Piercing the Future Fog

Four Perceived Robotic Futures

A Major Qualifying Project Report:

Submitted to the faculty

of the

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Bachelor of Science

by

____

Michael Brauckmann

Date: 11, March 2013

____

Professor John M, Wilkes, Advisor

Abstract

It cannot be denied that emerging technologies can rapidly transform the way we do business, fight wars, and live our lives. The social implications of these technologies and the changes they bring about are far reaching and rarely predicted in advance of the technological development. Hence, a better means of forecasting socio-technical implications is needed. Some experts suggest that a technology could be so transformative as to create a singularity i.e. a point beyond which all the rules change and forecasting is impossible.

In the book *Wired for War* P.W. Singer suggests that robotics is such a transformative technology and we are already headed toward a singularity. Inspired by the challenge of seeing through the fog of a period where all the major trends change, I designed a modified Delphi study to investigate the perceptions of about 200 students pursuing careers in robotics and other technical and non-technical fields at four universities. These students represent the voices of three professional groups expected to take part in the intensifying robotics debate. The respondents assess the likelihood, desirability, and ethical implications of four possible scenarios for the development of robotic technology.

It is beyond the scope of this study to address the question of whether the future of robotics is in principle predictable. However, it is not difficult to answer simpler questions with methodological implications for the field of technology assessment. How much consensus exists among students in different fields about where the technology is going? If there is strong consensus on some direction, how desirable is that future perceived to be? A consensus in perceptions may mean the technology is not entirely unpredictable. Common expectations can even become a self-fulfilling prophesy. Such a consensus poses a significant sociopolitical issue whether or not it proves to be accurate if it results in an attempt to control the direction of the field or resist undesirable outcomes.

Another important question is how important the goals and mindset of the institution driving the technological development are perceived to be in determining the socio-technical consequences of the technology. Each of the four scenarios posits a different institution driving the development of robotics and each takes a different stance on Isaac Asimov's three Laws of Robotics, and differs in the number of them it violates. Using these responses we investigate the concerns of students in different fields about the perceived direction of robotics technology.

Table of Contents

Abstract i

Introduction 1

Literature Review 4

Methodology 12

Developing the Scenarios 14

Developing the Survey 15

The Sample 24

Results 28

Likelihood 28

Desirability 30

Ethical Concerns 33

Worcester State and Boston University Study 36

Discussion 46

Conclusion 49

Bibliography 52

Appendix A – Scenarios 53

Appendix B – Survey 57

Introduction

In July 2002, a Pentium III microprocessor powers up, bringing an iRobot PackBot Scout to life (NASA, 2011). This robot is one of the first ground-based robots ever used in a war-zone. Four years later, in 2006, more than five thousand robots would be deployed in Iraq and Afghanistan (Singer, 2010). The capability of this technology to rapidly transform the battlefield is apparent. But what is the result of this transformation? Will robots prove to be an unstoppable force that will end human warfare? Or will robots be a new weapon which nations will race to develop and with which even more wars will be fought? Will robots serve as discriminating protectors; reducing casualty's rates and avoiding collateral damage on a still human battle field? Or will they lead to wars of ruthless carnage carried out by machines with no empathy or pity to constrain them? In any case the action has moved out of cyber space to the real world of objects and lives. Most observers expect the sociopolitical implications to be far reaching. Singer documents what he calls an incipient "revolution in military affairs". Singer actually goes further than this, suggesting that this technological revolution could be the one that changes everything, starting with the man-machine relationship. He notes that in the field of robotics there is talk of a coming "singularity" beyond which the future becomes so turbulent and the possibilities so numerous that it is in principle unpredictable. As thought provoking as this claim is to a roboticist, it is an even greater challenge to the social scientists working on the problem of technological forecasting.

It cannot be denied that emerging technologies can rapidly transform the way we do business, fight wars, and live our lives, and one can hardly control technology if its' most likely path and social implications cannot be predicted. This is not the first revolution in military affairs. One can see the changes in prevailing trends caused by previous RMAs and the extent to which these changes made their implications unpredictable. Certainly the introduction of gunpowder, which led to the development of firearms and cannons rendering a knight's heavy armor and city walls useless, qualifies as one RMA. Another example is aviation. Just 40 years after the Wright Brother's first flight the skies above WWII battle fields filled with sophisticated propeller powered fighters and bombers while the first jet fighters were being developed. Most military plans centered on a strategy for air superiority and aircraft carriers became the most important ships in the Pacific theater. In the last half century atomic bombs and missile technology caused another RMA. In retrospect, the socio-political implications of these developments are pretty well understood and many were predictable, some features of the ensuing arms race were actually predicted at the outset by a group working on the Manhattan Project. (Feuer)

Are we witnessing something on this scale with robotics applied to military operations? It may not be an insuperable problem to predict which direction this institution is likely to take the field. The ability of military institutions in the USA to control technology (even those it considers to be sensitive and secret) is subject to question, but not the likely technological direction it will try to foster. A harder question is what other institutions will try to do with the same technological capabilities once they are developed, assuming that they become available in a market place of ideas and devices.

It is beyond the scope of this study to address the question of whether the future of robotics is in principle predictable. However, it is not difficult to answer simpler questions with methodological implications in the field of technology assessment. How much consensus exists among students in different fields about where the technology is going? If there is strong consensus on some direction, how desirable is that future perceived to be? In short, is there a sufficient consensus to attempt to control the direction in which the field is headed, and specifically to resist some of the outcomes currently perceived to be likely? My partners (doing an IQP) and I (doing an MQP)also intend to ask the question of how important the educated public thinks the beliefs, mindsets and goals of the institution controlling the development of a technology tend to be in determining the socio-technical consequences of an emerging technology. However, my focus is on the methodological issues in technology assessment and their focus on the development of the instrument, field testing it, assessing the degree of consensus and looking at the question of generalizability of findings.

These questions will be investigated through a study of perceived robotic futures. The study will be built around gathering data on the reactions of respondents to a series of four alternative futures, developed into scenarios of about a page in length. In each scenario, a different institution drives the development of the field of robotics. Respondents will be asked to answer a questionnaire on likely hood, desirability and social concerns for each scenario.

Research in the forecasting the impacts of the telephone indicates that non-technical professionals in affected fields have made more accurate predictions of a technologies impact than the engineers who developed the technology. (Pool) Additionally it is clear that any attempt to control the direction of the technology would involve a public debate and many social interest groups with political voices. Hence, the respondents are aspiring experts in the field of robotics, as well other technologists and non technical majors. The robotics majors represent the experts and leaders in the field driving the technological development. Those students majoring in other technical fields may contribute to the field or even hold management positions in the companies developing robotics technologies. Their voice represents that of the roboticists' peers in the technological community. The students in the non-technical sample are drawn from Clark University and represent public opinion leaders, the media, and political activists with expertise in journalism, law, medicine and business. They are also more likely to become teachers and political figures than WPI graduates. A consensus in the perceptions of so broad an array of aspiring professional may mean the technology is not entirely unpredictable. Common expectations can even become a self-fulfilling prophesy. Such a consensus poses a significant sociopolitical issue whether or not it proves to be accurate as it captures the voices in the coming robotics debate. This may provide a better methodology for forecasting socio-technical implications. If the expected outcome of the technology can be determined then an attempt to control the direction of the field or resist undesirable outcomes is moe likely to be undertaken.

Literature Review

A robot is a machine built upon the "sense-think-act" paradigm—that is, they are man-made devices that sense their environment, process data, and respond based on what they've perceived (Singer, 67). The PackBots, which have been deployed in Afghanistan and Iraq, are far from the only robots out there. iRobot also makes the Roomba, small disk-shaped vacuum cleaner robot. Predator drones armed with missiles patrol foreign skies. Industrial robots tirelessly work on the production lines of factories across the globe. The field of robotics is developing extremely quickly.

In Wired for War P.W. Singer tells the story of this emerging technology and its impact on society. The vast majority of research in this field in the United States comes from military funding programs such as the defense advanced research project agency DARPA. According to Singer (2010), some 80% of what is spent in this country comes from defense dept. sources. Programs for developing a single robot frequently have budgets in excess of several million dollars. The first section of the book covers the current robotic technologies employed by US troops in Iraq and Afghanistan. From clearing improvised explosives and roadside bombs, to flying surveillance missions in Iraq, to taking out insurgents with Hellfire missiles; these early robotic warriors have paved the way for robotics in the military. Some of these robots are designed as scouts, made to go into places people don't want to. Others, Foster-Miller's SWORDS platform and the predator drone, are intended to hunt down and kill humans.

While the original PackBot and Talon platforms included robotic arms, Foster-Miller's SWORDS version of the Talon is a prototype designed to carry and fire weapons. Capable of carrying anything from an M-16 to a .50-caliber machine gun to a rocket launcher, the SWORDS robots are amazingly accurate (Singer 30). iRobot is also developing a shotgun-wielding version of their PackBot. Singer interviews the scientists and engineers developing these robots, as well as the soldiers who use them. Through these interviews the argument is made that these technologies are the building blocks to a much greater change in the way we fight wars. The possible developments that come from combining these technologies with things like communication networks and artificial intelligence sound like scenes from a movie, and indeed many of them draw their inspiration from science fiction.

In 1998, Vice Admiral Arthur Cebrowski predicted that the introduction of computers and near-instant communication would produce something he called "Network Centric Warfare." He predicted that this change would be a paradigm shift called a "revolution in military affairs," or RMA. "RMAs typically involve the introduction of a new technology or organization, which in turn creates a whole new model of fighting and winning wars. A new weapon is introduced that makes obsolete all the previous best weapons (Singer 2004)." Just as the introduction of guns made highly trained knights nearly worthless, Cebrowski predicted that near-instant communication would create a similar change in warfare. Unfortunately, network-centric warfare introduced a sort of information overload, proving Cebrowski wrong. Singer predicts that robotics will be the technology that revolutionizes military affairs, "perhaps even leading to the rise and fall of global powers (Singer 204)." The combination of what Cebrowski envisioned with robotics incorporating artificial intelligence, and thus increasing autonomy, is what is truly far reaching and might really be an RMA.

Singer however looks beyond the RMA. He sees robotics causing a cascade of interdependent and complicated changes in society more generally. The social implications of these technologies and the changes they bring about are far reaching, and unforeseen effects may be even greater than the predicted outcomes. Singer cites Futurist Ray Kurzweil, whose company focuses on predicting trends in technology to “catch the train at the right moment.” Kurzweil believes that we are on the verge of such technological breakthroughs that they will change all the rules in an event he calls the Singularity (Singer 2004). Singer, and Kurzweil, are not alone in their belief that robotics is bringing about the Singularity which will turn the system we know upside down. Bill Joy is the cofounder of Sun Microsystems and author of a short article entitled “Why the Future Does Not Need Us”, in which he explains why he is uneasy about the danger we face in the 21 century (Joy). Joy’s anxiety started when he read a preprint of Kurzweil’s book The Age of Spiritual Machines, a story of a utopian future where man becomes one with robotics gaining near immortality. But Joy did not see this as a likely path of the technology Kurzweil described; instead he saw a future in which mankind all but obsoletes itself.

Joy urges us to consider the consequence of allowing more and more decisions to be made for us by machines. He warns that no hostile takeover or willing surrendering of control will be needed. The technical system will simply become more and more complex until no human will be able to make intelligent decisions and we will become so dependent on the machines that flipping the power switch would be tantamount to suicide. (Joy) Joy compares robotics along with genetic engineering and nanotechnology to Pandora ’s Box and warns that we have nearly opened it, and what comes out will never be put back in a box. In his words:

> “We are being propelled into this new century with no plan, no control, no brakes. Have we already gone too far down the path to alter course? I don't believe so, but we aren't trying yet, and the last chance to assert control - the fail-safe point - is rapidly approaching” (Joy).

