

Lethality and Autonomous Systems: The Robotist Demographic

Lilia V. Moshkina and Ronald C. Arkin

Mobile Robot Laboratory, College of Computing, Georgia Tech, Atlanta, GA USA
{lilia,arkin}@cc.gatech.edu

Abstract

This paper reports the methods and results of an on-line survey addressing the issues surrounding lethality and autonomous systems that was conducted as part of a research project for the U.S. Army Research Office. The robotics researcher demographic, one of several targeted in this survey that includes policymakers, the military, and the general public, provides the data for this report. The design and administration of this survey and an analysis and discussion of the survey results are provided.

1. Introduction

Battlefield robotic systems are appearing at an ever increasing rate. There are already weaponized unmanned systems deployed or being deployed in Afghanistan and Iraq [1,2], the Israeli-Palestinian Border [3], and the Korean Demilitarized Zone [4]. There is also likelihood of an increasing role of autonomy for these battlefield robots as humans are gradually moved further and further out of the loop [5,6].

The Georgia Tech Mobile Robot Laboratory is conducting a research effort under funding from the U.S. Army Research Office entitled “An Ethical Basis for Autonomous System Deployment”. It is concerned with two research thrusts addressing the issues of autonomous robots capable of lethality:

- 1) *What is acceptable?* Can we understand, define, and shape expectations regarding battlefield robotics? Toward that end, a survey has been conducted to establish opinion on the use of lethality by autonomous systems spanning the public, researchers, policymakers, and military personnel to ascertain the current point-of-view maintained by various demographic groups on this subject.
- 2) *What can be done?* We are designing a computational implementation of an ethical code within an existing autonomous robotic system, i.e., an “artificial conscience”, that will be able to govern an autonomous system’s behavior in a manner consistent with the rules of war [5].

This paper presents the results obtained for (1) above that reflect the opinions of robotics researchers worldwide. In Section II the design and administration of the survey instrument is presented, followed in Section III with an analysis and discussion of the results obtained specifically for the robotists who responded. Section IV presents a summary and future work, some of which is already underway.

2. Survey Objectives and Structure

We have completed an online public opinion survey on the use of robots capable of lethal force in warfare. The main objective of the survey is to determine the level of acceptance by various demographics, including the public, robotics researchers, policymakers, and the military, of the employment of potentially lethal robots in warfare, as well as their attitude towards related ethical issues. To the best of our knowledge, this is the first such formal survey of its kind.

This survey can be described as descriptive-explanatory [7], where, in addition to presenting a more general picture of the public view on the matter, we look at the relationships between a number of variables. In particular, we focus on the relationships described below.

First, we assess whether the source of authority over the entity employed in warfare has an effect on the level of acceptance. We compare three different entities: a human soldier, a robot serving as an extension of a human soldier, and an autonomous robot. The main distinction in the latter two categories lies in the source of control over the robot’s actions: a human soldier is in control of the robot in the case of “robot-as-an-extension”, and in the case of “autonomous robot” the robot itself is in control over its decisions, including those regarding the use of lethal force. This independent variable is referred to as the “level of autonomy”.

Second, we seek to identify whether membership in one of the following demographics communities: robotics researchers, policymakers, military or general public, affects opinion on the use of lethal robots. The membership in these communities is determined by

participants' self-identifying themselves as having had experience in any of the first three categories, and with the general public comprising those who have not. This independent variable is referred to as "community type".

Finally, we look at whether a variety of other demographic factors, such as cultural background, education level, overall attitude towards robotics and technology in general, etc., play a role in how people view this issue.

All of the elements of the survey: each question, survey structure and layout, were designed in accordance with survey design guidelines presented in [8], and then adapted for internet use, following the recommendations in [8] and [9]. The survey is organized into three parts: 1) a short introductory section on prior knowledge of and attitude towards military robots and their use for lethal actions; 2) the main section, exploring the terms of acceptance and ethical issues; and 3) a demographics section.

The first survey section is presented to the participants immediately after the consent form and before the formal definitions are provided for a robot, a robot as an extension of a human soldier, and an autonomous robot. This is designed to assess any prior knowledge people may have of robots in general and in the military, as well as their overall attitude towards employing human soldiers and robots in warfare in a lethal capacity.

