

**On the Ethical Dimension of Artificial Moral Agents Following the
Principles of *Jus In Bello* (Justice in War)**

Emmanuel A. Gonzalez

Jardine Schindler Elevator Corporation, 8/F Pacific Star Bldg., Sen. Gil Puyat Ave. cor. Makati
Ave., Makati City 1209, Philippines

Department of Computer Technology, College of Computer Studies, De La Salle University
Manila, 2401 Taft Ave., Malate Manila 1004, Philippines.

E-mail: emm.gonzalez@delasalle.ph

Abstract

In the past decade, research on the philosophy and the development of artificial moral agents (AMAs) have increased in order to seek for the implementation of moral decision-making faculties basically in machines—computers, robots, etc. Because of this, ethical reasoning becomes one of the primary challenges computer scientists, engineers, and philosophers are trying to solve—by developing an ethical dimension that can be incorporated in a machine so that it can make decisions on its own subjected to ethical and morality issues. However, when this ethical dimension is implemented in autonomous weapon systems (AWSs) such as cruise missiles, torpedoes, uninhabited/unmanned aerial vehicles (UAVs) and uninhabited/unmanned combat aerial vehicles (UCAVs), difficulty and confusion arise due to the question of who should hold responsible when the AWS itself is involved in an atrocity that would normally be described as a war crime under the principle of *jus in bello*—rules concerning acceptable practices while engaged in war. Thus, this paper presents a solution to the problem which is by incorporating artificial morality into the AWSs making them AMAs. Literatures supporting the possibility of incorporating an ethical dimension that is able to consider moral principles to be applied in AWSs are also presented.

Introduction

Artificial moral agents (AMAs) are entities that can be employed in a system in order to have a functionality of being able to make decisions based on an ethical dimension—a dimension consisting of ethical rules that play side-by-side in order to make proper decisions given a certain set of stimuli. Basically, an AMA must be able to make decisions that must be considered ethical in nature¹. It must be able to honor privacy, protect civil rights and individual liberty, and the welfare of others (Allen *et al*, 2006a). However, there are still a lot of issues concerning the responsibility for the actions of AMAs. For example, the decision to send AMAs, like artificially intelligent robots, into war has raised a lot of skepticism to philosophers, scientists, and engineers, especially when the AMA committed a war crime under the principles of *jus in bello* (justice in war, to right conduct in the midst of battle). Someone might ask who is responsible for AMA war crimes: is it the persons who designed, engineered and programmed the system, is it the commanding officer who authorized and/or ordered to use it, or the machine itself? This is indeed a very difficult question to answer; however, according to the doctrine of *justum bellum* (just war theory), it is a necessary condition for fighting a just war that someone must be held responsible for the deaths that occur in the course of war. It has been argued in (Sparrow, 2007) that this condition cannot be met in relation to deaths, directed or due to collateral damage, of the civilian therefore deploying such AMAs are unethical. Discussions on the philosophical issues of *justum bellum* are not anymore considered, but the readers can refer to (Orend, 2001) for more information.

The application of AMAs to autonomous weapon systems (AWSs) is one of the foreseen overwhelming technological advancements in the field of military engineering and it is very

¹ In this paper, the kind of ethics is not considered; however, most researches deal with the utilitarian and Kantian way of ethical reasoning especially for its applications to autonomous weapon systems.

