

This is a preprint version of the following article:

Brey, P. (2012). Anticipating Ethical Issues in Emerging IT. *Ethics and Information Technology* 14, 305-317.

# Anticipating Ethical Issues in Emerging IT

## Abstract

In this essay, a new approach to the ethics of emerging information technology will be presented, called *anticipatory technology ethics* (ATE). The ethics of emerging technology is the study of ethical issues at the R&D and introduction stage of technology development through anticipation of possible future devices, applications, and social consequences. In the essay, I will first locate emerging technology in the technology development cycle, after which I will consider ethical approaches to emerging technologies, as well as obstacles in developing such approaches. I will argue that any sound approach must centrally include futures studies of technology. I then present ATE and some applications of it to emerging information technologies. In ATE, ethical analysis is performed at three levels, the technology, artifact and application levels, and at each level distinct types of ethical questions are asked. ATE analyses result in the identification and evaluation of a broad range of ethical issues that can be anticipated in relation to an emerging information technology. This ethical analysis can then be used for ethical recommendations for design or governance.

**Keywords:** Anticipatory technology ethics; emerging technologies; uncertainty; futures studies; forecasting; technology assessment; ethical impact assessment; design; governance

## 1. INTRODUCTION

Jim Moor (2005) has argued that we need better ethics for emerging technology. Current ethics, he holds, is insufficiently equipped to address the revolutionary changes that are being brought about with new and emerging technologies. He argues that we need ethical approaches that are better informed concerning new technologies and their social consequences and that are more proactive in identifying and addressing ethical issues in relation to them. I agree. Far too little work has been done on developing such approaches. The aim of this paper is to answer to Jim Moor's call by presenting a particular approach for the ethical analysis of emerging technology, with specific reference to information technology.

The paper is structured as follows. In the next section, the difference will be discussed between emerging and established technology, and various stages of emergence will be distinguished. In section 3, issues and problems for an ethics of emerging technology will be discussed, and recent ethical approaches to emerging technology will be discussed and critiqued. In section 4, my own approach will be presented, which is called *anticipatory technology ethics* (ATE). I will present ATE as a promising new approach that builds on previous approaches. In section 5, I will then discuss how this approach can be applied in the ethical analysis of emerging information technologies.

## 2. ETHICS AND STAGES OF TECHNOLOGICAL EMERGENCE

What are the characteristics of an emerging technology and how is it different from a technology that is fully realized and established in society? And how would an ethics of emerging technology be different from an ethics of established technology? I aim to answer these questions through a discussion of Jim Moor's three-stage model of the development of new technology. Moor distinguishes three stages: the introduction stage, the permeation stage and the power stage. In the *introduction stage*, few implementations of the technology exist, and the technology is still largely something that exists in labs and drawing boards. The few devices that exist are often seen as intellectual curiosities or playthings. The cost of the technology is high, few people know and use it, and the technology's integration into society is minor. In the *permeation stage*, the technological devices have become more conventional and standardized. The cost is dropping, and the number of users is growing. The integration into society is moderate. In the *power stage*, finally, the technology is firmly established. It is readily available, low-

cost, and widely used, and many people in society are affected by it directly or indirectly. The impact on society is therefore large.

Moor argues that ethical problems multiply as technology moves from the introduction to the power stage. This is because at the introduction stage, there are few users, few different types of devices, and few uses of them. There are therefore few new types of activities and situations that may evoke ethical problems. As the technology evolves, there will be many users, uses, and types of devices, and therefore more ethical problems to consider. He summarizes this progression in ethical complexity in “Moor’s law”: As technological revolutions increase their social impact, ethical problems increase. Moor advocates, however, that for emerging technologies, ethicists do not wait until new technological devices and uses manifest themselves, but rather that they become proactive. This means that ethicists should “learn about the technology as it is developing and to project and assess possible consequences of its various applications.” (p. 119)

Moor’s account is compelling, but I believe that it skips an important stage in the development and diffusion of new technology. This is what I will call the R&D (research & development) stage. This stage comes before the introduction stage. Whereas in the introduction stage, there are already early applications of a new technology, in the R&D stage these applications do not yet exist. Rather, research is directed at the development of basic techniques that may further down the road result in concrete applications. Research in nanorobotics, for example, is concerned with the developments with machines or robots with components whose size is at or near the nanolevel. Current research is focused on developing basic techniques for the development of such nanobots. No actual nanorobots currently exist, so the technology is still wholly in the R&D stage. Tissue engineering is an example of a technology that is currently moving from the R&D stage to the introduction stage. It is concerned with developing techniques for the generation of artificial tissues to restore, maintain, or improve biological functions. Such research may eventually result in various devices and applications, which are now starting to emerge. Emerging technology can be defined as technology at the R&D and introduction stages. As technology moves to the permeation and power stage, it becomes established technology that is entrenched in society.

Ethics at the R&D stage is likely to be different from ethics in the three later stages because in the R&D stage, applications are still largely speculative, whereas in the later stages, many concrete applications are already in development or are being considered. Ethics in the R&D stage will focus on general ethical issues in relation to new techniques and on speculative ethical analysis of possible future applications. Ethics in the introduction stage will focus mostly on present and future applications that are already being considered. Within the R&D stage, moreover, a further distinction can be made between the *research stage* and the *development stage*. Research focuses on basic techniques, principles and methods that can be used for later development of concrete devices or processes, whereas development focuses on the actual design and manufacture of devices and processes. In the research stage, no knowledge may yet exist about any possible devices or applications that may result from the research, so ethical reflection on future consequences may be wholly speculative at this stage. In the development stage, actual devices and processes are being designed and developed, so ethical reflection can have a more concrete focus. Ethics at the introduction stage is less speculative as more is known about possible devices and applications, but it is still speculative in part, as many devices that may be developed on the basis of the technology still do not exist, and there is still uncertainty about the ways in which newly developed devices will be used and what will be the social consequences.

Ultimately, ethical assessment of emerging technologies concerns the question of what is good and bad about the devices and processes that they may bring forth, and what is right and wrong about ways in which they may be used. Since at the R&D and introduction stages many devices, usage patterns and social consequences are not yet present, ethical assessment turns speculative, as it focuses on particular R&D activities and techniques and then projects possible devices and usage patterns which are then assessed ethically. Such assessments may then be used to make ethical recommendations for R&D practices themselves, so as to increase the likelihood that these practices yield morally desirable devices and uses. Or they may be used for policy.

