International® Baccalaureate

Extended essay cover

Candidates must complete this page and then give this cover and their final version of the extended essay to their supervisor.

Candidate session number	
Candidate name	
School number	
School name	

Examination session (May or November)	MAY	Year	2012

Diploma Programme subject in which this extended essay is registered: _Philosophy_ ✓

(For an extended essay in the area of languages, state the language and whether it is group 1 or group 2.)

Title of the extended essay: _The moral status of artificial intelligence._

Candidate's declaration

This declaration must be signed by the candidate; otherwise a grade may not be issued.

The extended essay I am submitting is my own work (apart from guidance allowed by the International Baccalaureate).

I have acknowledged each use of the words, graphics or ideas of another person, whether written, oral or visual.

I am aware that the word limit for all extended essays is 4000 words and that examiners are not required to read beyond this limit.

This is the final version of my extended essay.

Candidate's signature: _____ Date: _11th Jan. 2012_

Supervisor's report and declaration

The supervisor must complete this report, sign the declaration and then give the final version of the extended essay, with this cover attached, to the Diploma Programme coordinator.

Name of supervisor (CAPITAL letters) _____

Please comment, as appropriate, on the candidate's performance, the context in which the candidate undertook the research for the extended essay, any difficulties encountered and how these were overcome (see page 13 of the extended essay guide). The concluding interview (viva voce) may provide useful information. These comments can help the examiner award a level for criterion K (holistic judgment). Do not comment on any adverse personal circumstances that may have affected the candidate. If the amount of time spent with the candidate was zero, you must explain this, in particular how it was then possible to authenticate the essay as the candidate's own work. You may attach an additional sheet if there is insufficient space here.

Strong-minded and independent, the candidate accepted only the lightest of supervision on my part after our initial discussion with respect to possible parameters. He is a keen philosopher and a keen physicist making the subject of the essay one playing to his strengths. He showed excellent determination throughout, together with the willingness to explore avenues that were not immediately obvious — ones, indeed, that turned into dead ends (despite which, he kept up his positive approach).

This declaration must be signed by the supervisor; otherwise a grade may not be issued.

I have read the final version of the extended essay that will be submitted to the examiner.

To the best of my knowledge, the extended essay is the authentic work of the candidate.

I spent [2] hours with the candidate discussing the progress of the extended essay.

Supervisor's signature: _____ Date: 14th Jan. 2012

Assessment form (for examiner use only)

Candidate session number	

Achievement level

Criteria	Examiner 1	maximum	Examiner 2	maximum	Examiner 3
A research question	2	2		2	
B introduction	2	2		2	
C investigation	3	4		4	
D knowledge and understanding	2	4		4	
E reasoned argument	3	4		4	
F analysis and evaluation	2	4		4	
G use of subject language	3	4		4	
H conclusion	1	2		2	
I formal presentation	2	4		4	
J abstract	1	2		2	
K holistic judgment	2	4		4	
Total out of 36	23				

of examiner 1: _____
TAL letters)

Examiner number: _____

of examiner 2: _____
TAL letters)

Examiner number: _____

of examiner 3: _____
TAL letters)

Examiner number: _____

IB Cardiff use only: B: _____

IB Cardiff use only: A: 104836   Date: 24/04/19

The Moral Status of Artificial Intelligence

Research Question: Could an artificially intelligent entity be deserving of moral status?

Philosophy

Extended Essay

May 2012

Word Count: 3,989

Abstract

Could an artificially intelligent entity be deserving of moral status? This question *Rg.* arose to me when considering why we think of some things as worthy of moral treatment and others not. People are always treated as moral, but there are many grey areas when it comes to moral opinions, where there are many different opinions. These include animals, the environment, and most importantly to me, machines.

I carried out my investigation by reading several essays, books, and articles, relating to varied topics, from neuroscience to free will. From my reading I discovered early on that there were many different methods that people used to ascribe moral status. They were all based on the entity in question possessing certain attributes or qualities. In this essay I examined some of the most prominent of these asserted qualities: life, emotions, and consciousness. I also discussed the nature of personhood with regards to moral status, and whether it is exclusive to human beings, or might be able to be obtained by artificial intelligence.

