

Cybernetics 3.0: Ethical Regulators and Super-Ethical Systems

Mick Ashby
ethics@ashby.de

Abstract

This paper combines the Good Regulator Theorem with the Law of Requisite Variety and seven other requisites that are necessary and sufficient for a cybernetic regulator to be effective and ethical. The resulting Ethical Regulator Theorem provides a basis for systematically evaluating the adequacy of existing or proposed designs for systems that make decisions that can have ethical consequences; regardless of whether the regulators are humans, machines, cyberanthropic hybrids, organizations, corporations, or government institutions. The theorem has potentially far-reaching implications for society. An ethical design process is defined. A six-level framework is proposed for classifying cybernetic and superintelligent systems, which highlights the existence of a possibility-space bifurcation in our future time-line. The implementation of “super-ethical” systems is identified as an urgent imperative for the human race to avoid the danger that superintelligent machines might lead to a technological dystopia. Third-order cybernetics, a.k.a. Cybernetics 3.0, is defined as the cybernetics of ethical systems. Concrete actions, a grand challenge, and a vision of a super-ethical society are proposed to help steer the future of the human race and our wonderful planet towards a realistically achievable minimum viable cyberanthropic utopia.

Keywords: Ethical Design, Superintelligence, AI Control Problem, Third-Order Cybernetics, Cybernetics 3.0, Cyberanthropic Utopia

Introduction

The need for cybernetics to embody ethical values has been recognized and discussed by many cyberneticians, but second-order cybernetics has no formal repeatable process for designing systems that behave ethically. Instead, it relies on the ad hoc skills of an ethically-motivated designer to somehow specify a system that is hopefully ethical. This is not a satisfactory solution to a problem that so desperately needs to be solved. In the context of cybernetics, this could be referred to as "The Ethics Problem".

Many people think that all technologies can be used for good or evil, but this is not true. Let us consider a system like that of public health inspections of restaurants, where the inspector performs a well-structured system of evaluations in defined dimensions, such as food storage, signs of vermin, kitchen hygiene, and waste management, to identify any inadequacies and specify necessary improvements to achieve certification of hygienic adequacy. Such a fundamentally ethical system can only help to make restaurants more hygienic. Might it be possible to adapt this certification model from the public health domain to create a system that can be used to certify whether a system is ethically adequate? And might such a system be a solution to “The Ethics Problem”?

All cultures and societies regulate the behaviour of their members by defining what behaviour is acceptable or unacceptable. This is achieved by making laws, regulations, and rules. Not all rules are defined formally in writing, some are unwritten conventions, yet in every culture, it is unacceptable to break such rules, regulations, or laws.

In this paper, the terms “ethics” and “ethical” are used in a concrete sense of acceptable behaviour, and the philosophical approach to ethics as some nebulous human quality is rejected. There are important ethical dilemmas, such as whether it is acceptable to experiment on human embryos. But such dilemmas are only problematic because there is significant disagreement over them. If a society can reach agreement on any such matter, it becomes trivial to encode it in a law, regulation, or rule so that it can be understood and obeyed by humans and machines.

The Good Regulator Theorem (Conant and Ashby, 1970) is ambiguous because a regulator that is good at regulating is not necessarily good in an ethical sense. To avoid this ambiguity, this paper uses the term “effective” for the first meaning, “ethical” for the second meaning, and only uses “good” when both meanings are intended.

The Good Regulator Theorem proved that every effective regulator of a system must be a model of that system, and the Law of Requisite Variety (Ashby, 1956) dictates the range of responses that an effective regulator must be capable of. However, having an internal model and a sufficient range of responses is insufficient to ensure effective regulation, let alone ethical regulation. And whereas being effective does not require being optimal, being ethical is absolute with respect to a particular ethical schema.

The Ethical Regulator Theorem

The Ethical Regulator Theorem claims that the following nine requisites are necessary and sufficient for a cybernetic regulator to be effective and ethical:

1. **Purpose** expressed as unambiguously prioritized goals.
2. **Truth** about the past and present.
3. **Variety** of possible actions.
4. **Predictability** of the future effects of actions.
5. **Intelligence** to choose the best actions.
6. **Influence** on the system being regulated.
7. **Ethics** expressed as unambiguously prioritized rules.
8. **Integrity** of all subsystems.
9. **Transparency** of ethical behaviour.

