INFORMATION TECHNOLOGY AND MORAL VALUES

Information technology is now ubiquitous in the lives of people across the globe. These technologies take many forms such as personal computers, smart phones, the internet, web and mobile phone applications, digital assistants, and cloud computing.

In fact the list is growing constantly and new forms of these technologies are working their way into every aspect of daily life. In some cases, such as can be seen in massive multiplayer online games, these technologies are even opening up new ways of interacting with each other. Information technology at its basic level is technology that records, communicates, synthesizes or organizes information. Information can be understood as any useful data, instructions, or meaningful message content.

The word literally means to "give form to" or to shape one's thoughts. So a basic type of information technology might be the proverbial string tied around one's finger to remind or inform you that you have some specific task to accomplish today. Here the string stands in for a more complex proposition such as "buy groceries before you come home." The string itself is not the information, it merely

symbolizes the information and therefore this symbol must be correctly interpreted for it to be useful. Which raises the question, what is information itself?

Unfortunately, there is not a completely satisfying and philosophically rigorous definition available, though there are at least two very good starting points. For those troubled by the ontological questions regarding information, we might want to simply focus on the symbols and define information as any meaningfully ordered set of symbols. This move can be very useful and mathematicians and engineers prefer to focus on this aspect of information, which is called "syntax" and leave the meaningfulness of information or its "semantics" for others to figure out.

Claude E. Shannon working at Bell Labs produced a landmark mathematical theory of communication (1948), where he took his experiences in cryptography and telephone technologies and worked out a mathematical formulation describing how syntactical information can be turned into a signal that is transmitted in such a way as to mitigate noise or other extraneous signals which can then be decoded by the desired receiver of the message (Shannon 1948; Shannon and Weaver 1949). The concepts described by Shannon, along with additional important innovations made by others who are too many to list, explain the way that information technology works, but we still have the deeper issue to resolve if we want to thoroughly trace the impact of information technologies on moral values.

The second starting point is a bit more deeply philosophical in nature. Here we begin with the claim that information either constitutes or is closely correlated with what constitutes our existence and the existence of everything around us. This means that information plays an ontological role in the manner in which the universe operates. A standpoint such as this would place information at the center of concern for philosophy and this idea has given rise to the new fields of Information Philosophy and Information Ethics. Philosophy of Information will not be addressed in detail here but the interested reader can begin with Floridi (2010b, 2011b) for an introduction. Some of the most important aspects of Information Ethics will be outlined in more detail below.

Every action we take leaves a trail of information that could be recorded and stored for future use. For instance, you might use the simple technology of keeping a detailed diary listing all the things you did and thought during the day. But today you could augment that with even more detail gathered with advanced information technologies some examples include; all of your economic transactions, a GPS generated plot of where you traveled, a list of all the web addresses you visited and the details of each search you initiated online, a listing of all your vital signs such as blood pressure and heart rate, all of your dietary intakes for the day, and many other examples can be imagined. As you go through this thought experiment you begin to see the complex trail of data that you generate each and every day and how that same data might be collected and stored though the use of information technologies. Here we can begin to see how information technology can impact moral values. As this data gathering becomes more automated and ever-present, we must ask who is in control of this data, what is to be done with it, and who will insure its accuracy. For instance, which bits of information should be made public, which held private, and which should be allowed to become the property of third parties like corporations? Questions of the production, access and control of information will be at the heart of moral challenges surrounding the use of information technology.

One might argue that this situation is no different from the moral issues revolving around the production, access and control of any basic necessity of life. But there is one major difference, if one party controls the access of some natural resource, then that by necessity excludes others from using it. This is not necessarily so with digital information, it is non-exclusory, meaning we can all at least theoretically possess the same digital information because copying it from one digital source to another does not require eliminating the previous copy. Since there is no physical obstacle to the spread of all information, then there remain only appeals to morality, or economic justice, which might prevent distributing certain forms of information. Therefore, understanding the role of moral values in information technology is indispensable to the design and use of these technologies (Johnson 1985; Moor 1985; Nissenbaum 1998; Spinello 2001). It should be noted that this entry will not directly address the phenomenological approach to the ethics of information technology since there is a detailed entry on this subject).

1. The Moral Challenges of Information Technology

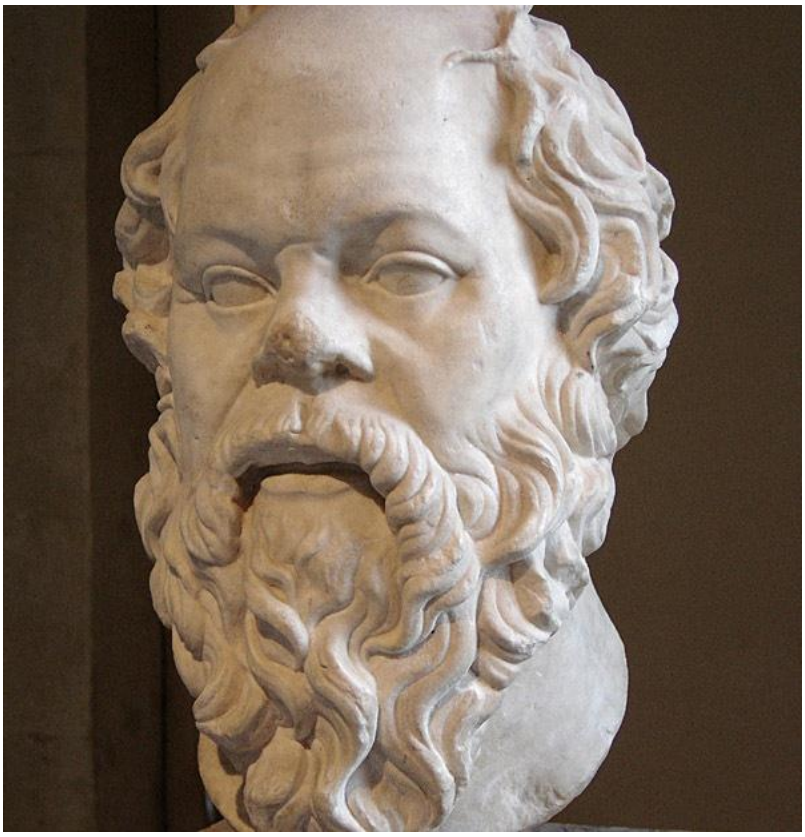

The move from one set of dominant information technologies to another is always morally contentious. Socrates lived during the long transition from a largely oral tradition to a newer information technology consisting of writing down words and information and collecting those writings into scrolls and books.

