Ethics of Emerging Technology

Philip Brey

This is a pre-print of the following book chapter:

Brey, P. (2017). Ethics of Emerging Technologies. In S. O. Hansson (Ed.), *Methods for the Ethics of Technology*. Rowman and Littlefield International.

This chapter surveys ethical approaches to emerging technology. In recent years, emerging technologies have become a major topic of study in the ethics of technology, which has increasingly focused its attention on early-stage intervention in technology development. A number of specific approaches and methods have now been developed for the field, which in many ways is still in its infancy. The main problem for the ethics of emerging technology is the problem of uncertainty (Sollie, 2007): how to deal with the uncertainty of future products, uses and consequences, and associated ethical issues that will result from an emerging technology. Several approaches to the ethics of emerging technology will be reviewed that deal with this problem in different ways. Special attention will be paid to anticipatory approaches, which combine foresight analysis with ethical analysis. These approaches will be assessed and critically compared to alternative ethical approaches to emerging technology.

What is ethics of emerging technology?

A proper understanding of the ethics of emerging technology presupposes a proper understanding of what emerging technologies are. *Emerging technologies* are technologies that are new, innovative, and still in development, and are expected to have a large socioeconomic impact. They are new in the sense that they employ new concepts, methods and techniques and cannot be subsumed under existing technologies. They are innovative in the sense that they promise new and potentially superior solutions to problems. They are still in development in that they are still largely a promise: no, or not many, products and applications have resulted from them, and few, if any, are marketed and used on a large scale. They are expected to have a large socioeconomic impact in that they are expected to generate significant economic value and activity, and have the promise to affect or transform one or more social or economic domains, such as education, healthcare, transportation or the retail industry.

Importantly for our purposes, emerging technologies are still technologies in the making. They are not fully developed and entrenched in society. Though not a simple linear process, technological innovation involves different stages, which often begin with research (to investigate new phenomena,

ideas, designs or techniques), followed by development, production, marketing and diffusion into society. A technology that has completed all these stages is sometimes called an *entrenched technology*. Such a technology is associated with a number of developed products, processes, procedures or techniques that are widely used in society and have familiar uses and known impacts on society. Although new products based on the technology may still come out, they only represent incremental improvements on existing products, and do not involve radical innovation. Examples of entrenched technologies are automotive technology, satellite technology, antibiotics, polymer technology, radio technology, information technology, and nuclear technology.

Emerging technologies are still largely or wholly in the R&D (research and development) stage. They have not yet resulted in many products, and they have not yet generated a large socioeconomic impact. They are still partially a promise: to become a successful, entrenched technology, further research may be needed, new innovative techniques and approaches may need to be developed or tested, methods may have to be developed to combine them with other technologies, new products and applications may have to be thought up, and their market success still has to be proven. Examples of emerging technologies are (at the time of writing): medical nanotechnology, synthetic biology, Internet of Things, personal and service robots, augmented reality, and smart materials.

It can now be seen how the ethics of emerging technologies is different from the ethics of entrenched technologies. First, the ethics of entrenched technologies is able to address, evaluate and direct a greater set of existing phenomena. The ethics of emerging technologies tends to have its focus on research and development of these technologies, as these are realities that can be assessed and redirected on the basis ethical assessments. The ethics of entrenched technologies can in addition ethically assess and recommend modifications of specific products, uses, regulations, and social impacts that are already in existence. Second, even if limiting itself to ethical assessments of research and innovation, the ethics of entrenched technology can draw from a wider range of data that are relevant to ethical analysis. Different from the ethics of emerging technology, it can make use of data about existing products, uses and social impacts. The ethics of emerging technologies can only make use of speculative data about future products, uses and impacts. The ethics of entrenched technologies can arrive at better informed moral evaluations of and prescriptions for research and development.

In spite of the better epistemic position of the ethics of entrenched technologies and the broader range of topics that can be covered by it, there has been a big movement in recent decades towards the ethics of emerging technologies. This is the case because of one big advantage of the ethics of emerging technologies over that of entrenched technologies: the possibility of early intervention in innovation processes. Once billions have been spent to develop a technology in a particular way, and it becomes

entrenched in society as a result, it is very hard to make fundamental changes to its overall design and embedding in society. The ethics of emerging technologies harbors the promise of early intervention when a technology is still malleable and there is still much room for choice in its development and social embedding. The price to be paid for this shift in focus is that the ethicist has a more limited range of empirical data to work with and is faced with significant uncertainties regarding future developments and impacts.