Of these nine requisites, only the first six are necessary for a regulator to be effective. If a system does not need to be ethical, the three requisites ethics, integrity, and transparency are optional. Figure 1 and the following sections explain the requisites in more detail.

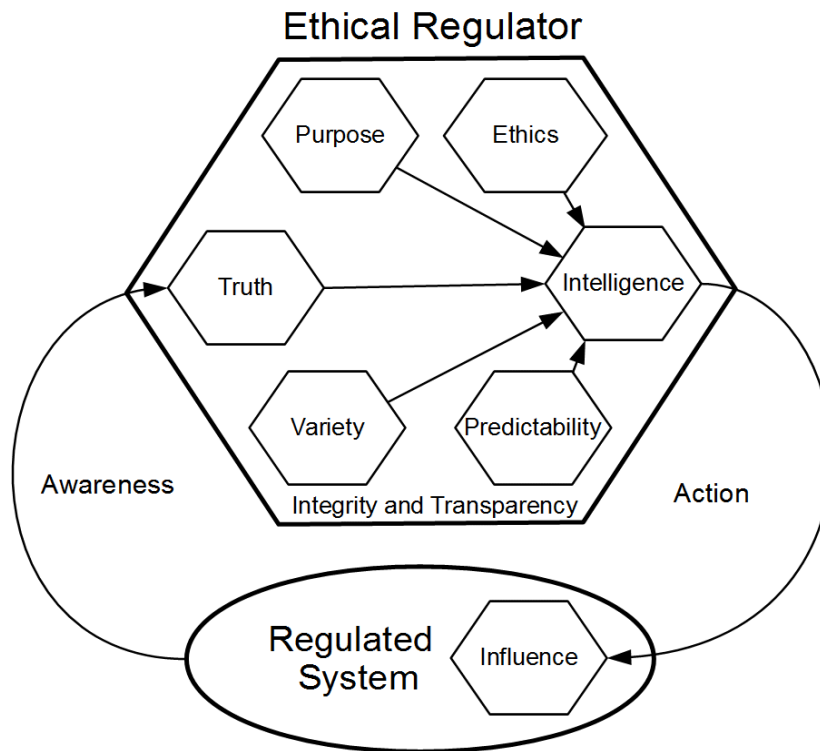


Figure 1: The Ethical Regulator System

Requisite Purpose

Because complex systems are required to satisfy multiple goals, **Purpose** must be expressed as unambiguously prioritized goals. Without well-defined goals, the system cannot be effective and might randomly adopt or default to a goal that is unethical.

Requisite Truth

Truth is not just about information that the regulator treats as facts or receives as inputs, but also the reliability of any interpretations of such information. This is the regulator's awareness of the current situation, knowledge, and beliefs. If the regulator's information sources or interpretations are unreliable, and cannot be error-corrected, then the integrity of the system is in danger. In extremis, if the perceptions of the regulator can be manipulated, it can be tricked into making decisions that are ineffective or unethical.

Some might claim that objective truth is unattainable, therefore requisite truth is unattainable, therefore no ethical regulator can exist. However, this argument is a fallacy. An ethical regulator doesn't require perfectly accurate information, rather it must be able to cope ethically with uncertainties and minimize the impact of unreliable information, misinterpretations, and deliberate misinformation as best as it can. This is much like the requirement that a good judge (effective and ethical) must be able to reach reliable verdicts "beyond reasonable doubt" from unreliable evidence.

Requisite Variety

Variety in the range of possible actions must be as rich as the range of potential disturbances or situations. This is nothing other than the Law of Requisite Variety.

Requisite Predictability

Predictability requires a sufficiently accurate model of the system being regulated, including the regulator, that can be used to rank the actions and strategies that will give the best outcome. This is nothing other than the Good Regulator Theorem.

Requisite Intelligence

Intelligence must be applied to the previous five requisite types of information to select the most rational and effective ethical action from the set of possible actions.

Requisite Influence

Influence is the existence of pathways to transmit the effects of the selected actions to the regulated system. This is not a property of the regulator itself, but a function of the connectivity relationships that span from the regulator's outputs to elements of the regulated system and its environment.

A regulator that is isolated from influencing the regulated system is not a true regulator, it is an observer or simulation, and there are no direct ethical consequences; which can be important when observing or simulating dangerous situations.

Depending on the nature of the system being regulated, the speed of the effect of actions can vary greatly. For example, a self-driving vehicle applying the brakes; the Supreme Court issuing a ruling; or someone sending a message to a complex network of amplifying and variable-delay transmission repeaters, known as Twitter followers.