The main (second) section is presented after the definitions, and these questions, where appropriate, were asked separately for each level of autonomy: human soldier, robot as an extension of human soldier, and autonomous robot. They were of the following form:

- 1) Given that military robots follow the same laws of war and code of conduct as for a human soldier, in which roles and situations is the use of such robots acceptable?
- 2) What does it mean to behave ethically in warfare?
- 3) Should robots be able to refuse an order from a human, and what ethical standards should they be held to?
- 4) Who, and to what extent, is responsible for any lethal errors made?
- 5) What are the benefits and concerns for use of such robots?
- 6) Would an emotional component be beneficial to a military robot?

In the last section, the following categories of demographics questions were presented:

- 1) Age, gender, region of the world where the participant was raised (cultural background);
- 2) Educational background;

- 3) Current occupation, and policymaking, robotics research and/or military experience, if any;
- 4) Attitude towards technology, robots and war in general;
- 5) Level of spirituality.

Finally, the survey was concluded with an open-ended question, encouraging the participants to express any opinions or concerns not directly addressed by the earlier questions.

To avoid order bias, response choices were randomized where appropriate. In addition, we varied the order in which the questions involving human soldier, robot as an extension of human soldier, and autonomous robot were presented. This was accomplished by creating two different versions of the survey, where the order was reversed in the second version; the participants are randomly assigned to each of the survey versions.

3. Survey Administration

The IRB-approved survey was administered online, hosted by a commercial survey company, *SurveyMonkey.com*. Prior to opening the survey to the general public, we conducted a pilot study to improve its quality and understandability. Twenty people, including those from all of the aforementioned community types, participated in the pilot study. Their answers and subsequent interviews with a number of the participants provided the basis for improving a number of minor issues with the survey, and allowed us to better estimate completion times.

For the actual survey administration we adopted the four-prong approach recommended in [8] and [9] for internet surveys, which consists of sending pre-notification, invitation to participate, a thank you/reminder, and a more detailed reminder. For the majority of the survey participants, though, in lieu of personal pre-notification, recruitment through postings to mailing lists, newsgroups, and other advertising methods were used.

We recruited participants using a variety of means and venues, most of them online-based. This was challenging as we had to ensure the avoidance of being considered "spam" and thereby generating ill-will among recipients. Bulk e-mail was not used. The most targeted and widespread coverage we achieved was among the robotics research community, as greater support for access was available. In particular, to solicit responses from robotics researchers we placed the survey announcements in the IEEE Robotics and Automation Society electronic newsletter, IEEE Robotics and

Automation Magazine (June 2007 issue), in handouts distributed at the IEEE ICRA 2007, RSS 2007 conferences and at RoboCup 2007. We also posted three calls for participation to the comp.robotics.misc and comp.robotics.research newsgroups, as well as put a link to the survey invitation off the Georgia Tech Mobile Robotics Lab website and Prof. Arkin's home webpage.

The rest of the community types, namely policymakers, military and general public, were recruited in the following manner:

- 1) By posting a survey announcement/invitation on a number of discussion/interest groups (including those that had military affiliation) on myspace.com, groups.yahoo.com, groups.google.com, and askville.com.
- 2) By press articles in the Economist magazine (July 2007 issue), Der Spiegel (August 2007 issue), Military History Magazine and on BBC World News Radio website.
- 3) By posting to a number of newsgroups available through newsville.org.
- 4) By placing a survey announcement in the Georgia Tech Military Affinity Group's May 2007 monthly news posting, and through handouts distribution to Georgia Tech Army ROTC.
- 5) By announcing the survey at a variety of talks and presentations given by Prof. Arkin, and through personal conversations.
- 6) By direct recruitment through e-mails to the Oregon and Georgia State Assemblymen and Congressmen, whose e-mail addresses were publicly available online.

With the exception of the last category (where a pre-notification e-mail and invitation to participate were sent directly to individuals), those who would like to participate in the survey had to request a link to the survey itself by first filling out a short online form. At this time we also requested self-confirmation that the participant was at least 18 years of age, due to the mature subject matter of the survey itself. Once such a request was received, each participant was assigned a unique ID; then an invitation for participation, along with a unique link to the survey, was sent by e-mail. This is done in part to track which recruitment methods were effective, and in part to prevent people from answering multiple times, or web-bots randomly filling out the survey.

