Writing Stories to Explore The Ethics of Technology

By John M. Artz, PhD
Associate Professor of
Information Systems
The George Washington University

Part One: The Context

Table of Contents

Introduction

Imagine that a new technology has been developed that allows you read other people's thoughts. Imagine further, that the technology is sufficiently inexpensive that nearly anyone could afford it. Is this a good thing? Should it be freely available? Should it be regulated in any way? Should it be banned? Should usage be controlled by social conventions or laws? Should its use be restricted to authorities or would that just make it worse? Most people would have no idea how to answer these questions, nor would they have any idea how find any answers.

Consider another scenario. Imagine that human cloning has been sufficiently advanced so that any given person could be replicated as many times as desired. There are some obvious ethical issues that arise. Is it OK to create clones for spare parts? Is it OK to create batches of clones for specific kinds of work, such as dangerous work or work requiring special skills or abilities? There are also some trickier questions. We currently believe in some basic principles such as "all men are created equal" and "one person, one vote". Would these ideas go out the window? Would all clones be considered equal to non-cloned humans? Would the original, from whom the clones were created, have no special status? What would keep someone from creating armies of identical clones to vote according to their beliefs? If clones are truly identical, should they still each get one vote each; or one vote for the whole lot; or some sort of weighted vote?

While we are considering farfetched scenarios, think about the problems that might arise if we could prevent aging. This is not quite as fanciful as the others since aging research is making impressive progress. Should we put a limit on how long people can hang around? Is it better to make new people even if we have to sacrifice older ones? Or are we better off keeping the older ones who may develop wisdom and perspective over time? If you have to decide between the wisdom of older people and the fresh ideas of younger people, how would you find the right balance? Can you live with the implications of your decision?

Here is just one more example to consider before we return to reality. What if machines could reason *perfectly*? Could we just turn over the government to them? This is also not as farfetched as it may first appear. How many other important functions have been turned over to computers? Could airplanes fly without computer systems? Could airports function? Could the subway run? What about the heating and air conditioning in large buildings? We feel some degree of comfort in the fact that these computerized functions do have human operators. But, could 'you' operate the system? If the system malfunctioned, how many 'operators' could really step in and fix it? In the past century, we have turned over an enormous amount of our daily operations to computers. In three or four more centuries, why not the government? It would be hard to argue that government of the people performs flawlessly.

These scenarios are all examples of ethical problems that arise as the result of new or evolving technology. And what they have in common is that our current understanding of ethics is inadequate for addressing them. The most prevalent ethical theory in modern times is called *consequentialism*. In this view, we evaluate the moral quality of an act or rule based upon its consequences. It is certainly reasonable to do this except in those cases in which we have no idea what the consequences may be.

And that seems to be an aspect that issues in the ethics of technology have in common. A new technology is developed or evolved and we have no idea what the consequences may be.

Marshall McLuhan is often attributed with the observation that looking to the past to understand the future is like driving while looking in the rear view mirror. While there is some doubt as to whether or not he actually said this, the observation stands on its own. In fact, the analogy is so apt that we can investigate it even further for additional insight. Think of yourself as driving a car and steering based upon what you saw through the rear view mirror. This might actually work in certain cases. If the road ahead was straight and you weren't driving too fast, you might actually get away with it. However, if the road ahead was curvy and you began to accelerate, then keeping your eyes on the rear view mirror might not be such a great idea. Now, let's turn back to the future with what we have learned here. If the future is pretty much like the past and does not come at you too quickly, then looking to the past to understand the future might work. However, if the future is markedly different and is coming at you more quickly, then looking to the past to understand the future might make no more sense than driving by looking in the rear view mirror. And this is exactly the problem we have in the ethics of technology. New technologies create new worlds that are very different from the old ones that they leave behind. And new technologies are coming at us at an increasing rate. So, what do we do? We need a way to get a 'headlight into the future' so that we can look ahead to make decisions about moral and ethical issues rather than relying on our past experience.

Granted the above scenarios are pretty farfetched. In fact, they were intended to be. But let us consider some less remote. Should personal privacy be protected on the World Wide Web or does this impede the free flow of information? Don't we use the term 'censorship' for impeding the free flow of information? And, doesn't that carry a negative connotation? How about intellectual products that are available in digital form? Should antiquated copyright laws apply to them? If not, what is appropriate? And what might the implications of other arrangements be? Should there be restrictions on the way people behave through social interaction technologies? Or is it all a case of "let the visitor beware"? And in virtual worlds is it better to protect the anonymity of residents or is it better for everyone to know who they are dealing with. How about video games? They are becoming increasingly more prevalent. Are there ethical consequences to the design and play of these games? If so, what are those consequences? This second set of scenarios is much more concrete and much more relevant than the first set. However, both sets have the same thing in common. They are ethical questions that arise as the result of new technologies.

The two sets of scenarios were presented in this order for a specific reason. That reason is that on issues that we are facing today, we often mistakenly believe that we already know the answers. This is because we have the tendency to evaluate new situations with current standards. The first scenarios were pushed far enough into the future to make any of today's standards irrelevant thus showing that we do not have all the answers. The second, more relevant, set are concerns that we are currently facing. However, we tend to evaluate them in terms of the past. Privacy should be protected in the future because it was important to protect it in the past. But is that really true? Intellectual property should be protected in the future because it was protected in the past. But, is that still true?

Questions in the ethics of technology pose some serious stumbling blocks. One of the dominant theories in ethics, mentioned earlier, suggests that we determine the ethical quality of an act based upon the consequences. And, yet, in the case of ethical issues that arise as the result of a new technology, we rarely know what the consequences might be. So, how do we determine ethical standards in an area where consequences are unknown? The answer is quite simple. We can use stories. Stories allow us to explore possible worlds with the idea of selecting the one that we feel most comfortable with. But, I am getting a little ahead of myself.