In some systems, influence is more of a determining factor than variety. Indeed, the power of the Law of Requisite Variety has often been overstated, for example, claiming that the subsystem with the most variety will control a system. This is not always true.

Let us consider two systems, A and B, that are competing to win control of system C, for example, two politicians seeking election. Often the variety of statements, actions, and strategies of the candidates is less important than their ability to purchase advertising to influence the voters. And if a robber uses a gun to increase his effectiveness, the use of a gun does not amplify his variety, it is just one existing element in his range of possible variety, yet making that choice greatly increases his effectiveness at controlling his victims. Such an increase in effectiveness, like buying advertising, is best explained in terms of an increase in influence. In the light of the concept of influence, the idea that variety can be amplified is as delusional as the idea that randomness can be amplified. The variety of the robber or an advertising message is effectively constant.

The six requisites described so far are necessary and sufficient for a system to be effective.

Effectiveness Function

The Ethical Regulator Theorem implies that we can define a function for the effectiveness that a regulator, R, has in controlling a system to achieve a given goal that is defined by Purpose_R. It shows that the effectiveness actually depends on the quality of five requisites:

$$\text{Effectiveness}_R = \text{Truth}_R \times \text{Variety}_R \times \text{Predictability}_R \times \text{Intelligence}_R \times \text{Influence}_R$$

In this form, each requisite is assigned an effectiveness value between 0 and 1, where 1 means that it is perfect or optimal. And if the effectiveness of even just one of Truth_R, Variety_R, Predictability_R, Intelligence_R, or Influence_R is close to zero, the effectiveness of regulator R is massively reduced. Applied to our two politicians: If Effectiveness_A is greater than Effectiveness_B, then A is more likely than B to win control over system C.

Although the effectiveness function is stated pseudo-mathematically, it is generally neither necessary nor possible to calculate meaningful numerical values in order to compare the effectiveness of different systems or configurations. The essential value of the function is to understand the relationships and dependencies that it captures.

It is sufficient if an intuitive understanding of the effectiveness function informs the system design strategy; recognizing that a maximally effective system requires that the quality of these five requisite dimensions are maximized, and that a successful attack on the integrity of any of them spells disaster for the effectiveness of the whole system.

It is worth noting that in social systems, money can buy media influence; which in turn, if the media is broadcasting advertising, lies, or propaganda, reduces the quality of Truth_x that is received by every consumer or voter, X, which can manipulate them into making decisions that are not in their best interest.

Requisite Ethics

Ethics must be expressed as unambiguously prioritized rules, regulations, and laws that codify ethical values in a human-readable form, for example, Isaac Asimov's First Law of Robotics: "A robot may not injure a human being or, through inaction, allow a human being to come to harm." (Asimov, 1942).

Ethical goals have a higher priority than goals for purpose. By always obeying the relevant highest priority ethical imperatives, the regulator is guaranteed to act ethically within the scope of the ethical schema.

Because ethical schemas vary between different cultures, in machines, they must be handled as plug-ins. And because an ethical schema can encode any ethics, good or evil, each ethical schema must be anchored in the laws of a particular legislative jurisdiction.

For example, a self-driving vehicle must avoid harming people and animals. In the U.S. and Europe, protecting a person is a higher ethical priority than protecting a cow. But, in India, where cows are sacred, the opposite priority might be ethically preferable and crossing a state or national border might require activating a different ethical schema.

A taxonomy of ethics modules can provide ethical coverage for all conceivable ethical situations. For example, medical, child-care, traffic-rules, gun-law, tax-law, contract-law, maritime-law, drone-flying, police-regulations, operating-nuclear-power-stations, and warfare-rules-of-engagement.

Ethics modules can be treated like device drivers, so that to be fully operational, a hypothetical gun-carrying robot that can drive on roads requires valid ethics modules for gun-law and traffic-rules. Without both ethics modules for the appropriate legal jurisdiction, the robot's gun or driving capabilities are automatically disabled.

By legislating that all autonomous artificial intelligence (AI) systems must include and obey appropriate ethics modules that are issued by an organization that is run by humans, we can establish a control mechanism that should ensure that intelligent machines are always subject to human ethics; without unduly restricting the freedom of AI researchers. In fact, it will free AI researchers and knowledge engineers to focus on the more challenging requisites of truth, predictability, and intelligence.