Famously Socrates was somewhat antagonistic to writing and he never wrote anything down himself. Ironically, we only know about Socrates' argument against writing because his student Plato ignored his teacher and wrote it down in a dialogue called"Phaedrus" (Plato). Towards the end of this dialogue Socrates discusses with his friend Phaedrus the"…conditions which make it (writing) proper or improper. Socrates tells a fable of an Egyptian God he names Theuth who gives the gift of writing to a king named Thamus. Thamus is not pleased with the gift and replies,

If men learn this, it will implant forgetfulness in their souls; they will cease to exercise memory because they rely on that which is written, calling things to remembrance no longer from within themselves, but by means of external marks.

Socrates, who was adept at quoting lines from poems and epics and placing them into his conversations, fears that those who rely on writing will never be able to truly understand and live by these words. For Socrates there is something immoral or false about writing. Books can provide information but they cannot, by themselves, give you the wisdom you need to use or deeply understand that information.

Conversely, in an oral tradition you do not simply consult a library, you are the library, you are a living manifestation of the information you know by heart. For Socrates, reading a book is nowhere near as insightful as talking with its author. Written words seem to talk to you as though they were intelligent,

but if you ask them anything about what they say, from a desire to be instructed, they go on telling you the same thing forever.

His criticism of writing at first glance may seem humorous but the temptation to use recall and call it memory is getting more and more prevalent in modern information technologies. Why learn anything when information is just an Internet search away? In order to avoid Socrates' worry, information technologies should do more than just provide access to information; they should also help foster wisdom and understanding as well.

1.1 The Fundamental Character of Information Technologies

Early in the information technology revolution Richard Mason suggested that the coming changes in information technologies would necessitate rethinking the social contract (Mason 1986). What he could not have known then was how often we would have to update the social contract as these technologies rapidly change. Information technologies change quickly and move in and out of fashion at a bewildering pace. This makes it difficult to try to list them all and catalog the moral impacts of each.

The very fact that this change is so rapid and momentous has caused some to argue that we need to deeply question the ethics of the process of developing emerging technologies (Moor 2008). It has also been argued that the ever morphing nature of information technology is changing our ability to even fully understand moral values as they change.

Lorenzo Magnani claims that acquiring knowledge of how that change confounds our ability to reason morally "…has become a duty in our technological world" (Magnani 2007, 93). The legal theorist Larry Lessig warns that the pace of change in information technology is so rapid that it leaves the slow and deliberative process of law and political policy behind and in effect these technologies become lawless, or extralegal. This is due to the fact that by the time a law is written to curtail, for instance, some

form of copyright infringement facilitated by a particular file sharing technology, that technology has become out of date and users are on to something else that facilitates copyright infringement (Lessig 1999). But even given this rapid pace of change it remains the case that information technologies or applications can all be categorized into at least three different types each of which we will look at below.

All information technologies record (store), transmit (communicate), organize and/or synthesize information. For example, a book is a record of information, a telephone is used to communicate information, and the Dewey decimal system organizes information. Many information technologies can accomplish more than one of the above functions and, most notably, the computer can accomplish all of them since it can be described as a universal machine, so it can be programmed to emulate any form of information technology.

1.1.1 Moral Values in Information Recording

We live in a world rich in data and the technology to record and store vast amounts of this data has grown rapidly. The primary moral concern here is that when we collect, store, and/or access information it is done in a just manner that anyone can see is fair and in the best interests of all parties involved. As was mentioned above, each of us produces a vast amount of information every day that could be recorded and stored as useful data to be accessed later when needed. But moral conundrums arise when that collection, storage and use of our information is done by third parties without our knowledge or done with only our tacit consent.

The control of information is power. The social institutions that have traditionally exercised this power are things like, religious organizations, universities, libraries, healthcare officials, government agencies, banks and corporations. These entities have access to stored information that gives them a certain amount of power over their customers and constituencies. Today each citizen has access to more and more of that stored information without the necessity of utilizing the traditional mediators of that information and therefore a greater individual share of social power.

One of the great values of modern information technology is that it makes the recording of information easy, and in some cases, it is done automatically. Today, a growing number of people enter biometric data such as blood pressure, calorie intake, exercise patterns, etc. into applications designed to help them achieve a healthier lifestyle. This type of data collection could become more automated in the near future.

There are already applications that use the GPS tracking available in many phones to track the length and duration of a user's walk or run. How long until a smartphone collects a running data stream of your blood pressure throughout the day perhaps tagged with geo-location markers of particularly high or low readings? In one sense this could be immensely powerful data that could lead to much healthier lifestyle choices. But it could also be a serious breach in privacy if the information got into the wrong hands which would be easily accomplished since third parties have access to information collected on smartphones and online applications.

When searching on the Internet, browser software records all manner of data about our visits to various websites which can, for example, make webpages load faster next time you visit them. Even the websites themselves use various means to record information when your computer has accessed them and they may leave bits of information on your computer which the site can use the next time you visit. Some websites are able to detect which other sites you have visited or which pages on the website you spend the most time on. If someone were following you around a library noting down this kind of information you might find it uncomfortable or hostile, but online this kind of behavior takes place behind the scenes and is barely noticed by the casual user.

According to some professionals, information technology has all but eliminated the private sphere. Scott McNealy of Sun Microsystems famously announced in 1999: "You have zero privacy anyway. Get over it" (Sprenger, 1999). Helen Nissenbaum observes that,

[w]here previously, physical barriers and inconvenience might have discouraged all but the most tenacious from ferreting out information, technology makes this available at the click of a button or for a few dollars (Nissenbaum 1997)

and since the time when she wrote this the gathering of data has become more automated and cheaper. Clearly, earlier theories of privacy that assumed the inviolability of physical walls no longer apply but as Nissenbaum argues, personal autonomy and intimacy require us to protect privacy nonetheless (Nissenbaum 1997).

A final concern in this section is that information technologies are now storing user data in "the cloud" meaning that the data is stored on a device remotely located from the user and not owned or operated by that user, but the data is then available from anywhere the user happens to be on any device he or she happens to be using. This ease of access has the result of also making the relationship one has to one's own data more tenuous because of the uncertainty about the physical location of that data. Since personal data is crucially important to protect, the third parties that offer"cloud" services need to understand the responsibility of the trust the user is placing in them. If you load all the photographs of your life to a service like Flickr and they were to somehow lose or delete them, this would be a tragic mistake that might not be repairable.