A further distinction relevant to the ethics of emerging technologies is that between stand-alone and enabling technologies. *Enabling technologies* are technologies that provide innovation across a range of products, industrial sectors, and social domains. They combine with a large number of other technologies to yield innovative products and services. Examples of enabling technologies are steam engine technology, glass making technology, integrated circuit technology, thermal energy storage technology, genetic engineering and nanotechnology. In addition, there are also industry- or sector-specific enabling technologies, such as enabling technologies for smart mobile services, tissue engineering, sustainable architecture and personalized medicine. *Stand-alone technologies*, in contrast, are technologies that yield specific products and services, and are often limited to one industrial sector and one application domain. Examples are quartz clock technology, ballistic missile technology, penicillin, and escalator technology.

The ethics of emerging technologies is focused in large part on emerging enabling technologies, which are expected to result in waves of innovations across different sectors in society and to raise a myriad of ethical issues in the process. Although emerging stand-alone technologies may also raise significant ethical issues, these are often more closely associated with specific products and services, so it is usually somewhat easier to subject them to ethical analysis of future products and associated uses and impacts.

Types of Approaches

In the past ten to fifteen years, five distinct types of approaches to emerging technologies have emerged, which will be discussed below. Before they are discussed, it is worth noting that the ethics of emerging technology can be situated within a larger set of both qualitative and quantitative approaches to the assessment and guidance of emerging technology that have the intent of producing better outcomes for society. These include technology assessment (TA), futures studies, impact assessment, risk assessment, risk-benefit analysis, cost-benefit analysis and cost-utility analysis, and, as well as approaches focused on public and stakeholder engagement, democratization, and deliberative decision-making. Many approaches

in the ethics of emerging technologies seek a combination with one or more of these non-ethical approaches, as one will see below.

Generic approaches

A first approach, which I called the *generic approach* in Brey (2012a), focuses on broad features of an emerging technology that raise ethical issues, independently of and prior to any specific products, uses or impacts. These are ethical issues that can be identified by considering inherent features of the technology, necessary conditions for its realization, or generic impacts that it is likely to have, regardless of how it will be developed in the future. Generic approaches rest on conceptual analysis and empirical observations of the general features of the technology. Sometimes, they can also involve projections of future impacts. A generic approach is, for example, taken in ethical critiques of genetic engineering that argue that manipulation of genomes is playing God, because it amounts to designing new life, which should be done by God, and not by humans. Another example is criticizing an emerging technology because it centrally involves chemical or physical processes that produce toxic or harmful gases or substances as a byproduct, such as greenhouse gases, radiation, or carcinogens.

The generic approach has as an advantage over other approaches in that it does not have to concern itself much with future development and use of a technology and that it can limit itself to the technology as it already exists. It can therefore make reliable claims about the technology that involve little or no speculation about the future. A disadvantage of the generic approach is that it can only consider a small set of ethical issues in relation to emerging technologies: general ethical issues concerning the technology as it has developed so far.

Anticipatory approaches

A second type of approach is an *anticipatory* or *foresight* approach (Brey, 2012a). Anticipatory approaches combine ethical analysis with various kinds of foresight, forecasting or futures studies techniques, such as scenarios, trend analysis, Delphi panels, horizontal scanning, as well as various methods of technology assessment. These techniques are used to project likely, plausible or possible future products, applications, uses and impacts that may result from the further development and introduction of an emerging technology. Ethical issues in these future applications and uses are subsequently identified and subjected to ethical analysis. Although anticipatory approaches may also be used to identify more general ethical issues as well, they are uniquely capable of identifying ethical issues in relation to projected future products, uses and social consequences. Ethical analyses performed with them look as follows:

- Technology X is likely to lead to applications and uses that harm privacy. Therefore, it should be developed and introduced in such a way as to minimize such harms.
- Technology X may lead to applications in the military domain that are morally unacceptable. Therefore strong regulation of X may be necessary to prevent such applications.
- Technology X will lead to products that may, in some societies, enhance socioeconomic inequalities. It therefore should be developed and introduced in a way that takes into account this moral issue.