important, at least in the philosophical and humanitarian sense that ethical decision-making must be taken into account. Since one of the aims of the United States Army's Future Combat Systems Project is to manufacture 'robot armies' for deployment in 2012, the inclusions of ethical and moral decision making for its application to AMAs in modern warfare is very important. Existing example of an AWS is the Boeing Corporation's Standoff Land Attached Missile-Expanded Response (SLAM-ER) cruise missile (Ray, 2003; Thompson & Nissan, 2004; Franke & Ari, 2004; Ray, 2005; Kuhlmann *et al*, 1998) that has a built-in automatic target recognition capability that allows it to choose a target once it reaches an area of operation. Another example is the US Air Force's Low Cost Autonomous Attack System (LOCAAS), which is a turbine-powered munition designed to autonomously search for targets of theatres containing missile defenses, air missile systems, and interdiction/armor targets as well (Niland *et al*, 2005; Lewis *et al*, 2006; Covault, 1997;). Other types of AWSs include torpedoes, submersibles, robots for urban reconnaissance, uninhabited/unmanned aerial vehicles (UAVs), and uninhabited/unmanned combat aerial vehicles (UCAVs). Even the "launch-on-warning" policy for strategic missiles advanced by the U.S. military during the late 1980s could be considered autonomous, in which military officials have supported believing that the use of onboard computers will make all necessary decisions. However, there is still a dangerous instability in such AWSs even if they are working out perfectly (Rauschenbakh, 1988; Martin & Schinzinger, 2005).

The incorporation of AMAs in AWSs, however, is difficult due to the issue of incorporating an ethical dimension which is practically a "bounded morality." The boundaries of this ethical dimension are actually studied under the field of *machine ethics*—a relatively new research area under the united fields of computer science, engineering, and philosophy—which

actually ignited further research as what has been presented in the AAAI 2005 Fall Symposium and in the special issue of the IEEE Intelligent Systems magazine in July/August 2006 (Anderson & Anderson, 2006). Machine ethics is concerned on how machines behave towards humans and other machines, which is guided by a set of rules following acceptable ethical principles. The courses of action to take by a machine must rely on a domain of ethical rules and principles—the *ethical dimension*. With this, the behavior of such machines becomes more acceptable in real-world compared to those machines without such a dimension. However, there is still a debatable question regarding the ability of a machine to be fully ethical in nature (Moor, 2006). According to Moor, an average adult is a full ethical agent. Furthermore, a fully ethical agent is able to make explicit judgments and is able to justify them reasonably, which according to Moor, is only applicable to an average adult human. Can a machine be a fully ethical agent? Applying and relating it in the scope of this paper—is it possible to have an AWS that is fully ethical such as an average adult person? Although the question is outside the scope of this paper, it is still a very important issue to discuss and debate, especially in the implementation process—the development of computational algorithms and its limitations due to hardware constraints.

All of these questions and issues are very important in the development of AMAs, which makes the field of machine ethics very important as well. The development of ethical agents or basically AMAs is based on the fact that we want machines to treat us (humans) well. Another thing is that since, in accordance to the scope of this paper, AWSs are becoming more sophisticated; it is more likely that these systems will exhibit autonomy. However, in designing AMAs for AWSs, it is very important to consider and understand the laws of war. It is very important that these AMAs observe just and lawful conducts in war, thus giving justice to those who are victims in any form during and after the war. *Justum bellum* (just war theory) actually

deals with the justification on why and how wars are fought, which is believed to be theoretical and historical (Moseley, 2006). Just war theory is theoretical in the sense that it is concerned with the ethics of justifying war and its forms of warfare. It is also considered historical due to the influence of past experiences and the historical body of rules of agreement employed. From *justum bellum*, philosophers and theorists distinguish between the rules that govern *jus ad bellum* (just of war) from *jus in bello*. According to Moseley, these two are by no means mutually exclusive; however, they present a set of moral guidelines in particular war and situations.

***Jus In Bello*—To a Just and Lawful Conduct in War**

Jus in bello is the term used by war theorists and philosophers referring to justice in war, essentially concerning the right conduct in the midst of a battle. It has been influential and subject to much debate due to the accounts developed by an American political theorist named Michael Walzer (Walzer, 1991)². Basically, in *jus in bello*, the objective is to hold the agents of war responsible for their actions, regardless if it is good or bad³.

The standards of *jus in bello* contains rules of just conduct that fall under two broad principles of discrimination and proportionality. However, Walzer added a third standard of banning means *mala in se*, that means “methods evil in themselves.” Moseley, on the other hand, added another principle, namely the principle of responsibility, which demands an examination of where responsibility lies in war.

