

Automation and Accountability in Decision Support System Interface Design

M.L. Cummings¹

Massachusetts Institute of Technology

ABSTRACT

When the human element is introduced into decision support system design, entirely new layers of social and ethical issues emerge but are not always recognized as such. This paper discusses those ethical and social impact issues specific to decision support systems and highlights areas that interface designers should consider during design with an emphasis on military applications. Because of the inherent complexity of socio-technical systems, decision support systems are particularly vulnerable to certain potential ethical pitfalls that encompass automation and accountability issues. If computer systems diminish a user's sense of moral agency and responsibility, an erosion of accountability could result. In addition, these problems are exacerbated when an interface is perceived as a legitimate authority. I argue that when developing human computer interfaces for decision support systems that have the ability to harm people, the possibility exists that a moral buffer, a form of psychological distancing, is created which allows people to ethically distance themselves from their actions.

1. INTRODUCTION

Understanding the impact of ethical and social dimensions in design is a topic that is receiving increasing attention both in academia and in practice. Designers of decision support systems (DSS's) embedded in computer interfaces have a number of additional

¹ 77 Massachusetts Avenue, Room 33-305, Cambridge, MA 02139, MissyC@mit.edu, (617)252-1512, (617)-253-4196(Fax)

ethical responsibilities beyond those of designers who only interact with the mechanical or physical world. When the human element is introduced into decision and control processes, entirely new layers of social and ethical issues (to include moral responsibility) emerge but are not always recognized as such. Ethical and social impact issues can arise during all phases of design, and identifying and addressing these issues as early as possible can help the designer to both analyze the domain more comprehensively as well as suggest specific design guidance. This paper discusses those accountability issues specific to DSS's that result from introducing automation and highlight areas that interface designers should take into consideration.

If a DSS is faulty or fails to take into account a critical social impact factor, the results will not only be expensive in terms of later redesigns and lost productivity, but possibly also the loss of life. Unfortunately, history is replete with examples of how failures to adequately understand decision support problems inherent in complex sociotechnical domains can lead to catastrophe. For example, in 1988, the USS *Vincennes*, a U.S. Navy warship accidentally shot down a commercial passenger Iranian airliner due to a poorly designed weapons control computer interface, killing all aboard. The accident investigation revealed nothing was wrong with the system software or hardware, but that the accident was caused by inadequate and overly complex display of information to the controllers (van den Hoven, 1994). Specifically, one of the primary factors leading to the decision to shoot down the airliner was the perception by the controllers that the airliner was descending towards the ship, when in fact it was climbing away from the ship. The display tracking the airliner was poorly designed and did not include the rate of target altitude change, which required controllers to “compare data

taken at different times and make the calculation in their heads, on scratch pads, or on a calculator - and all this during combat (Lerner, 1989).”

This lack of understanding the need for a human-centered interface design was again repeated by the military in the 2004 war with Iraq when the U.S. Army’s Patriot missile system engaged in fratricide, shooting down a British Tornado and an American F/A-18, killing three pilots. The displays were confusing and often incorrect, and operators, who only were given ten seconds to veto a computer solution, were admittedly lacking training in a highly complex management-by-exception system (32nd Army Air and Missile Defense Command, 2003). In both the *USS Vincennes* and Patriot missile cases, interface designers could say that usability was the core problem, but the problem is much deeper and more complex. While the manifestation of poor design decisions led to severe usability issues in these cases, there are underlying issues concerning responsibility, accountability, and social impact that deserve further analysis.

Beyond simply examining usability issues, there are many facets of decision support system design that have significant social and ethical implications, although often these can be subtle. The interaction between cognitive limitations, system capabilities, and ethical and social impact cannot be easily quantified using formulas and mathematical models. Often what may seem to be a straightforward design decision can carry with it ethical implications that may go unnoticed. One such design consideration is the degree of automation used in a decision support system. While the introduction of automation may seemingly be a technical issue, it is indeed one that has tremendous social and ethical implications that may not be fully understood in the design process. It

is critical that interface designers realize the inclusion of degrees of automation is not merely a technical issue, but one that also contains social and ethical implications.

