gotterbarn comments:

> "Pop" ethics might have had a place when computing was a remote and esoteric discipline, but I believe that in the current environment this approach is dangerous to the preservation and enhancement of values. This model of computer ethics does not forward any of the pedagogical objectives for teaching ethics [prescribed by PA] (Gotterbarn 1992).

Nonetheless, PE offers advantages. Some sensitisation to ethical problems is an important preliminary to CE. There is little point in providing a solution to someone unaware of the problem, particularly when the solution is not simple. Secondly, the variety of concerns is vital to CE (professional, legal, moral, social, political, etc.) and must be appreciated from the start. For this purpose a variety of case studies helps. For instance, Epstein 1997 provides an example of PE found by many lecturers to be useful as preliminary reading for a course in CE. The objection to PE is that it goes no further than cataloguing examples, and that it is frequently used to support NA. Methodologically, NA provides a useful point of reference because it repre-

sents an ideal lowest bound for the foundationalist debate, comparable to the role played by relativism in metaethics. In terms of logical analysis, any other approach can be seen as starting from the assumption that NA should be avoided, if possible. Positions can then be ranked depending on their distance from NA, whilst failure to defend any successful alternative confirms NA as the only negative conclusion.

**The professional approach: CE is a pedagogical methodology**

The first positive reaction to the policy vacuum has been an appeal to the social responsibility of computer professionals. This has meant, among other things, developing a professional-ethics approach (PA) to CE, which has stressed pedagogical need (for an overview see Gotterbarn 1991, 1992). According to PA, CE should:

> introduce the students to the responsibilities of their profession, articulate the standards and methods used to resolve non-technical ethics questions about their profession, develop some proactive skills to reduce the likelihood of future ethical problems, [. . .] indoctrinate the students to particular set of values [. . .] and teach the laws related to a particular profession to avoid malpractice suits (Gotterbarn 1992).

PA argues that there is no deep theoretical difference between CE and other professional ethics like business ethics, medical ethics or engineering ethics, only a variety of pedagogical contexts (Gotterbarn 1991, 1992). And since CE courses have the goal of creating ethically-minded professionals not ethicists, it is not necessary, and it may actually be better not to have philosophers teaching them. After all

> Philosophers are no more educated in morality than their colleagues in the dairy barn; they are trained in moral theory, which bears about the same relation to the moral life that fluid mechanics bears to milking a cow (Robert K. Fullinwider, cited in Gotterbarn 1992).

This argument is not uncommon in academe. Mathematics courses, for example, are taken by many faculties, from Engineering to Economics. But should the lectures be given by mathematicians, who are presumably masters of the material, or lecturers from the application area, who may appreciate better its particular application? Apart from political and financial arguments, the latter view is often seen as reinforcing a sense of "subject" in the application domain, whilst the less applied view of the former may be seen as broadening an established subject with new applications. The arguments concerning the lecturing of Applied Ethics appear similar. It may perhaps be argued that at university level such courses ought to enable participants to solve new problems as they arise (what are the fundamentals?), whilst in specialised professional institutions such courses are typically under pressure to be more prescriptive:

> In applied professional ethics courses, our aim is not the acquaintance with complex ethical theories, rather it is recognizing role responsibility and awareness of the nature of the profession (Gotterbarn 1992).

PA has a number of major advantages. It stresses the vital importance of CE education, taking seriously issues like technical standards and requirements, professional guidelines, specific legislation or regulations, levels of excellence and so forth. It thus exposes the risky and untenable nature of NA and places PE in perspective, revealing it insufficient by itself. PA defends the value and importance of a constructive PE, by developing a "proactive" professional ethics (standards, obligations, responsibilities, expectations etc.), favourable to value-supporting and welfare-enhancing (development and uses of) ICT products (Bynum 1992; Gotterbarn 1992). Furthermore, PA defends a realistic pedagogical attitude, pragmatically useful to sensitise and instruct students and professionals. Its ultimate aim is to ensure that:

> ethical values, rules and judgements [are] applied in a computing context based on professional standards and a concern for the user of the computing artefact (Gotterbarn 1991).