The principle of discrimination concerns who are legitimate targets in war and is considered the most important principle in *jus in bello*. The idea in this principle is that no individual can justly be attacked unless he has, through his own actions, lost his basic human

² A reflection on Michael Walzer is available presented in (Orend, 2001).

³ The words *good* and *bad* are actually subjective to the reader. The discussion on the goodness and badness of an outcome or a process is considered as a fuzzy expression and is out of the scope of this paper.

rights by being engaged in harming. Furthermore, it must be understood that this individual can either be a soldier or a civilian because a soldier engaged in harming is no different from a civilian supplying goods to assist a soldier engaged in harming. If the individual is not engaged in the military effort of their nation, then that individual must not be a target of lethal force. As what Walzer asserts, “legitimate act of war is one that does not violate the rights of the people against whom it is directed.” (Walzer, n.d.)

The principle of proportionality, on the other hand, concerns the amount of lethal force that must be given in a legitimate target during war times. It is mandated in *jus in bello* that soldiers must only deploy proportionate lethal forces such that excessive harm and purposeless and wanton violence are prohibited. In this principle, the means of attacking a legitimate target is of interest. A soldier using a lethal weapon such as a gun that has bullets that contain glass or metal shards does violate the principle of proportionality because the wounds and injuries inflicted by this weapon expresses suffering and cruelty.

The most general rule in *jus in bello* is that armies must not employ acts or weapons that can destroy and shock the moral conscience of mankind. These means *mala in se* acts are considered to be forbidden because of its terrible consequences and awfulness. The use of nuclear, chemical, and biological weapons fall under this category because their use will inflict unlawful results such as the killing of innocent civilians including their properties, and will inflict shock to the moral conscience of mankind. As what Orend said, “We don’t have to do a cost-benefit analysis to determine whether such acts are impermissible in warfare: we already judge such acts to be heinous crimes.” (Orend, 2001). Furthermore, he also said that, “The

intentional destruction, and/or forcible displacement, of whole peoples, as Walter⁴ suggests, is something we find ‘literally unbearable.’”.

Artificial Morality in Artificial Moral Agents

The ethics of the decision to send AWSs during war is discussed in (Sparrow, 2007). It is said that someone must be held responsible for the results and consequences in the course of war and Sparrow believed that it would be unethical to deploy such systems because of “these conditions cannot be met in relation to the deaths that occur in the course of war.” Sparrow argued that “it will be unethical to deploy autonomous systems involving sophisticated artificial intelligences in warfare unless someone can be held responsible for the decisions they make where these might threaten human life.” However, it can also be argued that incorporating sophisticated artificial intelligences may not be the only solution to solve this problem. In relation to the principles of discrimination, proportionality, and no means *mala in se* in *jus in bello*, it can be argued that applying an ethical dimension under the field of artificial morality in AWSs can solve this problem. Thus, if the AWS has its own ethical dimension that would ensure zero war crime, then the problem of determining who is responsible for the consequences of war actions will not be a problem at all.

The primary goal of artificial morality is to design artificial moral agents (AMAs) that act as if they have their own ethical dimension by including sensitivity to the values, ethics, and legality of activities into the system. Strategies for implementing artificial morality into systems may have a lot of variations (Allen *et al*, 2007; Allen *et al*, 2006a; Allen *et al*, 2006b; Allen *et al*,

⁴ As referenced to Walzer’s rape example in Walzer, *Wars*, 129-137; C. MacKinnon, “Crimes of War, Crimes of Peace” in S. Shute and S. Hurley, eds. *On Human Rights* (New York: Basic Books, 1993), 83-110.

2005; Wallach *et al*, 2005), but the exploration of the technological and philosophical issues in making AMAs into explicit moral reasoners are relatively the same.

Artificial morality can be developed using two primary methods: the bottom-up and top-down approaches. A hybrid approach of bottom-up and top-down methodologies is also possible. In bottom-up approaches, the AMA's development is based on learning and evolution. The AMA is exposed to the environment in which appropriate behavior is selected and/or rewarded. This can be done by either providing unconscious mechanistic trial and failure of evolution, or by providing educational development in the learning AMA. However, the bottom-up approach is difficult to implement in AWSs because of its costly, timely, and infeasible development cycle—the need for experience, i.e. performing destructive tests and simulations.