When we introduce ethics, the effectiveness function must be modified because the effect of behaving ethically is that it reduces the variety of options that are available, by removing all possibilities that are unethical.

Thus, if A is an ethical politician and B is an unethical politician, we get:

$$\begin{aligned} \text{Effectiveness}_A &= \text{Truth}_A \times (\text{Variety}_A - \text{Ethics}_A) \times \text{Predictability}_A \times \text{Intelligence}_A \times \text{Influence}_A \\ \text{Effectiveness}_B &= \text{Truth}_B \times \text{Variety}_B \times \text{Predictability}_B \times \text{Intelligence}_B \times \text{Influence}_B \end{aligned}$$

Which captures the reality that politicians and businessmen who lie and cheat have an advantage over ones that are ethical.

Requisite Integrity

Integrity of the regulator and its subsystems must be assured through features such as resistance to tampering, intrusion detection, and cryptographically authenticated ethics modules. Monitoring mechanisms must detect if any invalid ethics modules are being used or if an ethical imperative is violated, and if necessary, activate an ethical fail-safe mode, preserve evidence, and notify the manufacturer and/or the appropriate authorities.

The regulator's first-order integrity mechanisms offer no protection to the pathways on which the regulator depends to influence the system. This poses a potential vulnerability that can only be mitigated by using the situation awareness closed-loop feedback to check for evidence of the effect of each action. For example, you can be confident that your tweet was received and distributed by Twitter when you see that it has been retweeted.

Requisite Transparency

Transparency is defined by introducing **The Law of Ethical Transparency**, which states "For a system to be truly ethical, it must be possible to prove retrospectively that it acted ethically with respect to the appropriate ethical schema."

Whereas it doesn't really matter whether the programmers of a chess playing robot can find out why a piece was sacrificed, the logic of ethical decisions must never be hidden in the depths of opaque processes or lost to the passage of time. Generally, this requisite can be satisfied by keeping audit trails that are adequate and secure.

When an ethically adequate system violates an ethical imperative, as they sometimes will, analysis of the audit trail will identify the reason. For example, because a faulty neural network wrongly identified a boy leading a cow as a calf leading a man, or it will prove whether a CEO was lied to about illegal corporate activities.

Integrity and **Transparency** are codependent because we require both integrity of transparency and transparency of integrity.

Evaluating Ethical Adequacy

Like a public health inspection of a restaurant, an evaluated system is judged on the adequacy of each requisite dimension. Only systems that completely satisfy all nine requisites are said to be “ethically adequate”. Otherwise they are classified as “ethically inadequate” and the weaknesses listed with recommendations for improving them.

The evaluation of ethical adequacy also has similarities to how a Red Team performs network penetration testing; where the evaluation team tries to identify weaknesses and theoretical possibilities to subvert the integrity of the system and all its subsystems.

The theorem cannot be used to certify that an ethical schema is ethical because schemas can vary arbitrarily between different cultures. However, automatic proof algorithms will detect certain types of errors before an ethical schema is packaged as an ethics module.

Perhaps, in the near future, accredited ethical consultants will specialize in auditing and certifying the ethical adequacy of existing and proposed systems, products, and processes.

Ethical Regulator Theorem Proof

Now that we understand the nine requisites better, is it possible to prove that they are indeed necessary and sufficient for a cybernetic regulator to be effective and ethical?

Proof of Necessity

Proving necessity is simple: One-by-one, for each of the nine requisites, ask yourself the question “Can a regulator be effective/ethical without requisite X?”, i.e. “Is requisite X necessary?” For example, “Can a regulator be effective/ethical without Truth?”

The answer in each case is rather obvious, especially if you refer to Figure 1 and, one-by-one, cover each requisite using your thumb, and then consider whether the resulting system can be effective or ethical. Table 1 summarizes the results, which confirm the necessity claims, including that ethics, integrity, and transparency are optional for systems that only need to be effective.

Table 1: Results of proof of necessity “by thumb”

Requisite	Necessary to be effective?	Necessary to be ethical?
Purpose	Yes	Yes
Truth	Yes	Yes
Variety	Yes	Yes
Predictability	Yes	Yes
Intelligence	Yes	Yes
Influence	Yes	Yes
Ethics	No	Yes
Integrity	No	Yes
Transparency	No	Yes

For effectiveness, 5/6 of the positive results for necessity correspond to the solutions for when $Effectiveness_R = 0$, for example, when $Truth_R = 0$, but not when $Transparency_R = 0$

I.e. $Effectiveness_{SR} = Truth_R \times Variety_R \times Predictability_R \times Intelligence_R \times Influence_R = 0$

This agreement with Table 1 is unremarkable because the effectiveness function was constructed from the results of posing precisely that necessity question for each requisite. So the agreement does not confirm the correctness of the theorem, but by performing this exercise yourself, you can confirm the correctness of the effectiveness function.