One of the primary results of PA has been the elaboration and adoption of usage regulations and codes of conduct in ICT contexts (libraries, universities, offices etc.), within industry and in professional associations and organisations, as well as the promotion of certification of computer professionals. PA addresses mainly ICT practitioners, especially those involved in software development, where technical standards and specific legislation provide a reliable, if minimal, frame of reference. As we shall see, PA's goals are pedagogical not metaethical. Unfortunately, sometimes PA is interpreted as the only correct way to understand the whole field itself, as if CE could be reduced to a professional ethics in a strict sense:

> The only way to make sense of "Computer Ethics" is to narrow its focus to those actions that are within the horizon of control of the individual *moral*

computer professional (Gotterbarn 1991, see also Gotterbarn 1992; Gotterbarn 2001 presents a less radical view).

This strong view has further led to radically anti-philosophical positions (Langford 1995). However, a strong PA is far too restrictive, for at least three reasons. First, strong PA disregards the significant fact that, contrary to other purely professional issues, CE problems – e.g., privacy, accuracy, security, reliability, intellectual property and access – permeate contemporary life. Although strong PA can rightly argue that moral problems somehow involving ICT (e.g., theft using a computer) should not be vaguely confused with distinctively CE problems (e.g., software copyright issues), this restriction does not yet justify the reduction of all palpably CE problems to just professional issues. To be coherent, strong PA could reply that any citizen of an information society should be treated, to various degrees, as an ICT professional, to whom some corresponding professional guidelines should apply; but this would mean just accepting the fact that CE cannot be reduced to a specific professional ethics without the latter losing its perspicuous meaning. Strong PA becomes undefeatable but empty. Second, interpreting PA as providing a conceptual foundation for CE is to commit a mistake of levels. It is like attempting to define arithmetic on the basis only of what is taught in an introductory course. Without a theoretical approach PA is mere para-CE, to use an expression coined by Keith Miller and used by Bynum (1992) in analogy with paramedic, to describe a middle level between pop CE and theoretical CE. Theoretical CE underpins PA and requires a different approach from it. Finally, understanding CE as just a professional ethics, not in need of any further conceptual foundation, means running the risk of being at best critical but naive, and at worst dogmatic and conservative. On the one hand, focusing on case-based analyses and analogical reasoning, a critical PA will painfully and slowly attempt to rediscover inductively ethical distinctions, clarifications, theories and so forth already available and discussed in specialised literature. On the other hand, an uncritical PA will tend to treat ethical problems and solutions as misleadingly simple, non-conflicting, self-evident and uncontroversial, a matter of mere indoctrination, as exemplified in "The 10 Commandments of Computer Ethics" approach. Deferring to some contingent "normal ethics" currently accepted within the agent's society, to adapt a Kuhnian expression, is itself a very significant ethical decision at least because, when normal ethics is methodologically coherent, it limits itself to providing negative prescriptions, since lists of "don'ts" are easier to implement, and much less questionable, than positive recommendations. Moral standards, values and choices are always legitimised by ethical positions and arguments, at least implicitly. Applying normal ethics may then be sufficient in everyday life; but it is only the first step towards a mature approach that can uncover, evaluate, criticise and modify at least some of the accepted presuppositions working in CE, and thus hope to improve them. Thus, PA may be seen pragmatically as an historical first step towards a more mature CE.