Top-down approaches on the other hand are the most feasible development strategy for AMAs. In this approach, the AMA is given a rule-base of conditional propositions that can be programmed in a form of a non-volatile memory. The idea is that moral principles are used as rules for the selection of actions that are appropriately ethical. Although, often criticized to be insufficiently robust and unsuitable for providing a general theory of ethical action, there are still some instances where this method provides the best solution given available technology.

The most ideal approach in the incorporation of artificial morality into AMAs is to use the combination of both top-down and bottom-up approaches. This hybrid approach, however, may not be suitable for some AWS such as individual missiles and torpedoes because of the destructive learning actions during system evolution.

Towards an Ethical Autonomous Weapon System

In order to have a fully-ethical AWS, the principles of *jus in bello* must be incorporated into the system through the development approaches under the field of artificial morality. Since, in the previous section, it is evaluated that the top-down approach is the most feasible development strategy, the bottom-up and hybrid approaches may not be considered anymore, but it is still assumed the last two approaches may be useful in some way. Using this strategy, a rule-base is developed incorporating the principles of *jus in bello* as a set of conditional propositions. For example, if C is the set of all ethical codes, R is the regulated AWS behavior, and $D_{\text{ETH}} = \{C, R\}$ is the ethical dimension, then the conditional proposition would have a structure of “if $c \in C$ then $r \in R$.” However, one might think that the development process is time-consuming and requires an ample amount of memory. This is true, but will be dependent on the current technology in smart weapons.

As indicated before, the most important principle in *jus in bello* is the principle of discrimination which concerns who are legitimate targets in war. However, the big problem with this principle is the determination “who” are the legitimate targets during wartime. As Walzer said: a legitimate target is someone or something that is engaged in harming. This means that anyone or anything that is not engaged in harming is immune from direct or indirect attack from soldiers. However, another question arises on whether the human rights of soldiers are also violated, i.e. they must also be immune from direct or indirect lethal attacks. As what is said in (Orend, 2001): “...soldiers do something which causes them to forfeit their rights, much as an outlaw country forfeits its state rights to non-interference when it commits aggression.” In the case of Walzer: “No one can be forced to fight or to risk his life, no one can be threatened with war or warred against unless through some act of his own he has surrendered or lost his rights.”

Walzer also said that: "...all soldiers forfeit their right not to be targeted with lethal force, whether they be of just or unjust nations, whether they be tools of aggression or instruments of defence." Thus, if an AWS is designed to follow Walzer's philosophy, then any soldier that is detected by the AWS can be a lethal target. One thing to be considered, however, is that when attacking a legitimate target, it must be ensured that those who are immune to aggression must be protected from any harm, i.e. collateral damage must be prevented. Although a lot of AWSs nowadays have the capability of reducing collateral damage such as the Northrop Grumman BAT munitions and its derivatives, which were evaluated by the US Air Force and Army and currently deployed in Iraq (Eshel, 2007), the decision making and the accuracy of these systems are still based on the inputs given by its commanders. AWSs developed by Northrop Grumman, for example, are based on the concept of "excellent precision-kill success"—which still follows the inputs given by its commander. Thus, aside from artificial intelligence, if artificial morality is also incorporated in these systems, there is a chance that the accuracy and collateral damage problem is also reduced. If the commander gives a false input to an AWS in the sense that it will cause too much collateral damage, the AWS that will be sent to its target location will just make sure that the coordinates given by its commander are met without recognizing its effects and by just relying on its "excellent precision-kill success" concept. However, if the AWS was design to have its own artificial morality, i.e. it is able to discriminate legitimate targets, then the AWS must be able to make on-line decisions given the inputs by the commander and the current situation in the target area. It must be able to make some calculations whether the target given by the commander will cause severe collateral damage, or it must be able to decide based on the ethical dimension that is set inside the AWS.