### **3. ETHICAL APPROACHES TO EMERGING TECHNOLOGIES**

Because the ethical analysis of emerging technology involves consideration of future events, it has to consider an epistemological problem, the *problem of uncertainty* concerning future devices, applications, uses and social consequences of emerging technology (Sollie, 2007). The basic issue is: how can we do ethical analysis of technologies when we do not have reliable knowledge of their future applications and social consequences? It would seem undesirable that ethicists lose themselves in speculation over future ethical issues that might emerge but that may be based on completely inadequate projections of future developments. On the other hand, it would also seem undesirable that ethicists remain silent about emerging technologies because of this uncertainty. Remaining silent would have the implication that ethicists can only comment on technologies when they are fully developed and their

impacts on society become transparent. By then it has become much more costly and difficult to steer technology development into a more ethical direction.

Two ethical approaches are possible in response to this problem, one more conservative and reliable, the other more uncertain and speculative. The first approach is to focus only on ethical issues that can be known or predicted reliably. These are ethical issues that seem unavoidable in any future application or use of the technology. Such issues correlate with characteristics inherent to the technology, or characteristics that are likely to manifest themselves over a broad range of possible applications. For example, when nuclear energy technology was being developed it was known early on that whatever particular systems or devices would be built on the basis of it, there would be a problem of radioactive waste, which requires ethical deliberation. Likewise, when genetic technology was being developed it was known from the beginning that it would involve the modification of genetic material, which was considered to be intrinsically morally controversial. So even when particular applications or uses are not yet known, it is often possible to identify generic ethical issues that are likely to manifest themselves as the technology progresses, and these can be discussed at an early stage. Because this approach identifies generic aspects of technologies that pose ethical issues I will call this approach the *generic approach*.

A second approach is more speculative, in that it actually aims to predict or hypothesize devices, uses and social consequences of an emerging technology. Ethicists may either develop such forecasts themselves or rely on existing forecasting studies. The forecasts are then used to explore ethical issues that would present themselves if these forecasts were to come true. For example, ethicists may forecast that nanotechnology will yield applications for targeted drug delivery in the human body using nanoparticles, and that such applications will become widely available to both doctors and patients. On the basis of such a forecast, ethical issues may then be identified that may occur when such devices are being used, such as issues of privacy, confidentiality and informed consent. I will call this approach the *forecasting approach* to the ethics of emerging technology.

Ethicists will not always be well-equipped to engage in forecasting studies themselves, and may rely on forecasting studies performed by other scholars. Forecasting studies of technological devices, uses and social consequences are performed in two related fields: futures studies and technology assessment. *Futures studies* is a field that aims to study possible or probably futures (Bell, 1997). Futures research includes many different forecasting approaches, such as environmental scanning, causal layered analysis, the Delphi method and scenario methods. Some of these, like the Delphi method, rely on the consultation of experts in various fields. Others may rely on surveys, time series analysis, regression analysis, or simulations. Some of the work in futures studies focuses on technology forecasting. In such studies, future technologies are forecast, including the development spread of certain types artifacts, and in some cases their utilization and the social consequences that may result from their use. The other field that engages in forecasting of new technologies is *technology assessment* (TA), a field that studies the effects of new technologies on industry, the environment and society, evaluates such effects and develops instruments to steer technology development in more desired directions (Tran & Daim, 2008; Grunwald 2009). TA makes such assessments on the basis of known or potential applications of the technology. Thus, TA in part relies on, and in part engages in, futures studies. Both futures studies and TA can hence be useful for forecasting the development of emerging technologies.

The generic and forecasting approach each have their benefits, and they are not mutually exclusive. The forecasting approach has an advantage over the generic approach that it is able to consider more ethical issues, by including not only those that are generic to the technology but also those that are specific to projected future devices and their uses. Its potential disadvantage is that its ethical assessments are based on forecasts that are to some degree speculative and that may be incorrect. However, to the extent that forecasts can be reliable, a forecasting approach will be able to anticipate many more ethical issues than a mere generic approach would, and would therefore be more valuable. In this essay, therefore, I focus on ethical approaches that have a major emphasis on forecasting.

While there are many forecasting approaches to technology, there are hardly any ethical approaches to technology that involve forecasting. It is only in the past five or six years that such approaches have started to develop. One of the first is *ethical technology assessment* (eTA), proposed by Palm and Hansson (2006), which has as its purpose “to provide indicators of negative ethical implications at an early stage of technological development” (p. 543). This approach relies on studies in technology assessment and on close interactions with developers of technology. The goal of eTA is not to predict far into the future, but rather to continually assess current practices in technology development and provide feedback to designers and policy makers. The ethical analysis of an emerging technology takes place by confronting projected features of the technology or projected social consequences with ethical concepts and principles. Palm and Hansson propose an ethical checklist of nine issues to identify the most common ethical issues in emerging technologies. This list contains issues like privacy, sustainability, issues of control, influence and power and issues of gender, minorities and justice. Not all of these issues are ethical in a conventional sense, but all can be framed as ethical issues.

Palm and Hansson's approach does a good job at advocating the need for ethical TA, and then presents an original approach that seems workable and appears to cover a lot of different issues. However, eTA is rather vague in its methodology, as it does not specify in detail what kind of knowledge needs to be acquired from technology developers and from TA and how it should be acquired, and it also does not spell out in detail how ethical analysis can be performed on the basis of this knowledge. In addition, the ethical checklist of nine items seems somewhat limited. A final issue is that eTA only looks at the near future and does not seem suitable for ethical assessments over longer timeframes.

Another recent approach is *ethical impact assessment* approach by David Wright (2010). This is an approach for the ethical evaluation of new information technologies by developers to ensure that ethical issues are taken into account in their development. Wright's approach relies on an extended ethical checklist which contains ethical values and principles along with sets of questions raised by these values and principles, questions that must be answered during the ethical impact assessment. For example, his checklist contains a principle of universal service will include such questions as "Will the project or service be made available to all citizens? When and how will this be done?" and "Will training be provided to those who do not (yet) have computer skills or knowledge of the Internet?". For answering these questions, Wright relies on stakeholder discussions, including expert workshops, consultations and surveys.