## Theoretical CE and the uniqueness debate

Any applied or professional ethics must necessarily make room for critical theorising, even if it does not have to consider it one of its own tasks. PA at its best distinguishes between pedagogical problems and metatheoretical research, descriptive and normative questions, practical and theoretical issues, commonsensical applications and conceptual criticisms of some normal ethics. Among the fundamental questions that PA does not mean to address are: Why does ICT raise moral issues? Are CE issues unique (in the sense of requiring their own theoretical investigations, not entirely derivative from standard ethics)? Or are they simply moral issues that happen to involve ICT? What kind of ethics is CE? What justifies a certain methodology in CE, e.g., reasoning by analogy and case-based analysis? What is CE rationale? What is the contribution of CE to the ethical discourse? PA programmatically avoids entering into such investigations and coherently leaves them to theoretical CE. Theoretical CE can then be introduced as the logical stage following pop CE, NA and PA. Historically, it has developed along two lines, which can be usefully introduced through the "uniqueness debate". This has aimed to determine whether the moral issues confronting CE are unique, and hence whether CE should be developed as an independent field of research with a specific area of application and an autonomous, theoretical foundation. The debate arises from two different interpretations of the policy vacuum problem, one more radical, the other more conservative (Floridi 1999b is a collection of papers; Bynum 2000; Tavani 2000 are two overviews).

## The radical approach: CE as a unique discipline

According to the radical approach (RA), the presence of a policy and conceptual vacuum indicates that CE deals with absolutely unique issues, in need of a completely new approach (Mason 1986; Maner 1996, 1999). Thus, RA argues that

> [Computer Ethics] must exist as a field worthy of study in its own right and not because it can provide a useful means to certain socially noble ends. To exist and to endure as a separate field, there must be a unique domain for computer ethics distinct from the domain for moral education, distinct even from the domains of other kinds of professional and applied ethics. Like James Moor, I believe computers are special technology and raise special ethical issues, hence that computer ethics deserves special status (Maner 1999).

In terms of logical analysis, RA presents several advantages. It counteracts the risk run by NA of under-evaluating CE problems. Taking seriously their gravity and unprecedented novelty, RA improves on the various pop versions of CE, including PA, by stressing the methodological necessity of providing the field with a robust and autonomous theoretical rationale, if it wishes to deal with ICT-related moral issues successfully. Nevertheless, RA is confronted by at least four problems. First, to establish that CE is a unique field, the argument quoted above (Maner 1999) requires the explicit and uncontroversial identification of some unique area of study (actually, the argument appears to be of the form "uniqueness only if special domain" and "special domain" therefore "uniqueness". It is rectified if "only if" is replaced by "if"; in the original words, if "there must be" is replaced by "it suffices that there be", which at least renders the argument valid, if less plausible). Yet, RA seems unable to show the absolute uniqueness of any CE problem. None of the cases provided by Maner 1996 and 1999 is uncontroversially unique, for example. That is to be expected. It would be very surprising if any significant moral issue were to belong fully and exclusively only to a limited conceptual region, without interacting with the rest of the ethical context. It does not happen in any other special context such as business ethics, medical ethics and environmental ethics, and it remains to be shown why it should happen even in principle in CE. Second, in reply to the difficulty just seen, one could argue that CE problems could be made, or become, or discovered to be increasingly specific, until they justify the position defended by RA. This reply runs the risk of being safe but uninteresting because empirically unfalsifiable. It certainly keeps the burden of proof on the RA side. But let us suppose that a domain of unique ethical issues in CE were available in principle. The basic line of reasoning would still be unacceptable. The "uniqueness" of a certain topic is not simply inherited as a property by the discipline that studies it. On the one hand, specific moral problems – e.g., abortion, or the profit motive – may still require only some evolutionary adaptation of old macroethical solutions, that is theoretical, field-independent, applicable ethics, to e.g., medical or business ethics. On the other hand, specific disciplines, e.g., environmental ethics, are not necessarily so because they are involved with unique problems, for they may share their subjects – e.g., the value of life, the concept of welfare – with other disciplines, the difference resting, for example, in their methodologies, aims or perspectives. The other two problems encountered by RA are methodological. Given the interrelatedness of ethical issues and the untenability of the equation "unique topic = unique discipline", it is not surprising that RA is forced to leave unspecified what a mature CE could amount to in detail, as a unique discipline. Finally, by overstressing the uniqueness of CE, RA runs the risk of isolating the latter from the more general context of metaethical theories. This would mean missing the opportunity to enrich the ethical discourse.