One might think that it would be impossible to determine the current situation of the target area especially when there is no direct light-on-sight from the source of the aerial equipment to its destination. Fortunately, this problem can be solved in many ways. One way is to use spies, which is greatly emphasized on Sun Tzu's philosophy of war—The Art of War (Sun Tzu, 1983; Sun Tzu, 2007). As what Sun Tzu said: "Hence it is only the enlightened ruler and the wise general who will use the highest intelligence of the army for purposes of spying, and thereby they achieve great results. Spies are a most important element in war, because upon them depends an army's ability to move." He also wrote: "If you know the enemy and know yourself, you need not fear the result of a hundred battles." This concept of using spies can be used, however, not in the form of a human soldier or civilian, but in the form of technological advancement such as the use of satellites and Real-Time Video-Based Global Positioning System (RTVB-GPS) and Information Warfare (Infowar, IW) (Kopp, 2007). The use of satellites and RTVB-GPS can easily be used, however, the use of Infowar may not be that easy, because it also requires gathering of information from the opponent's facilities. Cyberwar can also be used for spying purposes, but it requires a lot of effort in order to hack or penetrate an opponent's computer network. Another problem with cyberwar is that it can only be used if the coordinate of the target of attack by the incoming AWS contains sufficient visual information, which can be gathered with the use of video cameras for spying purposes. Technically, the aerial AWS must be able to make decisions on its own prior to its arrival, which is based on the visual information gathered through RTVB-GPS. Furthermore, the AWS must be able to disengage all of its explosives and/or diverge to another location if it is able to detect that the target is not legitimate.

The other two principles of *jus in bello*—proportionality and banning means *mala in se*—can be greatly achieved in AWSs by ensuring that the amount of damage given by the aerial

equipment to the legitimate target is enough such that excessive harm is prevented, and the means itself of attacking the legitimate target doesn't cause moral disturbance to the people surrounding it. The amount of damage that is given to the legitimate target can be controlled again, based on the visual information gathered through RTVB-GPS prior to the arrival of the aerial equipment. Furthermore, the aerial equipment must be able to turn-on enough amount of explosives that is proportional to the volume of the legitimate target as long as the principle of discrimination in *jus in bello* is satisfied.

Ethical Dimension Techniques for AWSs

One of the most computationally-efficient approaches used in computational ethics is Jeremy Bentham's theory of hedonistic act Utilitarianism (McLaren, 2006; Bentham, 1948), which has been demonstrated in a prototype named JEREMY developed by Anderson *et al* (Anderson, Anderson, & Armen 2005). In Jeremy Bentham's theory, the idea is to use "moral arithmetic" in which one calculates the amount of pleasure and displeasure of those affected by every possible outcome in an ethical dilemma. In this theory, the "total net pleasure" is computed for each alternative action using the formula:

$$N = \sum_{i \in \mathcal{A}} IDP$$

where N is the total net pleasure, I is the intensity of the action, D is the duration of the action, P is the probability of the action to occur, i is a counter which is an element of \mathcal{A} , which is the set of all alternative actions given the dilemma. It has also been indicated in (Anderson, Anderson, & Armen, 2005) that the values of I , D , and, P are relatively small such that the product of the three values will not exceed unity.

The utilitarian approach in the design of an ethical dimension, however, seems to be premature because there are possibilities of having actions that are violent and would violate human rights, even though it may have the highest “total net pleasure.” A good example of such situation is when VIKI (Virtual Interactive Kinetic Intelligence), a humanoid character in the famous movie *I, Robot*, sees the laws of robotics as applying to humanity as a whole and thus justifies injuring some humans for the greater good—possibly due to having the highest “total net pleasure.” As what VIKI explains: “The Three Laws (of Robotics) are all that guide me. To protect humanity, some humans must be sacrificed. To ensure your future, some freedoms must be surrendered. We robots will ensure mankind’s continued existence. You are so like children. We must save you from yourselves.” In order to solve this problem, Anderson *et al* made their second prototype ,W.D. (Anderson, Anderson, & Armen, 2005), based on William D. Ross’s seven *prima facie* duties (Ross, 1930), and John Rawls’ theory of reflective equilibrium (Rawls, 1940). The general idea of the second prototype is that the theory of Ross provides a comprehensive set of duties or principles relevant to ethical cases, while Rawls’ approach provides the foundation for a procedure to make ethical decisions given those duties. This procedure defines cases as an evaluation of a set of duties having integral values between -2 to 2, which depend on how strongly each duty was violated or satisfied. In W.D., a value of -2 represents serious violation, while +2 represents maximal satisfaction of duties.

