

This is a preprint version of the following article:

Brey, P. (2012). Anticipatory Ethics for Emerging Technologies. *Nanoethics* 6(1), 1-13.

Anticipatory Ethics for Emerging Technologies

Philip A. E. Brey

Abstract

In this essay, a new approach for the ethical study of emerging technology ethics 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. I will argue that a major problem for its development is the problem of uncertainty, which can only be overcome through methodologically sound forecasting and futures studies. I will then consider three contemporary approaches to the ethics of emerging technologies that use forecasting: ethical technology assessment, the techno-ethical scenarios approach and the ETICA approach, and I considered their strengths and weaknesses. Based on this critical study, I then present my own approach: ATE. ATE is a conceptually and methodologically rich approach for the ethical analysis of emerging technologies that incorporates a large variety of ethical principles, issues, objects and levels of analysis, and research aims. It is ready to be applied to contemporary and future emerging technologies.

Keywords: Anticipatory technology ethics; emerging technologies; uncertainty; futures studies; forecasting; technology assessment

1. INTRODUCTION

Technologies are at different stages of development and of societal uptake. Some technologies have yielded many concrete devices and applications and are used by a many different people in a variety of contexts. For such technologies, ethical analysis has the benefit that many of the ethical issues have already been identified in society. For instance, a large variety of ethical issues in relation to the Internet have been identified not only by ethicists, but also by users and other stakeholders who run into them as they use or deliberate on the technology. Other technologies, however, are still emergent: they are at an early stage of development and have not yielded many applications and societal consequences. They are still largely, or fully, at the research and development (R&D) stage, meaning that they are still at the stage of research into basic techniques, or at an early stage of development which at most has resulted in lab prototypes and experimental applications but little or no serious products that are being used by ordinary users.

For technologies at the R&D stage, ethical issues relating to their use in society cannot be known reliably, as their impact on society lies in the uncertain future. At the stage of fundamental research, the focus is 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. At this stage, no knowledge may yet exist about possible devices or applications that may result from the research, so ethical reflection on future consequences may be wholly speculative at this stage. At the development stage, the focus is on the design and manufacture of actual devices and processes. At this stage, more information is known about possible designs, but there is still considerable uncertainty about the devices and systems that will eventually gain societal acceptance, the ways in which these may ultimately be used, and the societal consequences that their use will bring. So at this stage, also, there is much uncertainty regarding ethical issues and ways in which these may be approached.

The question that is the focus of this essay is how we can identify and evaluate ethical issues for technologies that are still emerging because they are still at the R&D stage. With the accelerated pace of technological change in contemporary society, and the major impact that technology has on people's lives, early identification and evaluation of ethical issues is an important aim. Early identification can help users and other societal actors better prepare for future moral dilemmas, and can also help steer R&D or usage practices so as to

avoid or minimize ethically undesirable consequences. Yet, so far very little research has been directed at developing sound approaches and methods for ethical analysis of emerging technologies. It is only in recent years that such research has seriously gotten underway. My aim in this essay is to review some of this recent work and to present a new, integrative approach for the ethical study of emerging technologies.

Emerging technology can be defined as technology at the R&D and introduction stages. Rather, research is directed at the development of basic techniques that may further down the road result in concrete applications. 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?

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.

The paper is structured as follows. In the next section, two approaches within the ethics of emerging technology will be distinguished, based on how they deal with the problem of uncertainty about the future. In section 3, three recent ethical approaches to emerging technology will be discussed and critiqued, and it will be concluded that neither is fully satisfactory. 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, and I will provide examples throughout its discussion how it can be applied in practice.

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.

2. ETHICS, UNCERTAINTY AND FORECASTING

The central problem for an ethics of emerging technologies is that we do not know the future, and therefore do not know which actual issues will play out once the technology is fully developed and entrenched in society. Because emerging technology is technology in the making, many questions about its nature, its future use and its social consequences are still undecided. For this reason, many ethical issues in relation to it cannot yet be identified or analyzed reliably. We can speculate about future applications and uses, but as history has shown, speculations about future technology are often way off the mark, meaning that we may end up exploring an irrelevant set of ethical issues.

The ethics of emerging technology therefore has to deal with an epistemological problem, the *problem of uncertainty* concerning future devices, applications, uses and social consequences (Sollie, 2007). The question is how it can deal with this problem in a responsible manner. On the one hand, it is to be avoided that ethicists lose themselves in idle speculation on future ethical issues in technology that in most cases turn out to rest on mistaken projections on how the technology will actually evolve. On the other hand, it is to be avoided that ethicists feel that they can say nothing about emerging technologies because they do not know which devices and uses will result from them. So the question is how ethicists can come to assessments of emerging technologies that are based on somewhat reliable knowledge of the future.