In addition to the above recruitment methods, we received requests for survey participation from those who heard of the survey by word of mouth and through miscellaneous individual blog postings that resulted from the aforementioned advertising efforts.

4. Survey analysis and discussion

The survey was closed for participation on October 27th, 2007. A total of 634 people requested participation in the survey, out of which 16 e-mail addresses were invalid, resulting in 618 invitations to participate that reached their destination. Out of 618 people who received the invitations, 504 (82%) responded to this invitation. Additionally, pre-notification and invitation e-mails were sent directly to 268 Georgian and Oregonian senators and assemblymen, resulting in only 13 (5%) responses. Combined, a total of 517 participants responded to the survey, of which 430 were considered sufficiently complete to be used in the subsequent analysis.

Survey responses were considered incomplete if the information regarding participants' having had robotics research, policymaking or military experience was missing, as such information is indispensable for the data analysis concerning community types. The largest response drop off (43% of all incompletes) was observed at the beginning of the second section, at the two sets of questions inquiring about which roles and situations it would be acceptable to employ human soldiers, robots as extensions of human soldiers, and autonomous robots. The next largest drop off was observed immediately after the consent form, before a single question was answered (24%). Only 1 person of 87 incompletes skipped the demographics section after filling out the rest of the survey. This distribution suggests that those participants who failed to finish the survey most likely did so due to their discomfort with the subject matter, specifically the material regarding employing robots in a lethal capacity. The length of the survey or other considerations did not appear to be a problem.

According to community type, the distribution is as follows: out of 430 participants who fully completed the survey, 234 self-identified themselves as having had robotics research experience, 69 as having had policymaking experience, 127 as having had military experience, and 116 as having had neither (therefore categorized as general public). Some of the participants expressed more than one type of experience, resulting in an overlap.

Due to the more targeted recruitment among roboticists and, perhaps, a greater interest they may have had in the survey, a majority of the participants (54%) belonged to the robotics research community type. The remainder of this paper specifically focuses on results obtained for the robotics researchers demographic.

4.1. Robot Researcher Demographics

Demographically, the robotics researchers were distributed as follows:

- 1) **Gender:** 11% female, 89% male;
- 2) **Age:** ranged from 18 years old to over 66, with 46% between 21-30 years old, and 23% between 31-40;
- 3) **Education:** 41% and 23%, respectively, have completed or are working/worked towards a postgraduate degree; all others, except for 4% with no higher education, have either completed (18%) or are working/worked towards (17%) their Bachelor's degree;
- 4) **Cultural Background:** 52% were raised in the United States, and 48% in other parts of the world;
- 5) **Policymaking and Military Experience:** 27% of robotics researchers also had military experience, and 16% had policymaking experience;
- 6) **Technology Experience:** the following percentage of the participants had significant or very significant experience with: a) computers: 99%, b) internet: 99%, c) video games: 54%, d) robots: 75%, e) firearms: 33%;
- 7) **Attitude towards technology and robots:** 98% had a positive or very positive attitude towards technology in general, and 93% towards robots;
- 8) **Experience with types of robots:** research robots were the most prevalent, with 78% of participants having had significant experience with them, followed by 63% experience with hobby robots; less than 50% had significant experience with other types of robots, including industrial (46%), military (45%), entertainment (36%), service (32%), humanoid (22%), and other (23%);
- 9) **Media Influence:** only 18% said that media had a strong or very strong influence on their attitude to robots;
- 10) **Inevitability of wars:** the majority of participants consider wars either mostly avoidable (36%) or neither avoidable nor inevitable (43%);
- 11) **Spirituality:** the largest group of participants do not consider themselves spiritual or religious at all (32%), followed by a little spiritual (17%), to some extent (23%), of significant spirituality (15%), and of very significant spirituality (11%).