Another author Kevin Kelly writes in his book What Technology Wants that something entirely new has emerged which he calls the technium. He views technology analogous to a biological organism evolving as much by internal processes as by human choice. He claims it is “whispering to itself” becoming increasingly autonomous, has “wants” and urges and a direction in which it wants to go. Kelly claims this technium has become “as great a force in our world as nature” and it would be unreasonable to expect it to obey us. Rather than even attempt to control it he guides us to learn what it wants, and where it will go, to listen to it, and decide how to “optimize technologies blessings while minimizing the costs”. The increasing trend toward autonomy is evident in his work and commented upon extensively. Von Neumann, the inventor of the first useful computer, whose architecture is still prominent in many microprocessors, noted that technology was a process of increasing “structure, organization information, and control.” Kelly called it “a vital force that throws us forward or pushes against us.”

This study was inspired by Singer’s book Wired for War and his concerns and warning about the current trends in the field of robotics. One of the four scenarios is drawn loosely from his description but projected ahead to 2050. When confronted with Ideas like those expressed by Joy and Kelly, Singers warning may even come across as a moderate voice. The shape of the future lies in the balance of the policies, social changes, and socio-economic decisions made in the present. We can consider the lessons of the scientists working on the atomic bomb

> “The danger that things will move to fast, and in a way in which the process can take on a life of its own. We can as they did create insurmountable problems in no time flat. We must do more thinking up front if we are not to by similarly surprised and shocked by the consequences of our inventions.” (Singer 2004)

Even Bill Joy in his pessimistic view of the future and near certainty that we were creating a dystopian future believed that this was the moment to take a stand.

> “Have we already gone too far down the path to alter course? I don't believe so, but we aren't trying yet, and the last chance to assert control - the fail-safe point - is rapidly approaching… If we could agree, as a species, what we wanted, where we were headed, and why, then we would make our future much less dangerous - then we might understand what we can and should relinquish. I believe that we all wish our course could be determined by our collective values, ethics, and morals. If we had gained more collective wisdom over the past few thousand years, then a dialogue to this end would be more practical, and the incredible powers we are about to unleash would not be nearly so troubling.” (Joy 9)

It may not be possible to predict the course of technology and, even if it is possible, controlling what that direction will be might still prove to be an insurmountable problem. However, if no attempt is made then we are certainly left to whims and urges of Kelly's technium whatever they may be. It is the author's opinion that it would be foolish not to make every effort to understand and direct the path of these technologies. If we make it our goal to understand where this technology is taking us, and shape our socio-technical policies so as to guide it in a favorable direction, that at least improves the odds that the future will be the result of our deliberations rather than technological inertia. "Wired for war" gives the reader an idea of the changes to come and raises many important questions about robotics and human nature itself that must be answered. Surely, then, it is a good idea to take a look at what futures are possible and ask how people perceive them.

To this end, we developed our four scenarios, each outlining a different possible future for the field of robotics. Each scenario varies in that the institution driving technology has different goals and ambitions which lead to a different path of development. Hence, responses will reveal the perceived effect of the institutional goals and mindset. Singer seems concerned that the US Military has ill-advisedly crossed an ethical line in the man-machine relationship, and will one day regret having done so when the USA is no longer the technology leader in their field. It's just a matter of time before the USA's current military capabilities are widely available to other nations and hostile political groups. According to Singer, it is possible to be short sighted and act in this way because the military avoids looking at the ethical implications of the technologies they work with. As Michael Goldblatt, DARPA's defense sciences office director, puts it "You can't let the fear of the future inhibit exploring the future." In the words of another DARPA program manager, "That [considering ethics] is above my pay grade." Hence, we thought to incorporate an ethical dimension into our scenarios.

The iRobot Corporation takes its name from Isaac Asimov's book *I, Robot*. Considering that iRobot is developing killer robots, this association is rather peculiar. Asimov was a science fiction writer and published a series of short stories known as *I, Robot* during the late 1940's. The book describes how, over the course of a lifetime, robotics begin as simple mechanics and develop into complex entities containing "positronic brains" somewhat more like the human brain than microcontrollers. In this alternate future, all robots follow the Laws of Robotics:

1. A robot may not injure a human being, or through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law

With these ethical laws in place, humanity thrives in the company of these intelligent machines. Robopsychologist Dr. Susan Calvin explains that strict adherence to these laws prevents robots from performing act or undertaking tasks that are immoral, dangerous, or generally undesirable.

iRobot's machines clearly violate all three of Asimov's laws. The military, in fact, "explicitly wants robots that can kill, won't take orders from just any human, and don't care about their own lives." So much for Laws One, Two, and Three (Singer 432)." The people at iRobot, however, believe that Asimov would approve of what they have created and "think it's cool as hell (Singer 25)." Actually it is doubtful that he would be swept away by technological optimism. In our scenarios each institution driving the development of robotics takes a different stance on robotic ethics. We chose to adopt Asimov's three laws as our basis of ethics thus every institution varies in its ability to accept Asimov's laws given its goals. By gathering people's perceptions of these scenarios, we hope to see if acceptance of Asimov's three laws reduces concern about ethical issues and stand as a guide for ethics in the field of robotics.

The design of this study is inspired by the Delphi study technique which traditionally includes a panel of experts in the field being assessed. Our study differs from this format in two basic ways, and there are precedents for these changes in the Technology Assessment and Public Understanding of Science and Technology literatures. 1) We have chosen to use a college student sample instead of experts, and 2) we chose to sample from aspiring experts both in and outside the field of robotics. It is the robotics majors who will stand in for our panel of experts. Previous research in the field of public perceptions of nuclear technology by Jon Miller Et al. showed that the views of college students could approximate the literate college educated portion of the US population called the "attentive public". Another study in Aerospace innovation showed that student opinions were a rough approximation for expert opinions. This study included WPI students, WPI graduates and NIAC experts, and found that the technological breakthrough that 80% of experts found most likely about half of the students also found most likely. The other 50% of students tended to be a random scattering of other responses- noise but not enough to obscure the signal which was in line with the overall pattern of agreement with the experts. Hence, the more attainable student sample of aspiring experts is an acceptable proxy for experts and to be preferred where a high response rate is important. High response rates make it possible to claim greater representativeness and generalizability of findings.

Another reason that the college student sample is appropriate is that the scenarios are set in the timeline of the careers of current college students so the technological developments discussed represent the contribution of the next generation to the field. This makes current students a more appealing sample than current experts in several respects. As for the students outside the field of robotics (which includes other technical majors and non-technical majors) research on the accuracy of predictions about the social impact of the telephone suggests that those affected by a technology may provide more accurate predictions than the engineers involved in the development of the technology. Similarly there will clearly be other voices, in the public debate over robotics and involved the process of making polices, which are not those of technical experts in the field.

Ellul describes a “Technological Mentality” which is employed by engineers and scientists. It is largely based on a efficiency criteria and involves a narrowly focused preference for objective criteria and short term implications. While this mentality has certainly aided in the development of increasingly useful and efficient technologies, it results in short term thinking and can lead people to miss side effects with unintended consequences. The non technological members of society tend to think more broadly and over the long term, and more easily consider the effect of a technology outside of the domain of application it was designed for. These predictions are based on hunches, experiences, and judgments about what new capabilities the technology might provide. This subjective thinking is rarely convincing to the efficiency oriented expert engineers. None the less these predictions often prove accurate once a technology is developed and its unintended effects begin to trickle into society.

For example, when the telephone and telecommunication equipment first became available most of the engineers working on it pictured a mass media communications system like today’s radio broadcasts. They claimed there would be phone lines fanning out from the opera halls so all the world could listen in. Obviously this was not the most substantial effect of the new technology, which was more suitable for point to point communication than broadcasts. They were distracted by the prior existence of radio being used for point to point applications such as ship to shore, and considered the niche filled. It was later that the logic of radio for broadcast and telephone for point to point asserted itself. Meanwhile, it was businessmen wanting to move production to cheaper land outside the city who saw the technology as a way to keep the office in the city center while moving the factory without loss of communication. Firemen and doctors, pharmacists could all see whay this capability would be important to them. Hoteliers predicted the development of the sky scrapers as troops of telegram delivery boys were rendered obsolete. . Clearly it makes sense for one attempting to assess the future of a technology (despite the possibility of a singularity) to consider the voices of those outside the field that might better understand how the technology is likely to be applied in their own professions and fields.

Methodology

Robotics technology is expected to enable rapid changes in the way we live our lives, and we may be heading toward a singularity which changes all the rules and trends making the future impossible to predict. Yet Singer, Joy and others warn we may not like where robotic technology is taking us. The social implications of these technologies and the changes they bring about are certainly likely to be far reaching, but to call them unpredictable is more than saying they are profound transformations. It implies that the future is so in spin that anything is equally possible. That seems unlikely.

This study was inspired by the methodological challenge of seeing through (or forecasting beyond) a period where all the major trends change to do a technology assessment of robotics technology. There are two parts to a technology assessment, the forecast and then an examination of the implications of things coming out as one has predicted. This is worth doing even if the forecast is not fulfilled, especially if the message of the analysis is a cautionary tale of a future one hopes can be avoided. This study will not address the question of whether the future of robotics is predictable. Instead two questions will be addressed. How much consensus exists among students in different fields about where the technology is going? If there is strong consensus on some direction, how desirable is that future perceived to be? These questions have methodological implications in the field of technology assessment. Such a consensus poses a significant sociopolitical issue whether or not it is accurate, especially if it results in an attempt to control the direction of the field.

Hence, a modified Delphi study was designed. This study investigates the perceptions of about 200 students pursuing careers in robotics and other technical and non-technical fields at four universities. As was the case with the telephone, the "experts" (in this case the robotics majors) may not have the most accurate predictions of the future of the field, thus these students represent the equally interesting and valid voices of three professional groups expected to take part in the coming round of the robotics and society debate. Another important question is how important the goals and mindset of the institution driving the technological development are perceived to be in determining the socio-technical consequences of the technology.

Four scenarios describing possible future in the development of robotics technology have been developed. Each one posits a different lead institution providing the bulk of the developmental funding for the field of robotics. Our questions about the perceived importance of institutional influence shaping the field are answered indirectly by examining changes in the perceptions of those likely to be affected in different ways. The underlying question to be addressed is whether people aspiring to enter the field of robotics and their peers aspiring to other technical and non-technical fields are equally concerned about where the field of robotics is headed, and if those concerns are mitigated by the values prominent in the mindset of the institution leading the field. This ethics question was embedded in each of the four scenarios. By this I mean that references were made directly or indirectly made to Asimov's laws in each case, but it was done in a parallel and smoothly flowing way that kept it from being obtrusive. Differing reactions to the scenarios imply that it really does matter what institution is playing the lead role. The respondents assess the likelihood, desirability, and ethical implications of four possible scenarios for the future of robotic technology. If there is a consensus that it greatly matters who develops this powerful technology, the stage is set for further investigation using social methods that get beyond perceptions data. The following sections detail the development of the scenarios, and the survey, and explain the methods for distributing and collecting the instrument, and tools used for analysis of the data.

Developing the Scenarios

In order to determine students' perceptions of robotics technology being developed under different institutions four scenarios were developed. Each scenario posits a different institution driving the development of robotics and each takes a different ethical stance on Isaac Asimov's 3 Laws of Robotics. Currently the vast majority of funding for research and development of robotics technology in the USA comes from the Department of Defense. In one scenario this trend was continued in the other three scenarios another institution replaced the military as the institution driving the development of the field. Each institution has a different goal for the technology; to explore and take advantage of lunar resources, to aid in meeting a major global food and environmental crisis, to take advantage of eldercare opportunities in the commercial sector, and to gain an advantage on the battlefield.