1.1.2 Moral Values in Communicating and Accessing Information

Information technology has forced us to rethink a simple notion of privacy into more complex theories that recognize both the benefits and risks of communicating all manner of information. The primary moral values of concern are privacy, ownership, trust and the veracity of the information being communicated.

Who has the final say whether or not some information about a user is communicated or not? Who is allowed to sell your medical records, your financial records, your friend list, your browser history, etc.? If you do not have control over this process, then how can you claim a right to privacy? For instance Alan Westin argued in the very early decades of digital information technology that control of access to one's personal information was the key to maintaining privacy (Westin 1967). It follows that if we care about privacy, then we should give all the control of access to personal information to the individual.

Most corporate entities resist this notion as information about users has become a primary commodity in the digital world boosting the fortunes of corporations like Google or Facebook. There is a great deal of utility each of us gains from the services of internet search companies. It might actually be a fair exchange that they provide search results for free based on collecting data from individual user behavior that helps them rank the results. This service comes with advertising that is directed at the user based on his or her search history. That is, each user tacitly agrees to give up some privacy whenever they use the service.

If we follow the argument raised above that privacy is equivalent to information control then we do seem to be ceding our privacy away little by little. Herman Tavani and James Moor (2004) argue that in some cases giving the user more control of their information may actually result in greater loss of privacy. Their primary argument is that no one can actually control all of the information about oneself that is produced each day.

If we focus only on the little bit we can control, we lose site of the vast mountains of data we cannot (Tavani and Moor 2004). Tavani and Moor argue that privacy must be recognized by the third parties that do control your information and only if those parties have a commitment to protecting user privacy will we actually have any real privacy and towards this end they suggest that we think in terms of restricted access to information rather than strict control of personal information (Tavani and Moor 2004).

Information security is also an important moral value that impacts the communication and access of user information. If we grant the control of our information to third parties in exchange for the services they provide, then these entities must also be responsible for restricting the access to that information by others who might use it to harm us (see Epstein 2007; Magnani 2007; Tavani 2007). With enough information, a person's entire identity might be stolen and used to facilitate fraud and larceny.

The victims of these crimes can have their lives ruined as they try to rebuild such things as their credit rating and bank accounts. This has led to the design of computer systems that are more difficult to access and the growth of a new industry dedicated to securing computer systems.

The difficulty in obtaining complete digital security rests in the fact that security is antithetical to the moral values of sharing and openness that guided many of the early builders of information technology.

Steven Levy (1984) describes in his book,"Hackers: Heroes of the Computer Revolution," a kind of"Hacker ethic," that includes the idea that computers should be freely accessible and decentralized in order to facilitate"world improvement" and further social justice (Levy 1984; see also Markoff 2005). So it seems that information technology has a strong dissonance created in the competing values of security and openness based on the competing values of the people designing the technologies themselves.

This conflict in values has been debated by philosophers. While many of the hackers interviewed by Levy argue that hacking is not as dangerous as it seems and that it is mostly about gaining knowledge

of how systems work, Eugene Spafford counters that no computer break-in is entirely harmless and that the harm precludes the possibility of ethical hacking except in the most extreme cases (Spafford 2007).

Kenneth Himma largely agrees that hacking is largely unethical but that politically motivated hacking or "Hacktivism" may have some moral justification though he is hesitant to give his complete endorsement of the practice due to the largely anonymous nature of the speech entailed by the hacktivist protests (Himma 2007b). Mark Manion and Abby Goodrum agree that hacktivism could be a special case of ethical hacking but warn that it should proceed in accordance to the moral norms set by the acts of civil disobedience that marked the twentieth century or risk being classified as online terrorism (Manion and Goodrum 2007).

A very similar value split plays out in other areas as well, particularly in intellectual property rights. What information technology adds to these long standing moral debates is the nearly effortless access to information that others might want to control such as intellectual property, dangerous information and pornography (Floridi 1999), along with the anonymity of both the user and those providing access to the information (Nissenbaum 1999; Sullins 2010).

For example, even though cases of bullying and stalking occur regularly, the anonymous and remote actions of cyber-bullying and cyberstalking make these behaviors much easier and the perpetrator less likely to be caught. Arguably, this makes these unethical behaviors on cyberspace more likely that the design of cyberspace itself tacitly promotes unethical behavior (Adams 2002; Grodzinsky and Tavani 2002).

Since the very design capabilities of information technology influence the lives of their users, the moral commitments of the designers of these technologies may dictate the course society will take and our commitments to certain moral values (Brey 2010; Bynum 2000; Ess 2009; Johnson 1985; Magnani 2007; Moor 1985; Spinello 2001; Sullins 2010).

Assuming we are justified in granting access to some store of information that we may be in control of, there is a duty to ensure that that information is useful and accurate. If you use a number of different search engines to try to find some bit of information, each of these searches will vary from one another. This shows that not all searches are equal and it matters which search provider you use.

All searches are filtered to some degree in order to ensure that the information the search provider believes is most important to the user is listed first. A great deal of trust is placed in this filtering process and the actual formulas used by search providers are closely held trade secrets. The hope is that these decisions are morally justifiable but it is difficult to know. If we are told a link will take us to one location on the web yet when we click it we are taken to some other place, the user may feel that this is a breach of trust.

This is often called"clickjacking" and malicious software can clickjack a browser by taking the user to some other site than is expected; it will usually be rife with other links that will further infect your machine or sites that pay the clickjacker for bringing traffic to them (Hansen and Grossman, 2008).

Again the anonymity and ease of use that information technology provides can facilitate deceitful practices. Pettit (2009) suggests that this should cause us to reevaluate the role that moral values such as trust and reliance play in a world of information technology.

Lastly in this section we must address the impact that the access to information has on social justice. Information technology was largely developed in the Western industrial societies during the twentieth century. But even today the benefits of this technology have not spread evenly around the world and to all socioeconomic demographics. Certain societies and social classes have little to no access to the information easily available to those in more well off and in developed nations, and some of those who have some access have that access heavily censored by their own governments.

This situation has come to be called the "digital divide," and despite efforts to address this gap it may be growing wider. While much of this gap is driven by economics (see Warschauer 2003), Charles Ess notes that there is also a problem with the forces of a new kind of cyber enabled colonialism and ethnocentrism that can limit the desire of those outside the industrial West to participate in this new"Global Metropolis" (Ess 2009). John Weckert also notes that cultural differences in giving and taking offence play a role in the design of more egalitarian information technologies (Weckert 2007). Others argue that basic moral concerns like privacy are weighed differently in Asian cultures (Hongladarom 2008; Lü 2005).