Back in the early 1990's I attended a conference in Washington DC related to personal privacy as it was affected by computer systems and the Internet. At that conference, there was a panel discussion on the issue of personal privacy. All of the panelists had the same basic position – we must protect personal privacy. I began to wonder how we could have a debate on an ethical issue when everybody was on the same side. So I began to ask around – what was the other side of this issue? What was the downside of protecting personal privacy? As I asked around, I was astonished to find that not only did nobody have the slightest idea what the other side of the issue might be, but that many people looked askance at me for suggesting that there might be another side. This is a problem because ethics is always a trade off of competing interests. When everybody is on one side of an ethical issue, it is not ethics. It is merely reinforcing prevailing beliefs. I felt that if we wanted to protect personal privacy, we should choose to do that over the alternatives and know the costs. And yet, among those people that I talked to at the conference, and over the next few years, I could not find any alternatives. Further, there was little, if any, concern for the costs. So, I set out to discover the alternatives and costs. I took a position that all personal information should be freely available over the Internet and anyone should be able to access it by going to a local library. I actually think I came up with some compelling arguments for this side of the issue. But my arguments were not at all well received. Most people had a very negative visceral reaction to my position disliking me for suggesting the position as much as they disliked the position.

I won't go into my arguments, here, for freely available personal information. However, I will say that I was astonished by the unwillingness of the various people I spoke to, to even consider other sides of the argument. I began to wonder why this was the case and how I might get people to consider other perspectives. I began to see this same phenomenon occurring in other areas of computer ethics. In fact, most of computer ethics consisted of little more than one sided opinions that were held dearly by a few and ignored by many. You should not copy software. You must protect copyrights. People are dehumanized by technology. And on, and on. Over time, I began to see that the problem with computer ethics was not just that it was one sided but that the conclusions that were drawn were based on past experiences from when the key technologies did not exist. In other words, if you freeze things the way they are today and add a new technology you get one set of conclusions. If you consider the new world created by the new technology you may come to a different set of conclusions.

That solved part of the problem which is why does everyone fall on one side of this issue. But it raised a bigger problem which was how do you get people to consider questions in computer ethics within the context of a future world? That requires a healthy and disciplined imagination, attributes that are typically in rather short supply. Around the same time, and, as far as I knew, completely unrelated, I was trying to teach myself how to write creative fiction. This is another long story from which I will spare the

reader. But, I will say that at some point I saw a connection between these two endeavors. On one hand, I was trying to figure out how to get people to envision a future different from today in order to evaluate ethical problems in the appropriate context. On the other hand I was creating different worlds through the use of creative fiction. Bingo! The way to get people to envision a different future is to write a story about it. Stories can be used as an exploratory mechanism for ethical and moral issues. In fact, stories have long been used in this way. But, I just wanted to turn the focus more clearly on the future as brought about by technological changes.

Over the next few years, I wrote a series of papers on the role of stories in computer ethics. I laid down some pretty good arguments for the use of stories and those arguments will be covered later in this book. However, each hurdle we get over just brings us to the next. While it may be true that stories play an important epistemological role in computer ethics, we still have the problem of getting people to write stories. People, who write software, typically do not write creative fiction. So, my next hurdle was to figure out how to teach people to write stories.

My target audience for this new talent was, at the time, people involved in computer ethics. For them, I felt, that finding the writer within was a little too close to a Zen experience and I would need a much more structured approach. So, I looked at the top down software design technique that most computer people would be familiar with. In the top down approach to designing software, you specify what the software has to do and then through a process of decomposition and stepwise refinement, you identify the software modules needed to achieve the top level functionality. I thought this could be used to design stories as well. You identify what you are trying to achieve in the story and then identify the components needed to achieve that. The components, in this case, are the elements of a story (characters, settings, etc) rather than software modules. But the basic design idea is the same.

Then in another unrelated serendipitous event the university in which I work embarked on a new writing program called Writing in the Disciplines. Under this approach, writing classes would be offered, not by the English Department, but rather by the individual schools who would develop classes that would teach students how to write within that particular discipline. Each school was on the hook to develop writing classes within their own disciplines. I saw this as an opportunity to test out my technique for writing stories. So, I seized the moment and offered a writing course.

In yet another unrelated event (it is lovely how unrelated events fall together so nicely sometimes) the academic department which I was in was reorganized. As a result of this reorganization my focus was expanded from information technology to include technology in general. I had been considering, for some time, expanding the scope of my research and this provided a great opportunity. Computer ethics is just one instance of the ethics of technology and I figured the techniques that I was working on had broader applicability. So, I offered a course in Writing Stories to Explore the Ethics of Technology.

I employed the writing techniques that I had developed and refined them over a half a dozen or more offering of this course. There was much tweaking and much refinement along the way. Those techniques make up the bulk of this book along with the philosophical justification of the technique. My

goal, in this book, is to present what I have developed in as coherent a manner as possible despite the serendipitous events and circuitous route that really led up to it.

Chapter 1 – The Ethics of Technology

What is Ethics?

Ethics can be frustrating and problematic but, in most cases, does not have to be nearly as bad as most people think. This is because - when people encounter ethical problems they are usually the trickiest and thorniest of all ethical issues. For example, the question of whether or not it is acceptable to steal a loaf of bread to feed your starving family has no satisfactory resolution. On the other hand, the question of whether or not it is acceptable to shoot your neighbor's dog because his barking is keeping you awake is not so difficult to figure out.

Years ago, I was attempting to engage an undergraduate class in a discussion of computer ethics. One student, somewhat grumpily, point out that ethics is just "a bunch of rules that tell you what not to do". I still remember this remark many years later because it reflects an all too common misunderstanding of ethics – a misunderstanding that is held as often by people involved in computer ethics as by people outside of the field.

This misunderstanding is further complicated by morality and religious concepts of good. Since being ethical, in the minds of many people, equates with being good or being a good person, exploration of ethical issues is often frustrated as people may not be willing to consider sides of an issue that may suggest that they are not a good person. This makes discussion of ethical issues difficult as some may have preconceived notions of what is right or what is good and others may be reluctant to challenge those views. So, in order to counter these unfortunate misconceptions, I would like to offer some insights and some ground rules on ethics and the exploration of ethical issues.

Ethics is a collection of guidelines for appropriate behavior within a social group that considers both the well being of the individuals and the need for harmony within that group. You can't just have individuals running around doing whatever they please because that will impact the well being of other individuals. On the other hand, you can't just force all individuals to adhere to a monastic discipline as that would have a negative impact on their sense of well being. In order for society to thrive, the individuals must thrive. And in order for individuals to thrive, they must have a harmonious environment in which to thrive. So, ethics attempts to find a workable balance between the two.