Each scenario was designed to expand the current state of robotics technology for approximately 50 years, and to picture similarly advanced robotics systems. This time frame was chosen so that the scenario would represent the contribution of the current generation of students to field at the end of their careers. Each scenario then represents a perceived future of robotics under the leadership of varying institutions trying to address different real world problems. Each scenario is designed to raise ethical questions about the direction of robotics technology and its social and technical implications. Differing views on these implications between scenarios will reveal the effect of the driving institution. Additionally, each scenario represents a different relationship to Asimov's three laws of robotics.

Although they come from works of science fiction, Asimov's Laws are the best known statement in the literature on ethics in robotics and keeping the technology under social control. We adopt Asimov's framework with care. Asimov wrote the laws before the first transistor was developed, his positronic brains and our microprocessors share next to nothing in common. As one roboticist put it, "People ask me about whether or not our robots follow Asimov's laws. There is a simple reason [that they don't]. I can't build Asimov's laws into them (Singer 432)." Furthermore, the entire premise of Asimov's short stories is that the three laws do not entirely prevent robots from behaving in undesirable ways.

We have been very careful in adopting his framework in that it is the corporations in control of the development of robotics that are the ones following the ethical code not the robots per se. . The institutions in control in each scenario vary in their willingness and ability to accept Asimov's laws, from complete acceptance in the lunar scenario to complete rejection in the military scenario with the others falling somewhere in between. At this point of our research there were concerns about the clarity, and readability of the scenarios as well as how long it would take respondents to read through all four of them. A pilot study was conducted in a single WPI class containing about 80% robotics majors in order to obtain initial responses to the scenarios. Feedback from this class allowed for critiques that were grounded in experience and set the stage for editorial adaptation of the stimulus and response items. Following this pilot study, the scenarios were also modified to avoid confusion and to shift attention to the social implications of the technology itself, downplaying the many feasibility concerns coming from the robotics majors about how such a thing might be implemented.

Developing the Survey

A questionnaire was attached to each scenario in order to collect data on the direction and strength of participant reactions to the scenarios. In the end, the hope was to produce a rank order from most to least likely and most to least desirable, though ties were possible. The same Indicator questions were used on each scenario to enhance comparability between scenarios and make such a rank ordering possible. The questionnaire consists of five variable indicator items: one designed to assess the likelihood question; two to address the desirability of the scenario in general and as a economic and technical stimulus; two more to get at the severity of ethical issues raised by the technology. One of these ethics items left open the nature of the ethical concerns that concerned the respondent and the other picked up on the man machine relationship specifically to tie into the extensive literature about technology becoming autonomous and getting out of control.

Each response is intended to reveal a different aspect of the participant's perceptions of a possible direction in which robotics could develop and gives one an idea of what they expect to see from the technology. The study is simplified by treating the scenarios as alternatives, though in fact they are not mutually exclusive and in fact are likely to co-exist and interact. The four scenarios do not represent the only possibilities for robotics and the respondents' actual best prediction of what will really happen is not directly assessed. Instead, this is a search for consensus on the direction of the technology and whether the social implications associated with most likely directions are reassuring or disquieting.

It was decided to keep the number of differing response categories to a minimum to avoid confusion and improve the appearance of the survey. Each question was worded such that it could be answered on either a likelihood or desirability scale. Four response categories were chosen so that there would be no middle ground. Hence, participants would be encouraged to think about the question enough to choose a side. The two response scales used on the questionnaire are as follows:

Unlikely	**Somewhat Unlikely**	**Somewhat Likely**	**Very Likely**
Undesirable	**Somewhat Undesirable**	**Somewhat Desirable**	**Very Desirable**

At this point a walkthrough of the five items in the order they were asked after each scenario is in order so that comments can be made about what variable the indicator is supposed to tap and what the logic was for addressing each key variable in this fashion.

How likely is it that this scenario could come about?

This question was used to support a comparison of the four scenarios to reveal which scenario's application area (space, the seas, personal service or warfare) was perceived as the most probable direction of application and hence have funds for technology development in the field. It was important to allow for ties, so a forced rank ordering item was avoided. It is only of passing interest what the majority of the whole stratified sample considers to be most likely as the study is designed to be internally comparative. Each of the three strata, in the sample will first be considered separately in this regard. This study is designed to reveal the level of consensus between our three sample strata (robotics majors, other technical majors and non-technical (liberal arts) majors). Thus it is primarily the level of agreement within and between these groupings is of interest. One wants to see If there is a significant consensus among these people with different academic backgrounds and literacy on the subject at hand. Then a comparison can be made with desirability to determine if the perceived most likely direction of the technology is also the most desirable.

If the scenario came about, would the resulting technology be likely to spin-off many applications that significantly advance the field of robotics?

This question was developed to determine the amount of influence the technology described in scenario would have in terms of stimulating robotics and possibly other related fields. High responses on this question are intended to indicate socio economic impact potential. However, on its face it also means that the participants see this as a promising direction of technology development that will spread outside the scope of the scenario. If a development is perceived as likely to spin-off and stimulate secondary effects on society and the economy it is especially interesting from the standpoint of the coming singularity argument. A “singularity” is a complex, interaction of explosive technological developments to the point that predicting where it is going and what effect it will have is likely to be impossible. While many spinoffs would not be enough by itself to support the notion of a coming singularity, as proposed in the literature and noted by Singer, many spin- off applications would be part of a singularity pattern. If a robotics advance is highly transferrable to other ends, it might usher in a dynamic and volatile period in which robotics technology could be involved in a technological revolution evocative of the singularity idea. So, perceived spinoff potential raises two questions of interest to this study. Is the technology particularly likely to get out of control and does it matter who funds the development of the technology in terms of provoking an upheaval one might call a singularity after which developments are unpredictable? One theoretical premise of this study is that it does matter which institution develops the technology and for what purpose. However, this is in principle an empirical question subject to testing. However, since the data being gathered here cannot directly address that question, it is for now a theoretical assumption, and will be tested only in the world of perceptions. I can only address the question of whether the sample believes that it matters which institution is in charge. I can also see if the respondents perceived the scenarios to have different likelihoods of generating spinoffs or not.

There are those who claim, with some justification, that technology will be applied to war whatever its initial area of development and application was. Vice Versa may also apply, ie. that military capabilities will soon be turned to other ends. For example, the internet was a DARPA project aimed at robust communications that could survive a nuclear war. Clearly that has not been its most significant application and it is increasingly considered a socially transformative communications medium. On the other hand, this could be an exceptional case. Most technology developed by the military is classified and subject to secrecy requirements that limit its spinoff potential, or at least delays it.

At this point we are not collecting data intended to (or in principle able to) resolve the questions of what the future will really be and whether it really matters what organizations fund and execute the initial development of a robotic capability. It is still interesting to find out if those affected believe it matters what the lead agent is and what their expectations about the future are. The perceptions of those in the field of robotics clearly matter and I would contend that the perceptions of their other technical and non-technical peers likely to be affected by these developments about whether "the who and why" of robotic development matters just as much. Note that we are asking only about the 50 year period which their careers will span- and they will be acting on these perceptions at least initially. If all four scenarios are considered to have massive and essentially equal potential for spinoff, the sample is saying that it does not matter who does what and why- robotic technology is intrinsically revolutionary and possibly uncontrollable- ie., they expect that the singularity is coming.

If the scenario came about, how desirable or undesirable would the resulting changes in the quality of life be?

The change in quality of life is used as a general and non-specific indicator of the effect the technology change would have on the society it is introduced into. It was important to get beyond narrow efficiency and economic implications of robotics and get into disruption and displacement issues, if they concerned the respondent. A broader than economics intent had to be clear, hence "quality of life" for people. A desirable effect on the quality of life indicates that the technology improves society in some way or at least alleviates the social issue it was designed to address. Undesirable responses indicate the technology may create worse problems than it solves, upset the balance in the system, displace workers or even get out of control. The key is a perception that it does not seem likely to solve problems, or that in solving one problem it might have unintended consequences that were negative side-effects and create even worse problems. By comparing these responses across scenarios and across the three groups in our sample one can determine if a consensus exists on the scenarios most likely to have desirable outcomes and compare them to the perceived most likely scenarios.

If the scenario came about, how desirable or undesirable would the resulting changes in the man machine relationship be?

Having two parts to the desirability question was an effort to separate out the major theme of dependency of people on machines and inversions in the man-machine control relationship from the many other questions raised by the movement of automation into a robotics phase (and the creation of artificial intelligence) that one could consider undesirable trends. Having two questions which could easily be combined into a composite item was a modest recognition that this was a multidimensional variable. Similar to the quality of life question, this question is intended to measure the social desirability of the scenario. Whether it is dependence on machines to meet some basic need or the formation of a caretaker relationship, the way in which machines interact with humans is inevitably changed by the kind of advancements in robotics technology under discussion. Questions of subordination and autonomy are bound to come up and thus impact the man-machine relationship that we are accustomed to seeing.

From a man -machine partnership to explore and mine the moon under lunar surface conditions hazardous to humans, to reshaping the ecology of the seas to feed humans, to directly putting vulnerable humans under robotic care, the stakes are rising. In the end, creating machines designed to hunt, ambush and kill humans raises the ultimate question of who is in control here especially if there seems to be a trend from human in the loop to increasing autonomy in these killer bots. But all along the way to this “terminator” extreme, the man-machine relationship is one thing you want to watch, and the control issue it raises is the focus of Asimov’s laws.

Whether the acceptability and rated desirability of the scenarios tracks with the degree to which the scenario violates these laws is one of the questions under study. Responses to this question will also be checked for consensus among robotics majors, Other technical majors, and non-technical majors. It is not clear that WPI and Clark University students will see things the same way, as they did not in the case of nuclear power during the late 1970's. This is a matter where trust and confidence in the technology and the institutions creating and managing it become increasingly important to public acceptance.

The perception of who was in charge and public confidence in that institution (be it "science", "government" or "private industry") greatly affected public attitudes toward nuclear power in the 1970's prior to the Three Mile Island (TMI) incident. At both WPI and Clark University there was high confidence in science as an institution, but only the WPI students perceived scientists to be in charge of the nuclear industry via the Nuclear Regulatory Agency. The Clark University students viewed the nuclear industry as a venture of the private sector, known for cost cutting in areas related to public safety. After the TMI (Three Mile Island) incident in 1979 and the Chernobyl accident in 1986 (Ukraine in the then USSR) the dynamics changed, in part due to the discrediting of all the organizations in charge of the technology. The nuclear establishment seemed not to have been worthy of public trust and the charges of institutional failure were now specific rather than possibilities derived by analogy.

In the case of the nuclear debate, the release of the film "The China Syndrome", shortly before the TMI incident, had already presented the possibility of a nuclear meltdown disaster due to corporate evasion of safety regulations during the construction of a nuclear power plant. At Clark University the TMI incident moved campus opinion from 60% anti-nuclear to 80% anti-nuclear. At WPI it went from 55% pro-nuclear to 75% pro-nuclear. Hence, there was an incident associated with the polarization of opinion about this technology in that case. We seem to be in the pre-polarization period of public attitudes toward robotics as there is not yet a famous incident to interpret as evidence of how safe the technology is and why. However, a debate about the legality of UAV attacks on civilians who may or may not be terrorists, as there has been no trial, is coming. It will probably be at the UN with some nations tryingto brand the US as guilty of war crimes in its use of this technology. This could be a polarizing event that makes robotic technology controversial.