2. AUTOMATION IN DECISION SUPPORT SYSTEMS

In general, automation does not replace the need for humans; rather it changes the nature of the work of humans (Parasuraman & Riley, 1997). One of the primary design dilemmas engineers and designers face is determining what level of automation should be introduced into a system that requires human intervention. For rigid tasks that require no flexibility in decision-making and with a low probability of system failure, full automation often provides the best solution (Endsley & Kaber, 1999). However, in systems like those that deal with decision-making in dynamic environments with many external and changing constraints, higher levels of automation are not advisable because of the risks and the inability of an automated decision aid to be perfectly reliable (Sarter & Schroeder, 2001).

Various levels of automation can be introduced in decision support systems, from fully automated where the operator is completely left out of the decision process to minimal levels of automation where the automation only presents the relevant data. The application of automation for decision support systems is effective when decisions can be accurately and quickly reached based on a correct and comprehensive algorithm that considers all known constraints. However, the inability of automation models to account for all potential conditions or relevant factors results in brittle-decision algorithms, which possibly make erroneous or misleading suggestions (Guerlain et al., 1996; Smith, McCoy, & C. Layton, 1997). The unpredictability of future situations and unanticipated responses from both systems and human operators, what Parasuraman et al. (2000) term

the “noisiness” of the world makes it impossible for any automation algorithm to always provide the correct response. In addition, as in the *USS Vincennes* and Patriot missile examples, automated solutions and recommendations can be confusing or misleading, causing operators to make sub-optimal decisions, which in the case of a weapons control interface, can be lethal.

In addition to problems with automation brittleness, significant research has shown that there are many drawbacks to higher levels of automation that relegate the operator to a primarily monitoring role. Parasuraman (2000) contends that over-automation causes skill degradation, reduced situational awareness, unbalanced workload, and an over-reliance on automation. There have been many incidents in other domains, such as nuclear power plants and medical device applications, where confusing automation representations have led to lethal consequences. For example, in perhaps one of the most well-known engineering accidents in the United States, the 1979 cooling malfunction of one of the Three Mile Island nuclear reactors, problems with information representation in the control room and human cognitive limitations were primary contributors to the accident. Automation of system components and subsequent representation on the instrument panels were overly complex and overwhelmed the controllers with information that was difficult to synthesize, misleading, and confusing (NRC, 2004)

The medical domain is replete with examples of problematic interfaces and ethical dilemmas. For example, in the Therac-25 cases that occurred between 1985-1987, it was discovered too late for several patients that the human-computer interface for the Therac-25, which was designed for cancer radiation therapy, was poorly designed. It was

possible for a technician to enter erroneous data, correct it on the display so that the data appeared accurate, and then begin radiation treatments unknowingly with lethal levels of radiation. Other than an ambiguous “Malfunction 54” error code, there was no indication that the machine was delivering fatal doses of radiation (Leveson & Turner, 1995).

Many researchers assert that keeping the operator engaged in decisions supported by automation, otherwise known as the human-centered approach to the application of automation, will help to prevent confusion and erroneous decisions which could cause potentially fatal problems (Billings, 1997; Parasuraman, Masalonis, & Hancock, 2000; Parasuraman & Riley, 1997). Reducing automation levels can cause higher workloads for operators; however, the reduction can keep operators cognitively engaged and actively a part of the decision-making process, which promotes critical function performance as well as situation awareness (Endsley, 1997). Higher workloads can be seen as a less-than-optimal and inefficient design approach, but efficiency should not necessarily be the primary consideration when designing a DSS. Keen and Scott-Morton (1978) assert that using a computer aid to improve the effectiveness of decision making is more important than improving the efficiency. Automation can indeed make a system highly efficient but ineffective, especially if knowledge needed for a correct decision is not available in a predetermined algorithm. Thus higher, more “efficient” levels of automation are not always the best selection for an effective DSS.

While it is well established that the use of automation in human computer interfaces should be investigated fully from a design standpoint, there are also ethical considerations, especially for interfaces that impact human life such as weapon and medical interfaces. What might seem to be the most effective level of automation from a

design viewpoint may not be the most ethical. The focus on the impact of automation on the user's actions is a critical design consideration; however, another important point is how automation can impact a user's sense of responsibility and accountability. In one of the few references in the technical literature on humans and automation that considers the relationship between automation and moral responsibility, Sheridan (1996) is wary of individuals "blissfully trusting the technology and abandoning responsibility for one's own actions."