A strong point of Wright's approach are that it contains an elaborate ethical checklist with dozens of questions to be answered regarding the ethical aspects of new technologies, and that it contains elaborate procedures for involving stakeholders in the ethical analysis. A weak point, however, is that Wright does not make clear how forecasting takes place of expected artifact, applications and social consequences. Apparently, the technology developers and other stakeholders are believed to be able to come up with this information somehow. In addition, it is not explained how participants are capable of answering questions in the ethical checklist on the basis of the knowledge that they have. All in all, Wright's approach seems to be more suited for the ethical assessment of concrete IT design projects, in which design specifications already exist and the context of use is already known, than for the broad ethical assessment of emerging information technologies about which many uncertainties still exist.

A third recent approach, the *techno-ethical scenarios approach* of Boenink, Swierstra and Stemerding (2010) aims at ethical assessments of emerging technologies that are intended to help policy makers to anticipate ethical controversies regarding emerging technologies. It relies on scenario analysis, which is a well-established approach within futures studies. A unique feature of the approach is that it aims to anticipate the mutual interaction between technology and morality, and changes in morality that may result from this interaction. They want to take such changes into account when ethically assessing new technologies, so that new technologies are not evaluated from within a moral system that may not have the same validity by the time an emerging technology has become entrenched in society. This approach has several benefits over other approaches, such as a focus on detailed scenarios and attention to moral change. A major limitation, however, is that the approach is descriptive and predictive, rather than normative and prescriptive. It describes moral issues that are likely to emerge as the technology progresses, not ones that ought to emerge from an ethical point of view. And it considers how these are likely to be resolved, not necessarily how they ought to be resolved. What this approach may miss, as a result, are ethical issues that are unlikely to collect much public attention but that are nevertheless important (cf. Brey, 2000). Conversely, it may identify moral controversies that may emerge in public debate that are based on a false or misguided understanding of the technology or its social consequences.

A fourth and final approach, the *ETICA* approach (Stahl et al., 2010) is a method for the ethical assessment of emerging information and communication technologies (ICTs).<sup>1</sup> It is so general in scope, however, that nothing prevents its application to other types of technology as well. Thus conceived, the aim of the ETICA approach is to provide comprehensive overviews of ethical issues for emerging technologies that are likely to play out in the future, with an emphasis on the medium term. The ETICA approach makes use of projections of the future which it derives from futures research. It aims to arrive at *foresight analyses*, which are forecasting analyses that consider multiple possible futures, out of which one is chosen as most desirable or important to consider. The ETICA approach relies on multiple futures methods and studies, which are used to identify a range of projected artifacts and applications for particular emerging technologies, along with capabilities, constraints and social impacts. These data form the basis for ethical analysis, which consists of three stages. In the first stage, the identification stage, ethical issues are identified for particular applications, artifacts or technological properties.<sup>2</sup> Most of the ethical values and principles that are used at this stage are derived from a prior checklist. In a second stage, the evaluation stage, the ethical issues

---

<sup>1</sup> See also <http://www.etica-project.eu/>, especially the deliverables.

<sup>2</sup> The ETICA project also uses these data to perform social and legal analyses. However, in my discussion I will focus on its use for ethical analysis.

of the identification stage are subjected to ethical evaluation and are ranked and ordered in relation to each other. In a third and final stage, the governance stage, governance recommendations are developed for policy makers for dealing with the ethical issues described in the earlier stages.

The ETICA approach is possibly the most elaborate and promising ethical approach to emerging technologies that has been developed to date. It aims at sound ethical analysis as well as at thoroughness by considering a wide range of technological properties, artifacts, applications, and ethical issues. And it aims to make use of state-of-the-art work in futures studies. My own approach, presented below, derives inspiration from it, and adopts its distinction between its three stages of analysis: ethical identification, ethical evaluation and governance. Yet, the approach also has weaknesses. First, its claim to adopt a futures studies approach is somewhat dubious, as the main sources for locating ethical issues that have been used in the ETICA approach are texts which are not based on futures research. A second weakness is that many of the ethical analysis undertaken in the ETICA project appear to refer to generic properties of the technologies that are studied. In the project these are called “ethical issues stemming from the defining features of the technology” (Heersmink, van den Hoven & Timmermans, 2010, p. 27). The range of artifacts and implications that is considered is often somewhat limited, and elaborate descriptions of possible artifacts and applications are often missing.

## 5. ANTICIPATORY TECHNOLOGY ETHICS

What we have seen in the previous section is that an ethical approach to emerging technologies has to overcome various obstacles: it has to engage in forecasting without becoming too speculative, it has to integrate forecasting analyses with ethical analysis in which a normative stance is maintained, and it has to attain a certain breadth in doing so. Existing approaches succeed in some, but not all of these challenges. Based on the previous discussion, I will now present an ethical approach of my own, which I will call *anticipatory technology ethics (ATE)*. I will argue that ATE has the potential to meet all the criteria that a sound approach to ethical analysis of emerging technologies should have. ATE distinguishes itself from other approaches in its definition of objects of analysis, its particular approach to forecasting, and its methods of ethical analysis. I will now discuss these in turn.

### Levels and objects of ethical analysis

A first characteristic of ATE is that it contains *three levels of ethical analysis*: the technology, artifact and application level. At each of these levels, various *objects of ethical analysis* are defined, which are things, properties or processes that raise ethical issues. Its three levels of analysis are similar to those of the ETICA approach, which distinguishes defining features of a technology, artifacts and applications. However, in ATE a more refined conceptual apparatus is developed through which a larger variety of objects of ethical analysis is defined.