The WPI response in the case of nuclear power after TMI may seem surprising, but it depends on how the facts were interpreted. At Clark the key fact was that the experts said this kind of accident was highly improbable (1 in 1,000,000) and would probably never happen and yet it did. At WPI the prevailing view was that even with idiots and incompetents abusing a nuclear reactor they had not been able to make it meltdown to the point of breaching containment and harming the public. It was an economic disaster for the industry to be sure, but human error had been mitigated by built in automatic safety systems. Indeed, had all the operators taken a coffee break when the first alarm went off and left the system alone it would have shut down safely and the emergency core cooling system would have kept the system acceptably stable. The real problems began when the operators, confused about what had happened, shut down the ECCS (Emergency Core Cooling System).

Note the temptation by technologists to design humans out of the system and make them peripheral rather than create a transparent and fault tolerant man-machine interface and depend on well trained and highly paid operators. This issue is returning in the robotics debate as the “human in the loop” question about whether one really wants to seek fully autonomous systems or not? Economics push one to reduce the caliber and number of operators if possible. Other considerations such as economic liability and political or legal accountability push back the other way.

The robotics debate is still in its pre disaster phase and analogy based perceptions of the institutions in charge are likely to be very important, hence the scenarios we designed move the lead role from government to various forms of public- private or private ventures. On the other hand there has been a very active science fiction literature raising concerns about this technology's development.

The bulk of the nuclear power referents in science fiction tended to be fairly optimistic by comparison to those about robotics, but the first nuclear application was not a power plant, but an atom bomb that destroyed two whole cities. That history of surprise, dread and the strong reassurances given the public that "Atoms for Peace " had been tamed may have contributed to the public reaction of shock when nuclear technology finally did get out of control. The experts really were not on top of things and the unthinkable nearly happened at TMI and then did at Chernobyl. So, the issues of autonomy, subordination and control, highlighted by Asimov, are the focal point of this part of the perceived desirability variable tapped by this item.

If this scenario came about, how likely would it be to raise severe or challenging ethical concerns?

This item serves as a crosscheck item for the ethical concerns raised by the man-machine relationship. Major ethical concerns may be indicated by the man and machine relationship, but it is also possible that other values, especially an environmental ethic, and possibly issues having to do with the meaning of work from various religious perspectives, have significant bearing on reactions to the questions that robotics raises for humanity. An item that asked about the level of concern provoked by each scenario that was not specific to what those concerns were, seemed appropriate. This question is an estimate of the odds that severe ethical concerns would be raised by the technology developing for the purposes indicated under the control of the given institution in each scenario.

A consensus on high levels of ethical concern would be a very significant "red flag" even if the respondents did not see the ethical stakes rising with each violation of one of Asimov's laws, as we expected. Each scenario was designed to vary in its acceptance of Asimov's laws. Responses to this question will be used to determine a relationship between Asimov's laws and perceived ethical concerns. This serves as hypothetical test of Asimov's laws as ethical guidelines for robotics technology. The results of this question will also be compared with the scenarios deemed most likely to come about. In this case the two likelihood items will indicate whether the most likely scenarios are also the ones most likely to raise ethical concerns and challenges. If the current direction of the technology is deemed problematic on grounds of the emerging man machine relationship and those involved will likely be faced with ethical dilemmas, it is time to examine whether this is the direction the field or the funding agencies want to go.

The Sample

Student assessor panels were initially drawn from Worcester Polytechnic Institute (WPI, a technical college, and Clark University a liberal arts college. Later, there was additional data collection at Worcester State University in a non-technical class and at Boston University in a technical class.

Our initial objective at WPI and Clark University was to collect a stratified sample broken into three segments: Future Experts in Robotics, Future Experts in other Technical Fields, and Future Experts in Non-Technical Fields. These students represent the voices three professional groups expected to take part in the intensifying robotics debate. The robotics majors represent the leaders in the field driving the technological development. These are the ones holding the jobs and designing the products. They are also the ones whose jobs are most affected by the funding institution and its policies. Those students majoring in other technical fields may contribute to the field or even hold management positions in the companies developing robotics technologies. Their voice represents that of the roboticists' peers in the general technological community but they are not experts. The students in the non-technical sample are drawn from Clark University and represent public opinion leaders, the media, and political leaders and social activists. Theirs is the least expert, but not necessarily the least informed voice on the subject of technology and society issues.

In a traditional Delphi study one would normally draw on a small panel of experts. It would be difficult to know if they were representative of the full range of expert opinion on the subject but the credentials of the participants would provide credibility to the result. One often has to contact many experts to get a few to participate, so Delphi panels are probably not representative, and may actually be skewed toward the experts with the strongest opinions on the subject. However, past studies on other technical subjects (aerospace technologies) have revealed that Delphi panels and student samples tend to be in sync with one another one average, the major difference being more moderate opinions on the part of the students. Students are also more likely to respond, so one gets more intermediate positions represented.

An aerospace technology study (Climis, Learned, Bussey)) of the likelihood of various breakthroughs compared the ratings of WPI Alumni to those of WPI students and expert panels. What is of concern here is not the results but the response rate. The Alumni study produced about a 25% response rate, compared to about 65% response among the students. Yet, the distributions of response were comparable in both studies. The Expert Delphi study did not precisely report a response rate, but many experts were contacted to get about a dozen to respond. I estimate about a 10-15% response rate, half of that documented by the Alumni study. So, the students are more accessible and their response distributions lead to the same conclusions about majority views as those based on more expert respondents.

These findings led us to have confidence that using a student sample to calibrate a new instrument and get an initial sample distribution would be meaningful and an efficient use of our resources. Available time and resources would not allow us to generate a sample large enough to statistically represent all experts. The choice was a small Delphi panel or a student sample of hundreds stratified to allow us to see if the more expert students (aspiring experts) stand out from those in other courses of study. By substituting students for real experts one can get a statistically respectable sample size which is more likely to be representative of the universe of students from which it was drawn.

The preferred approach was to sample at the level of classes rather than individual students despite the problem that would create later in assessing the randomness of the sample. As a practical matter, the adequacy of the sample would be determined by the willingness of the instructor to devote class time to the study more than the selection of classes, though both would be a factor. The classes were not randomly selected. Some were selected to maximize the chance of finding robotics majors in them, others were selected so as to represent a typical distribution of majors at the college and others were selected in the hopes that the instructor would be interested in the subject and cooperate with the study by offering class time for administration.

Classes were selected from Worcester Polytechnic Institute (WPI) and Clark University based on distribution of the students we needed and the relevance of our study to the class topic. If the course instructor agreed to participate, they were given the option of having us distribute the survey with a five minute speech, or for us to take over the class and administer the survey followed by a presentation on the project. Three robotics courses and three other technical courses were selected at WPI, and three humanities courses at Clark University. There were massive differences in response rate and in effect, the courses with cooperative faculty members willing to devote class time to the study came to dominate the sample.

In total there were about 45 Robotics Majors from WPI, about 60 other Technical from WPI and 54 Non-Technical from Clark University, about 150 in all. The critical part of the study was getting a very high proportion of the total pool of robotics majors and that was done in part by getting strong cooperation from one key professor and good cooperation from the instructors in two other key courses. There was a near complete failure to obtain data from a related CS course in which the instructor provided modest access but no endorsement. Hence, just over half of the RBE major data came from one key course with 98% participation. This gives us some confidence that the data in that stratum is not self-selected since it is a required course and there was full participation. The other smaller courses involved were over 80% participation as well. The class data collection effort that was a bust had little to no impact on the study. Data collection at Clark University was more successful since the instructors were carefully selected so as to be likely to be interested, and 75% were. Again, one was exceptionally interested and her class participated fully but another allowed class time to be used for data collection as well. If there is any self-selection in the Clark University sample it is at the level of deciding to take a course dealing with science fiction or computer science, and not at the individual level of deciding whether to fill out the questionnaire and return it. Thus, the sample may not represent the whole student body, but is comparable to the kind of subset of the campus represented by the robotics majors at WPI, the most literate portion of the non-technical campus on this subject.

Where an honest effort to be representative of the campus as a whole was made was in selecting the "other technical" sample at WPI. This was an effort to go to the classes that fulfilled distribution requirements for the campus as a whole rather than to select any given major. Though first choice was the Engineering Science (ES) classes taken by most majors, the faculty members offering those courses could not make a course connection to the study and response rates were distressingly low (10-15%). As a result, we shifted to social science classes that fulfilled campus wide requirements in another way. Two Sociology classes, without much if any connection to the study other than as the basis for a social research methods illustration, ended up providing the bulk of the subsample, with 80-90% participation rates in each class.

Though a case could be made that the kind of student who would fulfill their social science requirement with sociology rather than economics or psychology classes might differ from the typical

WPI student, the likelihood of self-selection is certainly no greater than it was in the Clark University sample and less likely to be directly relevant to the subject at hand.

Two more classes in other colleges were sampled in order to determine if the results of this study generalized to other student populations. The technical sample was cross checked by going to a computer science class at Boston University where there is no robotics major, but this class spent one week out of 14 on the subject. The Clark University sample was cross checked with a sample from a Sociology class at Worcester State College that served a similar role for majors there as the Sociology class at WPI. The results of this comparison and their implication on the generalizability of this study will be discussed in another chapter.

The four scenarios and the question response items provide an instrument for capturing the perceptions of students on possible directions for the future of robotics. The stratified sample of 200 students from four colleges represents three major voices in the robotics debate. Even if the direction of robotics is in principle unpredictable, this study provides insight into the public debate that will take place in the future. If the current direction of the technology is deemed problematic by a wide consensus covering all three of these groups, the result may well be a serious attempt to control the technology.

Results

Likelihood

The ANOVA test between scenarios and post-hoc Turkey test showed the Military scenario was considered to be significantly more likely than all other scenarios with a mean 3.2. An ANOVA test between strata showed no significant difference between the non-technical, technical and expert strata on the military scenario. Hence the Military scenario is considered most likely by all sample groups

Figure 1 shows the distribution of means across scenarios. The Lunar Scenario was second and the others essentially tied for third place.

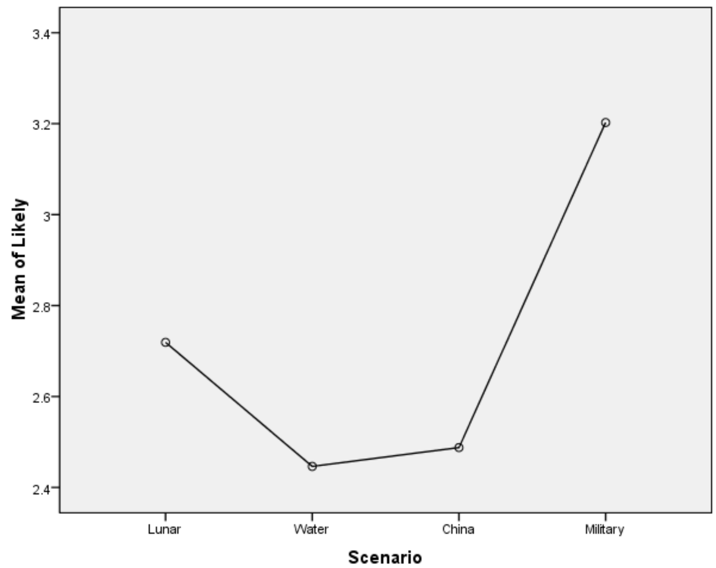

Figure 1: Distribution of averaged likelihood across all scenarios.

Differences did arise between strata on the Lunar and Water Scenarios. These differences were examined with crosstabs and a Chi Square test. For the lunar scenario the Tech sample was significantly different than both the RBE and the Non-Tech Group, with the Tech Group more likely to say the scenario is very unlikely. There was no significant difference between RBE and Non-Tech.