Two approaches are possible at this point, one more conservative and reliable, the other more uncertain and speculative. The first approach is to restrict oneself to ethical analysis of generic qualities of the new technology that are likely to manifest themselves in all or most future applications of the technology and that are likely to present ethical challenges. For example, when nuclear energy technology was being developed it was known early on that however it were to be developed, there would be a problem of radioactive waste, which requires ethical deliberation. 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. I will call this approach the *generic approach*.

A second approach is to speculate on future devices, uses and social consequences. This requires that ethicists either rely on existing forecasting studies or do such studies themselves. They can then use the forecasts to explore ethical issues. For example, ethicists can 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. They can then analyze ethical issues that are likely to occur when such devices are being used. I will call this the *forecasting approach* to the ethics of emerging technology.

The forecasting approach relies on predictive studies of future technological devices, uses and social consequences. Such studies are undertaken in two related fields. *Futures studies* is a field that aims to study what possible or probably futures may look like (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, whereas others may rely on surveys, time series analysis, regression analysis, or simulations. Some work in futures studies focuses on technology forecasting. It forecasts future technologies, including the development spread of certain types artifacts, and optionally their utilization and social consequences that may result from their use. *Technology assessment (TA)* is 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; Grünwald 2009). It 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 forecasting approach has as 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 is 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 preferable. In the next two sections, therefore, I will focus on forecasting approaches. I will first look at three contemporary forecasting approaches to the ethics of emerging technology, which I will critically evaluate. In the section thereafter, I will then present my own approach.

4. CRITIQUE OF EXISTING APPROACHES

In recent years, forecasting approaches to technology ethics have been gaining attention, although few mature approaches currently exist. In what follows I will consider three promising approaches that have been formulated in recent years: ethical technology assessment, the techno-ethical scenarios approach, and the ETICA approach. For each, I will consider their strengths and weaknesses, after which I will draw a general conclusion.

Ethical Technology Assessment

Ethical technology assessment (eTA), proposed by Palm and Hansson (2006), has as its purpose “to provide indicators of negative ethical implications at an early stage of technological development” (p. 543). Such indicators can subsequently be used to guide design or technology policy. The focus of eTA is on the whole life-cycle of technology development, from initial R&D to ultimate impacts on society. To attain an adequate understanding of future developments, eTA relies on studies in technology assessment (TA) and on close interactions with developers of technology. The interactions with technology developers are to guarantee an adequate understanding of the technology in question. Studies in TA are to provide insight into both the technology in question and its social consequences, and are also used to organize interactions with technology developers in which eTA is made relevant for the development process. 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. This yields areas in which a conflict may emerge between the technology and one or more accepted moral principles. This ethical knowledge may then be used to adjust design processes to avoid ethical concerns or to steer decision-making on an emerging technology. Palm and Hansson go on to 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 is one of the first ethical approaches explicitly targeted at emerging technologies. It 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. Still, the approach has a few limitations. Most importantly, it 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, as many recognized moral values and principles are not found on the list, such as autonomy, human dignity, informed consent, distributive justice, and so on. So it would seem one would need a much longer list to be able to do comprehensive ethical assessments of new technologies. And even then, moral issues could be into play for a new technology that are not included in the list. To identify such issues, it would seem that exploring moral intuitions of either stakeholders or the analyst would be in order.

The techno-ethical scenarios approach

The *techno-ethical scenarios approach* of Boenink, Swierstra and Stermerding (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. Boenink et al. argue that technology may change the way we interpret moral values and may also affect the relative importance of particular moral principles. For example, privacy may become a less important principle in an information society where personal information is ubiquitous, and the concept of human responsibility may change in a society in which human decision-making is supported by expert systems. 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.

The techno-ethical scenarios approach involves three steps. The first step, "sketching the moral landscape," aims to describe the new technology in question, as well as current moral beliefs, practices and regulations that are directly or indirectly relevant to the technology, and may optionally provide some historical background on the evolution of these beliefs and practices. The second step, "generating potential moral controversies, using NEST-ethics," aims to identify ethical issues and arguments regarding the new technology. This is done using the approach of NEST-ethics (Swierstra and Rip, 2007), which is an approach for identifying ethical issues and arguments in a new technology using a taxonomy of issues and arguments that have been used in past ethical controversies on technology. The NEST-ethics approach performs three tasks. First, it identifies promises and expectations concerning a new technology. Second, it identifies critical objections that may be raised against these promises, for example regarding efficiency and effectiveness, as well as many conventionally ethical objections, regarding rights, harms and obligations, just distribution, the good life, and others. Third, it identifies chains of arguments and counter-arguments regarding the positive and negative aspects of the technology, which can be used to anticipate how the moral debate on the new technology may develop. During this step, effects of the moral debate on the development of the technology may also be considered. These different steps may involve literature reviews of technologies, promises and expectations, literature reviews of ethical issues, as well as workshops with policy makers and TA experts.