### The conservative approach: CE as applied ethics

Some of the problems just seen are neatly avoided by the conservative approach (CA). CA defends two theses:

a) classic macroethics – e.g., Consequentialism, Deontologism, Virtue Ethics, Contractualism – are sufficient to cope with the policy vacuum. These theories might need to be adapted, enriched and extended, but they have all the conceptual resources required to deal with CE questions successfully and satisfactorily;
b) certain ethical issues are transformed by the use of ICT, but they represent only new species of traditional moral issues, to which already available metaethical theories need to, and can successfully, be applied. They are not and cannot be a source of a new, macroethical theory.

From (a) and (b) it follows that CE is a microethics, that is a practical, field-dependent, applied and professional ethics. Thesis (a) is weaker and hence less controversial than (b). To explain both, Johnson [2000] introduces an evolutionary metaphor (see also Naresh 1999 for a similar approach):

> Extending the idea that computer technology creates new possibilities, in a seminal article, Moor [1985] suggested that we think of the ethical questions surrounding computer and information technology as policy vacuums. Computer and information technology creates innumerable opportunities. This means that we are confronted with choices about whether and how to pursue these opportunities, and we find a vacuum of policies on how to

> make these choices. [. . .] I propose that we think of the ethical issues surrounding computer and information technology as new species of traditional moral issues. On this account, the idea is that computer-ethical issues can be classified into traditional ethical categories. They always involve familiar moral ideas such as personal privacy, harm, taking responsibility for the consequences of one's action, putting people at risk, and so on. On the other hand, the presence of computer technology often means that the issues arise with a new twist, a new feature, a new possibility. The new feature makes it difficult to draw on traditional moral concepts and norms. [. . .] The genus-species account emphasizes the idea that the ethical issues surrounding computer technology are first and foremost ethical. This is the best way to understand computer-ethical issues because ethical issues are always about human beings.

Since CA presents CE as an interface between ICT-related moral problems and standard macroethics, it enjoys all the advantages associated with a strong theoretical position. CA rejects NA. It accepts pop CE's recommendation: CE problems are important and significant, so much so that, for CA, they deserve to be approached both pragmatically and theoretically. It is compatible with, and indeed reinforces, PA since, for CA, CE is an ethics for the citizen of the information society, not just for the ICT professional. Being based on a macroethic perspective, CA can both promote a constructive attitude, like PA, and hope to adopt an evaluative stance, thus avoiding a naive or uncritical reliance on some contingent normal ethics. Finally, CA avoids RA's untenable equation and corresponding "isolationism", because the development of an evolutionary rather than a revolutionary interpretation of CE problems allows it to integrate them well within the broader context of the ethical discourse. Is then the CA position devoid of difficulties? Not yet, for CA is still faced by four shortcomings.

First, CA's weaker thesis (a) is controversial. It is at least questionable whether standard macroethics do indeed have all the necessary resources to deal with CE problems satisfactorily, without reducing them to their own conceptual frames and thus erasing their true novelty (Floridi 1999a provides arguments against (a)). It may be argued that precisely the fact that CE problems were unpredicted and are perceived as radically new casts doubts on the possibility of merely adapting old ethical theories to the new context.