Overly trusting automation in complex system operation is a well-recognized decision support problem. Known as automation bias, humans have a tendency to disregard or not search for contradictory information in light of a computer-generated solution that is accepted as correct (Mosier & Skitka, 1996; Parasuraman & Riley, 1997). Automation bias is particularly problematic when intelligent decision support is needed in large problem spaces with time pressure like what is needed in command and control domains such as emergency path planning and resource allocation (Cummings, 2004). Moreover, automated decision aids designed to reduce human error can actually cause new errors in the operation of a system. In an experiment in which subjects were required to both monitor low fidelity gauges and participate in a tracking task, 39 out of 40 subjects committed errors of commission, i.e. these subjects almost always followed incorrect automated directives or recommendations, despite the fact that contraindications existed and verification was possible (Skitka et al., 1999). Automation bias is an important consideration from a design perspective, but as will be demonstrated in the next section, it is also one that has ethical implications as well.

2.1 Automation and Accountability

While automation bias can be addressed through training intervention techniques (Ahlstrom et al., 2003, however see Skitka, et al., 1999 for conflicting evidence), the degradation of accountability and abandonment of responsibility when using automated computer interfaces are much more difficult and ambiguous questions to address.

Automated decision support tools are designed to improve decision effectiveness and reduce human error, but they can cause operators to relinquish a sense of responsibility and subsequently accountability because of a perception that the automation is in charge. Sheridan (1983) maintains that even in the information-processing role, “individuals using the system may feel that the machine is in complete control, disclaiming personal accountability for any error or performance degradation.”

Some research on social accountability suggests that increasing social accountability reduces primacy effect (the tendency to best remember the salient cues that are seen first) (Tetlock, 1983), which is akin to automation bias. Social accountability is defined as people having to explain and justify their social judgments about others. In theory, increased accountability motivates subjects to employ more self-critical and cognitively complex decision making strategies (Tetlock & Boettger, 1989). However, previous studies on social accountability focused on human judgments about other humans and did not incorporate technology, specifically automation, so they are somewhat limited in the application of social accountability to the discussion of computers and accountability.

Skitka, Mosier, and Burdick (2000) attempted to bridge the gap in researching accountability from a purely social perspective to one that included technology in the

form of automation. The specific intent of this study was to determine the effects of social accountability on automation bias. Instead of being held accountable for their judgments about other people, subjects were required to justify strategies and outcomes in computerized flight simulation trials. The results showed that not only did increased social accountability lead to fewer instances of automation bias through decreased errors of omission and commission, but also improved overall task performance (Skitka, Mosier, & Burdick, 2000).

If increased accountability can reduce the effects of automation bias, how then could decision support systems be designed to promote accountability? For complex socio-technical systems, accountability will most likely come from an established organizational structure and policies put in place by higher-level management. However, one tangible design consideration for accountability would be the number of people required to interact with a given decision support system. Research indicates that responsibility for tasks is diffused when people work in collective groups as opposed to working alone, and this concept is known as “social loafing” (see Karau & Williams, 1993 for a review). By designing systems that require the fewest individuals in a decision-making component, it is possible that erosion in accountability through social loafing could be diminished. However, while research indicates that people experience degraded task responsibility through collective action, the potential loss of a sense of moral responsibility and agency for operators interacting collectively through human-computer interfaces is not as clearly understood. It is likely that the computer interface becomes another entity in the collective group so that responsibility, and hence accountability, can be cognitively offloaded not only to the group, but also to the

computer. This is one area in human-computer interaction and accountability research that deserves significantly more attention.

3. DESIGNING A MORAL BUFFER

Because of the diminishment of accountability that can result from interactions with computers and automation, I argue that when developing a human computer interface for any system that has the ability to harm people, such as interfaces for weapons and medical interfaces, the possibility exists that a moral buffer, a form of distancing and compartmentalization, is created which allows people to morally and ethically distance themselves from their actions. The concept of moral buffering is related to but not the same as Bandura's (2002) idea of moral disengagement, in which people disengage in moral self-censure in order to engage in reprehensible conduct. A moral buffer adds an additional layer of ambiguity and possible diminishment of accountability and responsibility through an artifact or process, such as a computer interface or automated recommendations. Moral buffers can be the conduits for moral disengagement, which is precisely the reason for the need to examine ethical issues in interface design.