The third step of the techno-ethical scenarios approach, finally, is "constructing closure by judging plausibility of resolutions". In this step, the multitude of views and arguments from step 2 is reduced by imagining which resolution of the debate is the most plausible. The intention is to use steps 1 through 3 to develop a scenario of how the new technology will develop in the future, how this affects the moral landscape (i.e., moral beliefs, practices and regulations), and how moral closure is eventually reached. The particular scenario they develop, for example, considers how developments in molecular medicine may affect existing moral practices concerning medical experiments with human beings. They project several changes in these practices, based on a scenario study set in Dutch society between 2010 and 2030.

The techno-ethical scenarios approach has some obvious advantages over the eTA approach. It takes into account moral change. It moreover takes on a larger time-frame than eTA, which seems to focus on incremental steps. And it identifies not only ethical issues but also complex patterns of argumentation regarding them. Yet, the techno-ethical scenarios approach has an important limitation as well. This is that it is a descriptive and predictive approach, rather than a normative and prescriptive one. 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. As I have argued, important moral controversies may remain hidden because of the complexity or opaqueness of technological artifacts or practices (Brey, 2000). Such controversies are not likely to be included in techno-ethical scenarios. Conversely, moral controversies may ensue that are based on a false or misguided understanding of the technology or its social consequences. Such moral controversies do not present moral issues that ought to be considered in assessing emerging technologies, because they are based on false premises. In addition, moral controversies may ensue that are based on parochial moral concerns that would not be considered in an ordinary ethical analysis. My point is hence that moral controversies that may emerge in public debate may be different from moral issues that may result from thorough ethical assessments, even though there may be a large overlap in practice between the two. The current approach focuses on the former type whereas I think an ethical analysis of emerging technology should primarily focus on the latter, as its aim should not be to predict moral debate but to identify normative ethical issues.

The ETICA approach

The ETICA approach (Stahl, 2011; Stahl et al., forthcoming) is a recent method for the ethical assessment of emerging information and communication technologies (ICTs).¹ It is so general in scope, however, that nothing prevents its application to other types of technology as well, and it will for this reason be considered as a general approach for the ethical assessment of emerging technology. 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 medium-term future. The ETICA approach makes use of projections of the future which it derives from futures research. It aims to arrive at a *foresight analysis*, which is a forecasting analysis that considers 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, under the assumption that while individual studies will contain biases and shortcomings, their aggregate use will tend to yield more reliable results.

Ideally the ETICA approach would include doing one's own future studies, as its researchers say. However, in their study of emerging ICT's, limitations in resources limit them to two methods for identifying ethical issues in emerging technology. The first is to extract ethical issues from texts about particular emerging technologies in which ethical issues are discussed. Such texts include governmental and political sources, scientific sources such as research reports and journal articles, and non-academic sources such as published future visions of companies. The second is to use bibliometric analysis that finds correlations between emerging technologies and ethical values and concepts in a database of texts on ethics of technology in the academic literature.

The results of multiple futures research studies 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. In the first stage of ethical analysis, the identification stage, ethical issues are identified for particular applications, artifacts or technological properties.² Most of the ethical values and principles used in this approach are derived from a prior list of ethical issues for ethical evaluation in a European context. The resulting ethical issues are summarized in a *normative issues matrix*, which specifies relevant normative issues in relation to particular emerging technologies and the artifacts and applications that are expected to result from them. For example, an analysis of robots, as an emerging technology, may focus on particular applications such as service robots in households, robots as companions and robots as soldiers, and discuss ethical aspects of each application. The normative issues matrix also contains more general ethical issues with particular technologies that are not bound to particular applications. For example, an analysis of robots may focus on privacy issues in relation to the sensory capabilities of robots, or responsibility issues in relation to the behavioral autonomy of robots, or ethical issues that are specific to humanoid robots.