Second, CA is metatheoretically underdetermined. The evolutionary metaphor incorporates the tension between a radical and a traditional approach but does not resolve it. New species of moral problem could conceivably be so revolutionarily different from their ancestors – the digital instrumentation of the world can create such entirely new moral issues, unique to CE and that do not surface in other areas – to require a "unique" approach, as suggested by RA. Or they may represent just minor changes, perfectly disregardable for any theoretical purpose, as the conservative approach wishes to argue. The trouble is that CA, left with the tension now hidden in the evolutionary perspective, opts for the conservative solution to treat CE as an applied ethics, but then it does not and cannot indicate which macroethics should be applied. At best, this leads to the adoption of some standard macroethics to deal with CE problems, e.g., Consequentialism and, when the choice is not arbitrary, this further justifies the claim that, in terms of ethical theorising, there is not much to be learnt philosophically from this applied field. If new ICT-related moral problems have any theoretical value, either by themselves or because embedded in original contexts, this is only insofar as they provide further evidence for the discussion of well-established ethical doctrines. In this way, CA approaches NA: there are no CE specific problems, only ethical problems involving ICT. At worst, CA's lack of commitment leads to a muddle of syncretic and eclectic positions, often acritical and overlooking the theoretical complexity of the problems involved. CA's lack of an explicit metaethical commitment generates a logical regress: having accepted CE as a microethics, one then needs a metatheoretical analysis to evaluate which macroethics is most suitable to deal with CE problems. This logical regress tends to be solved by appealing either to some common-sensical view or to pedagogical needs. The former solution leads back to the assumption of some contingent, normal ethics as providing CE rationale (to simplify: I do CE by using Habermas' dialogue ethics because this is what my society approves as normal ethics). It thus undermines the critical and normative advantage that CA hopes to have over other approaches. The latter solution (to simplify: I do CE by using Virtue Ethics because this is what my students find more intuitive), apart from being equally arbitrary, represents the kind of unnecessary intrusion of philosophy into professional matters so rightly criticised in PA literature. Software engineers should not be required to read the *Nicomachean Ethics*.

Third, CA is methodologically poor. This is a consequence of the first problem. Lacking a clear macroethical commitment, CA cannot provide an explicit methodology either. It then ends by relying on common-sense, case-based analysis and analogical reasonings, often insufficient means to understand what CA itself acknowledges to be new and complex issues in CE.

Fourth, CA is metaethically unidirectional. Arguing for (b), CA rejects a priori and without explicit arguments the possibility, envisaged by RA, that CE problems might enrich the ethical discourse by promoting a new macroethical perspective. It addresses the question "what can ethics do for CE?" but fails to ask the philosophically more interesting question "is there anything that CE can do for ethics?". It thus runs the risk of missing what is intrinsically new in CE, not at the level of problems and concepts, but at the level of contribution to metaethics. A mere extension of standard macroethics does not enable one to uncover new possibilities (Gorniak-Kocikowska 1996, for example, argues that computer ethics is the most important theoretical development in ethics since the Enlightenment).

## The innovative approach: Information ethics as the foundation of CE

There is a third possible approach to theoretical CE, which is neither conservative nor radical, but innovative (IA; Bynum 1998, 2000 outline the need for an innovative approach to CE). IA builds on CA's advantages, but it avoids its shortcomings by rejecting the conservative restriction made explicit in (b). According to IA, CE problems, the corresponding policy and conceptual vacuum, the uniqueness debate and the difficulties encountered by RA and CA in developing a cohesive metaethical approach strongly suggest that the monopoly exercised by standard macroethics in theoretical CE is unjustified. ICT, by transforming in a profound way the context in which moral issues arise, not only adds interesting new dimensions to old problems, but leads us to rethink, methodologically, the very grounds on which our ethical positions are based. Although the novelty of CE is not so dramatic as to require the development of an utterly new, separate and unrelated discipline, it certainly shows the limits of traditional approaches to the ethical discourse, and encourages a fruitful modification in the metatheoretical perspective. Rather than allowing standard macroethics to occupy the territory of CE arbitrarily, as happens with CA, or exiling CE in an impossibly isolated and independent position, as proposed by RA, IA argues that theoretical CE should be promoted to the level of another macroethics because it does have something distinctive and substantial to say on moral problems, and hence can enrich the metaethical discourse with a new and interesting approach of unquestionable philosophical value. In previous work, we have defined this macroethical perspective Information Ethics (Floridi 1998, 1999a; Floridi and Sanders 1999, 2001; Floridi and Sanders forthcoming-c). Information Ethics (IE), understood as the theoretical foundation of applied CE, is a non-standard, environmental macroethics, patient-oriented and ontocentric, based on the concepts of information object/infosphere/entropy rather than life/ecosystem/pain. The definition requires some comments.