1.1.3 Moral Values in Organizing and Synthesizing Information

In addition to storing and communicating information, many information technologies automate the organizing of information as well as synthesizing or mechanically authoring or acting on new information. Norbert Wiener first developed a theory of automated information synthesis which he called *Cybernetics* (Wiener 1961 [1948]). Wiener realized that a machine could be designed to gather information about the world, derive logical conclusions about that information which would imply certain actions, which the machine could then implement, all without any direct input form a human agent.

Wiener quickly saw that if his vision of cybernetics was realized, there would be tremendous moral concerns raised by such machines which he outlined in his book *the Human Use of Human Beings* (Wiener 1950). Wiener argued that, while this sort of technology could have drastic moral impacts, it was still possible to be proactive and guide the technology in ways that would increase the moral reasoning capabilities of both humans and machines (Bynum 2008).

Machines make decisions that have moral impacts. Wendell Wallach and Colin Allen tell an anecdote in their book"Moral Machines" (2008). One of the authors left on a vacation and when he arrived overseas his credit card stopped working, perplexed, he called the bank and learned that an automatic anti-theft program had decided that there was a high probability that the charges he was trying to make were from someone stealing his card and that in order to protect him the machine had denied his credit card transactions.

Here we have a situation where a piece of information technology was making decisions about the probability of nefarious activity happening that resulted in a small amount of harm to the person that it was trying to help. Increasingly, machines make important life changing financial decisions about people without much oversight from human agents. Whether or not you will be given a credit card, mortgage loan, the price you will have to pay for insurance, etc. is very often determined by a machine.

For instance if you apply for a credit card the machine will look for certain data points, like your salary, your credit record, the economic condition of the area you're in, etc., and then calculates a probability that you will default on your credit card, that probability will either pass a threshold of acceptance or not and determine whether or not you are given the card. The machine can typically learn as well to make better judgments given the results of earlier decisions it has made.

Machine learning and prediction is based on complex logic and mathematics (see for example Russell and Norvig 2010), this complexity may result in slightly humorous examples of mistaken prediction as told above, or it might interpret the data of someone's friends and acquaintances, his or her recent purchases, and other social data which might result in the mistaken classification of that person as a potential terrorist, thus altering that person's life in a powerfully negative way (Sullins 2010).

It all depends on the design of the learning and prediction algorithm, something that is typically kept secret.

1.2 The Moral Paradox of Information Technologies

Several of the issues raised above result from the moral paradox of Information technologies. Many users want information to be quickly accessible and easy to use and desire that it should come at as low a cost as possible, preferably free. But users also want important and sensitive information to be secure, stable and reliable. Maximizing our value of quick and low cost minimizes our ability to provide secure and high quality information and the reverse is true also. Thus the designers of information technologies are constantly faced with making uncomfortable compromises. The early web pioneer Stewart Brand sums this up well in his famous quote:

In fall 1984, at the first Hackers' Conference, I said in one discussion session: "On the one hand information wants to be expensive, because it's so valuable. The right information in the right place just changes your life. On the other hand, information wants to be free,because the cost of getting it out is getting lower and lower all the time. So you have these two fighting against each other"(Clarke 2000—see Other Internet Resources)[1]

Since these competing moral values are essentially impossible to reconcile, they are likely to continue to be at the heart of moral debates in the use and design of information technologies for the foreseeable future.

2. Specific Moral Challenges at the Cultural Level

In the section above, the focus was on the moral impacts of information technologies on the individual user. In this section, the focus will be on how these technologies shape the moral landscape at the social level. At the turn of the century the term "web 2.0" began to surface and it referred to the new way that the world wide web was being used as a medium for information sharing and collaboration as

well as a change in the mindset of web designers to include more interoperability and user-centered experiences on their websites.

This term has also become associated with "social media" and "social networking." While the original design of the web by its creator Tim Berners-Lee was always one that included notions of meeting others and collaboration, users were finally ready to fully exploit those capabilities by 2004 when the first Web 2.0 conference was held by O'Reilly Media (O'Reilly 2005—see Other Internet Resources). This change has meant that a growing number of people have begun to spend significant portions of their lives online with other users experiencing an unprecedentedly new kind of lifestyle. Social networking is an important part of many people's lives now where massive numbers of people congregate on sites like Facebook and interact with friends old and new, real and virtual.

The Internet offers the immersive experience of interacting with others in virtual worlds where environments constructed from information. Just now emerging onto the scene are technologies that will allow us to merge the real and the virtual. This new "augmented reality" is facilitated by the fact that many people now carry GPS enabled smart phones and other portable computers with them upon which they can run applications that let them interact with their surroundings and their computers at the same time, perhaps looking at an item though the camera in their device and the "app" calling up information about that entity and displaying it in a bubble above the item. Each of these technologies comes with their own suite of new moral challenges some of which will be discussed below.

2.1 Social Media and Networking

Social networking is a term given to sites and applications that facilitate online social interactions that typically focus on sharing information with other users referred to as "friends." The most famous of these sites today is Facebook. There are a number of moral values that these sites call into question. Shannon Vallor (2011) has reflected on how sites like Facebook change or even challenge our notion of friendship. Her analysis is based on the Aristotelian theory of friendship.

Aristotle argued that humans realize a good and true life though virtuous friendships. Valor notes that four key dimensions of Aristotle's'virtuous friendship,' namely: reciprocity, empathy, self-knowledge and the shared life, are found in online social media in ways that can actually strengthen friendship (Vallor 2011). Yet she argues that social media is not up to the task of facilitating what Aristotle calls 'the shared life,' and thus these media cannot fully support the Aristotelian notion of complete and virtuous friendship by themselves (Vallor 2011).

Vallor also has a similar analysis of other Aristotelian virtues such as patience, honesty and empathy as they are fostered in online media (Vallor 2010). Johnny Hartz Søraker (2012) argues for a nuanced understanding of online friendship rather than a rush to normative judgement on the virtues of virtual friends.

There are, of course, privacy issues that abound in the use of social media. James Parrish following Mason (1986) recommends four policies that a user of social media should follow to ensure proper ethical concern for other's privacy:

1. When sharing information on SNS (social network sites), it is not only necessary to consider the privacy of one's personal information, but the privacy of the information of others who may be tied to the information being shared.
2. When sharing information on SNS, it is the responsibility of the one desiring to share information to verify the accuracy of the information before sharing it.
3. A user of SNS should not post information about themselves that they feel they may want to retract at some future date. Furthermore, users of SNS should not post information that is the product of the mind of another individual unless they are given consent by that individual. In both cases, once the information is shared, it may be impossible to retract.
4. It is the responsibility of the SNS user to determine the authenticity of a person or program before allowing the person or program access to the shared information.