The *technology level*, to start with, is the level at which a particular technology is defined, independently of any artifacts or applications that may result from it. A *technology* is a collection of techniques that are related to each other because of a common purpose, domain, or formal or functional features. Nuclear technology, for example, is the collection of techniques for the fission and fusion of atomic nuclei. Nanotechnology is the collection of techniques for manipulating matter on an atomic and molecular scale. And biometric technology pertains to methods for the measurement and recognition of physical and behavioral traits of humans for identification and authentication purposes. A *technique* is a procedure to accomplish a specific activity or task. For example, nanotechnology embodies such techniques as solid state silicon methods, focused ion beams, and molecular scale electronics. Techniques may depend on technological methods, processes, tools, knowledge and skills that make them possible. Within a technology, it is often possible to distinguish *subclasses* that are distinguished by a more specific purpose, domain, or set of features than the parent class. For example, in nanotechnology, it is possible to distinguish bionanotechnology, optical nanotechnology and DNA nanotechnology. Such subclasses are also technologies themselves.<sup>3</sup>

At the technology level, ethical analysis focuses on features of the technology at large, particular subclasses of it, or techniques within it. It then considers generic ethical issues that are attached to these features. These are either ethical issues inherent to the character of the technology, issues that pertain to consequences that are likely to manifest themselves in any or nearly any artifact or application of the technology, or issues pertaining to risks that the technology will result in artifacts or applications that are morally problematic. Genetic engineering, for example,

---

<sup>3</sup> Some technologies are defined in terms of specific types of artifacts that they aim to develop and use. Examples are fuel cell technology and membrane technology. In such technologies, the technology and artifact level blend into each other.

involves the manipulation of DNA in cells and organisms. This is a defining feature of the technology. At the technology level, a generic ethical issue is whether such manipulation violates natural order or the dignity of life. When nuclear technology was being developed, a moral discussion emerged whether the technology should be developed at all because of the potential to build a nuclear bomb. So here the technology is ethically criticized because of its potential to lead to dangerous or morally problematic applications. And nuclear energy technology can be critiqued for developing energy solutions that inevitably generate a problem of nuclear waste.

Let us now turn to the *artifact level*. On the basis of a technology, functional artifacts, systems and procedures are developed. For example, nuclear technology has yielded artifacts like nuclear reactors, nuclear bombs, x-ray imaging systems and ionization smoke detectors. It has also yielded procedures such as food irradiation and nuclear well logging. An *artifact* is a physical configuration that, when operated in the proper manner and in the proper environment, produces a desired result.<sup>4</sup> A *procedure* is a sequence of actions that, when performed in the proper manner in the proper environment using the proper tools, produces a desired result. The useful products of technology are technological artifacts and procedures. They are often the result of combining novel techniques within a technological field with more conventional techniques of engineering to produce artifacts and procedures that can be used in practice. Within each class of artifacts and procedures, it is moreover possible to distinguish various *subclasses*. For example, within the class of robots, one can distinguish subclasses of humanoid, industrial, mobile, and service robots. Similarly, there are often subtypes within a particular class of procedures.

At the artifact level, ethical analysis focuses on types of artifacts and processes that have resulted or are likely to result from a particular technology. It considers features of them that present moral issues. As was at the case at the technology level, such moral issue may present themselves for three reasons: because of the inherent character of the artifact, because the artifact has certain unavoidable consequences in most or all of its uses, or because certain potential applications of the artifact are so risky or morally controversial that it warrants reflection on the ethical justification of its manufacture. Examples are xxx , automobiles that produce greenhouse gases, smartphones that store and disseminate location data of users, and nerve gas weapons that can cause horrible agony and disfigurement.

At the *application level*, finally, ethical analysis focuses on particular ways of using an artifact or procedure, or on particular ways of configuring it for use. An *application*, as I will define it, is the concrete use of a technological artifact or procedure for a particular purpose or in a particular context, or a specific configuration of an artifact to enable it to be used in a certain way. Put differently, an application is a way of using or configuring an artifact or procedure. For example, a particular service robot may be configured and used to perform household chores, to assist the disabled, or to perform industrial tasks. These are different applications of it. The term “application” is sometimes used in a different way. Technological artifacts and procedures are sometimes called applications themselves. For example, an electro-galvanic fuel cell may be called an application of fuel cell technology. I will not use the term in this way, but will only use it to refer to particular uses or configurations of technological artifacts and procedures.<sup>5</sup>

Another way to think of an application is a situation in which one or more aspects of the *context of use* of an artifact or procedure are fixed. Such aspects may include the particular purposes for which the artifact is used (e.g., industrial vs. domestic use; cleaning vs. carpenting), the manner in which it is used (e.g., manually vs. automatically; for short or long durations), the characteristics of its users (e.g., male vs. female, skilled vs. unskilled, Western vs. non-Western), aspects of the social or physical context of use, properties of the technological configuration within which the artifact functions, and so on. As more and more elements of the context of use are fixed, more specific ethical issues may emerge from the dynamic interplay between artifact and its contextual elements. The use of artifacts by specific groups of users for specific within specific social, cultural and institutional arrangements will give rise to all kinds of ethical issues that are specific to these users, purposes, and contextual elements. This is what is being considered at the application level.

Let us consider ethical issues that may play at the application level. A first group consists of moral issues that concern the morality of certain purposes for which an artifact or procedure may be used. For example, moral issues may be raised by the use of in vitro fertilization for impregnating older women, the use of morphine for mercy killing, or the unauthorized use or dissemination of proprietary software. A second group consists of moral issues concerning side-effects or unintended consequences that occur in certain uses, in certain contexts of use, or for certain user groups. For example, a drug may cause cancer at a disproportionate rate for certain user groups, when used in combination with certain other drugs, or when used for an extended period of time. And computer games

---

<sup>4</sup> Certain complex artifacts, like power plants and railroad systems, may involve human actors as well. In such cases, human actors playing predefined roles are part of the design of the artifact, and the artifact is hence not a completely physical entity but also, in part, a social one.

<sup>5</sup> Even then, there are terminological difficulties. Should we say that an industrial robot is application of service robots to the industrial domain? Or should we rather say that an industrial robot is a subtype of service robots? This may depend on whether industrial robots are designed differently in important ways than generic service robots or whether they are only used in a different way.

may exacerbate social isolation for those individuals who already have weak social ties. A third group consists of moral issues pertaining the rights and interests of stakeholders who may be affected by a particular use of an artifact. For example, the use of new medical procedures without informed consent violates the rights of patients, and the construction and use of a power plant in a way that does not take into account concerns about pollution and noise by members of the local community also presents moral issues.