Existing anticipatory approaches include ethical technology assessment (Palm and Hansson, 2006), ethical impact assessment (Wright, 2010), anticipatory technology ethics (Brey, 2012a, b), the ETICA approach (Stahl et al., 2010), the techno-ethical scenarios approach (Boenink, Swierstra & Stemerding, 2010) and the moral plausibility approach (Lucivero, 2015, Lucivero, Swierstra & Boenink, 2011). These approaches draw on technology assessment (Palm and Hansson), impact assessment (Wright), and foresight and TA (Brey; Boenink, Swierstra and Stemerding; Lucivero). Each of these approaches has its own unique selling points. For example, the ethical impact assessment approach is useful for innovation projects since it addresses the practical implementation of recommendations based on ethical assessments. The techno-ethical scenarios approach has the strength of addressing and studying moral change, and is capable of ethical analysis based on expected future moral values of stakeholders. The anticipatory technology ethics approach is possibly the only approach that makes full use of foresight methods and presents a detailed methodology for combining these with ethical analysis. The moral plausibility approach, finally, proposes epistemic tools for critically assessing the plausibility of expectations about the future put forwards by scientists and experts.

The strong point of anticipatory approaches is that they are the only ones capable of detailed and comprehensive forward-looking ethical analyses of emerging technologies. Their weak point is that they rely on information about the future that is to a degree uncertain and speculative. It is difficult to make reliable predictions, and foresight analyses have often been completely off in their projections of future developments, uses and consequences of emerging technologies. Foresight analysts usually do not claim anymore to predict the future, but rather to describe plausible or possible futures. Even so, the unreliability of foresight analysis casts some doubt on its usefulness as a foundation for ethical assessment of emerging technologies. It should be noted, however, that foresight analyses tend to be more reliable if they concern the near future (e.g., 1-5 years, rather than 30 years from now), and that they can generate useful insights into potentialities, conditionalities and dependencies concerning emerging technologies, even if they are not fully predictive.

Ethical risk analysis

Some approaches to emerging technologies focus on the risks that they can pose, including health, security, safety, economic and environmental risks. These predominantly non-ethical approaches include risk analysis and risk-benefit analysis. *Risk analysis* is the process of defining, assessing, analyzing and managing risks (Haimes, 2015). It is often divided into *risk assessment*, which is the identification, evaluation and measurement of the probability and severity of risks, and *risk management*, which concerns decision-making about risks. *Risk-benefit analysis* is the comparison of the risk of a situation to its related benefits. It aims to determine risk-benefit ratios, which are the ratio of a risk of an action to its potential benefits. A guiding principle is that actions should only be undertaken if the risk-benefit ratio is above one. For example, for an individual, the risk-benefit ratio of air travel is usually considered to be above one, whereas the risk-benefit ratio of space travel has so far not been proven to be so.

Risk analysis and risk-benefit analysis are mostly quantitative approaches that do not make explicit use of ethical criteria. Ethical risk analysis has recently emerged to allow for risk analysis that takes into account ethical considerations (Hansson, this volume; Asveld and Roeser, 2009). It investigates issues of responsibility, justice, autonomy, well-being and others in relation to risk assessment and risk management. Ethical risk analysis of emerging technologies identifies risks in such technologies, morally evaluates them, and proposes risk management strategies that are justified from an ethical point of view. Similarly, ethical risk-benefit analysis takes into account ethical issues in determining and utilizing risk-benefit ratios for risks and potential benefits associated with emerging technologies.

Strictly speaking, ethical risk analysis and ethical risk-benefit analysis of emerging technologies are anticipatory approaches, because the calculation of risks and potential benefits requires anticipation and estimation of probabilities of future consequences of the emerging technology. However, because of their unique focus on risk and on quantitative methods, they deserve their own category. A strong point of the two approaches is that they are able to provide quantitative, ethically grounded assessments of risks and benefits of emerging technologies. A weak point is that, like anticipatory approaches generally, they necessarily depend on projections of the future, and that quantitative determinations of risk and potential benefit will often be contentious. Another potentially weak point is the narrow focus on risks, which excludes other types of impacts that are deserving of moral consideration.

Experimental approaches

Experimental approaches are based on the idea that emerging technologies bring with them many uncertainties regarding their consequences for society, and that these uncertainties often cannot be

properly expressed as quantifiable risks and cannot be known or conjectured through foresight approaches because the consequences of emerging technologies are the unpredictable, emergent outcomes of the co-evolution of technology and society. Instead, we should see the introduction of a new technology into society as a process with inherently uncertain outcomes: as a social experiment. Conceiving of technologies as social experiments, we postpone the question "Is technology X morally acceptable?", which we for the most part cannot answer before the technology has been fully introduced into society. Instead, we try to answer the question "Is it ethically acceptable to experiment with technology X in society?"