A key element in the development of a moral buffer is the sense of distance and remoteness that computer interfaces create for their users. This sense of distance created by computer interfaces can best be illustrated through a military weapons interface example; although, as will be demonstrated, moral buffers can occur in other domains. The military is currently developing smart weapons such as cruise missiles and unmanned combat aerial vehicles (UCAV's), which once launched, can be redirected in-flight to a target of opportunity in a matter of minutes. While these weapons will provide

the military with unprecedented rapid battlefield response, developing technologies of this sort also have the potential to become moral buffers that allow humans to kill without adequately considering the consequences. In general, these types of weapons can be fired from remote distances; for example, the military recently used missiles in Iraq that can be fired from over 1000 miles from their intended target with pinpoint accuracy. While this distance is effective in protecting our own forces, it is also likely that increasing the distance from the battlefield diminishes a sense of accountability.

The desire to kill the enemy from afar, termed “distant punishment,” is deeply rooted in the military culture, and even using the term “distant punishment” is a euphemistic form of moral buffering. Military historian and psychologist Dave Grossman contends that military personnel have a deep-seated desire to avoid personal confrontation, and thus use distant punishment as a way to exert military will without having to face the consequences of combat (Grossman, 1998). Grossman depicts the level of resistance to firing a weapon as a function of proximity to the enemy in Figure 1. In addition, he reports that there have been virtually no instances of noncompliance in firing weapons from removed distances, while there are significant instances of refusal to fire for soldiers engaged in hand-to-hand combat (Grossman, 2000).

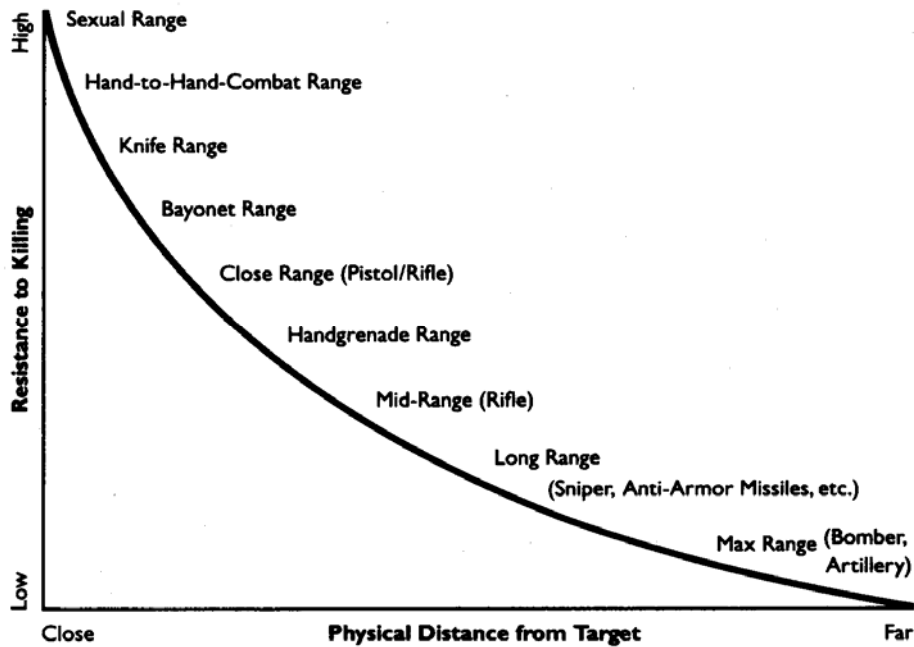


Figure 1: Resistance to Killing as a Function of Distance (Grossman, 1995)

In addition to the actual physical distance that makes it easier for people to kill, Grossman (1995) contends that emotional distance is a significant contributor as well. Emotional distancing in many domains is necessary for job performance, such as police work, the medical community, and in the military in general. However, there is a distinct difference in developing emotional distance for self or team preservation, and developing emotional distance through technology to make killing another human more palatable. Grossman contends that emotional distance in the context of killing can be obtained through social factors that cause one group to view a particular class of people as less than human, which include cultural elements such as racial and ethnic differences, as well as a sense of moral superiority. However, the primary emotional distancing element hypothesized by Grossman that should be of concern to interface designers is that of

mechanical distancing. In this form of emotional distancing, some technological devices provide the remote distance that makes it easier to kill. These devices can be TV and video screens, thermal sights, or some other mechanical apparatus that provides a psychological buffer, an element that Grossman terms “Nintendo[®] warfare (Grossman, 1995).” With the recent advancements in smart weapons that are controlled through computer interfaces that resemble popular video games, both the physical and emotional distancing that occur with remotely launching and controlling weapons provides an even greater sense of detachment than ever seen previously in modern warfare.