To conclude, I have identified three levels of analysis for an ethics of emerging technologies: (1) the *technology level*, at which morally relevant features of the technology at large are studied, as well as features of subclasses of the technology and particular techniques; (2) the *artifact level*, at which morally relevant features of artifacts and procedures are analyzed, as well as of particular subcategories of them; (3) the *application level*, at which morally relevant features of particular uses or configurations of artifacts and procedures are analyzed. The ethics of emerging technology should, I believe be aimed at all three levels. At the technology level, fundamental ethical issues pertaining to the technology are studied, whereas at the artifact and application levels, more specific and contingent issues are studied. It should not be concluded that the fundamental ethical issues are necessarily more important than the specific ones. They are more generic, but possibly of lesser importance than certain specific issues. For instance, any fundamental ethical issues with nuclear technology are probably of less ethical importance than specific issues relating to nuclear weapons and nuclear energy.

## Forecasting methods

I hold that different forecasting methods are required for the technology, artifact and application levels. I agree with Palm and Hansson that at the technological level, an understanding of the technology is best acquired from engineers. They are best positioned to describe the features that define the technology entails, the particular techniques and subclasses of technology that it contains and the techniques that may be developed in the future. Both for the present and future state of a technology, engineers are best positioned to inform ethicists, and we most likely need no consultation of experts from other fields or separate futures studies to get knowledgeable about the technology.

For the artifact and application levels, projections of the future are needed, this requires that ethicists either utilize or engage in futures studies. But how should they do this. First, I hold, they should utilize existing studies in forecasting and TA about the technology, to the extent that these are available. These provide ethicists with a first view of artifacts and applications that are likely to emerge in the future. Second, ethicists should initiate expert surveys and roundtable discussions with experts that yield expert predictions of possible or likely future artifacts and applications. Relevant experts would include engineers, technology forecasters and TA experts, as well as historians and sociologists of technology and marketing experts.

It would be useful if these experts would also reflect on the plausibility of projected futures in the forecasting studies that are being considered. Also, because the conjecturing of future artifacts and applications is an imaginative activity, it may be useful to consider policy documents, company studies, academic texts or even SF stories for ideas about possible future artifacts and applications, as long as these ideas are then subjected to scrutiny regarding their feasibility and plausibility. The consultation of existing futures studies and of relevant experts are both important steps to take, and may in many cases be sufficient. However, if these steps do not yield enough insight, it may be necessary for ethicists to do their own futures studies as well, possibly in tandem with futures studies researchers.

A thorough forecasting analysis of a new technology would consider how it is likely to evolve and mature over time, how it might be combined with other new and existing technologies to yield new artifacts and procedures and new application areas, and for which of these artifacts and procedures there is likely to be both significant demand and the possibility to realize a stable supply. It would do this systematically for different application domains, such as healthcare, food, transportation, entertainment, security and defense. And it would distinguish various types of applications by varying elements in the context of use, such as user types, use environments, and usage patterns. The result of such an analysis would be a systematic timed prediction of possible future artifacts and applications in various domains. This, of course, would be the ideal. Particular forecasting analyses may be less elaborate because of limitations in time and resources.

Out of all the possibilities, ethicists have a particular interest in those artifacts, applications and social consequences that may cause harm, violate rights, affect well-being, or cause unjust distributions of goods. This particular interest may imply that ethicists will sometimes have to develop their own forecasts and scenarios that focus on such matters. For instance, in studying future point-of-care testing devices that bring medical testing to the site of patient care, ethicists may want to consider specifically the potential impact on for different social groups, so as to be able to explore issues of distributive justice. In studying future deep brain implant techniques for psychiatric treatment, ethicists may want to explore in more detail the possibilities of abuse of such techniques, or potentially

negative side-effects on the well-being or autonomy of patients. Thus, ethicists will likely want to do extended futures studies of at least some artifacts and applications, in order to identify ethical issues that may not be transparent in the less specialized analyses from futures research.

## Methods of ethical analysis

Technological forecasting, as described above, results in descriptions of present and anticipated technologies, artifacts and applications. These descriptions constitute the input for ethical analysis. I agree with the ETICA approach that there are two stages to such ethical analysis: a first one in which ethical issues are identified (the identification stage) and a second one in which they are evaluated (the evaluation stage). Optionally, in a third stage the results of ethical analysis may be used to make ethical recommendations for technology development or for governance.

Let us now consider the *identification stage*. At this stage, descriptions of the technology are cross-referenced with ethical values and principles. It is investigated if features of the technology are likely to negatively impact moral values or principles. For instance, it is investigated if a future neurimaging system that makes cognitive processes visible may possibly harm privacy or autonomy. The question is how ethicists determine whether a particular technology, artifact or application may negatively impact moral values and principles. The general way in which this is done, I hold, is through an *operationalization* of the value or principle, which is a description of it that specifies real-world conditions for its realization or frustration. For instance, information privacy can be described as the right to control access to personal information about oneself. The real-world conditions that must be present for this value to be realized are hence that people have the ability to control access to such personal information. At the identification stage, it can then be ascertained whether particular information systems, as described at the forecasting stage, are likely to allow for such control, or whether there is a significant probability that such control will be absent.

Another question is how ethicists arrive at the values that they cross-reference with the technology. All three previously discussed ethical approaches in some way make use of an *ethical checklist* that contains ethical values, principles or arguments. I agree that such a checklist can be useful. It may help one to identify ethical issues that might otherwise have been missed. Such an ethical checklist should contain those ethical values and principles that are widely accepted in society and in ethics.<sup>6</sup> Table 1 represents an attempt at such a list.