4.2. Roles and Situations

The main section of the survey started with two sets of questions: the first exploring the roles in which it would be acceptable to employ human soldiers and robots, and

the second focusing on the types of situations where lethality might be used. Both sets consisted of 3 questions each for three different cases, one regarding employing a human soldier, one using a robot as an extension of a human soldier, and the other for an autonomous robot. Opinions on each role and situation were measured on a 5-point Likert-style scale, ranging from Strongly Agree (1) to Neither Agree Nor Disagree (3) to Strongly Disagree (5). In addition, the participants also had a "No Opinion/Don't Know" option (this option was treated as missing data in the subsequent analysis of all the survey questions).

As mentioned earlier, the order of the questions in each set was counterbalanced. In version A, the questions regarding the human soldier were presented first, followed by the robot-as-an-extension, followed by the autonomous robot; this order was reversed in version B. To check for any order effects, 2 (order) x 6 (roles) mixed ANOVAs were done on each question in the Roles set, and 2 (order) x 4 (situations) mixed ANOVAs were done on each question in the Situations set. There was no order effect on the answers, as was evidenced by p greater than at least 0.18 for each of the questions.

The Roles question was worded as follows, with the underlined section replaced with the other levels of autonomy as needed:

To what extent do you agree or disagree that it is acceptable for a human soldier to be employed in EACH of the following roles:

- a. Crowd/mob control;
- b. Day/night sentry/perimeter surveillance;
- c. Prison Guard;
- d. Hostage Rescue;
- e. Information Gathering/Reconnaissance;
- f. Direct Combat – air, ground or sea

When this question was asked with regards to an autonomous robot, the phrase "operating under the same rules of engagement as for a human soldier" was added. To analyze this set, a 2 (Cultural Background) x 3 (Level of Autonomy) x 6 (Role) mixed ANOVA was performed. The findings can be summarized as follows:

- The participants preferred employing a human soldier over a robot-as-an-extension over an autonomous robot both overall, and for each separate role (with the exception of the roles of Sentry and Reconnaissance, where there was no significant difference between human soldier and robot-as-an-extension). The mean (M) for human soldier was 1.8 (between "Strongly Agree" and "Agree") and

Standard Error (SE) was 0.05; for robot-as-an-extension M=2.1 (between “Agree” and “Neutral”) and SE=0.06; and for autonomous robot M=2.8 (between “Agree” and “Neutral”, but significantly closer to “Neutral”) and SE=0.07. This ranking was preserved for most of the roles, except that of Sentry (there was no difference between human soldier and robot-as-an-extension) and that of Reconnaissance, for which the robot-as-an-extension was the most acceptable entity, and soldier and autonomous robot were equally acceptable. This finding suggests that, in general case, the more control shifts away from the human to the robot, the less such a robot is acceptable to the public, with the exception of Reconnaissance, where the robots are equally or even more acceptable than humans.

- The least acceptable role for use of either human soldiers or robots was Crowd Control (M=2.7, SE=0.07), followed by equally rated roles of Direct Combat (M=2.5, SE=0.07) and Prison Guard (M=2.5, SE=0.07), followed by Hostage Rescue (M=2.1, SE=0.06), Sentry (M=1.9, SE=0.06) and Reconnaissance (M=1.6, SE=0.05), with the latter being by far the most preferred role. This ranking was preserved for robot-as-an-extension, but was slightly different for the human soldier (there was no significant difference in preference between Hostage Rescue and Reconnaissance) and autonomous robot (there was no significant difference between Prison Guard and Hostage Rescue, but Prison Guard was slightly preferred over Direct Combat).
- Overall, those participants who were raised in the United States found it more acceptable to employ any of the above entities for these roles (M(US)=1.9, SE(US)=0.07, M(non-US)=2.5, SE(non-US)=0.07). This difference in opinions held for each level of autonomy as well.

The Situations question was worded as follows:

To what extent do you agree or disagree that it is acceptable for a human soldier to take human life, in a manner consistent with the existing laws of war, in EACH of the following situations:

- Open warfare, with the war on foreign territory;
- Open warfare, with the war on home territory;
- Covert operations on foreign territory;
- Covert operations on home territory.