Lunar Scenario			sample groups			Total
			Non-tech	Tech	Rbe	
LLikely	Unlikely	Count	4	12	2	18
		% within sample groups	7.7%	22.6%	4.8%	12.2%
	Somewhat Unlikely	Count	10	14	13	37
		% within sample groups	19.2%	26.4%	31.0%	25.2%
	Somewhat Likely	Count	28	24	19	71
		% within sample groups	53.8%	45.3%	45.2%	48.3%
	Very Likely	Count	10	3	8	21

		% within sample groups	19.2%	5.7%	19.0%	14.3%
Total		Count	52	53	42	147
		% within sample groups	100.0%	100.0%	100.0%	100.0%

For the Water Scenario the Tech and Non tech samples were significantly different with the Tech Sample rating likelihood lower and the Non-Tech sample rating likelihood higher. Although half of the Non-Tech sample considered it very likely or somewhat likely and about a third of the Tech sample agreed, the striking difference was the 23% of the Tech sample found it very likely while only 6% of the Non Tech sample agreed. The RBE sample bridged the gap between them and was not significantly different than either of the other sample groups, close to the Tech group on very likely but closer to the Non-Tech group on somewhat likely.

Water Scenario			sample groups			Total
			Non-tech	Tech	Rbe	
WLikely	Unlikely	Count	6	11	5	22
		% within sample groups	11.5%	20.8%	11.6%	14.9%
	Somewhat Unlikely	Count	16	25	20	61
		% within sample groups	30.8%	47.2%	46.5%	41.2%
	Somewhat Likely	Count	18	14	15	47
		% within sample groups	34.6%	26.4%	34.9%	31.8%
	Very Likely	Count	12	3	3	18
		% within sample groups	23.1%	5.7%	7.0%	12.2%
Total		Count	52	53	43	148
		% within sample groups	100.0%	100.0%	100.0%	100.0%

Desirability

Initial test indicated substantial correlation between two variables used to measure desirability, the perceived changes in the man machine relationship and the perceived changes in quality of life. A desirability index was formed to combine these two variables by averaging each participant's response to the two questions. This allowed a direct comparison of means between scenario and strata on "desirability in general". An ANOVA test indicated differences in the military and lunar scenario between strata. A closer look indicated the RBE group stood out in both cases rating the military scenario less desirable than the other groups and the lunar scenario more desirable than the other groups. Figure 2 and show this variation of mean between strata, the tables below show the actual distributions for the lunar and military scenarios where the RBE majors stand out.

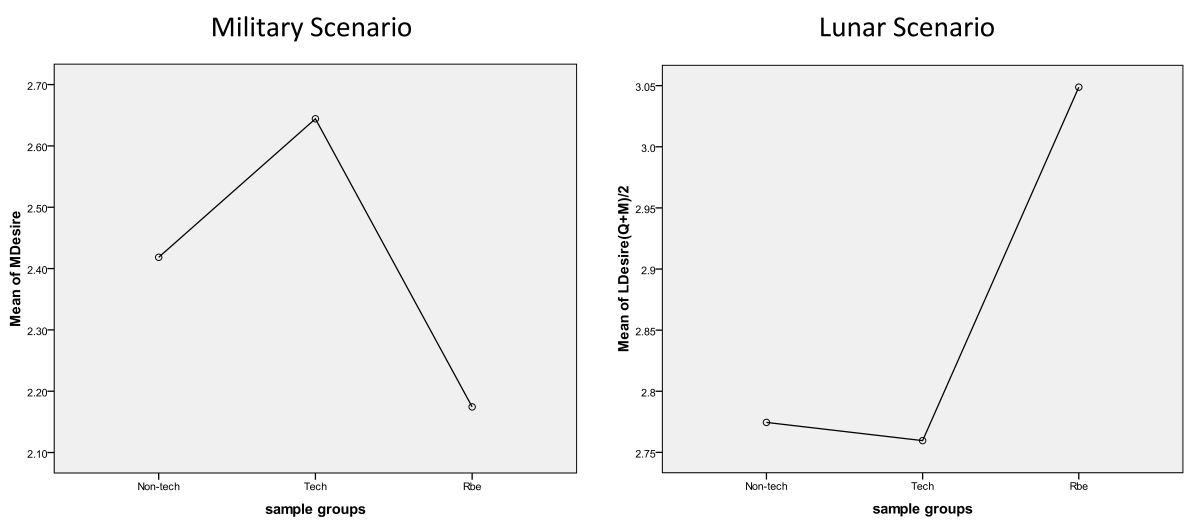

Figure 2: Variation of average desirability between strata

(Military scenario on left, Lunar scenario on right)

Clearly the robotics majors have reservations about working on military applications, and the notion of working on a lunar mining camp seems to appeal to them (2.15 vs. 3.05 mean desirability ratings). At 2.65 and 2.75 the other technical majors see little difference but tend to rate the lunar scenario a bit more desirable on average.

Military Scenario			sample groups			Total
			Non-tech	Tech	Rbe	
Desirability	Undesirable	Count	11	13	15	39
		% within sample groups	22.4%	25.0%	34.9%	27.1%
	somewhat undesirable	Count	18	10	17	45
		% within sample groups	36.7%	19.2%	39.5%	31.3%
	somewhat desirable	Count	9	13	5	27
		% within sample groups	18.4%	25.0%	11.6%	18.8%
	Very desirable	Count	11	16	6	33
		% within sample groups	22.4%	30.8%	14.0%	22.9%
Total		Count	49	52	43	144
		% within sample groups	100.0%	100.0%	100.0%	100.0%

Lunar Scenario			sample groups			Total
			Non-tech	Tech	Rbe	
Desirability	Undesirable	Count	12	12	2	26
		% within sample groups	23.5%	23.1%	4.9%	18.1%
	somewhat undesirable	Count	13	9	9	31
		% within sample groups	25.5%	17.3%	22.0%	21.5%
	somewhat desirable	Count	14	17	16	47
		% within sample groups	27.5%	32.7%	39.0%	32.6%
	Very desirable	Count	12	14	14	40
		% within sample groups	23.5%	26.9%	34.1%	27.8%
Total		Count	51	52	41	144
		`	100.0%	100.0%	100.0%	100.0%

Actual ordering varied in all three groups; however the Military and China scenarios were the lower half in desirability in every stratum. The Military scenario had the lowest desirability rating in both RBE and Non-Tech but The Tech sample rated the China scenario the least desirable. Only the RBE sample reported a significant difference between the China scenario and the Military scenario. Figure 3 shows the distribution of means by scenario. The table on the right shows the means for each sample strata.

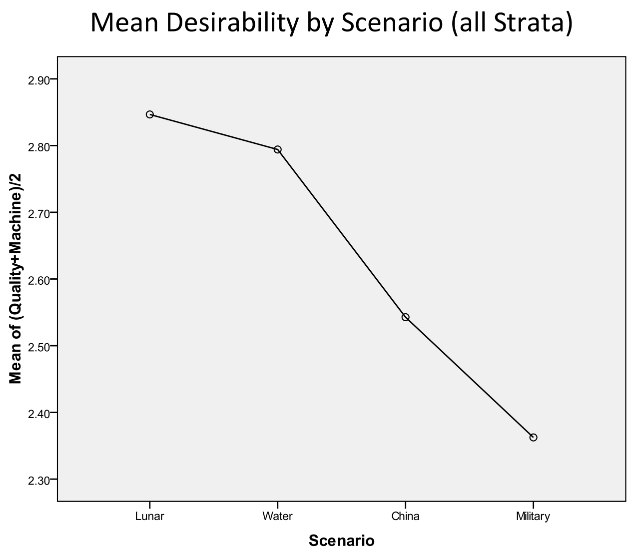

Mean Desirability Ratings by strata				
	Lunar	Water	China	Military
Non-tech	2.8	2.9	2.5	2.4
Tech	2.8	2.8	2.5	2.6
RBE	3.0	2.8	2.7	2.2
total	2.8	2.8	2.5	2.4

Figure 3: Distribution of average desirability across all scenarios

Ethical Concerns

The mean responses to the general ethical concerns question are compared in the following table. As expected, there are some interesting similarities to the composite desirability findings. For the sample as a whole the rating of the China Eldercare scenario and the Military scenario are nearly identical. The Lunar scenario raises the fewest ethical concerns, but the Waterworld scenario has moved to an intermediate position between these two extremes. However, this is overall. The rank ordering shifts depending on the stratum one is examining and the ethically charged scenarios got to their nearly tied positions in different ways. The non-technical majors from Clark U consider the China eldercare scenario more fraught with moral issues than the Military scenario.

The WPI students rank the two most controversial scenarios the other way around, but the Other technical majors at WPI see the two as about equally controversial while the WPI Robotics majors see both as more ethically concerning and also discern more of a more difference between them.

The mean ethical concerns ratings on the military scenario by the robotics majors at WPI and the Non-technical majors at Clark U are quite similar. It is the other technical majors at WPI that are most comfortable with the Military application rating it at a level lower level of ethical implications than the Robotics majors rated the China eldercare scenario.

Turning to the question of statistical significance, ANOVA analysis reveals that the only significant difference between the strata was on the China eldercare scenario. Post-hoc tests show that difference to be due to the Clark Non-Technical majors who rated the China scenario far more unethical than the RBE or Tech majors at WPI did. . However, it is also important to note that it is the mean 3.82 rating from the Non-Techs at Clark contrasting with the 3.50 rating from the Tech majors at WPI that is statistically significant. The 3.58 mean rating of the robotics majors is not significantly different from that of the other stratum involving WPI students. On the military scenario, which the Non-techs rated on average a more moderate 3.75 on ethical concerns, while the Tech majors at WPI moved up to 3.55; one no longer has a statistically significant difference with a sample of this size. In this case the RBE majors were intermediate with a mean of 3.69. These three ratings are not significantly different at the .05 level.

In terms of ethical issues, we have a theory to test. The hypothesis involving this item was that ethical concern would rise to the degree that the scenario violated one or more of Asimov's laws of robotics. Hence, we were expecting the Military Scenario to be most fraught with ethical concerns and the China eldercare scenario, which violated only one of these laws, to be more acceptable. Comparison of Means showed both RBE majors and Tech majors from WPI answered rating the scenarios in our predicted Asimovian order (which is the Lunar scenario as presenting the fewest ethical concerns, then the Waterworld scenario, then the China eldercare scenario and the Military scenario with the highest level of concern). The only problem with taking this position is that the differences in the ratings of the last two by the Techmajors is so small that they are statistically tied. The same situation holds when one compares the difference between the ratings of these two scenarios by the Non-Techs who flipped the order of the China and Military scenario but considered both to have higher levels of ethical concern compared to the other group that rated them as similar. Only the robotics majors seemed to really distinguish the two and put them in the expected Asimovian order. The Non-Tech majors from Clark are interesting since they not only rated the two scenarios as similar but also flipped the order of the China and Military scenarios to refute the Asimovian hypothesis explicitly.

Since we consider testing this hypothesis important and the refutation was hanging on a slender difference, a more powerful rank order analysis procedure was used to double check the finding. In a Guttman scale one examines the way all the different respondents individually rank ordered the scenarios. This is no longer based on the mean rating from the group, but involves finding out how predictable the rating order itself is, whether the differences are large or small. The Guttmann scale effectively tells us how often we would be wrong if we tried to predict each individual's rank order using the Asimov rules hypothesis.