In a second stage of ethical analysis, the evaluation stage, the ethical issues 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 ethical approach to emerging technologies that has been developed to date. It aims at thoroughness by considering a wide range of technological properties, artifacts,

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

² 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.

applications, and ethical issues. It also engages in ethical evaluation and develops recommendations for governance. And it aims to make use of state-of-the-art work in futures studies. Yet, the approach also has weaknesses. First, its claim to adopt a futures studies approach is somewhat dubious. The main sources of the ETICA approach for locating ethical issues are government and political texts, scientific texts, and non-academic texts. Many of the non-academic and government texts will not be based on scientific methods of futures research. Moreover, many of the scientific texts do not seem to be either. Judging from the literature references in the ETICA projects, many of these texts come from ethics and computer science journals, and most of them do not use methods of futures research.

Second, its assumption that “the overall discourse on future[s] technologies provides as good and reliable an understanding of the future as will be possible to achieve” (Stahl, Wakunuma & Rogerson, 2010, p. 9) is also dubious. Rather than merely aggregating predictions about new technologies, it would be better if the approach would provide independent critical assessments of such predictions and the methods used for arriving at them before such predictions are used as a basis for subsequent ethical analysis. It should be granted, though, that in the ETICA project some independent foresight research is undertaken to validate some of the predictions that are made. Third and finally, 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. For example, in the ethical analysis of robotics, most space goes to the consideration of generic ethical issues, and only a few types of robots and application areas of robotics are considered in detail.

Conclusion

My review of the three approaches has revealed strong and weak points in each approach. It has also brought forward various points to consider in an ethics of emerging technologies. A first point is through what approaches and methods technological forecasts are arrived at. The three approaches use various approaches from futures studies and technology assessment, including approaches developed as part of their own approach. A second point concerns the use of ethics and the identification and evaluation of ethical issues. How should this be done? Here, the three approaches also have different answers, though what they have in common is their drive to identify possible ethical issues or controversies and their heuristic use of ethical checklists in doing so. A final point, which has been more implicit in the discussion, concerns the question what an ethics of emerging technology actually studies: is it whole technologies and techniques, is it possible future artifacts, is it uses of artifacts, social consequences, or yet something else? To this question, also, the three approaches give different answers. These three points for an ethics of emerging technologies provide a good challenge to build and improve on the three approaches discussed above. That is what I will turn to in the next section.

5. ANTICIPATORY TECHNOLOGY ETHICS

Based on the previous discussion, I will now present an ethical approach of my own, which I will call *anticipatory technology ethics* (ATE). 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 distinguishes *three levels of ethical analysis*: the technology, artifact and application level (Fig. 1). At each of these levels, various *objects of ethical analysis* are defined: 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.

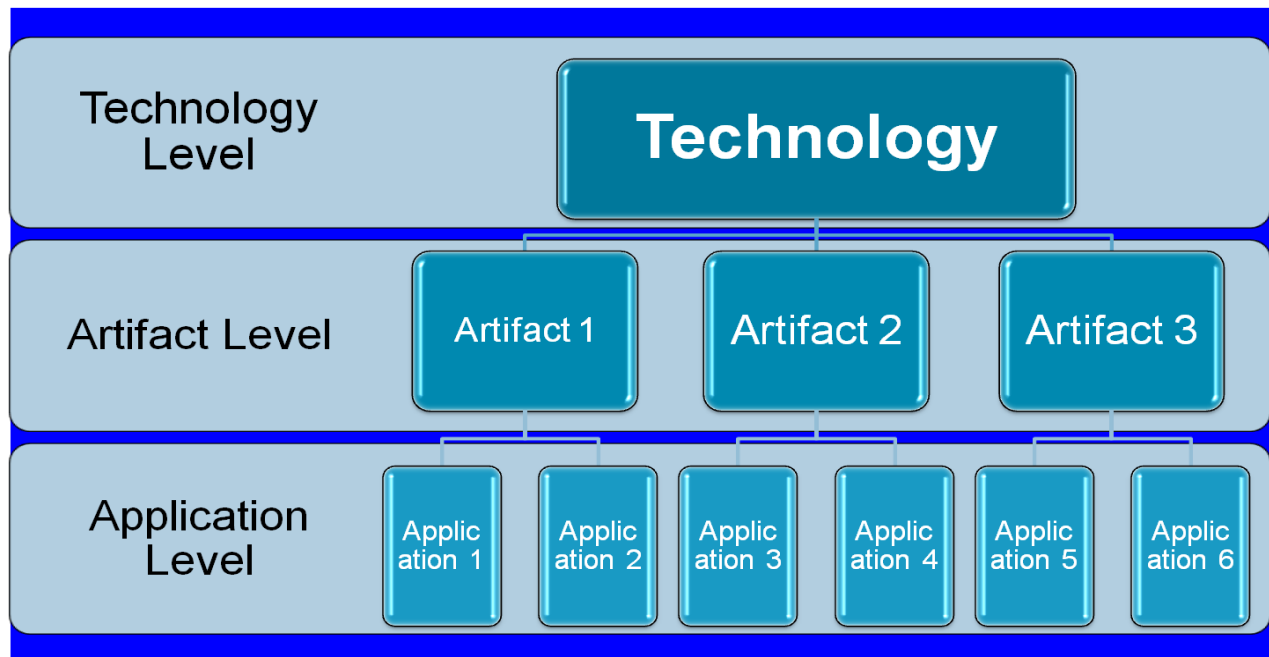


Fig. 1 Three levels of ethical analysis

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.³

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, 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