The interpretation of what is the primary object of the ethical discourse is a matter of philosophical orientation. Some macroethical positions (e.g., Virtue Ethics) concentrate their attention on the moral nature and development of the agent. They are properly described as agent-oriented, "subjective" ethics. Since the agent is usually assumed to be a single person, they tend to be individualistic. Some other positions (e.g., Consequentialism, Contractualism and Deontologism) concentrate their attention on the moral nature and value of the agent's actions. They are "relational" and action-oriented theories, intrinsically social in nature. Agent-oriented, intra-subjective theories and action-oriented, inter-subjective theories can be defined as "standard" or "classic", without necessarily associating any positive evaluation with either of these two adjectives. Standard macroethics tend to be anthropocentric and to take only a relative interest in the "patient", the third element in a moral relation, which is on the receiving end of the action, and endures its effects. Ontic power, however, brings with it new moral responsibilities. One can respect only what one no longer fears, yet knowledge is a process of increasing emancipation from reality and, in a world in which humanity can influence, control or manipulate practically every aspect of reality, philosophical attention is finally drawn to the importance of moral concerns that are not immediately agent/action-oriented and anthropocentric. Medical Ethics, Bioethics and Environmental Ethics are among the best known examples of this non-standard approach. They attempt to develop a patient-oriented ethics in which the "patient" may be not only a human being, but also any form of life. Indeed, Land Ethics extends the concept of patient to any component of the environment, thus coming close to the object-oriented approach defended by IE (Rowlands 2000). Capturing what is a pre-theoretical but very common intuition, non-standard ethics hold the broad view that any form of life has some essential proprieties or moral interests that deserve and demand to be respected, even if not absolutely but minimally, i.e., in a possibly overridable sense. They argue that the nature and well-being of the patient constitute its moral standing and that the latter makes important claims on the interacting agent and in principle ought to contribute to the guidance of the agent's ethical decisions and the constraint of the agent's moral beha-

viour. Non-standard macroethics place the "receiver" of the action at the centre of the ethical discourse, and displace its "source" to its periphery, and in so doing they help to widen further the anthropocentric view of who or what may qualify in principle as a focus of moral concern. The development of ethical theories just sketched provides a useful explanation as well as a further, metatheoretical, justification of IE. The various difficulties encountered by other approaches to CE can be reconnected to the fact that, far from being a classic, agent/action-oriented ethics, as it may deceptively seem at first sight, CE is primarily an ethics of being rather than conduct or becoming, and hence qualifies as non-standard. The fundamental difference, which sets IE apart from all other members of the same class of non-standard theories, is that in IE information objects as such, rather than just living systems in general, are raised to the role of universal patients of any action (see Floridi, forthcoming-b for a full defence of this view). Biocentric ethics usually ground their analyses of the moral standing of bioentities and ecological systems on the intrinsic worthiness of life and the intrinsically negative value of suffering. IE suggests that there is something even more elemental than life, namely being, understood as information; and something more fundamental than pain, namely entropy. According to IE, one should also evaluate the duty of any rational being in terms of contribution to the growth of the infosphere, and any process, action or event that negatively affects the whole infosphere – not just an information object – as an increase in its level of entropy and hence an instance of evil. Without information there is no moral action, but in IE information moves from being a necessary prerequisite for any morally responsible action to being its primary object. The crucial importance of this radical change in perspective cannot be overestimated. Typical non-standard ethics can reach their high level of universalisation of the ethical discourse only thanks to their biocentric nature. However, this also means that even Bioethics and Environmental Ethics fail to achieve a level of complete impartiality, because they are still biased against what is inanimate, lifeless, intangible or abstract (even Land Ethics is biased against technology and artefacts, for example). From their perspective, only what is intuitively alive deserves to be considered as a proper centre of moral claims, no matter how minimal, so a whole universe escapes their attention. Now this is precisely the fundamental limit overcome by IE, which further lowers the necessary condition that needs to be satisfied, in order to qualify as a centre of moral concern, to the minimal common factor shared by any entity, namely its information state. And since any form of being is in any case also a coherent body of information, to say that IE is infocentric is tantamount to interpreting it, correctly, as an ontocentric theory. The ethical question asked by IE is: "What is good for an information entity and the infosphere in general?" The answer is provided by a minimalist theory of deserts: any information entity is recognised to be the centre of some basic ethical claims, which deserve recognition and should help to regulate the implementation of any information process involving it. Approval or disapproval of any information process is then based on how the latter affects the essence of the information entities it involves and, more generally, the whole infosphere, i.e., on how successful or unsuccessful it is in respecting the ethical claims attributable to the information entities involved, and hence in improving or impoverishing the infosphere. IE brings to ultimate completion the process of enlarging the concept of what may count as a centre of minimal moral concern, which now includes every information entity. This is why it can present itself as a non-standard, patient-oriented and ontocentric macroethics.