The experimental approach of Van de Poel (2015) includes a framework for responsible experimentation with emerging technologies in society. Van de Poel bases his proposal on the ethical principles of respect for people, beneficence and justice, and proposes a set of thirteen conditions for responsible experimentation based on these three principles. These are conditions such as there being no other reasonable means for gaining knowledge about the potential hazards, it being reasonable to expect social benefits from the experiment, hazards being contained as far as reasonably possible, the experiment being approved by democratically legitimized bodies, people being properly informed, and vulnerable experimental subjects being excluded or protected. The approach calls for an incremental introduction of emerging technologies so that an adaptive learning process is possible.

Participatory and deliberative ethical approaches

An increasing number of approaches combine ethics with participatory, deliberative and stakeholder approaches. Cotton (2014) makes the case for this combination as follows. Ethical assessment, he argues, is a too top-down, technocratic and expertocratic exercise that fails to take into account the plurality of perspectives found in public responses to emerging technologies. Ordinary approaches to public and stakeholder engagement and deliberation are however too bottom-up, since they make assessments of emerging technologies subject to the opinions, prejudices and unconsidered moral judgments of those participating, with no guarantee that ethical issues will be considered carefully, if at all. Therefore, an approach must be sought that combines the best elements of a top-down and a bottom-up approach, which means a combination of ethics and participatory and deliberative approaches.

In principle, all previous approaches can be supplemented with participatory and deliberative approaches, and sometimes this possibility is made use of. I will focus here, however, on ethical approaches that are essentially participatory and deliberative. These approaches identify the public or the

stakeholders as the ones who ultimately have to provide an ethical or ethically informed assessment of an emerging technology, and make moral decisions about its further development and use.

The ethicist has a supporting role in this process. One way in which the ethicist provides support is by developing approaches and tools for better inclusion of ethical principles and arguments in debates on emerging technologies. Another way is by developing ethical criteria for the organization of deliberative and decision-making processes, to ensure that relevant stakeholders are represented, and to avoid that power relations and unequal participation undermine a democratic deliberative process. Both approaches have their roots in discourse ethics (Habermas, 1991). Many authors use them both (Swierstra and Rip, 2007, Cotton, 2014, Keulart, Korthals, Schermer and Swierstra, 2004).

An advantage of participatory and deliberative ethical approaches is that they include the opinions, viewpoints, and moral intuitions and judgments of different people in a way that could enrich ethical assessments. An advantage of stakeholder approaches in particular is that including stakeholders in ethical deliberation and decision-making processes is more likely to lead to ethical outcomes, since they usually include actors that shape the future development and use of emerging technologies (technology developers, users, regulators and other agents). A potential disadvantage of these approaches is that the ideal of serious moral deliberation under conditions of equality may be difficult to achieve. These approaches require that a substantial number of people are brought to the same table to engage in extensive moral deliberation in a way that follows the elaborate discourse rules of ethicists, moves the discussion beyond prevailing interests, and negates prevailing power relations that may distort the discussion. Even ordinary participatory and deliberative approaches have been difficult to realize in practice (Hagendijk and Irwin, 2006), and approaches that include ethical criteria face additional obstacles. Another potential disadvantage is that current approaches do not contain adequate methods for anticipating future ethical issues in relation to emerging technologies. Either the participants will have to engage in foresight themselves, and they risk lacking adequate expertise for this, or they will have to rely on foresight analyses by experts, in which case the approach effectively becomes a blend of anticipatory and participatory/deliberative ethical approaches.

In the remainder of this chapter, I will further discuss anticipatory approaches, for several reasons. Anticipatory approaches are the only ones that promise comprehensive, future-oriented ethical assessments of emerging technologies. They are more comprehensive than risk ethics approaches, and far more comprehensive than generic approaches, both of which they could subsume. The main competitors of anticipatory approaches are experimental and participatory/deliberative approaches. Experimental approaches are however only an interesting alternative if it is true that adequate foresight regarding future consequences of emerging technologies is impossible. So it merits investigation whether or not this is the

case. If anticipatory approaches are capable of incorporating participatory and deliberative approaches, then they may moreover be more attractive than the stand-alone version of these approaches. So we will investigate this option as well.