These systems are not typically designed to protect individual privacy, but since these services are typically free there is a strong economic drive for the service providers to harvest at least some information about their user's activities on the site in order to sell that information to advertisers for directed marketing.

2.1.1 Online Games and Worlds

The first moral impact one encounters when contemplating online games is the tendency for these games to portray violence. There are many news stories that claim a cause and effect relationship between violence in computer games and real violence. The claim that violence in video games has a causal connection to actual violence has been strongly critiqued by the social scientist Christopher J. Ferguson (Ferguson 2007).

However, Mark Coeckelbergh argues that since this relationship is tenuous at best and that the real issue at hand is the effect these games have on one's moral character (Coeckelbergh 2007). But Coeckelbergh goes on to claim that computer games can be designed to facilitate virtues like empathetic and cosmopolitan moral development so he is not arguing against all games just those where the violence inhibits moral growth (Coeckelbergh 2007). Marcus Schulzke (2010) holds a different opinion, suggesting that the violence in computer games is morally defensible.

Schulzke's main claim is that actions in a virtual world are very different from actions in the real world, though a player may “kill”another player in a virtual world, that player is instantly back in the game and the two will almost certainly remain friends in the real world thus virtual violence is very different from real violence, a distinction gamers are comfortable with (Schulzke 2010). While virtual

violence may seem palatable to some, Morgan Luck (2009) seeks a moral theory that might be able to allow the acceptance of virtual murder but that will not extend to other immoral acts such as pedophilia. Christopher Bartel (2011) is less worried about the distinction Luck attempts to draw; Bartel argues that virtual pedophilia is real child pornography, which is already morally reprehensible and illegal across the globe.

While violence is easy to see in online games, there is a much more substantial moral value at play and that is the politics of virtual worlds. Peter Ludlow and Mark Wallace describe the initial moves to online political culture in their book, *The Second Life Herald: The Virtual Tabloid that Witnessed the Dawn of the Metaverse*(2007). Ludlow and Wallace chronicle how the players in massive online worlds have begun to form groups and guilds that often confound the designers of the game and are at times in conflict with those that make the game.

Their contention is that designers rarely realize that they are creating a space where people intended to live large portions of their lives and engage in real economic and social activity and thus the designers have the moral duties somewhat equivalent to those who may write a political constitution (Ludlow and Wallace 2007). According to Purcell (2008), there is little commitment to democracy or egalitarianism in online games and this needs to change if more and more of us are going to spend time living in these virtual worlds.

2.1.2 The Lure of the Virtual Game Worlds

A persistent concern about the use of computers and especially computer games is that this could result in anti-social behavior and isolation. Yet studies might not support these hypotheses (Gibba, et al. 1983). With the advent of massively multiplayer games as well as video games designed for families the social isolation hypothesis is even harder to believe. These games do, however, raise gender equality issues.

James Ivory used online reviews of games to complete a study that shows that male characters outnumber female characters in games and those female images that are in games tend to be overly sexualized (Ivory 2006). Soukup (2007) suggests that gameplay in these virtual worlds is most often based on gameplay that is oriented to masculine styles of play thus potentially alienating women players. And those women that do participate in game play at the highest level play roles in gaming culture that are very different from those the largely heterosexual white male gamers, often leveraging their sexuality to gain acceptance (Taylor et al. 2009).

Additionally, Joan M. McMahon and Ronnie Cohen have studied how gender plays a role in the making of ethical decisions in the virtual online world, with women more likely to judge a questionable act as unethical then men (2009). Marcus Johansson suggests that we may be able to mitigate virtual immorality by punishing virtual crimes with virtual penalties in order to foster more ethical virtual communities (Johansson 2009).

The media has raised moral concerns about the way that childhood has been altered by the use of information technology (see for example Jones 2011). Many applications are now designed specifically for toddlers encouraging them to interact with computers from as early an age as possible. Since children may be susceptible to media manipulation such as advertising we have to ask if this practice is morally acceptable or not.

Depending on the particular application being used, it may encourage solitary play that may lead to isolation but others are more engaging with both the parents and the children playing (Siraj-Blatchford 2010). It should also be noted that pediatricians have advised that there are no known benefits to early media use amongst young children but there potential risks (Christakis 2009). Studies have shown that from 1998 to 2008, sedentary lifestyles amongst children in England have resulted in the first measured decline in strength since World War Two (Cohen et al. 2011).

It is not clear if this decline is directly attributable to information technology use but it may be a contributing factor.

2.3 Malware, Spyware and Informational Warfare

Malware and computer virus threats are growing at an astonishing rate. Security industry professionals report that while certain types of malware attacks such as spam are falling out of fashion, newer types of attacks focused on mobile computing devices and the hacking of cloud computing infrastructure are on the rise outstripping any small relief seen in the slowing down of older forms of attack (Cisco Systems 2011; Kaspersky Lab 2011). What is clear is that this type of activity will be with us for the foreseeable future. In addition to the largely criminal activity of malware production, we must also consider the related but more morally ambiguous activities of hacking, hacktivism, commercial spyware, and informational warfare. Each of these topics has its own suite of subtle moral ambiguities. We will now explore some of them here.

While there may be wide agreement that the conscious spreading of malware is of questionable morality there is an interesting question as to the morality of malware protection and anti-virus software. With the rise in malicious software there has been a corresponding growth in the security industry which is now a multi-billion dollar market.

Even with all the money spent on security software there seems to be no slowdown in virus production, in fact quite the opposite has occurred. This raises an interesting business ethics concern, what value are customers receiving for their money from the security industry? The massive proliferation of malware has been shown to be largely beyond the ability of anti-virus software to completely mitigate. There is an important lag in the time between when a new piece of malware is detected by the security community and the eventual release of the security patch and malware removal tools.

The anti-virus modus operandi of receiving a sample, analyzing the sample, adding detection for the sample, performing quality assurance, creating an update, and finally sending the update to their users leaves a huge window of opportunity for the adversary … even assuming that anti-virus users update regularly. (Aycock and Sullins 2010)

This lag is constantly exploited by malware producers and in this model there is an everpresent security hole that is impossible to fill. Thus it is important that security professionals do not overstate their ability to protect systems, by the time a new malicious program is discovered and patched, it has already done significant damage and there is currently no way to stop this (Aycock and Sullins 2010).