This is not the only balancing act that ethics attempts to pull off. In fact, as we shall see later, ethics can be seen as a collection of careful balancing acts between competing interests. For now let's simply say that as we try to find the balance between competing interests in society, no matter what we decide, there will be winners and there will be losers. When we decide that a particular business practice, for example, is unethical, we are deciding in favor of those who may be harmed by that practice and deciding against those who may benefit from it. So, it isn't just about right and wrong. It is more about winners and losers.

Where do these guidelines come from? The simple answer is that people decide on the guidelines. I don't want this to sound too frivolous. You can't just decide that anything is ethical and go with it. The guidelines arise out of a consensus, although that consensus may take time and may be heavily influenced by social, political or economic power. However, in those cases, when ethics are unduly influenced, the guidelines may well shift over time as the power shifts. If ethical principles seem to work in the long run, they may well be promoted to moral principles. So the restriction on accepting bribes in business transactions is largely an ethical principle. However, the prohibition against stealing would be considered a moral principle. That is, one can conceive of a business climate in which bribes were considered an acceptable practice in business transactions. However, it is hard to imagine a society in which stealing is a common and acceptable practice. When moral principles become embedded in a religious doctrine, they take on an even higher status and begin to define the differences between good and bad or right and wrong.

In the ethics of technology we are concerned mainly with determining guidelines for the acceptable uses of technology. We are not attempting to open the debate on moral principles or religious doctrine. Hence, discussions of these issues should be far less daunting than discussions of larger moral principles. The ethics of technology should be grounded, for the most part, in existing moral principles.

Sometimes this does not appear to be the case. For example, stealing violates a moral principle. And, in recent years many people claimed that copying digital information such as software, images, music or videos was immoral on that basis. However, the question is not whether stealing is now acceptable. The question is whether or not that act should be considered stealing. Similarly, we have a moral principle prohibiting murder. And yet, advances in technology that permit abortion and euthanasia are considered immoral, by some, because they permit murder. Again, we are not reopening the debate on whether or not murder is acceptable. We are asking, instead, if those acts should be considered as murder.

The point of all this is to say that the ethics of technology are open to debate and reasonable moral people may come to different conclusions. Too many people are afraid to challenge prevailing views in the ethics of technology because they feel that this has all been decided and all they have to do is learn the answers. But the answers are yet to be provided, and that is what this book is all about.

The Ethics of Technology

The ethics of technology can be even more frustrating and more problematic than traditional ethics, but that does not have to be the case either. The reason why it can be more frustrating is that you have all the problems associated with traditional ethics complicated by the fact that you are trying to develop guidelines of proper behavior for a moving target – an evolving technology. In this section we are going to define technology and the ethics of technology. I will explain why the ethics of technology is different and why it is important.

What is Technology?

We tend to think of technology as the various devices that both ease and complicate our lives: computers, cell phones, networks, DVD players, Cable TV, and the like. People in specialized fields may see technology as advances in transportation, communications or weapon systems. Some may see technology as advances in medicine. While this is not entirely incorrect, it is somewhat limited as it fails to capture the full significance of technology in our lives.

The word technology comes from an ancient Greek root word teché which can be translated somewhat loosely as 'craft' and a little more precisely as 'a reliable means for producing an end'. To provide a simple example, consider cooking. If you understand the mechanics of cooking and, based on that understanding, you can produce a desired result such as an edible meal, you are employing teché. If you throw stuff together and hope for the best you are not. Consider another example. Let's say you are trying to convince a friend of yours to vote for a political candidate that you favor. If you understand the rules of argumentation and the art of persuasion, you can construct a compelling argument and, hopefully, win them over. In this case, you have employed teché. If you just argue in circles with them and get frustrated you are not.

Teché is a body of knowledge about how to produce results. Consider how different this is from science. Science is a body of knowledge about the natural or social worlds. It can tell you 'what is', but cannot tell you how to bring about something different. Bringing about something different - whether it is producing a meal, convincing a friend, building a building, or creating a computer- is the domain of teché.

Now let's add another Ancient Greek root – logos. Logos is the root word we find in logic and in many sciences such as biology and psychology. Logos is a rigorous explanation of a phenomenon. So, biology is a rigorous understanding of life and psychology is a rigorous understanding of the psyche. When people hang around a bar and talk about an election, it is not logos. But when trained political scientists analyze the election it is logos.

If we put these two words together, we get teché logos or technology - a rigorous understanding of the means by which outcomes are produced. So, technology, in the large, is a body of knowledge about how to produce desired outcomes; just as science is a body of knowledge about the natural or social worlds. When these processes produce artifacts that impact our lives in the way that computers and networks do, we think of the artifacts as the technology rather than the result of technology.

This may seem like a picky point but it is important. Technology is the means by which we make changes to the world and bring about new worlds. When we create a new world, we have to take responsibility for the new world that we create. And the ethics of technology is all about creating new worlds in a responsible way.

Nonetheless, we tend to focus more on the ethics of artifacts. That is, we create computers; computers create problems; so we focus on the problems created by computers. But consider this. Let's say that a person is an expert in the teché of rhetoric and they persuade masses of people to do something differently, such as attack a neighbor or a subgroup. Isn't that bringing about a different world? Doesn't somebody have to take responsibility for that?

Before we get to abstract here, I should mention that, for the most part, this book will deal with more concrete, practical issues that arise as a result of specific implementations of technology such as computer ethics or bio ethics. However, the techniques employed to investigate these issues as described later are applicable not only to the wide variety of instances of the ethics of technology but to the large problems that arise from the more abstract view of technology. That is, we will be using stories to explore possible worlds for the purposes of gaining some moral insight.

What is the Ethics of Technology?

Ethics, as stated earlier, is a collection of guidelines for appropriate behavior within a social group that considers both the well being of the individual and the need for harmony in society. Technology is the means by which we bring about a different world, one that does not immediately follow from the present due to natural process. So, the ethics of technology addresses the guidelines for appropriate behavior in possible worlds brought about by the introduction of new technologies.