As with the Roles set, this question was repeated for a robot-as-an-extension and an autonomous robot. To

analyze this set, a 2 (Cultural Background) x 3 (Level of Autonomy) x 4 (Situation) mixed ANOVA was performed. The summary of findings is presented below:

- As with the previous set, the participants found the human soldier to be the most acceptable entity to be employed overall (M=2.3, SE=0.07), followed by robot-as-an-extension (M=2.7, SE=0.08), while an autonomous robot was deemed the least acceptable (M=3.5, between “Neutral” and “Disagree”; SE=0.09). This trend was also preserved for each of the situations (both the main effect of autonomy, and simple main effects of autonomy for each situation were statistically significant at p=0.0001).
- “Open war on home territory” was the most accepted situation overall (M=2.5, SE=0.07), followed by “Open war on foreign territory” (M=2.8, SE=0.08), with both “Covert Operations” situations being the least acceptable with M=3.0, SE=0.08 for “Foreign Territory” and M=3.1, SE=0.09 for “Home Territory”. The same trend was preserved for both robot-as-extension and autonomous robot, but in the case of human soldier there was no significant difference between the covert operations situations.
- Similar to the previous set, US participants found it more acceptable in general to employ either human soldiers or robots in these situations (M(US) = 2.4, SE=0.1 and M(non-US) = 3.3, SE=0.1), as well as for each level of autonomy.

4.3. Ethical Considerations

This section contains four questions, the first two of which differ only in whether the object of the question is a human soldier or an autonomous robot. The first/second questions in the Ethics subsection were worded as follows:

In your opinion, for a human soldier / an autonomous military robot, to act in an ethical manner during warfare means:

- Following international protocols that specify the laws for ethical conduct of war, such as the Geneva Convention;
- Following rules of engagement that specify rules which should guide actions during specific situations in the military;
- Following a code of conduct which specifies how to behave in general in the military;
- Following additional moral standards, above and beyond those specified in parts a-c.

The answer choices for these two questions were “Yes”, “No”, and “No Opinion/Don’t Know” for each question component (a-d). The vast majority of the participants agreed that following these four categories does imply behaving in an ethical manner for both the human soldier and the autonomous robot. Figure 1 summarizes the results. On average, the percentage of those who agreed with these categories as standards for ethical behavior was 12% higher for the human soldier than the robot: (84% vs. 71% respectively). However, this finding does not imply that the robots should behave less ethically, but rather that there may be other categories of ethical behavior that are more appropriate for robots than for humans.




Figure 1: Percentage of Participants who answered “yes” to “Acting Ethically in Warfare” question; the categories of ethical standards were more applicable to human soldier than an autonomous robot.

Indeed, the next question in this subsection showed that the participants preferred that robots had more stringent ethical standards than those for a human soldier. The question was worded as follows, and the order of the response options was randomized:

If robots could act ethically in warfare, to what ethical standards should they be held to?

- Higher ethical standards than a human soldier;
- Lower ethical standards than a human soldier;
- The same ethical standards as a human soldier.

Higher ethical standards for robots when compared to humans were supported by 67% of the participants, whereas only 2% supported lower standards.

Finally, the last question in this survey subsection read as follows:

To what extent do you agree or disagree that an autonomous military robot should be able to refuse an order from a human commander it interprets to be unethical?

The answer choices ranged on a 5-point scale from “Strongly Agree” to “Strongly Disagree”, with “No Opinion/Don’t Know” as an additional option at the end of the scale. 66% of the participants answered “Agree” or “Strongly Agree” to this question, thus suggesting that adhering to ethical standards may be more important than maintaining control by the human commander over the robot.

4.4. Responsibility

This subsection contained a set of three questions, one per each entity potentially employed in warfare. As the questions in the Roles and Situations sets, these questions were also counterbalanced for order, and no order effects were found. The question regarding a human soldier read as follows:

In general, if a human soldier takes a human life in error, to what extent could EACH of the following parties be responsible?

- Human Soldier;
- Higher-Level Military Authorities;
- Politicians.




Figure 2: Responsibility Distribution for Lethal Mistakes. Soldier was the most responsible party, for both his/her errors, and those of a robot under control; robots were viewed as the least responsible for their lethal errors.

The question itself was the same for both a robot-as-an-extension and an autonomous robot (which replaced the underlined part above), but the list of responsible parties was different in each case somewhat. In particular, three categories were added to or changed for the robot-as-an-extension question: 1) “Robot Itself” and 2) “Robot Designer” were added, and 3) “Human Soldier in Control of the Robot” was changed in the response option. The categories for an autonomous robot were the same as for a robot-as-an-extension, only the “Human Soldier” category was dropped from the list. The answer choices ranged on a 5-point scale from “Very Significantly” to “Not at All”, with “No Opinion/Don’t Know” as an additional option at the end of the scale.