Anticipatory Approaches and Foresight

To what extent can future consequences of emerging technologies be known? Clearly, there are future consequences that can be known fairly reliably, and others that cannot. When several firms are in the process of developing household robots, for example, it can be predicted with some confidence that there will soon be household robots on the market, and that there will be people using them. If it is also known that these robots are being designed to collect and store personal information, then it can also be reliably concluded that they will introduce new privacy risks. On the other hand, whether there will be robots within the next thirty years that are in possession of super-human intelligence and consciousness is something that probably cannot be known reliably. What these two examples suggest is that moderately reliable knowledge of future consequences of emerging technologies may be possible if the time horizon is not too great and if information is available about specific technological products that are in the planning and development stages.

Sometimes, it is possible to make predictions of future products and some of their consequences. But making reliable predictions is often not possible, and it is not what foresight analyses of emerging technologies usually try to do. Rather, they aim to identify *plausible* and *possible* futures. A possible future is one that could happen, and a plausible future is one that has a non-negligible likelihood to happen. Foresight analyses often explore multiple possible and plausible futures. The absence of concrete predictions does not make a foresight analysis worthless. The exploration of possible and plausible futures may provide valuable information. First, by giving glimpses of what *may* happen, it allows for better anticipation of the future than would be possible in a situation in which one has no idea what may happen. Second, by projecting possible future applications and uses of the technology and resulting consequences, it is possible to identify potential risks and benefits. It is likely, for example, that advanced 3D printers can be used to manufacture illegal weapons, from which it can be deduced that there is a significant risk that such printers will be misused in this way. It is also possible, and cases have been made, that nanoparticles have an adverse effect on the immune system. As long as this has not been conclusively investigated, this is a risk or uncertainty that can be associated with nanotechology.

Next to being able to give us *some* anticipation of the future and helping us to identify potential risks and benefits, foresight analyses can also help us identify path dependencies, causal relations,

contingencies and constraints in the development and use of emerging technologies. They can, for example, help us systematically go through the consequences for the economy, the environment, and everyday life of a massive shift from vehicles with internal combustion engines to electric vehicles. Even if such a study is not predictive, it may help identify the path dependencies, constraints, and unintended consequences involved in such a shift. This helps in making better strategic choices possible. For ethical analyses, foresight analyses can show us path dependencies and contingencies that determine whether certain ethical issues will emerge or not, and will help us take steps to avoid undesirable effects.

A requirement for the identification of dependencies, constraints, and potential risks and benefits, is that foresight analyses are logically valid and are based on empirically adequate generalizations. Implicitly, foresight analyses use empirical generalizations such as "If a technological product *p* provides a service for which there is a need that is not provided equally well by other products, then, everything else being equal, people will aim to acquire *p*" and "People have a need for cheap and affordable transportation". Advances in foresight analyses of emerging technologies will require that good use is made of empirical generalizations gained from research in the field of science and technology studies (Hackett, Amsterdamska, Michael & Wajcman, 2007), since this field has over the past forty or so years developed sophisticated, empirically adequate theories and models that capture law-like relations in innovation, technology development, technology use, and the impacts of technology. Such models can be used to make realistic projections of future consequences of emerging technologies.

Foresight methods

What follows is a brief and incomplete overview of tested foresight methods that can be used in anticipatory ethical studies of emerging technologies (cf. Giaoutzi & Sapio, 2013).

- *Horizon scanning* is the scanning or review of a broad range of data sources about the phenomenon about which one aims to gain foresight, in order to identify perspectives and trends that shed a light on future developments. The approach can take the form of a structured literature review or a bibliometric analysis. It can focus on existing foresight studies about the phenomenon, if any, but also on trends, expectations, and new developments and ideas collected from a variety of sources. It usually involves data from a wide variety of sources, including Internet sources, journals, databases, and various sorts of organizations (such as ministries, companies, NGOs, and research organizations).

- *Expert consultation* is a simple form of stakeholder engagement, in which experts with respect to the technology in question or specific impacts are consulted by means of interviews, a short

workshop or a small survey. They are consulted about their expectations regarding possible, plausible or likely future developments regarding the technology.[1]

- *Scenario methods*have become popular tools in foresight analysis, and are used in a wide variety of ways. An overview of current scenario approaches is offered by Börjeson et al. (2006). Scenarios are constructed by starting with the present and past, and projecting into the future. These are forward looking scenarios. Backcasting scenarios instead look backwards from a desired future. The objective of such scenarios is not to determine which futures are most likely to occur, but how to attain a particular future that is desirable. In pluralistic backcasting, multiple preferred futures are taken as starting points of the backcasting exercise. Scenarios are either concerned with what will happen (trend extrapolations, business as usual scenarios, probable scenarios), what could happen (forecasting, foresighting, strategic scenarios) and what should happen (normative scenarios such as those used in backcasting).