How is the ethics of technology different? When ethical issues arise out of a natural state of affairs, such as people living or working together in close proximity, we have ethical theories and approaches for examining these issues. One approach is to look to the past. What has been right in the past will, most likely, continue to be right into the future. For example, treating each other fairly and with dignity seems to have worked out well, in the past, and, and will probably continue to work well into the future. If a specific situation has not been encountered before, we can consider the consequences of various ethical positions.

Unfortunately, approaches to ethics that have served well in the past do not serve as well in the ethics of technology. Most notably is consequentialism which requires that you consider the consequences of various ethical decisions. This seems like a reasonable approach, but, what if you do not know the consequences? Many times, when a new technology is introduced, the consequences are unknown. So, how can we make ethical decisions based on consequences when we do not know the consequences? The ethics of technology is different from traditional ethics in that we are attempting to decide on guidelines for appropriate behavior in a world that does not yet exist – a world that will be brought about by the technology in question.

Why is the ethics of technology important? Increasingly, the ethical problems that we face are problems in the ethics of technology – computer ethics, medical ethics, the ethics of biotechnology, appropriate uses of the media, and so on. These technologies all create new worlds and we must attempt to make ethical decisions about appropriate behaviors in today's world for a world that does not yet exist.

Currently, we know how to look to the past or how to examine the present. But looking to the future presents some difficult problems. The ethics of technology is important because more and more of our ethical decisions are within this realm and fewer and fewer of our ethical decisions of the past provide us with the guidance that we need for addressing these issues.

Chapter 2 - Breaking with the Present

One of the problems that we encounter in the ethics of technology is that people steadfastly deny that the future will be any different from the past. It doesn't matter how much change a person has seen in the past. They, for reasons I completely fail to understand, will tend to see the future as an extension of the present. Consequently, when considering issues in the ethics of technology, there is a tendency to apply rules from the past that may not work so well in the future.

This is probably what happened with personal privacy on the Internet which is the issue that got me started on all this in the first place. In the past, there was a huge imbalance in access to information. The government held a lot of information. Corporations held a lot of information. And private citizens got crumbs. This put private citizens at a disadvantage in defending themselves in a wide variety of situations. Hence, an important component of privacy was to prevent the government, in particular, from getting information about you that put you in the position of having to defend yourself with very limited access to the same information.

Let's say, for example, that you were to rent a pornographic movie from a video store. (Remember when we used to have video stores?) To make the scenario easier to accept let's also say you did it by accident. That is, you thought you were renting a tour guide for a city in Texas and it turned out to be an adult movie. Further, let's say that someone who has a reason to discredit or defame you got this information and began to disseminate it. It would be difficult to defend yourself against such a disclosure. You would be pegged as some sort of shady immoral character who probably turned his collar up to avoid recognition as you went into a creepy store to rent filthy movies.

However, if you had some relevant information, defending yourself might be much easier. For example, if you knew how many people locally and nationwide had rented adult movies in general or that movie in particular, this might mitigate any claims against you. If you had the names of other people who had rented that movie (perhaps even the person making claims against you) that might provide some defense. If you had data on how many people inadvertently rented adult films each day, this might also make you look a little better. However, you would have none of that information. Historically, we are all used to having only crumbs of information.

And yet, with the Internet, all that has changed. Private citizens have access to staggering amounts of information, not only about themselves but in aggregate about other private citizens. In the past it was important to protect privacy and hence, many believe, privacy must be protected in the future at all costs. Granted many companies and many websites publish privacy policies that restrict the most flagrant violations of privacy. But, in general, privacy advocates of that era might be astonished to realize just how ineffective their efforts were. The problem here, in short, is that people evaluate ethical issues in technology using circumstance that often will no longer be the case when the technology is fully developed.

Writing Stories to Explore the Ethics of Technology

In this chapter, I want to try to dispel the notion that the future will be just like the past and that what we know about the past will serve us adequately in making decisions about the future. I am going to do this in two stages. First, I am going to bring up some very difficult issues that lay far enough into the future that it is obvious that today's standards will not apply. With this first set of scenarios, try to grant a "willing suspension of disbelief" and focus on how you would deal with them. If you have trouble suspending disbelief, imagine yourself going back 100 years and trying to explain radio, television, computers, cell phones, and the World Wide Web. A lot can happen in 100 years. So, as farfetched as these scenarios may seem, it is even more farfetched to claim that they definitely won't occur in the next 100 years. Once we have seen that the future could present us with some scenarios for which our past principles may not provide sufficient guidance, we can then consider some current scenarios which may not be well served by our current understanding. Hopefully, with these also we will see that the past may not provide sufficient guidance.

If Computers Could Reason Perfectly, Should We Turn Over Government To Them?

We will begin with a fairly easy one. We have turned over an enormous amount of responsibility to computers already. In fact, modern life would simply be impossible without the aid of computers. Airplanes could not fly. Subways and cars, even elevators, could not run. We could not efficiently maintain a controlled climate in a high rise building. The Internet would not be possible without reliance on computers and neither would cable TV. In fact, it only takes a minute to realize that you cannot think of very many things you do in our modern world that does not rely on a computer, somehow. Perhaps you could take a walk around the block without relying on a computer. But, if you check your watch to make sure you get home on time, you are relying on a digital device, if not a full blown computer, for that. If you get a call on your cell phone while on your walk, you are definitely relying on a computer. Take a second to think through a typical day and see how much of your normal daily life would be totally disrupted if you refused to rely on computers. In most cases we are very comfortable relying on computers as long as they work the way they are supposed to work.

We know that computers can process information much faster than we can. So we rely on them for a myriad of accounting and information processing tasks. We even rely on them routinely to make decisions for us. For example, when the fuel injector in your car decides to change the fuel mixture, it does not pop up a little window on the dash board to ask you if this is OK. We happily delegate countless little decisions to computers every day. What about the bigger decisions?

When you go to Amazon.com to buy a book, it recommends other books that you might like and you decide whether or not to order them. What if the algorithm for that became so refined that it knew, better than you, what books you will like? Should it go ahead and send those books? If it does indeed know what you would like isn't it silly to ask you and allow for human error? Wouldn't that be like having the fuel injector in your car ask if it is OK to change the fuel mixture?