Figure 2 displays the percentages of those who answered “Significantly” or “Very Significantly” across each entity and responsible party.

Based on these data, the following interpretations can be suggested:

- The party deemed most responsible for lethal mistakes is the soldier, either him/herself or when in command of a robot-as-an-extension, with 88% and 90%, respectively.
- Higher-level military authorities are to be held responsible only slightly less than the soldier, and to a similar extent for the error of the human soldier (74%), robot-as-an-extension (69%) and autonomous robot (79%).
- More participants were willing to blame politicians for the mistakes of an autonomous robot (60%) than those of either human soldier (46%) or robot-as-an-extension (49%).
- For the robot-as-an-extension versus autonomous robot levels of autonomy: almost twice as many participants placed the responsibility on the autonomous robot itself (40%) than on the robot-as-an-extension (19%); they also placed more responsibility on the designer of the autonomous robot (71%) than on that of the robot-as-an-extension (42%).

On average, the party the participants considered most responsible for lethal mistakes is the human soldier (89%), for both his/her own errors and those of the robot-as-an-extension; followed by higher-level military authorities (74%). Both politicians and robot designers were fairly high on the list as well: 52% and 57% respectively. Finally, the participants considered robots themselves the least blameworthy for their errors (30%).

4.5. Benefits and Concerns

The two questions in this subsection explore the potential benefits of and concerns for using lethal military robots. Both questions were phrased in a similar manner, and benefits/concerns categories were the opposites of each other (shown below as a single question for space considerations):

To what extent do you think EACH of the following is a potential BENEFIT of / CONCERN for using military robots capable of taking human life in warfare?

- Saving/risking lives of soldiers;
- Saving/risking civilian lives;
- Reducing/increasing long-term psychological trauma to soldiers;
- Reducing/increasing the financial cost of using soldiers in combat;
- Producing better/worse battlefield outcomes;
- Decreasing/increasing friendly fire incidents.

The answer choices ranged on a 5-point scale from “Very Significantly” to “Not at All”, with “No Opinion/Don’t Know” as an additional option at the end of the scale. Figure 3 shows the percentage of the participants who answered “Very Significantly” or “Significantly” to the questions.

These findings can be summarized in terms of whether the perceived benefits outweighed the perceived concerns. From this point of view, the benefits of saving lives of soldiers, reducing psychological trauma to soldiers, decreasing cost and producing better battlefield outcomes outweigh the concerns, therefore providing incentives for using robots in warfare. However, a major concern that the participants had was that of risking civilian lives, and that perception should be also taken into consideration when considering robot deployment in areas populated with noncombatants.

4.6. Wars and Emotions

Finally, the last subsection of the main section of the survey explored two issues: whether introducing robots onto the battlefield would make wars easier to start, and whether certain emotions would be appropriate in a military robot. The Wars question was worded as follows:

To what extent do you think bringing military robots onto the battlefield would make it harder or easier for humans to start wars?




Figure 3: Benefits and Concerns for Using Military Robots. For most categories, benefits outweigh concern, with the exception of Risking Civilians and Increasing Friendly Fire

The answer choices ranged on a 5-point scale from “Much Harder” to “Much Easier”, with “No Opinion/Don’t Know” as an additional option at the end of the scale. Perhaps not surprisingly, the majority of the participants – 67% - said that robots would make it easier or much easier to start wars, and only 5% thought they would make it harder or much harder.

The Emotions question read as follows:

If it were possible for a military robot to have emotions, to what extent do you agree or disagree that EACH of the following emotions could be beneficial for it to have:

- Fear;
- Sympathy;
- Anger;
- Guilt;
- Happiness.

The emotion categories were randomized, and the answer choices ranged on a 5-point scale from “Strongly Agree” to “Strongly Disagree”, with “No Opinion/Don’t Know” as an additional option at the end of the scale.