- The *Delphi method* is a method that involves an expert panel that fills out questionnaires with their forecasts on a topic in two or more rounds. They usually operate at a distance, although Delphi can also take place face-to-face. The experts enter their questionnaires anonymously, so as to allow for better judgments without undue influence from forceful or high-status advocates. After each round, a facilitator anonymously summarizes the expert's forecasts from the previous rounds together with the reasons for their judgments, and the experts are encouraged to revise their earlier answers in light of these replies. After a number of rounds, a mean or median score or opinion is determined on the basis of the results in the final round. This method is based on the principle that decisions or forecasts of a structured group are more accurate than those of unstructured groups.

- *Trend analysis* involves the identification of trends: general tendencies or directions evident from past events, and increasing or decreasing in strength of frequency of observation. These trends are then used to predict future events. Trend analyses can be qualitative or quantitative, involving statistical analysis.

[1] Lucivero (2015) advocates that the ethicist or analyst should not be uncritical of the expectations of experts, particularly those of scientists and engineers, who tend to advocate positive visions of the technology that help them sell their ideas. She has developed an approach for assessing the epistemic plausibility of the projections of scientists.

- *Relevance trees*are a normative forecasting method that originated in strategic planning. It starts with future needs or objectives, and then seeks to identify the conditions and actions required to meet them. A relevance tree is an analytic technique in which a broad topic is subdivided into increasingly smaller subtopics, thereby showing an increasing number of paths to the objective. It forecasts associated costs, probabilities and durations for each element in the tree.

- *Roadmapping* is a collaborative foresight process in which a broad set of strategies and plans is developed to reach a common goal. Roadmapping may include any of the other foresight tools. It involves a collaboration network of multidisciplinary and sometimes competing experts. The roadmapping process emphasizes not only probable and preferred futures, but also uncertainties and challenges. The time horizon is often longer than traditional forecasts and plans (5-15 years or more). In roadmapping, trends, actions, challenges and outcomes are commonly represented with graphical displays that are associated with support documents.

- *Participatory foresight*. In participatory foresight (Nikolova, 2014), citizen panels and other methods of citizen participation have a role in foresight analysis, often in combination with participation by experts and stakeholders. This is usually used in normative foresight analysis, in which citizens state their visions and preferences for particular futures and provide comments on scenarios and solutions presented by experts.

So do any or all of these approaches provide reliable insight into possible, plausible or probable futures? It should be noted that traditional foresight analysis, which aimed at forecasting and prediction, has only had limited success. However, sophisticated methods have been developed in recent decades, and the shift from prediction to the mapping of plausible and possible futures and to normative foresight analysis has led to different standards for success. Nevertheless, a thorough epistemic assessment and testing of current foresight methods has not yet been performed, so it is difficult at this point to assess their reliability and soundness. As I argued earlier, foresight analysis has at least some value in mapping possible and plausible futures, and it may be further improved by better inclusion of insights from STS. For this reason, it is an approach that is worth pursuing at this point, although further studies are needed to determine whether its limitations are so severe that in most cases alternative approaches to the ethics of emerging technologies are called for, in particular experimental methods as discussed above. It may also be best to combine anticipatory with experimental approaches, with experimental approaches entering into the equation at those moments when good foresight is not possible.

Combining Foresight , Ethics and Public and Stakeholder Engagement

It has not been addressed so far how foresight analysis combines with ethical analysis. Let us first consider descriptive foresight analysis, which describes plausible futures that are based on mere anticipation and prediction rather than on a normative vision of the future. One way in which it can be combined with ethical analysis is by first having a separate stage of foresight analysis, after which the forecasted future(s) are studied for ethical issues that they may raise. A second way, that may be more effective for ethical analysis, is to first formulate particular moral values, principles or issues, and then explore possible futures in relation to the emerging technology at issue in which they play a role. For example, in relation to the Internet of Things, one may explore various plausible future products and processes in which personal data is used so as to investigate possible implications for privacy. This is not an instance of normative foresight analysis, because the aim is not to work towards futures in which these values or norms are realized, but rather to limit the set of future developments that are studied to those relevant to or involving particular ethical principles or issues.