Not, precise enough description of how the investigation was undertaken

In the end, my investigation led me to the conclusion that there is nothing that inherently limits an artificially intelligent entity from moral status. They seem to be equally as capable of life, emotions, and consciousness as humans, even if they haven't fully achieved them yet. Perhaps more surprisingly they might even at some point in the future also be considered persons, and so attain all the moral status that comes with personhood.

Word Count: 242

Contents

Too general

Little information provided

Introduction

The question of whether artificially intelligent entities could be deserving of moral status is one that is becoming increasingly relevant in the modern age. As technology in the fields of computing and neuroscience advances ever more, the possibility of artificially creating human-like, or even conscious, beings draws closer[1]. How we would treat these entities then, is a question whose answer has the possibility to drastically alter our ethical attitudes and views in the near future.

Much of the current debate, both in the philosophy and science of artificial intelligence, circulates around the plausibility of artificial intelligence that would be indistinguishable from human intelligence, addressing the question 'Could machines think?[2,3] This question is a valid one, as whether anything non-biological or even non-human could be capable of consciousness or other cognitive processes is still in doubt. However, for the purposes of this essay I shall be almost avoiding this question entirely. I plan instead to examine the ethical implications that result from the assumption that machines will, at some time in the future, be capable of thought, feeling and consciousness comparable to that of humans

My answer will necessarily have to begin with the definition of the term 'Artificial Intelligence'. As discussed in this essay, artificial intelligence will be defined as manmade creations which receive input from their surroundings, and are capable of internally manipulating this or other data in a way similar to thought[4]. Whether this manipulation can be deemed 'conscious' is difficult to fathom and will form part of the discussion in this essay, but only insofar as it pertains to the question of in what way we might ascribe moral status to artificial intelligence.

Another definition required to answer this question is what it means to possess moral status. Its definition here will be: anything that can be held ethically accountable for its

---

[1] (Anderson 2001)
[2] (Churchland 1990) payes !.
[3] (Can Robots 'Think' like Humans? 2008)
[4] (Hauser 2007)

1

actions, would be worthy of certain rights, and that other moral beings would be

responsible for treating it in a way that respects its moral status. To answer this I will

investigate possible reasons for ascribing moral status.

Intuition

Instinctively, we have many preconceptions as to who and what is deserving of being

treated ethically. It seems clear to most that other intelligent human beings are deserving of

moral status and all that it entails, and those who disagree are often branded as

psychopaths. Also on an intuitive level, most would not accord moral status and respect to a

pebble, or piece of metal. These differences seem very clear to us. However a problem

arises in that there are many cases which are not so clear cut as to whether, and to what

extent certain 'borderline' cases can be ascribed any kind of morality.

For example, it is possible to conceive of a mechanical device, like a toy, which

appears like some small animal, whilst still being obviously artificial. At first glance it seems

intuitively that this should not be accorded moral rights, any more than another toy or

inanimate object should, and most people would have no qualms about harming or

destroying it. However, if it is then discovered that the function of this machine is to act in

such a way that it appears to be in pain when damaged, and be fearful of further harm, then

the immediate, instinctive decision becomes much more confused[5]. It may be much more

difficult to attack something that appears vulnerable, innocent, fearful, and capable of

suffering than something that is not, because it evokes a certain emotional empathy.

This implies that there are certain qualities which, if an entity possesses them, make

us more likely to view something as moral. To solve the original question of whether or not

an artificially intelligent entity could be deserving of moral status, one should first consider

the question: What, if any, attributes contribute to something being treated and considered

morally? In this system a sufficient condition for possessing moral status is one that, if

possessed by something, immediately qualifies them to be treated morally, regardless of

anything else. For example having three sides is a sufficient condition for being a triangle,

because any three sided object must be a triangle. A necessary condition is one that is

required for something to be treated morally, but by itself is not enough. For example,

---

[5] Example edited from (Miedaner 1977)   Not engh infmation

2

having four sides is necessary, but not sufficient, for being a square, because there are examples of four sided objects that are not squares, and yet every square must necessarily have four sides.[6]

It is very probable however, in the example above, that simply the outward appearance of the machine causes us to make assumptions about cognitive processes deeper within it, making us misjudge it in terms of morality. Another question then arises of how we can ascertain whether or not something truly possesses a certain quality, given that many of these attributes are probably cognitive and internal to a specific being.

Life, emotions and consciousness are some possible attributes which might be either necessary to, or at least contribute to, our attribution of moral status to a being.

why ?