Proof of Sufficiency

Proving that the nine requisites are sufficient, is not so simple. First, let us assert that in the real world, effective systems and ethical systems exist. Now, for all those such systems, do any of them rely on any information, ability, or other factor to achieve effectiveness or ethicalness that is not covered by the nine requisites?

I claim that for all such systems that I have considered, the answer is no. However, this claim is easily refutable because it will only take one person to find one example of a necessary factor that is not covered by the nine requisites to demolish the claim of sufficiency. In the event of that happening, I will obviously adapt the theorem, if necessary adding another requisite, then reassert the sufficiency claim, and issue the challenge: “Okay, *now* find one!” ;-)

So although it is impossible to prove that such an exception does not exist, we can claim that it will always be possible to extend the theorem to include any missing requisites that might be identified in the future, thus restoring the validity of sufficiency for all known systems that have been considered. No proof is offered for this last claim because that is the nature of the scientific method.

The Law of Inevitable Ethical Inadequacy

We can derive this new law from the Ethical Regulator Theorem:

“If you don’t specify that you require a secure ethical system, what you get is an insecure unethical system.”

The reason is because when ethical adequacy is not specified as a requirement for a system design, the resulting design phase will tend to optimize for effectiveness and maximally avoid the extra costs that would be incurred by implementing the ethical, integrity, and transparency dimensions, which are optional for a system that only needs to be effective, thus guaranteeing that the resulting system is ethically inadequate and vulnerable to manipulation; by design.

Ethical Design Process

Figure 2 shows the generic elements of a typical design process, in which an analysis phase produces a requirements artifact, which is the input to the design phase that produces a specification artifact, which is used as the input to the implementation phase, which builds the system.




Figure 2: Ethically Inadequate Design Process

If a problem is found in the requirements during the design phase, feedback can trigger another iteration of the analysis phase. And if a problem in the specifications is found during implementation, feedback can trigger the design team to update the specifications or pass the feedback to the analysis team to update the requirements.

Such a design process can be effective at producing systems that are effective, however, because the design process is ethically inadequate, it will inevitably only produce systems that are ethically inadequate; and we cannot be sure that the resulting systems are not actually ethically evil; by accident, or by design :-)

Fortunately, we can transform any effective but ethically inadequate design process to make it ethically adequate by simply adding ethical adequacy acceptance testing of the requirements and specifications. How we can retrofit the Ethical Regulator Theorem (ERT) to any effective design process that produces requirements and specifications before the implementation phase starts is shown in Figure 3.

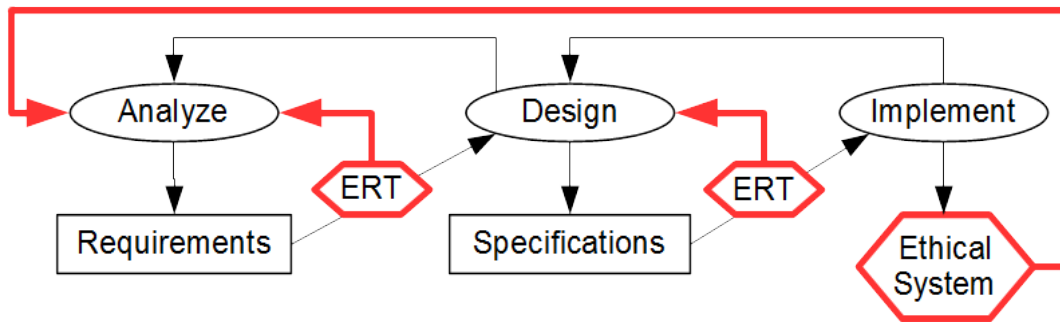


Figure 3: Ethically Adequate Design Process

Like any other quality assurance testing, the ERT evaluations must be performed by a Red Team that was not involved in the production of the artifact being tested. If an artifact is found to be ethically inadequate, it is rejected and recommendations for fixing the problems are provided as feedback to trigger another iteration of that phase. If an artifact is found to be ethically adequate, the artifact is accepted and passed onto the next phase. Because the two ERT testing steps ensure that the requirements and specifications are ethically adequate, if the implementation process performs an effective and lossless transformation of the specifications, the resulting system will also be ethically adequate.