Table 1 *The anticipatory technology ethics checklist*

- Harms and risks
  - o Health and bodily harm
  - o Pain and suffering
  - o Psychological harm
  - o Harm to human capabilities
  - o Environmental harm
  - o Harms to society
- Rights
  - o Freedom
    - Freedom of movement
    - Freedom of speech and expression
    - Freedom of assembly
  - o Autonomy
    - Ability to think one's own thoughts and form one's own opinions
    - Ability to make one's own choices
    - Responsibility and accountability
    - Informed consent
  - o Human dignity
  - o Privacy
    - Information privacy

---

<sup>6</sup> For particular purposes, it may be useful to employ more specific lists, e.g., lists that reflect European values, Asian values, conservative values or Christian values. In addition, it may be useful to develop specific checklists for specific types of technology. E.g., a checklist for information technology may focus on such values as privacy, security and accountability, whereas a checklist for medical technology may focus on such values as beneficence, nonmaleficence, human dignity and informed consent.



- Bodily privacy
    - Relational privacy
  - Property
    - Right to property
    - Intellectual property rights
  - Other basic human rights as specified in human rights declarations (e.g., to life, to have a fair trial, to vote, to receive an education, to pursue happiness, to seek asylum, to engage in peaceful protest, to practice one's religion, to work for anyone, to have a family, etc.)
  - Animal rights and animal welfare
- Justice (distributive)
  - Just distribution of primary goods, capabilities, risks and hazards
  - Nondiscrimination and equal treatment relative to age, gender, sexual orientation, social class, race, ethnicity, religion, disability, etc.
  - North-south justice
  - Intergenerational justice
  - Social inclusion
- Well-being and the common good
  - Supportive of happiness, health, knowledge, wisdom, virtue, friendship, trust, achievement, desire-fulfillment, and transcendent meaning
  - Supportive of vital social institutions and structures
  - Supportive of democracy and democratic institutions
  - Supportive of culture and cultural diversity

A disadvantage of ethical checklists is that they are necessarily incomplete, and may cause ethical issues that are specific to a particular technology or domain to be missed. For example, in the ethics of robotics it is sometimes proposed that advanced robots should have rights. Most ethical checklists will not recognize ethical values or principles granting rights to robots. In addition to employing an ethical checklist, it is therefore recommended that the technology ethics literature is also surveyed to identify ethical issues, and that the various artifacts and applications are also subjected to bottom-up ethical analyses. A bottom-up approach can either draw from moral values and principles expressed by stakeholders, or from moral intuitions of the analyst.

At the three technology levels, different kinds of ethical issues can be identified. At the technology level, ethical issues are either inherent, consequential, or pertaining to specific risks. *Inherent* ethical issues are issues relating to features or processes that are inherent to the technology and that are morally controversial. For example, manipulation of DNA is an inherent feature of genetic engineering, and it is also morally controversial. *Consequential* ethical issues are those that emerge in relation to consequences that are likely to manifest themselves in most or all artifacts or applications of the technology. For instance, nuclear energy technology tends to yield nuclear waste, in whatever manner it is developed, and nuclear waste raises ethical issues. Ethical issues pertaining to *specific risks*, finally, are ethical issues raised by foreseeable risks of the development of a particular technology. For example, the further development of nanorobotics has been argued to risk a “grey goo” scenario, in which we are faced with unstoppable self-replication nanobots.

At the artifact level, a similar distinction can be made between ethical issues that are inherent, consequential and pertaining to specific risks. Inherent issues are those that pertain to features or processes inherent to artifacts. For example, a video game that contains depictions that degrade human beings is inherently controversial, prior to any particular way of using it. *Consequential* issues pertain to morally controversial consequences that manifest themselves in most or all uses of an artifact or procedure. For instance, automobiles with gasoline-fueled engines have a difficult to avoid consequence of producing harmful greenhouse gases when used. Ethical issues pertaining to *specific risks*, finally, are issues raised by foreseeable risks of the use of a particular artifact. For example, smartphones are built in such a way that their use carries with it a significant risk that they are used by providers or others to track the location of their users.

At the application level, finally, three rather different types of ethical issues can be identified. First, there are ethical issues concerning the morality of the *intended use* of an artifact. These concern the purposes for which or way in which an artifact is used. For instance, nowadays the technique of in vitro fertilization (IVF) is no longer very controversial, but its use for the impregnation of older women is. A second type of issue pertains to *unintended consequences for users* of an artifact. These are consequences that violate rights or interests of users. They may occur for certain uses or for certain contexts of use or for certain user groups. For instance, violent computer games may have the morally undesirable consequence of increasing aggression in some users. However, such consequences may depend on the type of user and the type of use. A third type of issue pertains to *unintended*

*consequences of non-user stakeholders.* The use of artifacts may violate the rights of interests of non-user stakeholders just like it may do so for users. For example, a company may use toxic materials in its production processes that are emitted into the atmosphere, thereby harming members of the local community.

Thus, at the identification stage, descriptions of the technology are cross-referenced with ethical principles, including those from the ATE checklist, and with parameters that are specific to the three technology levels. After the identification stage, there is the *evaluation stage*, at which the potential importance of ethical issues is assessed, the likelihood that they will become a significant issue in society, as well as their relation to each other, including potential value conflicts. For instance, while in the identification stage it may have been determined that behavioral profiling by Ambient Intelligence systems presents privacy issues, in the evaluation stage it is determined how serious such privacy issues would be, what the arguments would be pro and con the permissibility allowing the violation of privacy by these technologies, and how privacy may conflict with other values in the use of Ambient Intelligence technology, such as autonomy, security and well-being.

After the evaluation stage, there are various optional stages in which the results of evaluation are applied for various purposes. One possible stage would be one on which the results of the evaluation stage are used to guide the development of the technology, and to ensure that designed artifacts conform to ethical standards. Let us call this the *design feedback stage*. This stage is prominent in Palm and Hansson's eTA. Methods like value-sensitive design (Friedman, Kahn & Borning, 2006) can be used to help implement the results of ethical evaluation into design processes. A second possible stage would be a *responsibility assignment stage*, at which moral responsibilities are assigned to different relevant actors for ethical outcomes at the artifact and application levels. For instance, if behavioral profiling by ambient intelligent systems present privacy issues, this stage may assign responsibilities to designers, retailers, government agents and users for handling these issues. A third possible stage would be a *governance stage*, as in the ETICA project. During this stage, governance recommendations are made for policy makers for dealing with the outcomes of the evaluation stage. This stage could make use of results of the responsibility assignment stage or the design feedback stage.