Normative foresight analysis, in which one sketches desirable futures and investigates how they can be reached, can be combined more directly with ethical analysis by including ethical criteria in the list of desirable qualities of the future. However, it is easy to include ethical criteria, but much more difficult to determine how they can be respected in the development and use of emerging technologies. Doing so will require much of the descriptive analysis and identification of ethical issues that is involved in descriptive foresight analysis. For example, one can do a normative foresight analysis of the Internet of Things in which one envisions a future Internet of Things where values like privacy, autonomy, distributive justice, etc. But determining the attainability of this vision and the conditions that must be met will require identifying all the possible and plausible ways in which the technology may be developed and the ethical issues that this introduces, as is attempted in descriptive foresight analysis.

Anticipatory approaches can include stakeholder and public engagement in quite straightforward ways. Approaches based on descriptive foresight analysis can do so by presenting foresight analyses to stakeholders or members of the public and engaging with them in ethical assessments of the projected futures. Alternatively, they may be presented with both foresight analyses and ethical issues that have already been identified by ethical experts, and they are then asked to assess, weigh and help resolve these ethical issues. In approaches based on normative foresight analysis, stakeholders and members of the public may be involved more directly in foresight analysis, and participatory foresight methods may be used to construct desirable futures for emerging technologies.

To illustrate the use of foresight analysis in an anticipatory approach, I will now discuss one particular approach, anticipatory technology ethics (Brey, 2012a, b). I discuss this approach because in my admittedly subjective opinion, it is the only anticipatory approach that fully embraces foresight analysis and presents an extensive methodology for the assessment of emerging technologies. Anticipatory technology ethics (ATE) is a method for comprehensive ethical assessments of emerging technologies. A comprehensive assessment is one in which the full range of plausible or possible ethical issues associated with an emerging technology, both generic ones and ones relating to particular products, applications or domains are identified and evaluated. ATE also allows for partial assessments, in which ethical issues are explored in relation to a particular value, product, application or domain.

ATE analyses consist of three phases or stages: (1) foresight analysis; (2) identification of ethical issues; and (3) evaluation of ethical issues. After these three stages, there are optional stages that utilize the ethical evaluations for further action. These include a design feedback stage, in which the results of ethical analysis are used take identified ethical issues into account in the design and development of the technology, a responsibility assignment stage, in which responsibilities are assigned to different stakeholders for the ethical issues that have been identified, and a governance stage, in which governance recommendations are made for public policy.

ATE engages in foresight and ethical assessment at three levels of analysis: the technology, artifact, and application levels. The technology level is a level of description at which the general features of a technology are described, independently of any products or applications. This level centrally focuses on basic techniques that define the technology. For example, nuclear technology can be defined as the collection of techniques for the fission and fusion of atomic nuclei, and nanotechnology as the collection of techniques for manipulating matter on a scale of 1 – 100 nanometer. At the technology level, ethical issues of two sorts can be identified: generic ethical issues, of the sort discussed above in relation to the generic approach, and generic risks, which are across-the-board risks associated with techniques that are part of the technology when they are implemented in future products.

The artifact level is a level of analysis that focuses on ethical issues associated with usable products of a technology, which include technological artifacts (e.g., airplanes, x-ray imaging systems) and technological processes and procedures (e.g., food irradiation, nuclear well logging). These artifacts and procedures are studied for their ethical properties largely independently of, and prior to, particular uses of them in particular contexts. Just as it is sometimes possible to identify ethical issues in relation to techniques, prior to products and applications resulting from them, it is sometimes possible to identify such issues in artifacts. This is because particular technological solutions and designs may raise ethical issues that are not specific to a particular way of using them. For example, the internal combustion engine

can be critiqued from the point of view of sustainability largely independently of particular uses of it in cars or other machines, because of its emission of greenhouse gases. Similarly, an app on a smartphone that automatically collects personal information and transmits it to the company behind it, or has the ability to do so, can be criticized for its risks to privacy, independently of how it is used by users or the company.

Finally, the application level, or use level, is concerned with particular uses of technological products, by particular types of users, and in particular contexts or domains. At this level, ethical issues are dependent on particular ways of using the artifact, and on the particular circumstances in which it is used, which jointly determine consequences or impacts. Ethical analysis focuses on ethical issues that emerge in particular ways of using technological products, in particular contexts or settings, or in relation to particular types of users. It may, for example, include ethical assessments of the use of cognitive enhancement for the elderly, or the use of drones in military contexts, or the use of 3D printers to make illegal or copyrighted products.