Let's say that a matchmaking website achieves 99% reliability in its matches. That is, when patrons of the website followed the recommendations of the website, it resulted in happy marriages 99% of the time. Would you rather go through the trial and error of the historical process of dating with its low probability of success? Even if the date is successful, the likelihood of a success marriage following this process is not very impressive. Wouldn't if better if the computer just accessed your personal information and sent over the love of your life?

Wouldn't this be just perfect? What rational person would not want to answer the doorbell and find the love of their life or the book they've always wanted to read? Yet, people considering these scenarios are likely to recoil and say "Noooooo! That isn't perfect at all!" It isn't that people are opposed to the idea. It is that they do not think it is possible. They think the computer will make mistakes. The doorbell will ring and your worst nightmare will be standing there. Or the book that showed up is sitting on a stack of a dozen other books that you will never get to.

Taking this notion, of computers making more and more of our decisions for us, a little further consider the lead in question - If Computers Could Reason Perfectly, Should We Turn Over the Government To Them? We have seen that computers already make a lot of decisions for us. We have also seen that they make better decisions in many cases. Our current form of government is as inefficient and flawed as our current approach to dating. So, if computers could do a better job, should we turn over decision making to them? Our first reaction to this suggestion is to recoil in horror. This is not because we don't want an efficient government that always makes the right decisions. It is because we don't think it will work.

The problem with denying the possibility of future scenarios like this is that if something like this does happen we will not have had time to seriously consider how to handle it. I am not advocating turning all of our decisions over to computers. What I am advocating is that if we begin to turn increasingly more of our decisions over to computers, which does appear to be happening, how do we sort out the possibilities and find the best way to do it? Denying the possibility while it occurs in fact does not seem to be a reasonable response.

If We Could Prevent Aging, Should We Live Forever?

Next we consider the possible impacts of long life spans. While the idea of living forever may test your imagination, the idea of living considerably longer should not be that difficult to imagine. First of all, the life span of the average person has almost doubled in the past couple centuries. So, having to deal with a much longer life span than one expected is not so farfetched. Second, it is not uncommon today for people to live well into their eighties or nineties. Medical science is continuing to make great advances. Preventative health care is making progress. And, finally, there seems to be some promising research in understanding the aging process. I don't want to speculate where all this is going, but it seems obvious that at some point we may have to consider the impact of people having dramatically longer life spans.

As the population continues to grow and the people continue to enjoy longer life spans, at some point, we are likely to face a problem of limited resources. When this happens we will have two choices. We let the scare resource problem sort itself out (i.e. the stronger take resources from the weaker). Or we can make social policies to keep the problem under control. At the broadest level, this would mean either limiting the number of new children or capping the life spans of older people. Initially, this sounds like a pretty easy choice. While it may seem distasteful to limit the number of children that people can have, it sounds much better than capping life spans which means forced euthanasia. But we are not off the hook so easily. If you limit children rather than life spans, you eventually wind up with an increasingly aged population. Is this a good thing? And without some restriction on life spans, the population will just become older and older. Clearly there is a role for younger people. But how do you decide the best way to maintain the balance? And, if you have limits on the number of young people, who gets to have them? Are some young people more desirable than others?

We have trouble today dealing with the concept of euthanasia. What will we do when we have to consider forced euthanasia for population control or to maintain population balance? How will we deal with that? The answer, I believe, is badly. We will apply our prejudices about euthanasia from the past to a future world that is markedly different. We will probably let many unfortunate epiphenomena occur to avoid coming to grips with some difficult moral decisions. Some will want to avoid this issue by claiming that over population won't be a problem in the future because we will have begun to populate other planets. The irony here is that people are not willing to accept the realistic possibility of dramatically extending life spans (which represents a problem) whereas they are willing to accept the less realistic possibility of colonizing other planets (which represents a solution).

If Human Cloning Were Possible, Should There Be Limits on the Uses of Clones?

In this next scenario we will push the boundaries of science even further to the day when human cloning is possible. In fact, let's push it to the point where you can create as many identical clones as you like. Some people might argue that you cannot create truly identical clones because the experience of growing up individuates people. However, if we can advance to the point of creating biologically identical humans, it is not much of a stretch to assume that we could control their experiences growing up as well.

Why would we want to do such a thing? There are lots of reasons most of which would offend modern sensibilities. However, that is not a reason to assume that we will never have to deal with them. First, you could make clones for spare parts. So, if you needed a hip transplant, you would not have to accept a mechanical implant. Better yet, if you needed a kidney transplant you would not have to worry about rejection if the kidney came from an identical clone. Another reason is to create armies of clones is for military purposes. Another reason might be for specialized workforces. Mind you, I am not advocating any of these possibilities. I am merely trying to create a plausibility argument. That is, no matter how offensive these possibilities may be to our modern sensibilities, it is certainly possible that people in the future may feel very differently on these issues.

Now, let us consider some of the potential moral and ethical issues. Do clones have the same status as individuals in our society? Should clones be allowed to vote? Initially, this seems like a shocking question. One of the deeply engrained tenets of our western world view is 'one person one vote'. However, this notion is based on a world of unique individuals that allows us to capture a diversity of opinions and gain the commitment of individuals for their governing bodies. However, an army of identical clones would not have a diversity of opinions. They would, presumably if identical, all share one opinion. If you wanted their opinion included in public discourse, you should only have to ask one of them. In addition, you would not gain their commitment by considering their opinions. Presumably they would be committed to whatever is good for the clone group. Now, consider the downside of allowing each one to vote. If somebody wanted to throw an election, all they would have to do is produce enough identical copies of a particular opinion to capture the election.

What about other rights? Should each copy of the clone have the same rights as every other person? Or should we begin to think about rights that accrue to groups rather than individuals? What if clone groups were created with special functions in mind. There may be laboring clones, teaching clones, professional clones, governing clones, clones who write books on ethics, and so on. If these groups are markedly different shouldn't their rights and responsibilities take this into consideration? Should our precept that 'all men are created equal' be revisited?

And finally, are there any boundaries on the usages of clones. Would it be OK to create some clones for spare parts? Could specialized clones be created for hazardous work? Could future wars be fought only by clones? Could a person who committed a crime have the sentence served by a clone? If we wished to populate other planets, would it be OK to breed clones specialized for the environment?