The famous Milgram studies of the early 1960’s help to illustrate how the concept of remoteness from the consequences of one’s actions can drastically alter human behavior. In these studies, the focal point of the research was to determine how “obedient” subjects would be to requests from someone they considered to be a legitimate authority. Under the impression that the real purpose of the study was to examine learning and memory, subjects, as the “teachers,” were told to administer increasing levels of electric shocks to another person, the learner, who was actually a confederate participant, when this person made mistakes on a memory test. While many different types of experimental conditions were examined, the one most pertinent to this discussion of moral buffers is the difference in subject behavior that was dependent on whether or not the teacher could see the learner. When the learner was in sight, 70% of the subjects refused to administer the shocks, as opposed to only 35% who resisted when the subject was located in a remote place, completely out of contact with the teacher (Milgram, 1975).

Milgram (1975) hypothesized that the increase in resistance to shocking another human when the human was in sight could be attributed to several factors. One important factor could be attributed to the idea of empathetic cues. When people are administering potentially painful stimuli to other humans in a remote location, they are only aware in a conceptual sense that suffering could result. Milgram had this to say about the lack of empathetic cues in military weapons delivery, “The bombardier can reasonably suppose that his weapons will inflict suffering and death, yet this knowledge is divested of affect and does not arouse in him an emotional response to the suffering he causes (Milgram, 1975).” Milgram proposed several other factors account for the distance/obedience effect to include narrowing of the cognitive field for subjects, which is essentially the “out of sight, out of mind” phenomenon. All of these factors are clearly present in the use of a weapons delivery computer interface, especially for one that controls weapons from over 1000 miles away.

In addition to physical and emotional distance and the sense of remoteness, and detachment from negative consequences that interfaces can provide, it is also possible that without consciously recognizing it, people assign moral agency to the computer, despite the fact that it is an inanimate object, which adds to the moral buffering effect. The human tendency to anthropomorphize computers has been well-established (Reeves & Nass, 1996). Furthermore, it has been established that automated decision support systems with ‘low observability’ can cause humans to view the automated system as an independent agent capable of willful action (Sarter & Woods, 1994). Low observability occurs in a complex system with high levels of automation authority (automation acts without human intervention) but little feedback for the human operator (Sarter & Woods,

1994). Viewing automation as an independent agent is also known as ‘perceived animacy’ and examples of this can be found in commercial airline cockpits where pilots will ask questions about flight management automation such as “What is it doing?” and “Why did it do that (Sarter & Woods, 1994)?”

In a research study designed to determine subject views about computer agency and moral responsibility, twenty-nine male computer science undergraduate students were interviewed concerning their views of computer agency and moral responsibility in delegation of decision making to the computer. Results suggested that these educated individuals with significant computer experience do hold computers at least partially responsible for computer error (Friedman & Millet, 1997). It follows then that if computer systems can diminish users’ senses of their own moral agency and responsibility, this would lead to erosion of accountability (Friedman & Kahn, 1997). In automated supervisory systems, human users can be isolated in a compartmentalized subsystem and detached from the overall system mission. This disengagement can cause them to have little understanding of the larger purpose or meaning of their individual actions. Because of this diminished sense of agency, when errors occur, computers can be seen as the culprits. When this diminished sense of agency occurs, “individuals may consider themselves to be largely unaccountable for the consequences of their computer use (Friedman & Kahn, 1997).”