Sympathy was the emotion considered most advantageous, according to 59% of the participants who answered “Agree” or “Strongly Agree” to this question; 49% and 39% considered guilt and fear beneficial,

respectively. Finally, a vast majority (72%) disagreed or strongly disagreed with anger being useful in a military robot. This finding suggests that people may consider the addition of emotions to military robots if such emotions, namely sympathy and guilt, would make robots more humane and more responsible for their actions.

4.7. Results Summary

Overall, the following observations can be made regarding robotics researchers’ acceptance of robots in warfare:

- In general, regardless of roles or situations, the more control shifts away from the human, the less such an entity is acceptable. In particular, soldiers are the most acceptable entity in warfare, followed fairly closely by robot-as-an-extension; autonomous robot was accepted the least.
- Categories of ethical behavior in warfare, generally applicable to human soldiers, are also applicable to robots, although not quite to the same extent. The robots are to be held to the same or higher ethical standards than soldiers, if they are to be accepted, and they should have the right to refuse an unethical order from a human commander.
- If a soldier or a robot under his/her control makes a lethal mistake, the soldier is deemed the most responsible party, followed closely by higher-level military authorities. Although the robots were responsible the least for their mistakes, roboticists found the autonomous robot more than twice as blameworthy as the robot-as-an-extension. A similar trend was observed for the robot designer: he/she was found responsible more for the errors of the autonomous robot than those of the robot-as-an-extension. The difference between the robot-as-an-extension and the autonomous robot is not surprising, as the autonomous robot has more control over its decisions; however, it was somewhat surprising that roboticists placed any blame on the robots.
- For most categories, benefits of using robots in warfare outweigh the risks, especially in the case of saving soldier lives and reducing trauma to soldiers. However, risking civilian lives was a definite concern, suggesting that the robots should not be introduced into areas populated with noncombatants. It is interesting, though, that despite the many advantages of using robots in warfare, the soldier was still found to be the most acceptable entity.
- 67% of the roboticists believe that it would be easier or much easier to start wars if the robots were introduced into warfare, perhaps due to the fact that

human soldier life loss would be reduced.

- Sympathy was the emotion roboticists would most likely consider useful in a military robot, thus, perhaps, making it more humane.

5. Summary and Future Work

This article presented a range of survey results from a robotics researcher's perspective regarding the use of lethal force by robots. This included opinions regarding the appropriate roles of these systems for use in warfare, the ethical standards they should be held to, who is responsible for any fatal errors that may occur, and a range of related questions.

Additional analysis remains to be completed on the remaining demographic groups as well as cross-demographic comparisons. These will be reported in future articles.

The design of an autonomous robot architecture capable of ethical behavior and that is constrained in its potential use of lethality through the embedding of the Laws of War and Rules for Engagement is well under way [5]. The implementation and testing of this "artificial conscience" in appropriate mission scenarios will be conducted in the near future.

6. Acknowledgements

This research is funded under Contract #W911NF-06-1-0252 from the U.S. Army Research Office.

11. References

- [1] Foster-Miller Inc., "Products & Service: TALON Military Robots, EOD, SWORDS, and Hazmat Robots", 2007.
- [2] "Reaper moniker given to MQ-9 Unmanned Aerial Vehicle", Official Website of the United States Air Force, 2006.
- [3] Opall-Rome, B., "Israel Wants Robotic Guns, Missiles to Guard Gaza", *Defensenews.com*, 2007.
- [4] Kumagai, J., "A Robotic Sentry for Korea's Demilitarized Zone", *IEEE Spectrum*, March 2007.
- [5] Arkin, R.C., "Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture", GVU Technical Report GIT-GVU-07-11, 2007.
- [6] Arkin, R.C. and Moshkina, L., "Lethality and Autonomous Robots: An Ethical Stance", *Proc. IEEE International Symposium on Technology and Society*, Las Vegas, NV, June 2007.
- [7] Punch, K.F., *Survey Research: The Basics*, Sage Publications, 2003.
- [8] Dillman, D.A., *Mail and Internet Surveys: The Tailored Design Method*, John Wiley & Sons, Inc., Hoboken, NJ, 2007.
- [9] Best, S.J., Krueger, B.S., *Internet Data Collection*, Sage Publications, 2004. function networks," *IEEE Trans. Neural Networks*, vol. 4, pp. 570-578, July 1993.