In the past most malware creation was motivated by hobbyists and amateurs, but this has changed and now much of this activity is criminal in nature (Cisco Systems 2011; Kaspersky Lab 2011). Aycock and Sullins (2010) argue that relying on a strong defense is not enough and the situation requires a counteroffensive reply as well and they propose an ethically motivated malware research and creation program.

This is not an entirely new idea and it was originally suggested by the Computer Scientist George Ledin in his editorial for the *Communications of the ACM*, "Not Teaching Viruses and Worms is Harmful" (2005). This idea does run counter to the majority opinion regarding the ethics of learning and deploying malware. Most computer scientists and researchers in information ethics agree that all malware is unethical (Edgar 2003; Himma 2007a; Neumann 2004; Spafford 1992; Spinello 2001). According to Aycock and Sullins, these worries can be mitigated by open research into understanding how malware is created in order to better fight this threat (2010).

When malware and spyware is created by state actors, we enter the world of informational warfare and a new set of moral concerns. Every developed country in the world experiences daily cyber-attacks, with the major target being the United States that experiences a purported 1.8 billion attacks a month (Lovely 2010).

The majority of these attacks seem to be just probing for weaknesses but they can devastate a countries internet such as the cyber attacks on Estonia in 2007 and those in Georgia which occured in 2008. While the Estonian and Georgian attacks were largely designed to obfuscate communication within the target countries more recently informational warfare has been used to facilitate remote sabotage. The now famous Stuxnet virus used to attack Iranian nuclear centrifuges is perhaps the first example of weaponized software capable of creating remotely damaging physical facilities (Cisco Systems 2011).

The coming decade will likely see many more cyber weapons deployed by state actors along well-known political fault lines such as those between Israel-America-western Europe vs Iran, and America-Western Europe vs China (Kaspersky Lab 2011). The moral challenge here is to determine when these attacks are considered a severe enough challenge to the sovereignty of a nation to justify military reactions and to react in a justified and ethical manner to them (Arquilla 2010; Denning 2008, Kaspersky Lab 2011).

The primary moral challenge of informational warfare is determining how to use weaponized information technologies in a way that honors our commitments to just and legal warfare. Since warfare is already a morally questionable endeavor it would be preferable if information technologies could be leveraged to lessen violent combat. For instance, one might argue that the Stuxnet virus did damage that in generations before might have been accomplished by an air raid incurring significant civilian casualties—and that so far there have been no reported human casualties resulting from Stuxnet.

The malware known as "Flame" seems to be designed to aid in espionage and one might argue that more accurate information given to decision makers during wartime should help them make better decisions on the battlefield. On the other hand, these new informational warfare capabilities might allow states to engage in continual low level conflict eschewing efforts for peacemaking which might require political compromise.

2.4 Future Concerns

As was mentioned in the introduction above, information technologies are in a constant state of change and innovation. The internet technologies that have brought about so much social change were scarcely imaginable just decades before they appeared. Even though we may not be able to foresee all possible future information technologies, it is important to try to imagine the changes we are likely to see in emerging technologies. James Moor argues that moral philosophers need to pay particular attention to emerging technologies and help influence the design of these technologies early on before they adversely affect moral change (Moor 2005). Some potential technological concerns now follow.

2.4.1 Acceleration of Change

An information technology has an interesting growth pattern that has been observed since the founding of the industry. Intel engineer Gordon E. Moore noticed that the number of components that could be installed on an integrated circuit doubled every year for a minimal economic cost and he thought it might continue that way for another decade or so from the time he noticed it in 1965 (Moore 1965). History has shown his predictions were rather conservative.

This doubling of speed and capabilities along with a halving of cost has proven to continue every 18 or so months since 1965 and shows little evidence of stopping. And this phenomenon is not limited to computer chips but is also present in all information technologies. The potential power of this accelerating change has captured the imagination of the noted inventor Ray Kurzweil who has famously predicted that if this doubling of capabilities continues and more and more technologies become information technologies, then there

will come a point in time where the change from one generation of information technology to the next will become so massive that it will change everything about what it means to be human, and at this moment which he calls "the Singularity" our technology will allow us to become a new post human

species (2006). If this is correct, there could be no more profound change to our moral values. There has been some support for this thesis from the technology community with institutes such as the Singularity Institute, the Acceleration Studies Foundation, Future of Humanity Institute, and H+.[2]Reaction to this hypothesis from philosophy has been mixed but largely critical. For example Mary Midgley (1992) argues that the belief that science and technology will bring us immortality and bodily transcendence is based on pseudoscientific beliefs and a deep fear of death.

In a similar vein Sullins (2000) argues that there is a quasi-religious aspect to the acceptance of transhumanism and the acceptance of the transhumanist hypothesis influences the values embedded in computer technologies that are dismissive or hostile to the human body. While many ethical systems place a primary moral value on preserving and protecting the natural, transhumanists do not see any value in defining what is natural and what is not and consider arguments to preserve some perceived natural state of the human body as an unthinking obstacle to progress.

Not all philosophers are critical of transhumanism, as an example Nick Bostrom (2008) of the Future of Humanity Institute at Oxford University argues that putting aside the feasibility argument, we must conclude that there are forms of posthumanism that would lead to long and worthwhile lives and that it would be overall a very good thing for humans to become posthuman if it is at all possible.

2.4.2 Artificial Intelligence and Artificial Life

Artificial Intelligence (AI) refers to the many longstanding research projects directed at building information technologies that exhibit some or all aspects of human level intelligence and problem solving. Artificial Life (ALife) is a project that is not as old as AI and is focused on developing information technologies and or synthetic biological technologies that exhibit life functions typically found only in biological entities.

A more complete description of logic and AI can be found in the entry on logic and artificial intelligence. ALife essentially sees biology as a kind of naturally occurring information technology that may be reverse engineered and synthesized in other kinds of technologies. Both AI and ALife are vast research projects that defy simple explanation. Instead the focus here is on the moral values that these technologies impact and the way some of these technologies are programmed to affect emotion and moral concern.

2.4.2.1 Artificial Intelligence

Alan Turing is credited with defining the research project that would come to be known as artificial Intelligence in his seminal 1950 paper "Computing Machinery and Intelligence." He described the "imitation game," where a computer attempts to fool a human interlocutor that it is not a computer but another human (Turing 1948, 1950). In 1950, he made the now famous claim that I believe that in about fifty years' time….one will be able to speak of machines thinking without expecting to be contradicted.