An example of how a computer decision support tool can become a moral buffer between the human and computer is that of the Acute Physiology and Chronic Health Evaluation (APACHE) system. The APACHE system is a quantitative tool used in hospitals to determine the stage of an illness where treatment would be futile. While it

could be seen as decision support tool to provide a recommendation as to when a person should be removed from life support systems, it is generally viewed as a highly predictive prognostic system for groups, not individuals (Helft, Siegler, & Lantos, 2000). The APACHE system could provide a moral buffer through allowing medical personnel to distance themselves from a very difficult decision (“I didn’t make the decision to turn off the life support systems, the computer did”). By allowing the APACHE system the authority to make a life and death decision, the moral burden could be seen as shifting from the human to the computer.

The designers of this system recommend that APACHE only be used as a consultation tool to aid in the decision of removing life support and should not be a “closed loop” system (Friedman & Kahn, 1997). The ethical difficulty arises when technologies like APACHE become entrenched in the culture. Since the system has consistently made accurate recommendations, the propensity for automation bias and over-reliance could allow medical personnel, who are already overwhelmed in the workplace, to increasingly rely upon this technology to make tough decisions. When systems like the APACHE system are deemed to be a legitimate authority for these types of decisions, the system could in effect become a closed-loop system, which was not originally intended. Instead of guidance, the automated recommendations could become a heuristic, a rule-of-thumb, which becomes the default condition, and hence a moral buffer.

The same psychological phenomenon that creates possible moral buffers in the use of computer interfaces for medical advice may apply to decision support systems for weapons delivery, and indeed, for any computerized system that can inflict harm upon

people. Acting through a seemingly innocuous apparatus like a computer interface and making potentially fatal decisions such as directing weapons through the click of a mouse can create a moral buffer and allow people to perceive themselves as not responsible for whatever consequences result. It could be argued that those people who actually control in-flight weapons are only following orders of superiors, and thus the actual operators are not responsible for their actions. In older military systems, a commander would make a weapons firing decision and then, for example, order an underling to push the button that actually launched a weapon. Unfortunately, command and control technology have outpaced both human reasoning capabilities and traditional command structures. In smart weapons control of the future, weapons will no longer be controlled by junior enlisted personnel with little training. Smart weapons control in the future will require complex problem solving and critical analysis in real-time, which will be accomplished by educated, highly trained personnel who have the ability to both approve and disapprove

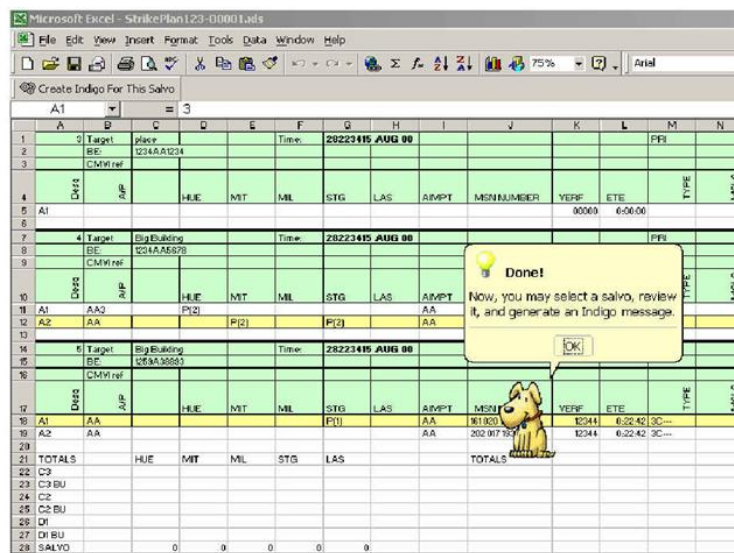


Figure 2: A Military Planning Tool

of a weapons launch (such as pilots who have the authority to not drop a bomb if the situation warrants.) It is precisely this group of decision makers who will be most affected by a moral buffer.

An example of how a particular design element could contribute to a moral buffer in the use of computer interfaces can be seen in Figure 2. This is a screenshot of an actual military missile planning computer program based on Microsoft's Excel[®] software package, which aids a military planner in planning an "optimal" mission (LoPresto, 2002). The user of this interface is likely to be a mid-career officer who is well educated and has the authority to choose between both resources and targets. The task of mission planning carries with it great responsibility, as millions of dollars in weapons, immeasurable hours in manpower, and scheduling of ships, planes, and troops are at the disposal of the planner. With users (the planners) bearing such serious responsibility, it is curious that the interface designers chose to represent the help feature using a happy, cute, and non-aggressive dog. A help feature is no doubt a useful tool for successful mission accomplishment, but adding such a cheerful, almost funny graphic could aid in the creation of a moral buffer by providing a greater sense of detachment in planning certain death through such an innocuous medium. It could be argued that in fact, this kind of interface is desirable as not to add to the already high stress of the mission planner; however, making the task seem more "fun" and less distasteful is not the way to reduce user stress.