Life

Moral status is not something that is granted uniquely to humans, by society. People will often feel bad about harming animals, although to varying degrees, depending on the person and the animal. Similarly people will often feel bad about harming machines, specifically those which most accurately simulate life, be it animal or human. This pathos is not usually felt towards non-living things, except perhaps when they are mistakenly attributed life-like qualities, such as stuffed toys or dolls, or perhaps when they are instrumental to effects on living things, such as the environment. This strongly suggests that there is something unique about life which brings with it an idea of morality.

However, this is not necessarily the case given that there are living things which are often ascribed no moral status at all. The most obvious examples of this are bacteria, or most plants, towards which people usually feel no remorse or empathy, even to the point of killing them with no remorse. This means that life by itself is not a sufficient condition by itself to receive moral treatment.

It is much more difficult however, to see how life could not be a necessary condition. Whilst it might be possible to conceive of non-living objects that we see as deserving of respect, such as the planet, i.e. the abiotic environment, it is probable that these are only viewed this way because they are instrumental to the good of other, living, things. For

---

[6] (Earl 2005)

example, we might only treat our non-living surroundings well, and not destroy them because they provides us with resources, because it is the habitat for many animals and plants, or because we regard them as beautiful, as opposed to because they have their own intrinsic value.[7]

Using life as the defining attribute that all moral beings must possess does raise the issue of why and how life distinguishes moral and non-moral things. There is a clear intuitive difference between the moral status of a bacterium and a human being, however if life were the only condition used to establish the morality of an entity then these two would be seen as the same, given that they are both living. It is possible then, that life is not genuinely the defining characteristic, but rather our feeling that it is derives from the fact there are many other characteristics, namely cognitive or emotional ones, which at the moment are only displayed by living organisms. These characteristics could then be used to distinguish between different living (or non-living, if it is possible for them to possess these attributes) moral entities.

On the other hand even if life were a necessary characteristic for moral status, no commonly held biological definition of life inherently excludes artificial intelligence. Usually definitions are based on something exhibiting certain characteristics. The most common of these are: adaptation, growth, homeostasis, metabolism, organization, reproduction, and response[8]. Metabolism is a simply a series of chemical reactions designed either to build or break down other chemicals or forms of energy and thus occurs in many non-living things, such as batteries, electrolytic cells, and even natural processes. Furthermore, there are examples of machines which can add parts to themselves and even create more copies of themselves as in robots in factories, showing growth and reproduction. There are also sophisticated robots which exhibit response, adaptation, homeostasis, and organisation; for example the robot TOPIO, which was designed to play table tennis, was able continuously adapting and responding to external stimuli, through a complex computing system[9]. It seems perfectly reasonable then, to imagine an artificially created being that by the current medical and biological definition is alive.

Some analysis and evolution in this techn

[7] (Waldron 2009)
[8] (Davison 2008)
[9] (Diginfo 2007)

4

This demonstrates that life is not sufficient by itself to generate a sense of morality. Instead it may be something else that is common to the certain specific animals, including humans, which accords them moral status. Such a quality could be cognitive or emotional states. Humans and the animals we feel pathos for all share thought or emotion in varying capacities. Seemingly at least one of these is required for something to be treated within the moral sphere.

Emotions

Emotions represent a psychological state, in reaction to surroundings or environment. They can reflect the impact that moral decisions have on other beings. This may be the reason they are so important in terms of ethics. Without anything to react to or be emotionally affected by moral actions (including the actor), they probably have no ethical meaning whatsoever. Certainly they would not register on a human scale of morality, which is so intrinsically linked with emotion as a reaction to moral events, guilt for doing something morally wrong, pride for doing something morally right, disappointment, anger, gratitude, etcetera for the moral actions of others[10]. However, emotions are complex cognitive concepts, and so require correspondingly complex neurological infrastructure to occur[11]. This makes them exclusive to relatively few organisms. Because of their similarity to life, humans often begin to imagine robots which simulate life with emotions, such as fear or shame. However, owing to a comprehensive understanding of the programming and electronic structure of these life-imitating machines, it can be said with some confidence that these robots do not yet display these emotions genuinely, but are only able to simulate the outward expression of certain emotions. The simple simulation of emotions is not sufficient for moral status.