The theory of hedonistic act Utilitarianism, the seven *prima facie* duties, and the theory of reflective equilibrium can be combined in order to form an ethical dimension for an AWS. To facilitate these theories, the procedure for making ethical decisions can follow the theory of reflective equilibrium utilizing the seven *prima facie* duties. Bentham’s Utilitarianism theory can then be used to determine the total net pleasure of each decision derived from Rawls and Ross.

However, in its application to AWSs, the computational complexity of having such systems is high such that the following assumptions should be initially made: 1) the cruise missile is autonomous that it is able to decide on its own given different environmental and societal parameters such as the type of location of target, the number of legitimate and non-legitimate targets involved, the exact location of each of the legitimate and non-legitimate targets, etc.; 2) the parameters transmitted to the cruise missile are in real-time; 3) the cruise missile has an air-time t ; the warhead of the cruise missile will detonate only when either commanded by the missile, or when the missile is within the proximity with its target with a radius r ; the warhead disengages when the air-time exceeds t ; and, the software works without error and the hardware is fail-safe.

Conclusion

It has been argued by Sparrow that “it will be unethical to deploy autonomous systems involving sophisticated artificial intelligences in warfare unless someone can be held responsible for the decisions they make where these might threaten human life.” However, with what is presented in this paper, the author believed that it will be ethical to have such autonomous systems if these systems are design in such a way that they are able to make decisions based on an ethical dimension, making them artificial moral agents. Sparrow also wrote that “...the more autonomous these systems become, the less it will be possible to properly hold those who designed them or ordered their use responsible for their actions.” But, if these systems have their own ethical dimensions and are able to make moral decisions, then the responsibility problem will then be eliminated. Finally, Sparrow ended his paper saying that “...the deployment of weapon systems controlled by artificial intelligences in warfare is therefore unfair either to

potential casualties in the theatre of war, or to the officer who will be held responsible for their use.” However, the authors believe that artificial intelligence is just one of the tools used to develop sophisticated AWS. Furthermore, artificial morality which is a relatively new field, must also be included into the development of such systems by cooperating with artificial intelligence in order to satisfy the principles of *jus in bello*.

Acknowledgment

The author would like to thank Dr. Colin Allen of the Department of History and Philosophy of Science and in the Cognitive Science Program at Indiana University, Bloomington, and Dr. Manuel Belino of the Mechanical Engineering Department at Mapua Institute of Technology for their invaluable support to the writing of this paper.

References

- Allen, C., Smit, I., & Wallach, W. (2005). Artificial morality: top-down, bottom-up, and hybrid approaches. In Brey, P., Grodzinsky, F., & Introna, L. (eds.) *Ethics of New Information Technology: Proceedings of the Sixth International Conference of Computer Ethics: Philosophical Enquiry (CEPE 2005)*, Enschede, Netherlands: Center for Telematics and Information Technology (CTIT), 21-28.
- Allen, C., Wallach, W., & Smit, I. (2006a). Why machine ethics?, *IEEE Intelligent Systems*, July/August, 12-17.
- Allen, C., Smit, I., & Wallach, W., (2006b). Artificial morality: top-down, bottom-up, and hybrid approaches. *Ethics and Information Technology*, 7, 149-155.
- Allen, C., Wallach, W., & Smit, I. (2007). Machine morality: bottom-up and top-down approaches for modelling human moral faculties, *A.I. and Society*: (in press)
- Anderson, M., Anderson, S., & Armen, C. (2005). Toward machine ethics: implementing two action-based ethical theories, *Proc. AAAI 2005 Fall Symp. Machine Ethics*, AAAI Press, 1-16.