A description of the test and its implications to philosophy outside of moral values can be found here. Turing's prediction may have been overly ambitious and in fact some have argued that we are nowhere near the completion of Turing's dream. For example, Luciano Floridi (2011a) argues that while AI has been very successful as a means of augmenting our own intelligence, but as a branch of cognitive science interested in intelligence production, AI has been a dismal disappointment.

For argument's sake, assume Turing is correct even if he is off in his estimation of when AI will succeed in creating a machine that can converse with you. Yale professor David Gelernter worries that that there would be certain uncomfortable moral issues raised. "You would have no grounds for treating it as a being toward which you have moral duties rather than as a tool to be used as you like" (Gelernter 2007). Gelernter suggests that consciousness is a requirement for moral agency and that we may treat anything without it in any way that we want without moral regard.

Sullins (2006) counters this argument by noting that consciousness is not required for moral agency. For instance, nonhuman animals and the other living and nonliving things in our environment must be accorded certain moral rights, and indeed, any Turing capable AI would also have moral duties as well as rights, regardless of its status as a conscious being (Sullins 2006).

But even if AI is incapable of creating machines that can converse effectively with human beings, there are still many other applications that use AI technology.

Many of the information technologies we discussed above such as, search, computer games, data mining, malware filtering, robotics, etc. all utilize AI programming techniques. Thus it may be premature to dismiss progress in the realm of AI.

2.4.2.2 Artificial Life

Artificial Life (ALife) is an outgrowth of AI and refers to the use of information technology to simulate or synthesize life functions. The problem of defining life has been an interest in philosophy since its founding. If scientists and technologists were to succeed in discovering the necessary and sufficient conditions for life and then successfully synthesize it in a machine or through synthetic biology, then we would be treading on territory that has significant moral impact.

Mark Bedau has been tracing the philosophical implications of ALife for some time now and argues that there are two distinct forms of ALife and each would thus have different moral effects if and when we succeed in realizing these separate research agendas (Bedau 2004; Bedau and Parke 2009).

One form of ALife is completely computational and is in fact the earliest form of ALife studied. ALife is inspired by the work of the mathematician John von Neumann on self-replicating cellular automata, which von Neumann believed would lead to a computational understanding of biology and the life

sciences (1966). The computer scientist Christopher Langton simplified von Neumann's model greatly and produced a simple cellular automata called "Loops" in the early eighties and helped get the field off the ground by organizing the first few conferences on Artificial Life (1989).

Artificial Life programs are quite different from AI programs. Where AI is intent on creating or enhancing intelligence, ALife is content with very simple minded programs that display life functions rather than intelligence. The primary moral concern here is that these programs are designed to self-reproduce and in that way resemble computer viruses and indeed successful ALife programs could become as malware vectors. The second form of ALife is much more morally charged. This form of ALife is based on manipulating actual biological and biochemical processes in such a way as to produce novel life forms not seen in nature.

Scientists at the J. Craig Venter institute were able to synthesize an artificial bacterium called JCVI-syn1.0 in May of 2010. While media paid attention to this breakthrough, they tended to focus on the potential ethical and social impacts of the creation of artificial bacteria. Craig Venter himself launched a public relations campaign trying to steer the conversation about issues relating to creating life. This first episode in the synthesis of life gives us a taste of the excitement and controversy that will be generated when more viable and robust artificial protocells are synthesized.

The ethical concerns raised by Wet ALife, as this kind of research is called, are more properly the jurisdiction of bioethics. But it does have some concern for us here in that Wet ALife is part of the process of turning theories from the life sciences into information technologies. This will tend to blur the boundaries between bioethics and information ethics. Just as software ALife might lead to dangerous malware, so too might Wet ALife lead to dangerous bacteria or other disease agents.

Critics suggest that there are strong moral arguments against pursuing this technology and that we should apply the precautionary principle here which states that if there is any chance at a technology causing catastrophic harm, and there is no scientific consensus suggesting that the harm will not occur, then those who wish to develop that technology or pursue that research must prove it to be harmless first.

Mark Bedau and Mark Traint argue against a too strong adherence to the precautionary principle by suggesting that instead we should opt for moral courage in pursuing such an important step in human understanding of life (2009). They appeal to the Aristotelian notion of courage, not a headlong and foolhardy rush into the unknown, but a resolute and careful step forward into the possibilities offered by this research.

2.4.3 Robotics and Moral Values

Information technologies have not been content to remain confined to virtual worlds and software implementations. These technologies are also interacting directly with us through robotics applications. Robotics is an emerging technology but it has already produced a number of applications that have important moral implications.

Technologies such as military robotics, medical robotics, personal robotics and the world of sex robots are just some of the already existent uses of robotics that impact on and express our moral commitments (see Capurro and Nagenborg 2009; Lin et al. 2011).

There have already been a number of valuable contributions to the growing field of robotic ethics (roboethics). For example, in Wallach and Allen's book *Moral Machines: Teaching Robots Right from Wrong* (2010), the authors present ideas for the design and programming of machines that can functionally reason on moral questions as well as examples from the field of robotics where engineers are trying to create machines that can behave in a morally defensible way.

The introduction of semi and fully autonomous machines into public life will not be simple. Towards this end, Wallach (2011) has also contributed to the discussion on the role of philosophy in helping to design public policy on the use and regulation of robotics.

Military robotics has proven to be one of the most ethically charged robotics applications. Today these machines are largely remotely operated (telerobots) or semi-autonomous, but over time these machines are likely to become more and more autonomous due to the necessities of modern warfare (Singer 2009). In the first decade of war in the 21st century robotic weaponry has been involved in numerous killings of both soldiers and noncombatants, and this fact alone is of deep moral concern.

Gerhard Dabringer has conducted numerous interviews with ethicists and technologists regarding the implications of automated warfare (Dabringer 2010). Many ethicists are cautious in their acceptance of automated warfare with the provision that the technology is used to enhance just warfare practices (see Lin et al. 2008; Sullins 2009b) but others have been highly skeptical of the prospects of a just autonomous war due to issues like the risk to civilians (Asaro 2008; Sharkey 2011).

3. Information Technologies of Morality

A key development in realm of information technologies is that they are not only the object of moral deliberations but they are also beginning to be used as a tool in moral deliberation itself. Since artificial intelligence technologies and applications are a kind of automated problem solvers, and moral deliberations are a kind of problem, it was only a matter of time before automated moral reasoning technologies would emerge.