It turns out that we would rarely be wrong. While there is some dither for the sample as a whole, relatively few Clark ratings registering large differences produced the small overall average difference in the wrong direction. For most people the response pattern is in the predicted order and even where there was divergence it involved only the order of those 2 items- ie., one error out of 4. Overall that converts to only 1/100 chances of being wrong when using Asimovian order to predict ethical order at the level of the individual respondent. Guttman scale procedures also reveal that there are two item groupings. The errors in order that do occur are within the Lunar and Waterworld pool of items or within the China and Military pool of items. There are very few if any Waterworld and China scenario reversals in rating. This is not surprising given the larger jump in means we had already observed, but it is also reassuring that there is nothing like random variation in the typical pattern of ratings producing those average figures. The Asimovian hypothesis has more statistical support across the strata than seemed to be the case at first blush.

Mean Ethics Rating By Scenario

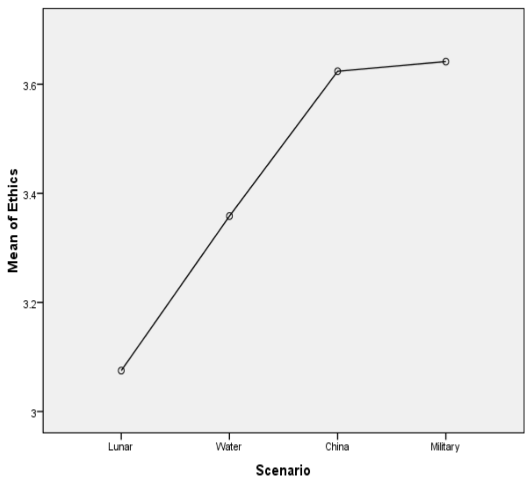

Mean Rateing of Ethical Concerns by Strata				
sample groups	Lunar	Water	China	Military
Non-tech				
Mean	3.21	3.39	3.82	3.75
Std. Deviation	.848	.874	.432	.483
Tech				
Mean	3.08	3.34	3.50	3.55
Std. Deviation	1.082	.919	.939	.845
Rbe				
Mean	3.00	3.33	3.58	3.69
Std. Deviation	.883	.969	.852	.780
Total				
Mean	3.10	3.35	3.64	3.66
Std. Deviation	.945	.913	.775	.719

Figure 4: Distribution of average ethical concerns across all scenarios

Worcester State and Boston University Study

There is always a question of how representative a sample is of the professions and voices one intended to capture. During the course of this project we collected data from WPI and Clark University. However WPI's robotics major was the first in the nation and remains the only accredited robotics program. The unique nature of the WPI robotics major raises questions about whether a study done anywhere else for the technical sample would yield the same result. There is even a possibility that non robotics students at WPI have been affected by the existence of the robotics program and would not have the same view as technical majors at another school. On the Clark University side there was also a generalizability issue, but it was of a different nature. The classes that agreed to participate in the study were rather unique. Were these students typical of even the rest of Clark University, much less non-technical majors more generally at other colleges and universities? Further, there was the question of whether these rather expensive and elite small private colleges were typical of the more general university student bodies, both public and private and at larger institutions.

Thus, a methodological concern would be examined before a final conclusion of the study was drawn. Given time constraints, nearby institutions would be approached. After some discussion, the project team decided that sampling students from a technical college at Boston University and a non-technical class at Worcester State University would give ample opportunity for the small elite college bias to be examined. Only one project team member could be spared for this activity so the effort could not be as large as the original data collection effort. Adam Vadala- Roth took this on as a special assignment in place of other writing and analysis assignments. It was clear that he would not be finishing with the rest of the team in order to carry out this spin off study.

Gathering data from WSU was a convenience to help Adam out , as he did not have access to a car. It seemed that a technical and a non-technical class from UMass Amherst would be more appropriate, but proximity and the existence of personal contacts at WSU prevailed. This smaller branch campus was local but did not include an engineering school. Still, these universities were different from Clark and WPI. One was not elite and the other was not small. They did not even offer the same array of majors as the colleges they were to mimic in the generalizability study.

Only one class at each school was studied as opposed to drawing a sample from more than one course as we had at WPI and Clark University. At BU the course sampled was Artificial Intelligence, our survey being given out during the one week in which they discuss robotics in their class. There was no robotics major at BU, but the other criteria we had considered using as the “expert” sample criterion could be used, as there were AI classes in Computer Science. That class at WPI had been approached, but was not a significant factor in the WPI sample due to a poor response rate.

At Worcester State University an Introductory Sociology class was recruited, on the grounds that a class of the same type had been recruited at WPI and had produced a distribution of majors typical of the college as a whole. On both campuses it fulfilled a distribution requirement. There were no classes comparable to those recruited at Clark University (one on science fiction, the other on computer programming) at WSU, but it seemed likely that the two classes that dominated the liberal arts sample were, in combination, typical of the Clark university student pool. To answer the questions about generalizability and whether a consensus was emerging between institutions the data was processed to produce cross tabulations and chi square calculations. A significance level of 0.05 was used for the chi squared test. However, we decided examine the distribution of any survey item with a significance level of .10 or below in order to identify a trend for one institution to be of more or less concern than another. A trend might emerge out of the not quite significant findings, but they cannot stand on their own. If the chi square significant figure fell below the decision threshold then our hypothesis that there was no significant difference is wrong and there is a difference to be explained.

Initially the distribution of non-technical major responses from WSU and Clark seemed very different. The initial conclusion was that the Clark findings were unlikely to generalize to other non-technical college populations. However, closer examination suggested that there was not only a pattern in the WSU findings, but also some interesting similarities to the Clark data as well as differences. The pattern made sense as well, though the hypothesis we brought to the study did not lead us to expect what was emerging at WSU. In retrospect, it made sense, and was even predictable. The following table represents the significance of each question on each scenario when comparing WSU to CU. Statistically significant findings and possible trends are denoted.

Pearson Chi Squared Significant Figure Chart

Legend:

** = Significant Difference (.05)

* = Not Significant but possibly part of a trend (.10)

Clark University vs. Worcester State University, Chi Square Significance Figures

Scenario	Likelihood	Spin Off	Quality of Life	Man / Machine	Ethics
NASA	0.014**	0.026**	0.636	0.511	0.598
Water World	0.969	0.530	0.763	0.289	0.015**
China	0.199	0.444	0.004**	0.002**	0.067*
Military	0.622	0.167	0.093*	0.011**	0.018**

Looking at the NASA scenario data for WSU compared to Clark yields some interesting findings. Examining the five items separately one finds that the hypothesis is supported and there are significant differences only for the likelihood and spin off questions. In the other 3 cases there were not significant differences. What they agree on is more important than what they do not agree about, from our perspective, so there is evidence of a general consensus emerging between the two institutions. The finding is still that one of the more desirable scenarios is among the least likely. Although the two schools still differed on how likely the Lunar scenario is.

			Nasa_Likely				Total
			Unlikley	Somewhat Unlikely	Somewhat Likely	Very Likely	
Institution	Clark	Count	4	10	28	10	52
		% within Institution	7.7%	19.2%	53.8%	19.2%	100.0%
	WSU	Count	0	11	14	0	25
		% within Institution	.0%	44.0%	56.0%	.0%	100.0%
Total		Count	4	21	42	10	77
		% within Institution	5.2%	27.3%	54.5%	13.0%	100.0%

			Nasa_spin_off			Total
			Somewhat Unlikely	Somewhat Likely	Very Likely	
Institution	Clark	Count	2	23	27	52
		% within Institution	3.8%	44.2%	51.9%	100.0%
	WSU	Count	1	19	5	25
		% within Institution	4.0%	76.0%	20.0%	100.0%
Total		Count	3	42	32	77
		% within Institution	3.9%	54.5%	41.6%	100.0%

For water world only the Ethics item produced a significant finding. This means that the two samples produced very similar distributions of responses on the four other items. Both likelihood and desirability rating are in alignment. However, it is more of an issue for the study when the students disagree about the ethical implications of a technological application than its likelihood. There is a real difference here, in that the Clark students are more concerned about this scenario coming to pass than the WSU students. That is an important matter that did not generalize, possibly reflecting a greater level of environment ethics concern at Clark than at WSU among non-technical majors.

			Water_Ethics				
			Unlikley	Somewhat Unlikely	Somewhat Likely	Very Likely	Total
Institution	Clark	Count	3	4	14	30	51
		% within Institution	5.9%	7.8%	27.5%	58.8%	100.0%
	WSU	Count	0	8	9	8	25
		% within Institution	.0%	32.0%	36.0%	32.0%	100.0%
Total		Count	3	12	23	38	76
		% within Institution	3.9%	15.8%	30.3%	50.0%	100.0%

The China scenario has significant differences on both of the desirability questions, quality of life, and the man machine relationship. Additionally the Ethical issues fell into the .10 significance category and closer examination showed a skew in the same direction as the others. Hence, all three items dealing with desirability issues are consistent is signaling a difference between the Clark and WSU students, and all in the same direction. This scenario was simply more questionable in the eyes of the Clark U students than those from WSU and the difference was real and consistent, though they agreed about its likelihood of occurring.

While there were some minor similarities between the two institutions on the questions dealing with likelihood and spinoff, WSU did not replicate Clark overall in applications dealing with robotics applied to medical assistance roles. The WSU students are much more okay with it than the Clark students. There is not a consensus and the findings did not replicate. This is an unexpected difference that has to be explained, though in retrospect, maybe we should have expected it. Two theories will be presented later in an effort to understand why this difference occurs.

			China_Quality				
			Undesirable	Somewhat Undesirable	Somewhat Desirable	Very Desirable	Total
Institution	Clark	Count	5	8	30	6	49
		% within Institution	10.2%	16.3%	61.2%	12.2%	100.0%
	WSU	Count	0	4	9	12	25
		% within Institution	.0%	16.0%	36.0%	48.0%	100.0%
Total		Count	5	12	39	18	74
		% within Institution	6.8%	16.2%	52.7%	24.3%	100.0%

			China_Machine				
			Undesirable	Somewhat Undesirable	Somewhat Desirable	Very Desirable	Total
Institution	Clark	Count	10	24	13	3	50
		% within Institution	20.0%	48.0%	26.0%	6.0%	100.0%
	WSU	Count	0	6	13	6	25
		% within Institution	.0%	24.0%	52.0%	24.0%	100.0%
Total		Count	10	30	26	9	75
		% within Institution	13.3%	40.0%	34.7%	12.0%	100.0%

The Military scenario followed a similar trend as the China scenario. Again all their questions dealing with desirability are in question. One finds statistically significant differences on the Man Machine relationship question, the Ethics question and Quality of life supporting the trend although not at a statistically significant level. In all 3 of these items the WSU students have less concern about the desirability of this application of robotics than the Clark students. Hence, it seems unlikely that this is an outlier rather than part of a real difference between the samples that needs to be explained.

			Military_Machine				
			Undesirable	Somewhat Undesirable	Somewhat Desirable	Very Desirable	Total
Institution	Clark	Count	14	15	13	7	49
		% within Institution	28.6%	30.6%	26.5%	14.3%	100.0%
	WSU	Count	0	6	10	8	24
		% within Institution	.0%	25.0%	41.7%	33.3%	100.0%
Total		Count	14	21	23	15	73
		% within Institution	19.2%	28.8%	31.5%	20.5%	100.0%

			Military_Ethics				
			Unlikley	Somewhat Unlikely	Somewhat Likely	Very Likely	Total
Institution	Clark	Count	0	1	11	39	51
		% within Institution	.0%	2.0%	21.6%	76.5%	100.0%
	WSU	Count	3	2	7	12	24
		% within Institution	12.5%	8.3%	29.2%	50.0%	100.0%
Total		Count	3	3	18	51	75
		% within Institution	4.0%	4.0%	24.0%	68.0%	100.0%

Both the China and Military scenarios presented some unpredicted differences that may affect the generalizeability of the study. Upon reviewing the dataset we found that just over a third of the students (9 out of 25) were majoring in either nursing or health education, fields clearly affected by the China scenario. Additionally Clark is a highly regarded liberal arts school that holds roughly comparable prestige in the liberal arts community as WPI does in the engineering community. It is a strong if not elite institution. WSU is a small state school that offers many different degrees and is not an elite school. WSU has a more vocational orientation, and may more generally represent workers whose jobs are directly affected by one or more of the proposed scenarios. In the retrospective technology assessment of the telephone the people most affected by the technology made better predictions then the technologists. Hence, our WSU sample may have been influenced by an additional voice that was not considered in the original design of the study.