In practice trying to find where the distinction lies between simulating emotion and genuinely expressing it is very difficult. It is not fully understood how the brain produces such complex processes such as emotion, but seeing as it is evidently possible to create this process in nature, it seems likely that it would be possible to do so in a machine in the

---

[10] (Strawson 2008)
[11] (John Crane 2009)

5

future. Given the current limited understanding of this however, if this were to be achieved in the near future, it would most likely be through experimental testing, trying to emulate the way the human brain operates[12]. Because of this, we might still not fully understand the cognition behind the emotions before we could actually create them in a machine. This would make distinguishing between fake and genuine emotion, as mentioned above, very difficult.

It is possible that emotions arise as an expression of mental states and consciousness, and so are dependent on them. In other words, it is possible that emotions could not exist without thought, or awareness. This could mean that although it seems like emotional state is deciding moral status, in reality it is something more basic that is the source of emotions, and hence is also common to all the animals that display emotions. Given this, a machine might not have to actually display emotion, but another cognitive process, which happens to also be the base of emotions[13]. As of yet, there is no certain way to test for emotions in a machine, however if emotions were simply emergent as a result of higher cognitive functions, then it could be more useful to discover these instead, and from that information infer either that specific machines do feel emotion, or that emotion is unnecessary and that the true necessary condition for morality is one of these processes, such as consciousness.

Consciousness

The most prominent reason for ascribing moral status would be level of consciousness. In common conceptions of morality there seems to be a spectrum of consciousness, which corresponds to a spectrum of moral status. For example, it is usually seen as more morally wrong to harm a dog than an ant, as the dog is more conscious than the ant, and is therefore more capable of feeling pain or suffering than the ant. This reflects a spectrum, in that we would ascribe more moral status to a dog than to a mouse, or to a mouse rather than an ant, or to an ant rather than a bacterium. These levels of morality correlate directly to their place in the spectrum of consciousness, in that, the more conscious a being is, the more 'moral' it is.

Because of this, if a machine could achieve some level of consciousness, then

---

[12] (Igor Aleksander 1999)
[13] (Steven Rose 1998)

6

seemingly it would have a place in the conscious, and therefore also moral, spectrum. It is however, very difficult to assess levels of consciousness and thought. One way, proposed by Alan Turing, is known as the 'Turing test'[14]. This test is a game, in which there are three players A, B, and C. A is a computer, B is a human contestant, and C is a human judge. The goal of the human contestant, B, is to convince C that it is the human. The goal of the computer, A, is to trick C into thinking that it is the human. The goal of C is to accurately determine which of the other two contestants is human, and which is a machine. So as to ensure that this is a test of intelligence and thought, and not of accurately simulating human appearance or voice, all of the players are in different rooms, and communicate with each other using typed messages which are transmitted across the rooms. The format of the game is that the judge asks a series of questions to each of the contestants, and they reply in the way in which they think is the most 'Human'. The aim of this test is to determine whether a machine could simulate human thought, and so be said to be thinking and 'conscious' like humans.

A common objection to this method of evaluating consciousness is that it doesn't actually demonstrate thought, only the ability to simulate it[15]. What the machine says and outwardly does might not reflect actual thought and consciousness on the inside. While this is clearly a legitimate objection, it is worth noting that people judging the outward expression and appearance of other human beings is exactly the method used every day by society to evaluate conscious thought. People are usually quite content to believe that the humans around them are conscious, simply on the basis of their outward appearance and what they say in conversation. However, in the case of humans their biological origins do also lend credence to their consciousness, in that they are much more likely to be similar to each other, and so one person more readily accepts that the other is conscious, given that they themselves are conscious, and believe the other person to be of a very similar makeup to them[16].

John Searle exemplifies the problem of judging thought simply from outward conversation with a thought experiment known as "The Chinese Room"[17]. This experiment

---

[14] (Turing 1950)
[15] (Igor Aleksander 1999)
[16] (Law 2003)
[17] (Cole 2009)