Asking these questions does not, by implication, answer them. I really have no idea what the right answers may be. However, I would encourage the reader to conduct the following mental experiment. First, answer the question as would be appropriate if these possibilities became realistic choices tomorrow. Then try to imagine a far distant future where these things are a reality. Imagine yourself as being in a position to make some hard decisions. Would you still decide the same way? And, finally, for whatever position you may be comfortable with try and imagine a situation in which you would reverse yourself. The purpose of this exercise is to show these issues would require some serious thought and a knee jerk application of unexamined prejudices is not likely to get us anywhere.

If You Could Read Minds, Would You Want To?

As we take on increasingly more improbable and unbelievable scenarios, let us assume, for our last one, that a technology has been developed that would allow you to read other people's minds. Would this be a good thing? I will not even attempt a plausibility argument for this one. However, I would point out that just a mere century and a half ago if you posed the possibility that a device were invented that would allow you to talk to someone anywhere on the globle, it would have been met with equal

skepticism. The prevalence of cell phones today reveals how unfounded such skepticism can be. So, who knows what the next century and a half will bring.

Let's say, for the sake of argument, that such a technology has been developed and that it is as inexpensive as a cell phone. Let's consider some of the thorny ethical questions that might arise. Would we do anything to control the usage of such a device? Would we simply allow anyone to read anyone else's thoughts? If we decided to restrict usage who would be allowed and who would not be permitted? Should use be limited to the government? That doesn't sound like a very good idea. Perhaps the device should be banned entirely. That doesn't seem to work either. If we restrict usage to the government that sounds a lot like 1984 and has some very negative connotations. And how can we keep a check on government usage if private citizens do not have access to the same information? Banning the technology does not seem to be a reasonable alternative because banning a technology does not prevent its use. It only prevents usage by honest citizens.

Perhaps social conventions will evolve. People have adjusted to a lot of things in the past and perhaps it is best to just rely on the natural evolution of conventions. The problem with this scenario is that modern guidelines provide little help and trying to imagine a future like this is also difficult. Perhaps the best thing is to just let the future unfold and deal with it. Perhaps, but I don't think so. The problem with "ethics as usual" is that, as the impact of technology increases, the time we have to sort out the ethical issues decreases.

Current Issues

The scenarios just presented are admittedly, and intentionally, rather farfetched. However, we have numerous issues today that are of a similar kind and not nearly as farfetched. Consider the use of cell phones. Cell phones present us with issues of privacy and protocol not unlike the issues that arise with mind reading equipment. Is it OK to talk on your cell phone on a crowded elevator? Are the other people on the elevator required to ignore you? Is appropriate to take a cell phone call from one person when you are having dinner with another? Is it acceptable to forward a text message? If you find somebody's cell phone can you look through the contents in an effort to find out whose phone it is? Should service vendors be allowed to collect and store text messages and conversations? Should they be allowed to process this information in order to figure out your interests so they can send you advertisements? What if those advertisements really do save you time in finding products and services that you like? Should people who are prone to buy from cell phone ads be given a cheaper rate for their cell phone service? Would it be OK to track your movement from you cell phone and send you ads from local vendors? Fortunately, most of these questions are of limited ethical significance and will get sorted out over time. But, they are a good example of problems presented by a new technology that we do not have easy answers to.

Consider another emerging technology – video games. Playing these games can be highly addictive and there are reports from time to time of some rather dire consequences. Some people see all the negative consequences of video games and feel that there usage should be limited. Others see them as a

promising new vehicle for work and education. Wouldn't it be great if work were as much fun as a video game? Wouldn't it be great if students enjoyed learning as much as they enjoy video games? Perhaps, but how do we know? Should there be any guidelines on the design and use of video games? Or should we just let it unfold and deal with the consequences? Remember, as things happen more quickly and have greater consequences, the "I'll deal with it when it happens" approach seems to be a little remiss.

One final example of a real modern issue is the problem of anonymity in virtual worlds. In virtual worlds like Sims Online, Habbo, and Second Life users can create an alternative online identity. (This is not limited to virtual worlds. You can create alternative identities in social interaction technologies such as Facebook, as well.) There is a benefit to this in that people often feel free to be more creative or expressive than they could be using their real world identity. This inspires imagination, creativity and personal exploration. On the other hand, when you are dealing with a virtual world persona, you have no connection to the real person. This can lead to negative emotional or financial consequences. So, is it better to maintain anonymity in virtual worlds for the benefits mentioned above? Or do the downsides dictate that Internet persona be clearly connected to a real person?

Summary

The point of this chapter was that we do not automatically have answers to the ethical issues that arise as a result of new technologies. We may think we have some answers but they are often knee jerk responses in which we apply rules from the past that may not apply to situations that we may encounter in the future. What we need is a way in which we can evaluate emerging ethical issues in the context of the future world that will be brought about by the technology rather than in the context of present. How do we create the context of a future world in which we can evaluate alternative answers to ethical problems? The answer, as we shall see in the next chapter is – stories. We can write stories to explore future possible worlds.

Chapter 3 – Writing Stories to Explore the Ethics of Technology

Why stories?

In the last chapter we considered some difficult (albeit hypothetical) ethical problems. There were three main difficulties in attempting to resolve these issues. First, attempting to apply known principles from the past within a present day context is not likely to work as the situations created by a technology must be considered within the context of the world created (or, at least, altered) by that technology. Second, it is difficult to 'imagine' the context of the future world within which the ethical decisions must be made. And, third, even if one could imagine what the future would be like, it would be very difficult to relate to it in the same way we can relate to the events in our daily lives. In this chapter, I will provide the basic justification for using stories to explore ethical issues that belong more to the future than they do to the present. In doing so, I hope to satisfactorily address these three difficult issues. However, before embarking on that explanation, a few preliminary remarks are in order.