Overall we are impressed with the similarities between the two schools, however several concerning differences have been noted above. The WSU sample seems to have captured a different voice in the non-technical community than that found at Clark University, a Liberal Arts College noted for its liberal political leanings, ecological ethics and environmental concern. If anything, the Clark sample seems likely to be more critical and skeptical of robotics applications than the typical non-technical sample is likely to be. The difference shows up especially in the man-machine relationship area and this raises ethical questions for the Clark students even when in the short run, elders benefit or their country gets a military advantage from a robotic application. Hence, the pattern of consensus emerging from the WPI-Clark comparison is all the more remarkable, given the sensitivities at Clark on the subject of robotics, and their attention to long run implications rather than just short term consequences.

Turning our attention to the WPI and BU sample, we find no significant difference on the surface, and only a few below the threshold we set for examining trends. The following table shows these the Chi square values.

** = Significantly lower than decision threshold

* = Greater decision threshold but close (gray area)

Worcester Polytechnic Institute vs. Boston University, Chi Square Significance Figures

Scenario	Likelihood	Spin Off	Quality of Life	Man / Machine	Ethics
NASA	0.469	0.796	0.578	0.111*	0.068*
Water World	0.274	0.640	0.166*	0.866	0.238
China	0.637	0.568	0.423	0.465	0.738
Military	0.566	0.267	0.428	0.465	0.467

Looking at the NASA scenario for WPI vs. BU, the similarities are still dominant and very apparent. Where there are potential or actual differences, it is the WPI students seeing more ethical and man-machine issues than the BU students. Again, if anything the main samples in the study (WPI and Clark) are from campuses with somewhat strong or extreme views on the robotics question, more sensitive, critical or skeptical than other places.

			NASA Man-_Machine relationship Item				
			Undesirable	Somewhat Undesirable	Somewhat Desirable	Very Desirable	Total
Institution	BU	Count	0	1	21	8	30
		% within Institution	.0%	3.3%	70.0%	26.7%	100.0%
	WPI	Count	4	12	38	14	68
		% within Institution	5.9%	17.6%	55.9%	20.6%	100.0%
Total		Count	4	13	59	22	98
		% within Institution	4.1%	13.3%	60.2%	22.4%	100.0%

			NASA Scenario -_Ethics item				
			Unlikely	Somewhat Unlikely	Somewhat Likely	Very Likely	Total
Institution	BU	Count	4	12	9	5	30
		% within Institution	13.3%	40.0%	30.0%	16.7%	100.0%
	WPI	Count	7	12	25	24	68
		% within Institution	10.3%	17.6%	36.8%	35.3%	100.0%
Total		Count	11	24	34	29	98
		% within Institution	11.2%	24.5%	34.7%	29.6%	100.0%

Looking at the water world scenario for WPI vs BU, the similarities are dominant, even more so than in the NASA scenario. We see a difference appearing on the quality of life question. However, we have not accepted differences at this level as evidence of a difference alone. Instead we have examined them to add evidence to general trends between the two data sets. Hence, the conservative decision is to say that the two samples are consistent.

In a prior section, a question was raised about whether the BU students were more like the expert or the other technical part of the WPI sample. When selecting classes at WPI an AI class was approached and would have been considered part of the expert sample. It is worthwhile to make the comparison with the AI class from BU. The pattern was that typically the BU students responded closer to the WPI robotics majors than the other majors. There were exceptions however, and about half of these exceptions had to do with the BU students seeing the likelihood of the robotic application's coming to pass more like the other technical students at WPI, and typically less likely than the WPI robotics majors. On 1-2 items they fell in between the two WPI samples and twice they took an extreme position compared to both parts of the WPI sample, but on the whole the BU students either saw fewer issues to be concerned about than the WPI sample as a whole or were more in agreement with the robotics majors at WPI than those in other majors. Hence, they could reasonable be considered an expert sample from another University.

The purpose of this chapter is to examine the ability of our sample to generalize to the three groups of aspiring experts we intended to capture: the robotics experts, the technological experts and the non-technological experts. Overall the samples the similarities between the original sample from WPI and Clark and the WSU or BU sample are promising. However the differences that do appear suggest that there will be additional voices taking part in the future robotics debate.

Discussion

The results obtained from this study have provided insight to the perceptions of college students about the future of an emerging technology. One of the most interesting aspects of our research is that regardless of academic preference or background, our participants answered 16 out of the 20 survey questions with consensus. Such a consensus has significant sociopolitical implications. Whether or not there prediction are entirely accurate the level of consensus indicates the direction of the debate over robotics technology as each sample strata is intended to represent a voice in this debate. All three strata found the military scenario to be the most likely scenario to come about. This again indicates a high level of consensus on the current direction of technology.

Turning to the desirability questions, the military scenario was rated lowest by the Robotics experts and tied with the china Scenario for least desirable in all other strata. These questions were designed to indicate the value to society of introducing the technology, and the effect of the technology on the man-machine relationship. Undesirable responses indicate the technology may create worse problems than it solves, upset the balance in the system, or even get out of control. Hence, there seems to be a strong degree of consensus that the most likely scenario to come about does not seem likely to solve problems, or that in solving one problem it might have unintended consequences which create even worse problems. It is also interesting to note the cases where our strata had differing opinions. In the Military scenario, WPI RBE majors marked the desirability notably lower than the other strata. This is an interesting discovery in itself. The future engineers who are likely to be designing robots capable of killing humans in fact perceive the technology itself as undesirable for mankind.

Another variance in perception discovered from our data analysis was that WPI RBE students viewed the Lunar scenario, where robots were used to mine resources from the moon was more desirable for mankind. Aspiring experts in the field of robotics see the Lunar scenario as a desirable achievement for their careers and a direction they would like to work in, however their perception of the likely direction of the field indicates they have resigned themselves to work on military applications. One explanation for this is that the current funding trends have made the military scenario the only financially viable one for the robotics engineer. This indicates that, the goals of the funding institution will play a large role in the development of the technology. The military scenario was designed with current funding patterns in mind. If our scenario is accurate, it will be very likely that the future will involve the robotic technologies depicted in the Military scenario. After analyzing the data we discovered that our participants by far and large believe the military scenario to be the most likely scenario to become reality. What is even more interesting is that there is tremendous consensus across all strata that advanced military robotics does not improve the quality of life and raises incredible ethical concerns.

Another important discovery was that Asimov's fictitious laws that were written in the mid nineteen hundreds actually hold much weight in the perceptions of humans on robotic ethics. Our data analysis revealed that both the desirability and ethical concerns accurately represented the degree of which Asimov's laws were followed in the four scenarios. The Lunar scenario had the highest rating for quality of life and lowest rating for ethical concerns while the Military scenario produced results with the lowest rating for quality of life and highest rating for ethical concerns. In all scenarios we discovered that students believed the technologies envisioned would have tremendous spin-off implications. The field of robotics is rapidly expanding and it is interesting to note that technological discoveries used in one application may very well serve useful in a myriad of other applications. It is very possible that the next technological revolution will stem from the field of robotics engineering.

In the beginning of this project we sought to find out if the funding patterns mattered when it came to the development of robotic technologies. After analyzing the data we found our answer: yes it does. In fact each scenario was perceived to have completely different results in the perceptions of college students. With this in mind, we believe that studies like this are important to help guide the development of robotics engineering so that it can improve the quality of life for humans while also reduce the amount of ethical concerns.

Additionally we propose that the strata represent three major voices that will shape the robotics debate. This approach to forecasting the direction of the technology captures ideas coming from inside the technical community and outside of it making more robust predictions possible. While an extreme amount of consensus on the direction of the technology suggests accurate predictions, differences between the strata such as those in the military scenario reveal likely topics of debate between these groups and the focus of future policy decisions. Our study suggests that there may be other voices contributing to this debate which were not captured by our stratified sample. Our generalizeability study at Worcester state university and Boston University showed a high degree of consensus among students at other colleges; however there were several differences that cannot be written off. We suggest the study be expanded further capturing management majors and students specializing in the fields affected by our scenarios. We found evidence of a unique voice rising from the health majors at Worcester State in response to the china scenario. We suspect there may be similar voices rising from those directly affected by any scenario. If this is the case it would make sense to expand the study to capture those voices especially for the scenarios considered to be most likely.

Conclusion

Like many emerging technologies Robotics can be expected to cause rapid changes in the way our society operates. The implications of these changes and the unintended side effects of the technology can cause a cascade of changes that is difficult to predict. Some experts including the P.W. Singer suggest that robotics technology is heading towards a point beyond which all the rules change and forecasting is impossible known as a singularity. This modified Delphi study was designed to provide a better methodology for forecasting the socio-technical implications in a period where all the major trends may be changing.

The perceptions of about 200 college students are investigated. Studying in robotics and other technical and non-technical fields, these students represent the voices of three professional groups expected to take part in the robotics debate. Four scenarios describing the future of robotics technology were assessed on the basis of likelihood, desirability, and ethical implications. It is not important whether the future of robotics technology is in principle predictable, instead we look for consensus about were robotics is going amongst students in different fields. If consensus exists on some direction we consider the perceived desirable of going in that direction. Such a consensus poses a significant sociopolitical issue whether or not it proves to be accurate if it results in an attempt to control the direction of the field or resist undesirable outcomes. Additionally we seek to understand how the goals of the of the institution driving technology affect its development and the perceived implications for the society it is introduced into. Each scenario posits a different institution driving the development of robotics and each takes a different ethical stance on Isaac Asimov's 3 Laws of Robotics.

Our findings reveal a high degree of consensus among students in the field of robotics and in Technical and Non-Technical fields that our military scenario is significantly more likely than any other scenario. Additionally there is tremendous consensus across all strata that advanced military robotics does not improve the quality of life and raises many ethical concerns. Asimov's fictitious laws proved to be a very good indicator of the ethical concerns that our participants perceived in each scenario. Our analysis revealed that both the desirability and ethical concerns accurately represented the degree of which Asimov's laws were followed in the four scenarios. Aspiring experts in robotics perceived even more issues with this scenario and still rated it most likely.

All this led us to conclude then the overwhelming percentage of funding currently devoted to military applications is controlling the direction of the field in spite of the growing concerns of the engineers developing the technology. If these perceptions do represent the voices that take part in the robotics debate, then we have identified a point of contention between the current trends and the direction in which the future robotics experts and their technical and non technical peers believe it should go. The stage is set for further studies into the social political implications of this consensus.

With this in mind, we believe that studies like this are important to pierce the future fog and help guide the development of robotics engineering. The implication of technology forecasts such as these is that technological development can and should be guided by moral values and careful thought.

Bibliography

Asimov, I. (1963). *I, Robot.* Garden City, NY: Doubleday.
Munroe, R. (n.d.). *Abstraction.* Retrieved 6 1, 2012, from XKCD: http://xkcd.com/676/
NASA. (2011, 5 1). *All-Terrain Intelligent Robot Braves Battlefront to Save Lives.* Retrieved 6 1, 2012, from NASA Spinoff: http://spinoff.nasa.gov/Spinoff2005/t_4.html
Singer, P. (2009). *Wired for War.* New York: Penguin Press.
Yamauchi, B. (n.d.). Packbot: A Versatile Platform for Military Robotics. Burlington, MA.