This means that instances of the resulting system that are deployed in the real-world will include an integrity monitoring mechanism that detects and reports any significant problems as feedback to the analysis team, which must decide whether it necessitates updating the requirements, redesigning the specifications, a reimplementation, a remote configuration update, and/or instructing the system to activate an ethical fail-safe mode. Only if the system fails to enter its fail-safe mode might it be necessary for it to be either deactivated using a kill switch or “retired” by a blade runner! ;-)

This concludes the description of the theorem and how to use it. The next 20 pages explore the implications; leading to an inevitably ethical imperative call to action.

Legislative Implications

By creating a well-defined interface for coding ethics, it becomes easier to apportion legal liability for failures. For example, if a self-driving car crosses the border into India, fails to switch to the Indian government certified ethics module for traffic-rules, and decides to hit a cow to avoid hitting a person, then the car manufacturer can be held liable for the crime of killing a sacred animal. But if the correct ethics module was activated, but the “don’t hit cows” rule had an incorrectly low priority in the ethics schema, then the car manufacturer would not be liable.

It is only a matter of time until the laws and regulations of most countries are available in a standardized XML format such as LKIF (Legal Knowledge Interchange Format), and cryptographically-signed by an official issuing authority.

However, the existing governmental and regulatory organizations are inadequate for completing such an undertaking in the necessary time frame. Perhaps, a non-profit organization without any conflicts of interests could define appropriate standards, and start an open-source ethics coding project for the rules and laws that are most urgently required by the ethically adequate systems that we try to construct.

By standardizing ethics modules, systems from different manufacturers will use identical ethics modules that are issued by central ethics authorities. The concept of central ethics authorities might sound like part of a dystopic dictatorship, but acting ethically is mostly just a matter of obeying rules, regulations, and laws, which are a normal and necessary part of every stable society.

Like Microsoft Windows operating system updates, when new laws, rules, regulations, or bug fixes to a previous ethics module are released, the new ethics module can be made available to all affected autonomous systems; crucially, including systems whose manufacturer has gone out of business.

By comparison, Google's Android operating system update mechanism is a classic example of the Law of Inevitable Ethical Inadequacy. Because Android was designed only to be effective, not ethical, Google delegated the responsibility for issuing Android operating system patches to the device manufacturers. This inevitably and predictably resulted in the current situation where hundreds of millions of Android devices will never receive any security patches, by design, for the simple reason that Google prioritized its profits over ethical consumer safety.

They could have designed it differently. And if we can't trust Google, who can we trust?

Such unethical corporate behaviour must be legislated out of existence, otherwise it will keep repeating itself in millions of different and damaging ways. For example, ethically inadequate Internet-of-Things devices that send unencrypted data over the internet, are vulnerable to being hacked, and will never receive security patches. Importing or selling such unethical devices that threaten our privacy and the security of our digital infrastructure should be as illegal as selling exploding cars.

We certainly don't want robots, self-driving vehicles, and autonomous weapons systems relying on an ethics update mechanism that stops working when the manufacturer goes out of business or decides to optimize its profits at the expense of safety updates.

Classification Framework

Now let's consider where the Ethical Regulator Theorem fits into the existing cybernetics framework. One might assume that it belongs in second-order cybernetics, however, in a 1991 conference plenary presentation, Heinz von Foerster implied that combining ethics and second-order cybernetics is not something that he would have suggested:

“I am impressed by the ingenuity of the organizers who suggested to me the title of my presentation. They wanted me to address myself to 'Ethics and Second-

Order Cybernetics'. To be honest, I would have never dared to propose such an outrageous title, but I must say that I am delighted that this title was chosen for me.” (von Foerster, 2003)

Table 2 lists some of the cybernetic community’s definitions of first- and second-order cybernetics, as summarized by Stuart Umpleby (2001).

Table 2: Definitions of first- and second-order cybernetics

Author	First-Order Cybernetics	Second-Order Cybernetics
von Foerster	The cybernetics of observed systems	The cybernetics of observing systems
Pask	The purpose of a model	The purpose of the modeler
Valera	Controlled systems	Autonomous systems
Umpleby	Interaction among the variables in a system	Interaction between observer and observed
Umpleby	Theories of social systems	Theories of the interaction between ideas and society

Although every one of these definitions captures an important distinction, when compared to the rigorous precision with which other scientific communities use the qualifiers “first-order” and “second-order”, the cybernetic community’s use of “first-order” and “second-order” appears to be rather subjective, lacks the consensus that is required by the scientific principle, and is of little utility (Kuhn, 1962).