## 5. APPLICATIONS OF ATE TO INFORMATION TECHNOLOGY

Emerging information technologies currently include, amongst others, Ambient Intelligence, neuroelectronics, radio-frequency identification, semantic web, quantum computing, affective computing and augmented reality. All these technologies will yield (and some are already yielding) technological devices and procedures that raise ethical issues. Using the methods of ATE, we can anticipate future developments in these technologies and ethical issues that result from their use by society. Ethical analysis of emerging information technologies can take several forms, depending on the object of analysis that is chosen and the range of ethical issues that is explored. In one type of analysis, the object of analysis is an emerging technology and a full range of ethical issues is explored in relation to it. For ethical analysis, a generic approach may be chosen, in which only general features of the technology are chosen, or a forecasting approach, which projects future artifacts and applications, and considers ethical issues for them as well. The emerging technology may also be studied in relation to a particular domain (e.g., healthcare, defense).

In a second type of analysis, the focus is not on a technology, but on a type of technological artifact or procedure. For instance, within robotics one may focus on future intelligent humanoid robots, and analyze the types of ethical issues that they may pose in the context of various applications. In a third type of analysis, the object of analysis is an emerging technology or artifact and ethical analysis is not directed at a large range of ethical issues, but on ethical issues in relation to a specific moral value or principle. For example, a study of present and future RFID applications may focus on privacy issues alone. A fourth type of ethical analysis takes place at the application level, and considers ethical issues in relation to a hypothesized future application. An example is an ethical analysis of the military use of wearable computers for detecting and photographing opposing forces. In such an analysis, it is not a specific type of artifact that is studied, but rather the use of an artifact or several related types of artifacts in a specific context of use. A final type is one in which not a technology but a moral value or principle is made central, and a wide range of emerging technologies is analyzed in relation to that value. For example, a study of how emerging technologies will affect privacy will consider possible and likely privacy implications of a wide range of emerging technologies.

To illustrate how ATE may work in practice, let us focus on a specific emerging technology, *Ambient Intelligence*, and let us see how a wide range of ethical issues may be considered in relation to it, both at the technology, artifact and application levels. Ambient Intelligence refers to electronic environments that are responsive to the presence of people and can react to their behaviors and inferred desires. Ambient Intelligence environments include embedded, networked and intelligent computer systems and chips, as well as sensors and actuators that are used to sense and respond to appropriate environmental stimuli. Ambient Intelligence, or AmI, is believed to be used in a variety of domains, including the home ("smart homes"), healthcare (hospital care, health

monitoring and assistance), transportation (intelligent vehicles), education (smart learning devices and classroom environments), and others (Cook, Augusto & Jakkula, 2009).

Ethical analysis at the technology level requires an identification of core features of AmI. On that basis, it is possible to identify ethical issues that are likely to persist throughout AmI applications. Langheinreich (2001) did this kind of analysis of privacy (see also Brey, 2005). He identified four core features of AmI that have implications for privacy, and then proceeded to analyze how these features are likely to impact privacy. These features are ubiquity, invisibility, sensing and memory amplification. For example, sensing is the ability of AmI systems to register aspects of the environment, including the presence, location, behaviors and verbal expressions of persons. Langheinreich points out that this feature may have specific impacts for personal privacy. Similarly, AmI systems are supposed to be largely invisible, by receding into the background. This will imply that users will not always know what they are present or what they are doing. If they are collecting and disseminating personal data, users will often not be aware of this. Similar analyses at the technology level may be made in relation to values other than privacy.

For ethical analysis at the artifact and application levels, prior forecasting is needed. Forecasting studies usually anticipate future artifacts and applications in three ways: technology-driven, domain-specific and need-based. *Technology-driven* studies emphasize how future technological developments in or combinations of an emerging technology with other technologies may yield new interesting devices. For example, studies may project how AmI devices equipped with cameras and motion tracking devices and with advanced AI may be capable of recognizing complex human emotions or intentions. *Domain-specific* studies consider various social domains, like the home, transport and communication, and consider how new devices and applications may be usefully adopted in such domains. *Need-based* studies, finally, focus on human needs or desires, like quality of life, nutrition or social status, and consider how new technological devices may contribute to satisfying such needs.

Existing futures studies of Ambient Intelligence most frequently consider a number of social domains, and then project possible artifacts and applications within those domains (e.g., Friedewald & Da Costa, 2003; Cook, Augusto & Jakkula, 2009). For instance, discussions of AmI in the home may include speculations of items like interactive bathroom mirrors, intelligent fridges capable of taking stock and making autonomous purchases, and stereo systems that based on contextual cues play music for which those present seem to be in the mood. Similarly, AmI devices in transportation may include intelligent navigation systems, intelligent airbags, smart traffic control systems, and sophisticated defect and environment surveillance systems within vehicles. The potential applications and social consequences of such devices may be explored in further analyses, including scenario studies.

Ethical analysis at the artifact level considers general ethical issues with these artifacts, issues that are likely to manifest themselves across different applications of them. For instance, an AmI device capable of recognizing human emotions or inferring human desires raises issues of privacy because of these very capabilities. At the application level, ethical issues are explored that pertain to the use of AmI devices for a particular purpose, in a particular context, in a particular manner, or by particular type of user. For instance, AmI devices capable of making intelligent decisions may ask their users for informed consent on their decisions. However, certain users, like young children or mentally incompetent adults, are not capable of informed consent. The use of such AmI devices by them thus raises ethical issues that are not present with other users. A systematic consideration of ethical issues at the application level requires a systematic consideration of a large variety of potential uses, user groups, and contexts of use.

After the identification of ethical issues with AmI, evaluation takes place, as described above. Optionally, then, design feedback, responsibility assignment and governance analyses may be performed. A design feedback analysis for AmI, for example, would consider how ethical issues at any of the three technology stages may be avoided or solved through appropriate technology development and design. For example, to address privacy issues with AmI, it may be opted to limit the storage and transmission of personal data by AmI devices, to encrypt any data that is stored or transmitted, to ensure that the processing of personal information is based on informed consent, and so on. A responsibility assignment analysis would identify responsibilities of designers, users and other stakeholders for ethical issues. It would, for example, propose responsibilities for designers to safeguard privacy in AmI, as well as for users to be aware of threats to their privacy and to respond appropriately. A governance analysis would include recommendations for policy makers regarding the regulation of the design and use of AmI devices that would minimize ethical problems.