This is still only an emerging technology but it has a number of very interesting moral implications which will be outlined below. The coming decades are likely to see a number of advances in this area and ethicists need to pay close attention to these developments as they happen. Susan and Michael Anderson have collected a number of articles regarding this topic in their book, *Machine Ethics* (2011), and Rocci Luppicini has a section of his anthology devoted to this topic in the *Handbook of Research on Technoethics* (2009).

3.1 Information Technology as a Model for Moral Discovery

Patrick Grim has been a longtime proponent of the idea that philosophy should utilize information technologies to automate and illustrate philosophical thought experiments (Grim et al. 1998; Grim 2004). Peter Danielson (1998) has also written extensively on this subject beginning with his book *Modeling Rationality, Morality, and Evolution* with much of the early research in the computational theory of morality centered on using computer models to elucidate the emergence of cooperation between simple software AI or ALife agents (Sullins 2005).

Luciano Floridi and J. W. Sanders argue that information as it is used in the theory of computation can serve as a powerful idea that can help resolve some of the famous moral conundrums in philosophy such as the nature of evil (1999, 2001). The propose that along with moral evil and natural evil, both concepts familiar to philosophy; we add a third concept they call artificial evil (2001).

Floridi and Sanders contend that if we do this then we can see that the actions of artificial agents …to be morally good or evil can be determined even in the absence of biologically sentient participants and thus allows artificial agents not only to perpetrate evil (and for that matter good) but conversely to 'receive'or 'suffer from' it. (Floridi and Sanders 2001)

Evil can then be equated with something like information dissolution, where the irretrievable loss of information is bad and the preservation of information is good (Floridi and Sanders 2001). This idea can move us closer to a way of measuring the moral impacts of any given action in an information environment.

3.2 Information Technology as a Moral System

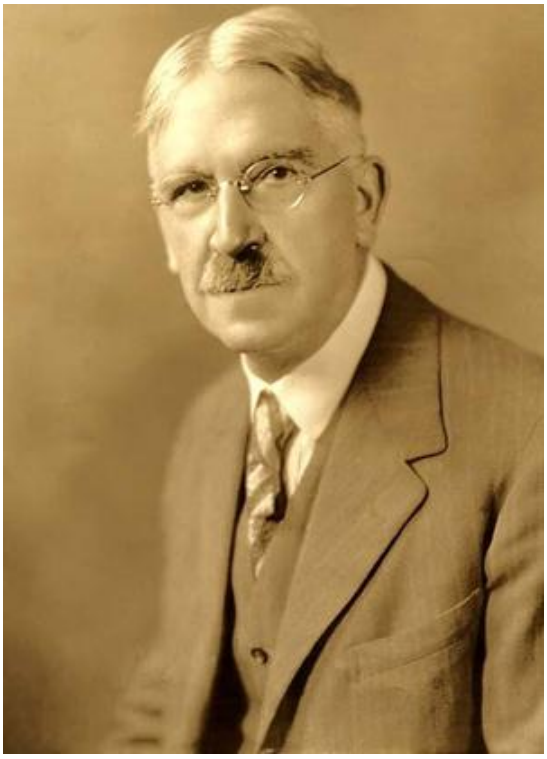

Early in the twentieth century the American philosopher John Dewey proposed a theory of inquiry

based on the instrumental uses of technology. Dewey had an expansive definition of technology which included not only common tools and machines but information systems such as logic, laws and even language as well (Hickman 1990). Dewey argued that we are in a 'transactional' relationship with all of these technologies within which we discover and construct our world (Hickman 1990). T

his is a helpful standpoint to take as it allows us to advance the idea that an information technology of morality and ethics is not impossible. As well as allowing us to take seriously the idea that the relations and transactions between human agents and those that exist between humans and their artifacts have important ontological similarities.

While Dewey could only dimly perceive the coming revolutions in information technologies, his theory is useful to us still because he proposed that ethics was not only a theory but a practice and solving problems in ethics is like solving problems in algebra (Hickman 1990). If he is right, then an interesting possibility arises, namely the possibility that ethics and morality are computable problems and therefore it should be possible to create an information technology that can embody moral systems of thought.

In 1974 the philosopher Mario Bunge proposed that we take the notion of a 'technoethics' seriously arguing that moral philosophers should emulate the way engineers approach a problem. Engineers do not argue in terms of reasoning by categorical imperatives but instead they use:

… the forms If *A* produces *B*, and you value *B*, chose to do *A*, and If *A*produces *B* and *C* produces *D*, and you prefer *B* to *D*, choose *A* rather than *C*. In short, the rules he comes up with are based on fact and value, I submit that this is the way moral rules ought to be fashioned, namely as rules of conduct deriving from scientific statements and value judgments. In short ethics could be conceived as a branch of technology. (Bunge 1977, 103)

Taking this view seriously implies that the very act of building information technologies is also the act of creating specific moral systems within which human and artificial agents will, at least occasionally, interact through moral transactions. Information technologists may therefore be in the business of creating moral systems whether they know it or not and whether or not they want that responsibility.

3.4 Informational Organisms as Moral Agents

The most comprehensive literature that argues in favor of the prospect of using information technology to create artificial moral agents is that of Luciano Floridi, and Floridi with Jeff W. Sanders. Floridi (1999) recognizes that issues raised by the ethical impacts of information technologies strain our traditional moral theories. To relieve this friction he argues that what is needed is a broader philosophy of information (2002).

After making this move, Floridi (2003) claims that information is a legitimate environment of its own and that has its own intrinsic value that is in some ways similar to the natural environment and in other ways radically foreign but either way the result is that information is on its own a thing that is worthy of ethical concern. Floridi (2003) uses these ideas to create a theoretical model of moral action using the logic of object oriented programming.

His model has seven components; 1) the moral agent a, 2) the moral patient *p* (or more appropriately, reagent), 3) the interactions of these agents, 4) the agent's frame of information, 5) the factual information available to the agent concerning the situation that agent is attempting to navigate, 6) the environment the interaction is occurring in, and 7) the situation in which the interaction occurs (Floridi 2003, 3). Note that there is no assumption of the ontology of the agents concerned in the moral relationship modeled (Sullins 2009a)

There is additional literature which critiques and expands the idea of automated moral reasoning (Adam 2008; Anderson and Anderson 2011; Johnson and Powers 2008; Schmidt 2007; Wallach and Allen 2010).

While scholars recognize that we are still some time from creating information technology that would be unequivocally recognized as an artificial moral agent, there are strong theoretical arguments in favor of the eventual possibility and therefore they are an appropriate concern for those interested in the moral impacts of information technologies.

Stamford University
Internet, BJA