A weapons control interface, even with the most elegant and thoughtful user design, may become a moral buffer, allowing users, who will be decision makers with authority and not subordinates "just following orders", to distance themselves from the

lethality of their decisions. Interface designers should be cognizant of the buffering effect when designing interfaces that require a very quick human decision, and be careful when adding elements such as the happy dog in Figure 2 that make a computer interface more like a leisure video game than an interface that will be responsible for lost lives. If computers are seen as the moral agents (i.e. I was only following the recommendations of the automation), military commanders may be tempted to use remotely operated weapons in real-time retargeting scenarios without the careful deliberation that occurred with older versions of weapons that required months of advance planning, and that once launched, cannot be redirected. Likewise, the same elements apply for users of any interface that impact human life such as medical devices and emergency response resources.

4. CONCLUSION

Because of the inherent complexity of socio-technical systems, decision support systems that integrate higher levels of automation can possibly allow users to perceive the computer as a legitimate authority, diminish moral agency, and shift accountability to the computer, thus creating a moral buffering effect. This effect can be particularly exacerbated by large organizations and the physical distancing that occurs with remote operation of devices such as weapons. For interface designs that require significant human cognitive contribution, especially in decision support arenas that directly impact human life such as weapons and medical systems, it is paramount that designers understand their unique role and responsibilities in the design process. The need for careful reflection on ethical issues should be a concern for the development of decision support systems for weapons; however, all domains in which computers have the

potential to impact human life deserve the same level of ethical and social impact analysis.

ACKNOWLEDGMENTS

I would like to thank Dr. Deborah Johnson, the University of Virginia Anne Shirley Carter Professor of Applied Ethics for her insights, framing suggestions, and most of all patience.

REFERENCES

32nd Army Air and Missile Defense Command (2003). *“Patriot Missile Defense Operations during Operation Iraqi Freedom.”* Washington DC: U.S. Army.

Ahlstrom, V., & Longo, K. (2003). *“The Human Factors Design Standard.”* Washington DC: FAA.

Bandura, A. (2002). “Selective Moral Disengagement in the Exercise of Moral Agency.” *Journal of Moral Education*, 31(2), 101-119.

Billings, C. E. (1997). Aviation Automation: The Search For A Human-Centered Approach. Hillsdale, N.J: Lawrence Erlbaum Associates.

Cummings, M. L. (2004). *“Automation Bias in Intelligent Time Critical Decision Support Systems.”* Paper presented at the AIAA Intelligent Systems Conference.

Endsley, M. (1997, June). Situation Awareness, Automation, and Free Flight. Paper presented at the FAA/Eurocontrol Air Traffic Management R&D Seminar, Saclay, France.

Endsley, M. R., & Kaber, D. B. (1999). “Level of automation effects on performance, situation awareness and workload in a dynamic control task.” *Ergonomics*, 42(3), 462-492.

Friedman, B., & Kahn, P. H. (1997). Human Agency and Responsible Computing: Implications for Computer System Design. In B. Friedman (Ed.), Human Values and the Design of Computer Technology (pp. 221-235). Stanford, CA: CSLI Publications.

Friedman, B., & Millet, L. I. (1997). Reasoning About Computers As Moral Agents: A Research Note. In B. Friedman (Ed.), Human Values and the Design of Computer Technology (pp. 205). Stanford, CA: CSLI Publications.