It may be objected that, as the theoretical foundation of CE, IE places the latter at a level of abstraction too philosophical to make it of any direct utility for immediate needs. Yet, this is the inevitable price to be paid for any attempt to provide CE with an autonomous rationale. One must polarise theory and practice to strengthen both (on IE as the ecological ethics of the new information environment see Floridi 2001 and forthcoming-a). IE is not immediately useful to solve specific CE problems but it provides the conceptual grounds that can guide problem-solving procedures in CE. Through IE, CE can develop its own methodological foundation, and hence support autonomous theoretical analyses of domain-specific issues, including pressing practical problems, which in turn can be used to test its methodology.

IE's position, like that of any other macroethics, is not devoid of problems. But it can interact with other metaethical theories and it contributes an important new perspective: a process or action may be morally good or bad irrespective of its consequences, motives, universality, or virtuous nature, but because it affects the infosphere positively or negatively. This is a major advantage. Without IE's contribution our understanding of moral facts in general, not just of CE problems in particular, would be less complete. The foundationalist debate in CE has lead to the shaping of a new ethical view.

## References


T.W. Bynum, editor. *Computers and Ethics*. Oxford: Blackwell, published as the October 1985 issue of *Metaphilosophy*, 1985.

T.W. Bynum. *Human Values and the Computer Science Curriculum*, http://www.southernct.edu/organizations/rccs/resources/teaching/teaching_mono/bynum/bynum_human_values.html, 1992.

T.W. Bynum. *Global Information Ethics and the Information Revolution*. Bynum and Moor, 274–289, 1998,

T.W. Bynum and J.H. Moor, editors. *The Digital Phoenix: How Computers are Changing Philosophy*. Oxford: Blackwell, 1998.

T.W. Bynum. A Very Short History of Computer Ethics. *APA Newsletters on Philosophy and Computers*: Spring February 1999, http://www.southernct.edu/organizations/rccs/resources/research/introduction/bynum_shrt_hist.html, 2000.

R. Epstein. *The Case of the Killer Robot*. New York: John Wiley and Sons, 1997.

L. Floridi. Does Information have a Moral Worth in Itself?. *Computer Ethics: Philosophical Enquiry (CEPE 98)*. London School of Economics and Political Science, London, December 14–15, preprint available at http://www.wolfson.ox.ac.uk/floridi/papers.htm, 1998.

L. Floridi. Information Ethics: On the Philosophical Foundation of Computer Ethics. *Ethics and Information Technology*, 1(1): 37–56, preprint available at http://www.wolfson.ox.ac.uk/floridi/papers.htm, 1999a.

L. Floridi, editor. *Etica and Politica*, special issue on *Computer Ethics*, 2, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html., 1999b.

L. Floridi. Ethics in the Infosphere. *The Philosophers' Magazine*, 6: pp. 18–19, 2001.

L. Floridi (forthcoming-a). Information Ethics: An Environmental Approach to the Digital Divide, *Philosophy in the Contemporary World* preprint available at http://www.wolfson.ox.ac.uk/floridi/papers.htm.