The Importance of Seeing Multiple Perspectives

Ethical issues are ALWAYS a balance between competing interests. In order to make good ethical decisions one must be able to see the merits of each position and make an informed decision regarding who or what gains and who or what looses. In order to solidify this point, let's consider a simple example. We have a rather deeply rooted prohibition against stealing. Most people would just say that stealing is simply wrong. However, like all moral principles it is a tradeoff. Consider who gains and who loses with this prohibition. First of all, those who gain are those who own things, especially those who own things that they cannot easily protect. It also benefits those who can acquire the things they need through available channels such as working and earning income. Those who lose are those who do not have the things they need or want and cannot readily acquire those things through available channels. What if somebody is lazy and doesn't want to work. Wouldn't it be easier for them to just steal what they need? Perhaps it would. What if someone is much larger and stronger than the people around him? Wouldn't it be easier for him to just take what he needs from others? Yes, it probably would. But the moral principle prohibiting stealing says that these people lose. In fact, most people would say that these people deserve to lose. And I would agree. I am not trying to undermine our moral prohibition against stealing. It provides an enormous benefit in terms or protecting individuals and increasing harmony in society. However, as with any moral or ethical principle there are winners and there are losers. We tend to see things as simply right when we don't care much about the losers, or when we think they deserve to lose.

Since ethical issues are always a tradeoff between competing interests, it is important to be able to see all sides of an issue before making an ethical decision. I tell my students (usually much to their dismay) that if you cannot argue more than one side of an issue then you have no right to an opinion. This is because, if you can only argue one side of an issue, you do not 'have' an opinion. You are merely repeating somebody else's opinion. In order to have an opinion of your own, you must fairly consider all sides of an issue and come to a conclusion regarding which side you support. In order to fairly consider

the other sides, you should be able to convincingly argue their positions and then explain why you did not choose that position. It is unfortunate, these days, that we see so many pundits in the media shouting at each other and talking past each other; refusing to acknowledge good points offered by the other side and refusing to give ground when it would be appropriate. This does not contribute to the debate. It merely contributes to the noise and confusion.

The following joke, adapted from Cathcart and Klein's delightful book Plato and a Platypus Walk into a Bar, makes this point, of not seeing the perspective of others, quite clear.

I was walking down the C&O Canal towpath one day when I saw a guy on the other side hollering and waving his arms trying to get my attention. I walked over to the edge of the canal so I could hear what he was saying.

"How do I get to the other side?" he inquired.

"Silly, you are on the other side", I replied.

This somewhat lame joke points to the fact that we often do not, or cannot, see things from another person's perspective. We are inclined to see things from our perspective and think that is the only legitimate perspective. So how can we make appropriate ethical decisions given the blinders we have for other perspectives? Here is where stories play an important role.

We interpret the events of the world based on the experiences we have had in the world. The reason why it is so difficult to see things from somebody else's perspective is that we do not have the lenses of their experiences from which to see the world. Sometimes, with a little effort we can put ourselves, 'in their shoes' and speculate on how they may see things. However, this is chancy and only works to the extent that we really do have experiences in common.

The benefit of reading a story is that it gives you an experience that you may not have had. Then, based upon that experience, you can see the world a little differently. Once you have seen the world from another perspective, it becomes yet easier to see it from yet other perspectives. If you have experienced only one perspective, it is very difficult to believe that there are other legitimate perspectives. As you see other legitimate perspectives you develop a more balanced view. Since stories provide us with a means of seeing a situation from multiple perspectives with competing interests, they help us to understand the richness and complexity in an ethical issue. This in turn helps us to make better informed ethical decisions. But helping us to see multiple perspectives is not the only benefit that stories provide. They allow us to see issues from future perspectives as well.

Writing Stories to Explore the Ethics of Technology

The Importance of Future Perspectives

As difficult as it may be to see an issue from multiple current perspectives, it is even more difficult to see an issue from the perspective of people in a possible future. As was mentioned earlier, we look at the protection of privacy from the perspective of the past where there was an incredible imbalance in who had access to information and who did not. So, when we consider the possible transgressions that could occur with personal information available on the internet, it is reasonable to assume that individuals need to be protected. However, if we consider this in the context of a future world where individuals are empowered by virtually limitless access to information, it does not seem quite the same. But, how do we put ourselves within the context of a future world. The answer, of course, is through stories. In stories we create fictional worlds and then experience those fictional worlds. From those experiences we can make decisions about what we like or don't like about that possible world.

This is not a new idea. From the very beginning, stories were used to create fictional worlds to provide us with vicarious experiences. Charles Dickens began a long line of social commentators who used fiction to shine a light on undesirable present and future consequences of various social policies and trends. In fact, it might be fair to say that stories have been a bit overused to explore future negative consequences and a bit under used to reveal future positive consequences. Following on McLuhan's analogy of driving while looking in the rear view mirror, it may be appropriate to point out that what we need is a headlight into the future. And, perhaps, stories can provide us with that headlight.

Some Conceptual Grounding

What I am proposing here is, hopefully, intuitively appealing. We can create stories which allow us to vicariously experience possible future worlds in hopes of making better ethical decisions regarding emerging technologies. But, 'intuitively appealing' is not enough. There are some compelling conceptual arguments for using stories as well. Later on in this book, I will explore these issues in greater depth. However, for now we will get a quick introduction to them as they will help with our 'willing suspension of disbelief'.

Narrative versus Logical Arguments

People like to make sense out of their experiences. It gives them a feeling that life makes sense and is, to some extent, predictable. So, we organize our experiences in terms of causal arguments, that is, A occurred and, so, then B occurred. And these causal arguments make us feel as though we have some control over our lives and that we understand what we are doing and why. Most people are aware of logical arguments of this form, but overwhelmingly narrative arguments are more common. A logical argument might proclaim some like the following: All conservatives prefer small government. Bob is a conservative. Therefore Bob prefers small government. However, as it turns out, people rarely actually use logical arguments in their daily lives. In fact, people rarely use logic outside of classes in logic. It is much more likely that people reason through an issue using a narrative argument. Narrative arguments

are causal sequences organized not according to the formal rules of logic but in terms of our understanding of human experience. For example, consider the following: Bob had a difficult childhood. There was never enough money and they were always struggling to make ends meet. Today Bob is very conservative and very careful with his money. Although this is not in a formal logical argument, people will look at this narrative argument and decide whether or not it makes sense in terms of their understanding of human nature. Narrative arguments are not exclusive like logical arguments. In logic, if A is true, 'not A' cannot be true as well. However, in narrative arguments you do not get this precision. Consider the following: Bob had a difficult childhood. There was never enough money and they were always struggling to make ends meet. Today Bob is a big advocate of government spending on social program to help the less fortunate. Both are reasonable claims given our understanding of human nature. However, consider what happens if we change it slightly. Bob had a difficult childhood. There was never enough money and they were always struggling to make ends meet. So, Bob decided to go out today and buy a big screen high definition television set. This would not be considered a valid narrative argument. While it may certain have been the case that Bob went out and bought a new TV, the premises do not really explain why.