- Grossman, D. (1995). On Killing. Boston: Little Brown & Co.
- Grossman, D. (1998). The Morality of Bombing: Psychological Responses to "Distant Punishment". Paper presented at the Center for Strategic and International Studies, Dueling Doctrines and the New American Way of War Symposium, Washington DC.
- Grossman, D. (2000). Evolution of Weaponry, Encyclopedia of Violence, Peace, and Conflict: Academic Press.
- Guerlain, S., Smith, P., Obradovich, J., Rudmann, S., Strohm, P., Smith, J., & Svirbely, J. (1996). "Dealing with brittleness in the design of expert systems for immunohematology." Immunohematology, 12, 101-107.
- Helft, P. R., Siegler, M., & Lantos, J. (2000). The Rise and Fall of the Futility Movement. New England Journal of Medicine, 343(4), 293-296.
- Karau, S. J., & Williams, K. D. (1993). Social loafing: a meta-analytic review and theoretical integration. Journal of Personality and Social Psychology, 65(4), 681-706.
- Keen, P. G. W., & Scott-Morton, M. S. (1978). Decision Support Systems: An Organizational Perspective. Reading, MA: Addison-Wesley.
- Lerner, E. J. (1989, April). "Lessons of Flight 655." Aerospace America, 18-26.
- Leveson, N. G., & Turner, C. S. (1995). "An Investigation of the Therac-25 Accidents." In H. Nissenbaum (Ed.), Computers, Ethics & Social Values (pp. 474-514). Upper Saddle River, NJ: Prentice Hall.
- LoPresto, L. M. (2002, February 14). PC-based Mission Distribution System (PC-MDS). Paper presented at the Cruise Missile Seminar, Norfolk, VA.
- Milgram, S. (1975). Obedience to Authority. New York: Harper and Row.
- Mosier, K. L., & Skitka, L. J. (1996). Human Decision Makers and Automated Decision Aids: Made for Each Other? In R. Parasuraman & M. Mouloua (Eds.), Automation and Human Performance: Theory and Applications (pp. 201-220). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.
- NRC Office of Public Affairs. (2004). "*The Accident at Three Mile Island*." Washington DC: United States Nuclear Regulatory Commission.
- Parasuraman, R. (2000). "Designing automation for human use: empirical studies and quantitative models." Ergonomics, 43(7), 931-951.
- Parasuraman, R., Masalonis, A. J., & Hancock, P. A. (2000). "Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance." Human Factors, 42(4), 636-659.

Parasuraman, R., & Riley, V. (1997). "Humans and Automation: Use, Misuse, Disuse, Abuse." Human Factors, 39(2), 230-253.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). "A Model for Types and Levels of Human Interaction with Automation." IEEE Transactions on Systems, Man, and Cybernetics, 30(3), 286-297.

Reeves B., & Nass, C. (1996). *The media equation: How people treat computers, television and new media like real people and places*. Stanford, CA: CSLI Publications, New York : Cambridge University Press.

Sarter, N. B., & Schroeder, B. (2001). "Supporting decision making and action selection under time pressure and uncertainty: The case of in-flight icing." Human Factors, 43, 573-583.

Sarter, N. B., & Woods, D. D. (1994, April). "Decomposing Automation: Autonomy, Authority, Observability, and Perceived Animacy" Paper presented at the First Automation Technology and Human Performance Conference.

Sheridan, T. B. (1996). "Speculations on Future Relations Between Humans and Automation." In M. Mouloua (Ed.), Automation and Human Performance (pp. 449-460). Mahwah, New Jersey: Lawrence Erlbaum Associates, Inc.

Sheridan, T. B., Vamos, T., & Aida, S. (1983). "Adapting Automation to Man, Culture and Society." Automatica, 19(6), 605-612.

Skitka, L. J., Mosier, K. L., & Burdick, M. D. (1999). "Does automation bias decision-making?" International Journal of Human-Computer Studies, 51(5), 991-1006.

Skitka, L. J., Mosier, K. L., & Burdick, M. D. (2000). "Accountability and automation bias." International Journal of Human-Computer Studies, 52, 701-717.

Smith, P., McCoy, E., & C. Layton. (1997). "Brittleness in the design of cooperative problem-solving systems: The effects on user performance." IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 27, 360-371.

Tetlock, P. E. (1983). "Accountability and the perseverance of first impressions." Social Psychology Quarterly, 46(4), 285-292.

Tetlock, P. E., & Boettger, R. (1989). "Accountability: A Social Magnifier of the Dilution Effect." Journal of Personality and Social Psychology, 57(3), 388-398.

van den Hoven, M. J. (1994). "Towards ethical principles for designing politico-administrative information systems." Informatization and the Public Sector, 3, 353-373.