Appendix A – Scenarios

NASA Races to Mine Lunar Resources

In 2030 various nations begin to compete for valuable resources on the moon. NASA contracts LunaCorp, a promising start-up, to develop technology for a lunar extraction industry. LunaCorp specializes in advanced robotics and telecommunications and competes against ArianeSpace and the China Great Wall Corporation. Lunacorp's vision is to use semi-autonomous robots to mine the lunar surface and to support a small number of on-site human operators, thus maintaining the majority of the human workforce Earth-side. This allows for a safer working environment and saves the expense of setting up a large lunar colony. By 2050 LunaCorp's mission control center in Massachusetts supports 950 operators working in three shifts, overseeing 50 people and 300 robots operating on the moon. Today in 2069, LunaCorp is expanding their mining operation by building more small colonies across the lunar surface. These new colonies will eventually allow LunaCorp to harvest resources from an area the size of North America--half the size of the moon.

LunaCorp's highly flexible, modular robotic fleet consists of four classes of robots: miners that collect oxygen, iron, titanium, silicon and helium from the regolith found near the base; hunter-gatherers that seek out rarer and more valuable gas and mineral resources like hydrogen, aluminum, chromium, platinum, calcium, and nitrogen; worker bots that aid in the construction of bases and shelters for human operators and transport liquid oxygen for delivery to low earth orbit; and an assembly robot that functions as a "queen bee," using materials gathered by other robots to manufacture new units. Only highly sophisticated electronics need be imported from Earth. LunaCorp is using this current setup as a stepping-stone to create increasingly autonomous robots. The goal is for one operator to control several units simultaneously, thereby increasing the size of the lunar robotic fleet and freeing up the Earth-side workforce to focus on further research and development.

If a company could be said to have a patron saint, LunaCorp's would be Isaac Asimov. LunaCorp researches exactly how to implement his three laws[1] regarding human/computer interaction. This is particularly necessary for the companion robots that live in close proximity to humans, operating greenhouses and completing routine tasks on the lunar base. This is also crucial for the training robots used by 5th-8th graders on the lunar base: each class gets one robot that they learn how to operate. Preparations are also afoot for the inevitable meeting of LunaCorp and China Great Wall robots. These robots eventually will compete for resources and might inadvertently prevent one another from completing their missions—a potentially hostile act.

Naval Robotic Technology Used to Colonize the Seas as Rising Sea Levels Engulf US Coastlines

As global warming continues, rising sea levels threaten to engulf more coastal cities. The inhabitants of these cities must either relocate inland and compete for scarce terrestrial resources or expand into the ocean and learn to support themselves with resources from it. The Marshal Islands, New Orleans, and Miami have chosen the latter solution and are developing technologies that allow humans to live in, on, and under the oceans. In a sense humans are reversing evolution and returning to the seas. In California, Louisiana, and Florida companies like WaterWorld and Atlantis compete to develop these sites.

The Marshallese have a specific problem: the sea is becoming more acidic as it rises, so much so that it is dissolving the coral reefs on which their islands rest. They have recruited the help of Woods Hole Oceanographic Institute to "float" their islands. Woods Hole has access to the US Navy's deep sea robotic technology used to locate the Titanic—this includes robots that can function on the ocean floor, far beyond the depth at which manned craft can survive. The Marshallese have also encountered food shortages. Since they are already working with deep-sea robotics technology in order to float their islands, they decide to enlist the aid of the robotics corporation Atlantis to create an agricultural economy in the ocean. Atlantis builds the Marshallese a new type of robot, a shark bio-mimic, designed to apparently replace sharks as the most dangerous fish in seas. Unlike sharks, though, these hunter-farmer robots use electrical power from charging stations scattered along the coasts rather than biological fuel from consuming fish, allowing these robots to harvest plant life and farm schools of fish in order to provide food for humans--effectively placing humans at the top of the oceanic food chain. These robots function almost as a new species, yet they follow Asimov's laws[2] and exist purely to re-purpose the ocean's resources for human consumption.

Louisiana and Florida, on the other hand, have enlisted the services of WaterWorld to extend their cities under the waves. WaterWorld has adapted old Navy technology to develop specialized construction and service robots that can work in shallow coastal waters. These robots develop the ocean floor, leveling it to make way for other robots that construct underwater cities. These construction robots work in teams, largely autonomous but programmed by humans on land. After a city is established, a second class of purely autonomous service robots take over the maintenance of the city and ensure the safety of its inhabitants. These robots are designed to be harmless to humans and they, too, are consistent with Asimov's laws.

Now, two types of robots rule the seas: the first, Atlantis' hunter-farmer robots that balance the oceanic ecosystem while harvesting fish and plant life for human consumption; and the second, WaterWorld's construction and maintenance robots that spend their hours looking for new places to develop underwater cities while maintaining the ones that already exist. While both these robots are true to the letter of Asimov's laws, they are responsible for the balance of an entire ecosystem--one humans depend on.

Robotic Care-Takers Bridge Gap between Humans and Machines

In 1978, China introduced a law that limited families to only one child in order to slow population growth. With this law in effect, the youth of China were torn by the Chinese culture which honors both professional achievement and familial values. In years prior, siblings shared the responsibility of tending to their aging parents while pursuing their professional ambitions. With the passing of the one-child law, however, unprecedented stresses were placed on Chinese children. In many cases, these children worked diligently to obtain a suitable career, generally in the field of science and technology, but had to relocate to cities in order to pursue their chosen profession. Unfortunately, this left the children's aging parents languish at home without care, an incredible dishonor in Chinese society. On the other hand, children who stayed to care for and assist their parents dishonored themselves by failing to supply money or gain prestige for their family.

In 2015, a small group of companies developed by unemployed college students began to work on a robotic solution to the problem. Within a few years, they founded Ant Farm, a corporation that developed semi-autonomous robotic surrogates that could be used to assist aging parents. These care-taking robots aided elderly parents by completing physically demanding house-hold labor and providing young adults with the ability to tend to their parents while still pursuing their own professional careers. These robots performed only simple functions on their own, primarily acting to notify their owners when parents needed assistance or extra care. In addition, these robots were designed to follow Asimov's laws very strictly, acting as simple-minded aids to elders. They followed orders from the elders and the youth without question, yet were incapable of many tasks more complex than sweeping floors.

In the current year, 2052, Auntie, the successor of Ant Farm, has grown tremendously. Robot surrogates have become extremely popular, allowing children to leave home and pursue college and professional careers without dishonoring their parents. Thousands of young people in China have entrusted the care of their elders to these robots. The robots themselves are now almost entirely autonomous. They are humanoid in form and have been designed to help seniors clean houses, cook meals, and maintain their lifestyle without the need of human care-takers. The robots are even capable of assisting seniors as they age and need to be bathed and constantly watched over. The young owners of the robots can check in and view status updates on their parents conditions at any time, but trust the robot to notify them if any accidents befall their parents. Demand is rising for robot care-takers with the ability to watch over parents after they have developed dementia and cannot be trusted to issue meaningful orders. The robots artificial intelligence software is incredibly advanced, differentiating between commands that are harmless and should be followed and commands that may endanger human life and should be ignored.

In order to expand its business to foreign markets, Auntie plans to market a nanny-robot to Americans. Anticipation for the release of Auntie's nanny-robot has peaked. Auntie has recently made claims that its newest line of nanny-robots will have the most advanced artificial intelligence ever seen. Auntie has advertised that its nanny-robot may in fact be a safer, more capable baby-sitter than an actual human being.

Certain scientists have begun to question the wisdom of placing these robots in positions of authority over humans. Some have noted that Auntie's robots now break Asimov's second law, allowing robots to effectively manipulate the very young and very old against their will in certain situations.

Robotic Warriors Revolutionize the Military

In 1998, Vice Admiral Aurthur Cebrowski, president of the Naval War College in Newport, predicted that the military was about to undergo a massive paradigm shift. He argued that this revolution of military affairs (RMA) would come about as new technology made near-instantaneous communication a reality. Secure, instant data transmission would clear away the fog of war and allow the generals and politicians to actually understand what was happening on the battlefield and make decisions in real-time.

At first, this new technology produced too much information. No human can sort through live video feeds from thousands of sources, understand it, and turn it into useful data that can be presented to commanders in real-time. Field commanders then received conflicting instructions from different levels of command. This new network-centric method of warfare could not alleviate the need for quick, intelligent decisions or allow people to process information faster.

While the implications of surveillance and instant communication were overestimated, the military vastly underestimated the importance of robots. When simple robots were introduced to the battlefield in 2005, soldiers used them in many ways they weren't intended. Many soldiers also became emotionally attached to their robots, demanding that their specific unit be repaired rather than accepting a replacement. Upon observing how dramatically simple robots changed warfare, the military invested heavily in the development of robotics.

Now, the year is 2050. Surveillance drones fly high above battle fields, keeping watch over troops, convoys, and bases. They collect images and videos that are transmitted directly to immensely powerful computers. These use advanced algorithms that process the data in real time, detect threats, and propose action plans to commanders. Computers and artificial intelligence entities have realized Cebrowski's dream of total battlefield awareness. Commanders are presented with everything they need to direct their troops, complete their missions, and minimize loss of life-- assuming the battalion is actually human. The military has also replaced most human troops with robots. Putting robots on the battlefield takes humans out of harm's way. Military combat robots are also stronger, faster, and more powerful than any human soldier. They can take direct fire from most weapons and are easily repairable when damaged, thus robot soldiers can take greater risks. They can afford to confirm targets as hostile before attacking. They can even use non-lethal weapons and capture enemies alive. This combination of robotic intelligence gathering and robotic warriors keeps humans out of harm's way, allowing for faster, more accurate decisions while protecting human lives.

These military robots, though, are designed for destruction. The military views Asimov's [3]three laws of robotics as short-sighted and ill-advised. Robots are programmed and must follow their instructions. Robots are immensely powerful weapons when allowed to kill, refuse instructions from outside the military chain of command, and engage in dangerous situations. In fact, they are so powerful that they may take more risks, reducing civilian casualties and collateral damage.

A significant question, though, is how opponents will respond when faced with robots instead of humans. When archery and later gunpowder were introduced to Europe, soldiers could kill from a distance instead of fighting hand-to-hand. These soldiers were perceived as dishonorable cowards rather than heroes, to be casually killed when an honorable knight would be held for ransom. Robotic soldiers raise this question again, except that the operator is a continent away. When a country is bombed, targets have little chance for survival. Robots, though, can seek out survivors. Humans with conventional weapons have little to no chance against the robots. Will this revolution of military affairs bring about a new age in protecting lives, where robots will take fire before firing and avoid collateral damage, or will it hail a new low in barbarism, where robots bring about an era of mechanized manslaughter?

Appendix B – Survey

How likely is it that this scenario could come about?

Unlikely Somewhat Unlikely Somewhat Likely Very Likely

If the scenario came about, the resulting technological development would have many applications, attract investment, and generate lots of spinoff technologies that would significantly advance the field of robotics.

Strongly Disagree Disagree Agree Strongly Agree

If the scenario came about, the resulting technological development would have what affect on the quality of life of the average person?

Undesirable Somewhat Undesirable Somewhat Desirable Desirable

If the scenario came about, the resulting technological development would have what affect on the man machine relationship?

Undesirable Somewhat Undesirable Somewhat Desirable Desirable

What does this scenario's relation to Asimov's laws say about the wisdom of pursuing technological development in this manner?

Undesirable Somewhat Undesirable Somewhat Desirable Desirable

Please comment on the scenario. (if you had any trouble with the questions above, please note it here as well)