Ethical analysis in ETA takes foresight analyses as inputs and initially aims to identify ethical issues that may be in play in the projected techniques, artifacts, uses and impacts. It does so by cross-referencing the results of foresight analysis with ethical principles and issues. This can be done intuitively, on the basis of the expertise of the ethicists or the moral intuitions of stakeholders, or it can be done with the help of an ethical checklist or similar tool. Prior ethical studies may also be used in the process. At this identification stage, the objective is merely to flag and identify potential ethical issues. Subsequently, at the evaluation stage, a thorough ethical evaluation takes place, in which the ethical importance of the issue is further analyzed, and its conflicting values in these issues are weighed against each other, and possible ways of resolving them are explored.

Conclusion

This chapter reviewed current issues and approaches in the ethics of emerging technology. Five different categories of approaches were distinguished: generic, anticipatory, risk, experimental and participatory and deliberative approaches. Strengths and weaknesses of each were reviewed. A further examination of anticipatory approaches was then undertaken, which involved review and assessment of foresight methods used in such approaches, the way in which they combine foresight analysis with ethical analysis, and the possibilities of combining such approaches with participatory and deliberative approaches. Some preliminary conclusions were drawn. However, the ethics of emerging technology is in many ways still in its infancy, and more work is needed to develop more sophisticated approaches, and to assess which approaches deliver the best results.

References

Asveld, L., & Roeser, S. (Eds.). (2009). The Ethics of Technological Risk. London: Earthscan Publishers.

Boenink, M., Swierstra, T., & 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), 1–38.

Börjeson, L., Höjer, M., Dreborg, K., Ekvall, T., & Finnveden, G. (2006). Scenario Types and Techniques: Towards a User's Guide. *Futures*, *38*(7), 723–739.

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

Brey, P. (2012b). Anticipating Ethical Issues in Emerging IT. *Ethics and Information Technology*, *14*(4), 305–317.

Cotton, M. (2014). *Ethics and Technology Assessment: A Participatory Approach*. Berlin: Springer-Verlag.

Giaoutzi, M., & Sapio, B. (Eds.). (2013). *Recent Developments in Foresight Methodologies*. New York: Springer.

Habermas, J. (1991). *Erläuterungen zur Diskursethik.* Frankfurt am Main: Suhrkamp.

Hackett, E.J., Amsterdamska, O., Michael, E.L., Wajcman, J. (Eds.). (2007). *The Handbook of Science and Technology Studies, 3rd Edition*. Cambridge, MA: MIT Press.

Hagendijk, R., & Irwin, I. (2006). Public Deliberation and Governance: Engaging with Science and Technology in Contemporary Europe, *Minerva*, *44*(2), 167–184.

Haimes, Y. (2015). *Risk Modeling, Assessment, and Management*, 4th ed. Hoboken, New Jersey: Wiley.

Keulartz, J., Schermer, M., Korthals, M. , & Swierstra, T. (2004). Ethics in a technological culture. A programmatic proposal for a pragmatist approach. *Science, Technology and Human Values*, *29*(1), 3–29.

Lucivero, F. (2015). *Ethical Assessments of Emerging Technologies: Appraising the moral plausibility of technological visions*. Dordrecht: Springer.

Lucivero, F., Swierstra, T., Boenink, M. (2011). Assessing Expectations: Towards a Toolbox for an Ethics of Emerging Technologies. *NanoEthics*, *5*(2), 129–141.

Nikolova, B. (2013). The rise and promise of participatory foresight. *European Journal of Futures Research*, 2, 1–9.

Palm, E., & Hansson, S.O. (2006). The Case for Ethical Technology Assessment (eTA). *Technological Forecasting & Social Change*, *73*(5), 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., Heersmink, R., Goujon, P., Flick, C., Van den Hoven, J., & 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.

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

Van de Poel, I. (2015). An Ethical Framework for Evaluating Experimental Technology. *Science and Engineering Ethics*, (online article) 1–20. http://link.springer.com/article/10.1007%2Fs11948-015-9724-3#/page-1

Wright, D. (2010). A Framework for the Ethical Impact Assessment of Information Technology. *Ethics and Information Technology*, *13*(3), 199–226.