L. Floridi (forthcoming-b). On the Intrinsic Value of Information Objects and the Infosphere, preprint available at http://www.wolfson.ox.ac.uk/floridi/papers.htm.

L. Floridi and J.W. Sanders. Entropy as Evil in Information Ethics. Floridi, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html, 1999b.

L. Floridi and J.W. Sanders. Artificial Evil and the Foundation of Computer Ethics. *Ethics and Information Technology*, 3(1): 55–66, 2001; and also *Etica and Politica*, 2(2), 2000, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/2000_2/index.html, 2001.

L. Floridi and J.W. Sanders (forthcoming-c). On the Morality of Artificial Agents, preprint available from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm, in A. Marturano and L. Introna editors, *Ethics of Virtualities. Essays on the limits of the bio-power technologies*, to be published for the series Culture Machine, Athlone Press, London.

T. Forester and P. Morrison. *Computer Ethics: Cautionary Tales and Ethical Dilemmas in Computing*, 2nd ed. 1994. Cambridge, Mass: MIT Press, 1990.

K. Gorniak-Kocikowska. The Computer Revolution and the Problem of Global Ethics. *Science and Engineering Ethics*, 2(2), 1996.

D.W. Gotterbarn. Computer Ethics: Responsibility Regained, first published in the National Forum, rep. in *Business Legal and Ethical Issues*. Australian Computer Society August 1993 and in Johnson and Nissenbaum 1995, http://www-cs.etsu-tn.edu/gotterbarn/artpp1.htm, 1991.

D.W. Gotterbarn. The Use and Abuse of Computer Ethics, special ethics issue of *The Journal of Systems and Software*, 17(1), http://www.southernct.edu/organizations/rccs/resources/teaching/teaching_mono/gotterbarn02/gotterbarn02_intro.html, 1992.

D.W. Gotterbarn. Software Engineering Ethics. In J. Marciniak, editor, *Encyclopedia of Software Engineering*, 2nd ed. New York: Wiley-Interscience, 2001.

D.G. Johnson. Sorting Out the Uniqueness of Computer-Ethical Issues. Floridi, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html, 1999b.

D.G. Johnson and H. Nissenbaum, editors. *Computers, Ethics, and Social Values*. Englewood Cliffs, NJ: Prentice Hall, 1995.

D. Langford. *Practical Computer Ethics*. London: McGraw-Hill, 1995.

W. Maner. Is Computer Ethics Unique?. Floridi, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html, 1999b.

W. Maner. Unique Ethical Problems in Information Technology. *Science and Engineering Ethics*, 2(2): 137–154. Revised version in Maner (1999), 1996.

R. Mason. Four Ethical Issues of the Information Age. *MIS Quarterly*, 10(1): 5–12, 1986.

J.H. Moor. What is Computer Ethics?. *Metaphilosophy*, 16(4): 266–275, http://www.southernct.edu/organizations/rccs/resources/teaching/teaching_mono/moor/moor_definition.html, 1985.

S. Naresh. Ethical Norms for the Information Society. *Proceedings of the First Session of UNESCO's COMEST*, Oslo April 1999, pp. 169–177. Paris: UNESCO, 1999.

D.B. Parker. *Ethical Conflicts in Computer Science and Technology*. Arlington, VA: AFIPS Press, 1981.

D.B. Parker. Ethical Dilemmas in Computer Technology. In W.M. Hoffman and J.M. Moore, editors, *Ethics and the Management of Computer Technology*. Cambridge, MA: Oelgeschlager, Gunn and Hain, 1982.

D.B. Parker. *Ethical Conflicts in Information and Computer Science, Technology, and Business*. Wellesley, MA: QED Information Sciences, 1990.

M. Rowlands. *The Environmental Crisis – Understanding the Value of Nature*. New York: St. Martin's Press, 2000.

H.T. Tavani. Computer Ethics: Current Perspectives and Resources. *APA Newsletters on Philosophy and Computers*, Spring February, 1999, http://www.apa.udel.edu/apa/publications/newsletters/v99n2/computers/feature-tavani.asp, 2000.