## 6. CONCLUSION

In this essay, I have presented a new approach to the ethics of emerging information technologies, called *anticipatory technology ethics* (ATE). I started my discussion with a discussion of stages of technology development, in which I placed emerging technology at the R&D and introduction stage of technology development. I then considered

ethical approaches to emerging technologies, and made a distinction between generic approaches, which ethically reflect on known features of technologies, and forecasting approaches, which speculate on possible future devices and applications. I argued that an ethics of emerging information technologies should centrally include forecasting in its approach.

I then presented my own approach: ATE. I discussed how levels and objects of ethical analysis are defined in ATE, how forecasting is used, and how ethical analysis takes place. ATE employs three levels of ethical analysis, the technology, artifact and application level, which each contain various objects of analysis. Knowledge of them is acquired through forecasting, including the use of existing forecasting studies, expert panels and surveys, and self-performed futures studies. Ethical analysis, finally, is performed at two initial stages, the identification and evaluation stage. At the identification stage, moral values and principles are operationalized and cross-referenced with technology descriptions resulting from the forecasting stage. The values and issues are derived from an ethical checklist as well as from the technology ethics literature and bottom-up analyses. At the evaluation stage, the potential importance of identified ethical issues is evaluated and these issues are elaborated. Evaluations may subsequently be used for improving technology development, for better governance of technology, or for other purposes. Finally, I provided a view of how ATE may be applied to particular emerging information technologies, using Ambient Intelligence as an example.

ATE constitutes a general approach for ethical analysis for new and emerging information technologies that is comprehensive yet flexible enough to be used and tailored in different ways. It can be used for the ethical analysis of any new and emerging information technology, as well as for the analysis of particular IT systems, software and devices, and their use in particular contexts and by particular users. ATE also incorporates stages of analysis in which ethical analysis can inform technology development and governance. ATE still needs to be developed further. It is my hope that it will prove a fruitful approach for the ethical analysis of emerging information technologies.

## REFERENCES

- Bell, W. (1997). *Foundations of Futures Studies, vol. 1: Human Science for a New Era: History, Purposes, Knowledge*. New Brunswick and London: Transaction Publishers.
- Boenink, M., Swierstra, T. and Stemerding, D. (2010). Anticipating the Interaction between Technology and Morality: A Scenario Study of Experimenting with Humans in Bionanotechnology, *Studies in Ethics, Law, and Technology*, 4(2), article 4.
- Brey, P. (2000). Disclosive Computer Ethics: The Exposure and Evaluation of Embedded Normativity in Computer Technology, *Computers and Society*, 30(4): 10-16.
- Brey, P. (2005). Freedom and Privacy in Ambient Intelligence, *Ethics and Information Technology* 7(3), 157-166.
- Cook, D., Augusto, J., Jakkula, V. (2009). Ambient Intelligence: Technologies, applications, and opportunities, *Pervasive and Mobile Computing*, 5(4), 277-298.
- Friedewald, M., and Da Costa, O. (eds.) (2003). *Science and Technology Roadmapping: Ambient Intelligence in Everyday Life (AmI@Life)*. Working Paper. Seville: Institute for Prospective Technology Studies IPTS. Retrieved online on 12-2011 at [http://foresight.jrc.ec.europa.eu/documents/SandT\\_roadmapping.pdf](http://foresight.jrc.ec.europa.eu/documents/SandT_roadmapping.pdf).
- Friedman, B., Kahn, P. and Borning, A. (2006). Value Sensitive Design and Information Systems, in *Human-Computer Interaction in Management Information Systems: Foundations*, eds. P. Zhang and D. Galletta. Armonk, NY: M.E. Sharpe.
- Grunwald, A. (2009). Technology Assessment: Concepts and Methods. In A. Meijers, *Philosophy of Technology and Engineering Sciences. Handbook of the Philosophy of Science vol. 9*. Amsterdam: Elsevier.
- Heersmink, R., van den Hoven, J. & Timmermans, J. (2010). *Normative Issues Report*. Deliverable D.2.2. Located at <http://www.etica-project.eu/>.

- Langheinrich, M. (2001). Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. *Lecture Notes In Computer Science; Vol. 2201 Archive. Proceedings of the 3rd international conference on Ubiquitous Computing*, Atlanta, Georgia, USA, Springer-Verlag London, UK, Pages: 273 - 291.
- Moor, J. H. (2005). Why we need better ethics for emerging technologies. *Ethics and Information Technology*, 7(3), 111–119.
- Palm, E. and Hansson, S. O. (2006). The case for ethical technology assessment (eTA). *Technological Forecasting & Social Change*, 73, 543–558.
- Stahl, B., Heersmink, R., Goujon, P., Flick, C., van den Hoven, J. and Wakunuma, K. (2010). Identifying the Ethics of Emerging Information and Communication Technologies: An Essay on Issues, Concepts and Method. *International Journal of Technoethics* 1(4), 20-38.
- Sollie, P. (2007). Ethics, technology development and uncertainty: an outline for any future ethics of technology, *Journal of Information, Communications & Ethics in Society*, 5(4), 293–306.
- Stahl, B., Wakunuma, K. and Rogerson, S. (2009). *Theoretical, Methodological, and Collaborative Issues of Stage I of the ETICA Project*. Deliverable D.5.4. Located at <http://www.etica-project.eu/>.
- Tran, T. and Daim, T. (2008). A taxonomic review of methods and tools applied in technology assessment, *Technological Forecasting and Social Change* 75(9): 1396-1405.
- Wright, D. (2010). A framework for the ethical impact assessment of information technology, *Ethics and Information Technology* (7 July 2010), 1-28.

## **Bio**

Philip Brey is full professor of philosophy of technology and chair of the department of philosophy of the University of Twente, the Netherlands. He is the past president of the Society for Philosophy and Technology, member of the board of INSEIT, and conference series director of the International Association for Computing and Philosophy (IACAP). He previously taught at Delft University of Technology and the University of California, San Diego. His research is focused on technology ethics and philosophy and ethics of computing and IT.