Overwhelmingly, people reason in terms of narrative arguments more often than using logical arguments. Since stories employ narrative arguments, they provide us with a means of reasoning through difficult ethical issues in a way that is much more natural for us and closer to the way in which we actually reason.

The Role of Emotions in Reasoning

A conventional view of the role of emotion in reasoning is that emotion has no role. In fact, we believe that emotions cloud reasoning. This is not only misleading and unfortunate, it is entirely false. Granted that emotion can sometimes high jack our reasoning and we should avoid such extremes. However, practical reasoning, without emotion, is simply not possible. If our reasoning is limited to logical syllogisms, then perhaps it may be possible to reason purely without emotion. However, as soon as our reasoning involves decision or values, it is simply not possible to reason in a purely objective manner. Consider a simple decision where we have a choice between three alternatives. However do we decide? Well, we choose the outcome that is the most desirable. How do we know which is most desirable? We choose the alternative that provides us with the greatest value. In fact, we use the term evaluate in the context of selecting alternatives. And value is not logical it is emotional. We value things if we feel good about them.

One might try to salvage the purity of reasoning by pointing to objective criteria for a decision. For example, we want the cheapest alternative or the alternative with the highest probability of success. But, those numbers are only measures of value and we are making the decision based on what we value.

Stories provide us with a superior vehicle for ethical reasoning because they allow us to connect with the argument emotionally. In fact, we care about the outcome of the story because we care about the

characters. Since ethical decisions are always value decisions, it behooves us to reason about them in a way that keeps us in touch with our values.

Case studies are a popular way of approaching ethical problems. And while they have merit for teasing out the subtleties of ethical reasoning, they fail completely to put those ethical issues in a human context. The characters in case studies are, at best, archetypes and at worst two dimensional cardboard characters. In order to understand the complexities of ethical dilemmas we need to understand them in their full human context.

A great example is an overused criticism of utilitarian thinking. Utilitarians believe that we should always act in such a way as to produce the greatest amount of happiness for the greatest number of people. This is a bit simplistic but my point here is not to poke fun at utilitarianism. Rather it is to reveal how emotional responses to human circumstances can trump purely logical ethical reasoning. Nonetheless, to get on with the story, there is a hospital in which six patients lay dying. One is in need of a heart transplant; one needs a kidney; one, a new set of lungs; another, a liver. And so on. A healthy person walks into the hospital. Should you carve up the healthy person in order to save the lives of the others? Or should you let the healthy person live and the others die? Utilitarianism would seem to suggest that carving up the healthy patient is the right answer because it will create great happiness for the six people who get to live and great unhappiness for the healthy person who got carved up. Any person with even the tiniest amount of moral development would recoil in horror at the idea of sacrificing the healthy person. Apparently there is an important value that is not captured in the prescription to produce the greatest amount of happiness for the greatest number of people. One might try to salvage the argument by saying that the unhappiness of the sacrificed patience outweighs the aggregate happiness of the saved patients. But that is not a legitimate argument since happiness cannot be measured with such precision. Further, the carved up healthy person is no longer unhappy as they have been sacrificed. Once again we could try to salvage this by point out that there may be great unhappiness among potential visitors to hospitals. But, would this be offset by the happiness of potential transplant recipients? The point is that ethical problems cannot be decided by pure logic alone. There is an important emotional component. That is, we must feel good about the decision. And stories help us consider the emotional implications of these decisions.

The Role of Imagination

We have a very ambiguous relationship with our imaginations. On one hand we are very distrustful of our imaginings. You often hear a child chided with the question “is that true or are you just imagining it?” The underlying claim is that products of our imagination should be dismissed as unreal. On the other hand, most of what we view as reality is a product of our imagination. And many important cognitive activities such as planning, forecasting, anticipating or mentally rehearsing rely on the imagination. Later on in this book, I will go into some detail about the importance of the imagination. But, for now, suffice it to say that our ability to analyze helps us understand the present, whereas our ability to imagine helps us to understand the future. Our imagination is no more or less worthy than our analytical

abilities. When used properly and in a disciplined manner we achieve positive results in the form of greater understanding. When used improperly or in an undisciplined manner we can easily mislead ourselves. But, that shortcoming is true for analysis as well.

The connection between stories and the imagination should be obvious. Stories are a product of the imagination. They also exercise and develop the imagination. In a world in which the events of the future are becoming more important than the events of the past, our ability to imagine becomes increasingly more important. And stories provide us with a vehicle for developing and exercising the imagination.

The Role of Stories in Moral Development

It is one thing to proclaim ethical principles. It is another thing entirely to get people to follow them. You can have laws and rules and codes of ethics. But, ultimately, people behave ethically when they want to and they want to when they have achieved the appropriate level of moral development. But, how do you get people to develop morally? Usually, this occurs as a reaction to the experiences that one has in their lives. But, it is not reasonable to expect each person to have the full set of personal experiences that would produce moral development. It is simply impractical to arrange all those experiences for all those people. However, one can have vicarious experiences through the stories they read. And in that way, stories can be used to promote moral development as well.

Conclusion

Several points have been raised in this chapter and, lest we lose the thread, it is important to refocus on the main points before closing. Stories allow us to see things from multiple perspectives. They allow us to see things from the perspective of a future world which may be a more appropriate context for evaluating an ethical issue. Stories allow us to engage our emotions in the reasoning process so that we can feel good about our decisions. And stories both engage and develop the imagination which allows us to understand future possibilities more fully. In the past we have used disciplined techniques such as logic and scientific method to understand the natural world and our history in it. But these techniques do not help us with problems of value or anticipating what the future may bring. As the future comes at us faster and faster, especially in the case of new technologies, it becomes incumbent upon us to understand the future before it gets here. And stories, used in a disciplined manner, may be just the vehicle to help us do that.