involves a closed room, in which there is a computer program capable of manipulating symbols. It receives an input of certain Chinese characters, then follows a set of logical rules in its programming that virtually transform the characters into new ones as output. It is programmed so that the symbols which are the output seem like a conversational response to the input symbols. This hypothetical programme can do this so well that it could 'pass' the Turing test, by being able to fool a human that it too was human based on its responses. There is a fluent Chinese speaker outside the room who is providing the input of characters into the room, by providing his half of a conversation. This Chinese speaker, without prior knowledge, would be fooled into thinking he was conversing with another, conscious, thinking, Chinese speaker on the inside of the room, given the output he receives from it. According to the Turing Test, this would be sufficient to say that the computer programme could speak and think in Chinese. Now Searle proposes to replace the computer programme with a human in the room who could not speak Chinese, who would receive the input of symbols, and produce an output by manually following the same set of rules which the computer programme followed. This process too would convince the Chinese speaker outside the room that he/she were conversing with someone inside the room who could understand Chinese. However, the human in the room clearly does not understand Chinese, and is simply manipulating characters. This Searle argues, is exactly the same as what the computer programme does. It does not actually understand Chinese, it is simply able to imitate understanding of Chinese. Searle calls this kind of machine "weak AI"[18,19]. He does however think that it is still possible to create a machine that displays 'genuine' understanding, but that the Turing test would not be sufficient to determine this. This machine, if suitably programmed, would be able to semantically understand natural language and be capable of conscious states as opposed to simply simulating them. He calls this kind of machine "strong AI". If this kind of machine did exist then it would literally have a 'mind', and be capable of consciousness, understanding, and thought in the same way that human minds are.

This is in short, the difference between simulating consciousness and genuinely having it. Whilst there are many computer programs that exist at the moment, most notably

---

[18] (Igor Aleksander 1999)
[19] (Searle 1999)

8

'chat bots'[20], which can hold a conversation by searching through a database of preset responses and so simulate consciousness, they do not yet appear to have reached the stage of authentic cognition.

Personhood

In the human moral framework there is a clear distinction made between 'people' and non-people. This affects all the ethical implications of their existence, their responsibilities to other people, and others to them. The terms 'Human' and 'Person' are not interchangeable. For example, an embryo or extremely brain damaged individual might be described as a human, but not as a person. It is therefore possible that being a human (that is, a member of the species, Homo sapiens) is not even a necessary condition for being a person. If a machine could be described as a person, then they would by definition gain all the ethical aspects that personhood entails.

To evaluate whether or not someone is a person it must be seen whether they express certain characteristics that make people, people. A possible list of these conditions could be: consciousness, language, emotion, and reason[21]. These conditions may not all be necessary to be a person, but at least some of these must be present in some capacity. For example, a human without the power of language may be said to still be a person, just as a sociopath who could not feel emotion could still be a person. However it would be difficult to conceive of a person without language, emotion, or the ability to reason. If a machine could demonstrate all of these characteristics, it would be very likely to be considered a person itself.

It has already been discussed that 'strong AI' would hypothetically be capable of consciousness and language[22]. Emotion and reason are both complicated and not fully understood mental phenomena. However, given the level of understanding required for consciousness, it wouldn't seem too much of a leap to presume that such machines could be programmed with feelings and logic, given that they would have genuine 'minds'. No such machine exists at present, and it is possible that consciousness, language, emotion, and

---

[20] (Carpenter 2011)
[21] (Stuart 2009)
[22] (Searle 1999)

9

reason are so complicated and abstract, that strong AI is altogether impossible to create, which would render the whole question irrelevant. However, the hypothetical machine I am proposing would be capable of all of these traits, and given that this is a thought experiment, it could be adapted to possess any other trait involved in being a person, so long as it did not contradict either logic, or the definition of itself as a machine.

These criteria might be seen as unsatisfactory to award personhood to certain machines to some people. They might argue for other reasons to discriminate between machines and human persons. These might include the origin, or material makeup of the machine[23]. The former is trying to suggest that because machines were not conceived or born in the same way as humans, they cannot ever be viewed on the same level morally. The latter says that because a machine's circuitry is not made up of the same stuff a human brain is, it cannot be capable of the same level of moral agency. Both seem to me to be very anthropocentric arguments. There are 'test-tube babies' created through IVF (In Vitro Fertilisation) which are neither conceived nor born the same way as 'normal' humans, and yet it is impossible to distinguish them from 'natural' humans, even from close physical examination[24]. There is also nothing yet discovered to be scientifically special about individual neurones or biological cells which could not hypothetically be reproduced with another material. Neurones are simply used to carry electrical and chemical signals at varying strengths and in varying directions depending on their input, in exactly the same way as circuits do. For example, if a single neurone in a human's brain were replaced with an electrical component which did exactly the same task, then it would not affect the outward actions of that person, and hypothetically, would have no impact on their thoughts or emotions either. By extension, if all the other brain cells were replaced by circuits one by one, then the effect would be the same, and there would be no change[25]. If this were true, then there would be nothing special about the human brain that makes it unique from anything else in the universe, except the order in which the particles are arranged, so as to carry certain electrical signals at certain times.