³ 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.

manner and in the proper environment, produces a desired result.⁴ 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 video games that contain depictions degrading to human beings, 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.⁵

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. cooking), 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 relating to the intended use of the artifact. They 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 for users. These are consequences that arise 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 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 non-user stakeholders who may be affected by a particular use of an artifact. For example the use of a new medical procedures by a surgeon 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

⁴ 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.

⁵ 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.

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 future 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 future 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 future studies as well, possibly in tandem with future 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.⁶ Table 1 represents an attempt at such a list.

Table 1 *The anticipatory technology ethics checklist*

- Harms and risks
 - Health and bodily harm
 - Pain and suffering
 - Psychological harm
 - Harm to human capabilities
 - Environmental harm
 - Harms to society
- Rights
 - Freedom
 - Freedom of movement
 - Freedom of speech and expression
 - Freedom of assembly
 - 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
 - Human dignity
 - Privacy
 - Information privacy
 - 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

⁶ 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.

- Justice (distributive)
 - o Just distribution of primary goods, capabilities, risks and hazards
 - o Nondiscrimination and equal treatment relative to age, gender, sexual orientation, social class, race, ethnicity, religion, disability, etc.
 - o North-south justice
 - o Intergenerational justice
 - o Social inclusion
- Well-being and the common good
 - o Supportive of happiness, health, knowledge, wisdom, virtue, friendship, trust, achievement, desire-fulfillment, and transcendent meaning
 - o Supportive of vital social institutions and structures
 - o Supportive of democracy and democratic institutions
 - o 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.

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.

6. CONCLUSION

In this essay, I have presented a new approach to the ethics of emerging technology ethics, called *anticipatory technology ethics* (ATE). I started my discussion with an appeal for the importance of ethical analysis of emerging technology, after which I described its nature and scope, defining it as 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. I then considered the problem of uncertainty and two approaches to it in the ethics of emerging technology, which I called the generic approach and the forecasting approach. I argued that a large part of the focus should be on forecasting, and I identified futures studies and technology assessments as two fields that engage in it.

I then considered three contemporary approaches to the ethics of emerging technologies: ethical technology assessment, the techno-ethical scenarios approach and the ETICA approach, and I considered their strengths and weaknesses. Based on this critical study, 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 is performed. ATE employs three levels of ethical analysis, the technology, artifact and application level, which each contain various objects of analysis. Knowledge of these objects of analysis 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.

ATE constitutes a general approach for ethical analysis for new and emerging 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 technology, as well as for the analysis of any particular artifacts and applications that this technology may yield. It also specifies how these ethical analyses may be used in technology development and in technology governance. Admittedly, some parts of the approach are still sketchy. It is my hope that a further development of ATE may advance the field of technology ethics and may afford more ethical development and governance of technology.⁷

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), art. 4.
- Brey, P. (2000). 'Disclosive Computer Ethics: The Exposure and Evaluation of Embedded Normativity in Computer Technology,' *Computers and Society*, 30(4): 10-16.
- 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.
- Grünwald, 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/>.
- Palm, E. and Hansson, S. O. (2006). The case for ethical technology assessment (eTA). *Technological Forecasting & Social Change*, 73, 543–558.
- 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. (2011). IT for a better future: hot to integrate ethics, politics and innovation. *Journal of Information, Communication & Ethics in Society*, 9(3), 140-156.
- Stahl, B., Heersmink, R., Goujon, P., Flick, C., van den Hoven, J. and Wakunuma, K. (forthcoming). Identifying the Ethics of Emerging Information and Communication Technologies: An Essay on Issues, Concepts and Method. *International Journal of Technoethics*.
- 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/>.

⁷ This paper was presented as a presidential address to the 17th conference of the *Society for Philosophy and Technology*, Denton, Texas, May 27th, 2011.

Swierstra, T., and Rip, A. (2007). Nano-ethics as NEST-ethics: Patterns of Moral Argumentation About New and Emerging Science and Technology, *Nanoethics* 1, 3-20.

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.