- Anderson, M., & Anderson, S. L. (2006). Machine ethics, *IEEE Intelligent Systems*, July/August, 10-11.
- Bentham, J. (1948). *Introduction to the principles of morals and legislation*, W. Harrison, ed., Hafner Press.
- Covault, C. (1997). LOCAAS sensor tests advance, *Aviation Week & Space Technology*, 147(17), October 27, 72-73.
- Eshel, D. (2007). Smart weapons for UAVs, *Defence Today*, September/October, 50-53.
- Franke, M. E., & Ari, G. (2004) Use of cargo aircraft for launching precision-guided munitions, *42nd AIAA Aerospace Sciences Meeting and Exhibit*, January 5-8, Reno, Nevada
- Khulmann, R., Gerth, D., Heckroth, K., & Nielson, J. (1998). SLAM-ER – an innovative update to a premier GPS guided weapon, *Proceedings of the 11th International Technical Meeting of the Satellite Division of the Institute of Navigation, Nashville, TN*, September 15-18, 583-589.
- Kopp, C. (2007). Fundamentals of information warfare, *Defence Today*, September/October, 71-73.
- Lewis, M., Polvichai, J., Sycara, K., & Scerri, P. (2006). Scaling-up human control for large UAV teams. Retrieved August 2, 2007. Available online: <http://www.cs.cmu.edu/~pscerri/papers/HFRPV05.pdf>
- Martin, M. W., & Schinzinger, R. (2005). *Ethics in Engineering*, 4th ed., New York: McGraw-Hill, Inc.
- McLaren, B. M. (2006). Computational models of ethical reasoning: challenges, initial steps, and future directions, *IEEE Intelligent Systems*, July/August, 12-17. 29-37
- Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics, *IEEE Intelligent Systems*, July/August, 18-21.
- Moseley, A. (2006). Just war theory, *The Internet Encyclopedia of Philosophy*. Retrieved August 12, 2007. Available online: <http://www.iep.utm.edu/j/justwar.htm>
- Niland, W., Stolarik, B., & Harman, J. (2005). The inclusion of a supplementary mission scenario into the multi UAV research tool, *AIAA Guidance, Navigation, and Control Conference and Exhibit*, San Francisco California, August 15-18.
- Orend, B. (2001). Just and lawful conduct in war: reflections on Michael Walzer, *Law and Philosophy*, 20, 1-30.

- Rauschenbakh, B. V. (1988). Computer war, in Anatoly Gromyko and Martin Hellman (eds.), *Breakthrough: Emerging New Thinking*, New York: Walker.
- Rawls, J. (1999). *A Theory of Justice*, 2nd ed., Harvard Univ. Press.
- Ray, E. (2003) CFP applied to separation of SLAM-ER from the S-3B, *21st Applied Aerodynamics Conference*, June 23-24, Orlando, Florida.
- Ray, E. (2005). Authorizing SLAM-ER use from the P-3C with CFD, *43rd AIAA Aerospace Science Meeting and Exhibit*, January 10-13, Reno, Nevada.
- Ross, W. D. (1930). *The Right and the Good*, Oxford Univ. Press.
- Sparrow, R. (2007). Killer robots, *Journal of Applied Philosophy*, 24(1), 62-77.
- Sun Tzu (1983). *The Art of War*, Clavell, J. (ed.), New York: Dell Publishing.
- Sun Tzu (2007). *The Art of War*, Giles, Lionel (ed.), Ann Arbor, Milwaukee: Borders Group, Inc.
- Thompson, A. M., & Nissan, R. (2004). Missile systems, *Aerospace America*, December, 94-95.
- Wallach, W., Allen, C., & Smit, I., (2005). Machine morality: bottom-up and top-down approaches for modeling human moral faculties. *Proceedings of AAAI Workshop on Machine Ethics*, Washington, D.C., November.
- Walzer, M. (1991). *Just an Unjust Wars*, 2nd ed., New York: Basic Books.