---

[23] (Law 2003)
[24] (Anonymous n.d.)
[25] (Law 2003)

This seriously undermines the position that humans are uniquely worthy of personhood. Thus it leaves open the possibility of an artificially intelligent entity reaching the point that it too could be treated as a person, and be treated as morally equal.

Conclusion

Given all of the arguments analysed it seems that the answer to my research question of "Could an artificially intelligent entity be deserving of moral status?" is yes. There do not seem to be any inherent limitations in machines that prevent them from one day reaching the same or similar levels of moral status as humans.

However, this conclusion has been reached on the basis that any methods for distinguishing morally between artificial intelligence and other, moral, entities are seemingly insufficient. Therefore the question as to what criteria should be used to decide whether something is moral or not has not been answered.

The most prominent obstacle to my investigation into the research question was the limited scientific understanding as to the complex nature of consciousness and how the human brain produces it, as well as other psychological states like emotions. It is possible that there will someday be discovered something as of yet unknown that will prove the uniqueness of the human brain, or the limitations of artificial intelligence, rendering them incapable of moral status.

11

The essay points a rif and shows a way of arguing in relation to it. It is relevant in general but the conceptual analysis of the issues and ideas pointed is only partially developed is an in-depth and critical philosophical treatment of them.

Bibliography

Anderson, Kevin. *Predicting AI's Future.* 2001.
http://news.bbc.co.uk/1/hi/in_depth/sci_tech/2001/artificial_intelligence/1555742.stm (accessed July 30,
2011).

Anonymous. *In Vitro Fertilisation.* http://en.wikipedia.org/wiki/IVF (accessed July 30, 2011).

Can Robots 'Think' like Humans? October 13, 2008. http://news.bbc.co.uk/1/hi/technology/7666836.stm
(accessed July 30, 2011).

Carpenter, Rollo. *Cleverbot.* 2011. cleverbot.com (accessed July 29, 2011).

Churchland, Paul M. Churchland and Patricia Smith. "Could a Machine Think?" *Scientific American*, January
1990.

Cole, Davild. *The Chinese Room Argument.* September 22, 2009. http://www.seop.leeds.ac.uk/entries/chinese-
room/ (accessed july 30, 2011).

Davison, Paul. *How to Define Life.* October 17, 2008.
http://www2.una.edu/pdavis/BI%20101/Overview%20Fall%202004.htm (accessed October 21, 2011).

Diginfo. *Tosy TOPIO, Table Tennis Playing Robot.* December 5, 2007. http://www.diginfo.tv/2007/12/05/07-
0601-d.php (accessed October 21, 2011).

Earl, Dennis. *The Classical Theory of Concepts.* October 15, 2005. http://www.iep.utm.edu/conc-cl/ (accessed
October 21, 2011).

Hauser, Larry. *Artificial Intelligence.* June 8, 2007. http://www.iep.utm.edu/art-inte/ (accessed July 2, 2011).

Igor Aleksander, John Searle, interview by Melvyn Bragg. *Artificial Intelligence* BBC Radio 4. April 29, 1999.

John Crane, Jette Hannibal. *Psychology Course Companion.* Oxford University Press, 2009.

Law, Stephen. *The Philosophy Gym p58-67.* Headline Book Publishing, 2003.

Miedaner, Terrel. *The Soul of Anna Klane.* Church of Physical Theology, 1977.

Searle, John. *Mind, Language, and society.* New York: Basic Books, 1999.

Steven Rose, Dan Robinson, interview by Melvyn Bragg. *The Brain and Consciousness* (November 19, 1998).

Strawson, P. F. *Freedom and Resentment and other essays.* Routledge, 2008.

Stuart, Tony. *You and Your Life.* 2009. http://www.sevenoaksphilosophy.org/downloads/youandyourlife.pdf
(accessed July 29, 2011).

Turing, A. M. *Computing Machinery and Intellignece.* Mind Vol. LIX, 1950.

Waldron. *Environmental Ethics.* 2009. http://www.sevenoaksphilosophy.org/ethics/environmental-ethics.html
(accessed April 2011).