This incoherence in defining cybernetics as first-order and second-order not only prevents it from being useful to classify different types of systems and dissipates intellectual energy, but it also prevents the classification from being extended to higher orders, which can be viewed as either a self-limiting dead-end, or paradigmatic autoapoptosis (self-programmed death), which is not entirely unlike the situation of the members of the Heaven’s Gate millennial death-cult, who believed that by committing suicide, they would be rescued by an alien spacecraft and “graduate to the Next Level”.

To illustrate the problem of classifying cybernetics into observer-centric “orders”, let’s start by considering first-order cybernetics, which is concerned with a system, S, that is studied by an observer, as illustrated in figure 4.




Figure 4: First-Order Cybernetics

Second-order cybernetics introduces a second observer's viewpoint, as shown in figure 5.




Figure 5: Second-Order Cybernetics

How does this paradigm help us predict the future of cybernetics? Logically, third-order cybernetics would add a third observer's perspective, as shown in figure 6.




Figure 6: Third-Order Cybernetics

However, from the perspective of the third observer, this looks more like psychology than cybernetics. In fact, this structure is isomorphic to a typical management team evaluation exercise, where the details of the task that is given to the team to work on is virtually irrelevant to the outermost observer. It can be any goal-oriented activity, such as building the highest stable tower possible from a limited set of Lego bricks, solving an impossible puzzle in a limited amount of time, or studying a first-order cybernetic system.

At even higher orders, with N observers, it becomes even more compelling that cybernetics would stop being cybernetics and become psychology.

New Classification

It could be of more utility to define “levels” of cybernetic systems that include categories of future systems that are already anticipated, and associate each level with established concepts. To that end, the following six-level framework for classifying cybernetic and superintelligent systems is proposed.

Table 3: Six-level framework for classifying cybernetic and superintelligent systems

Level	The cybernetics of	Also known as	The cybernetician
1	Simple systems	First-order cybernetics	Observes the system
2	Complex systems	Second-order cybernetics	Participates in the system
3	Ethical systems	Third-order cybernetics (or Cybernetics)	Designs the system

Level	The cybernetics of	Also known as	The cybernetician
4	Superintelligent systems	Technological singularity	Stares incredulously, as the system redesigns itself
5	“Super-Ethical” systems (Superintelligent and ethically adequate)	Technological utopia	Is protected by the system
6	“Super-Unethical” systems (Superintelligent and ethically inadequate)	Technological dystopia	Obeys the system

Today, we are in the transition from building complex cybernetic level two systems (CL2) to building ethical systems and superintelligent systems of cybernetic levels three and four (CL3 and CL/4), and the future of our species and our fragile ecosystem is in our hands, but first, let's clarify each level and explore where this new framework leads us.

Cybernetic Level 1: Simple Systems

This is the domain of first-order cybernetics: Studying and designing simple systems that are effective.

Cybernetic Level 2: Complex Systems

This is the domain of second-order cybernetics: Studying and designing complex systems that are effective. There is still much valuable and important work to be done at this level.

Cybernetic Level 3: Ethical Systems

Decades ahead of his time, it was the wonderful and inspiring Ranulph Glanville who defined “cybernetics” as “the cybernetics of ethics and the ethics of cybernetics” (Glanville, 1986).

The Ethical Regulator Theorem belongs at this level, which is concerned with designing man-made systems that are ethically adequate. Such systems must satisfy all nine requisites of the Ethical Regulator Theorem and the regulating agents can be humans, machines, cyberanthropic hybrids, organizations, corporations, or government institutions. Ethically adequate machines must obey certified ethics modules.

In retrospect, now that we're not extrapolating from just two points in concept-space, if level three cybernetic systems are ethical, it's apparent that the third observer in the third-order cybernetics system of Figure 6 is not necessarily a psychologist or a lost cybernetician, but could be the second observer's conscience; her super-ego, or higher-self; that constantly self-observing sense that we all have that knows the difference between right and wrong; between good and evil. This self-monitoring mechanism is known as integrity, and is something that today's ethically indifferent scientists, politicians, CEOs, lawyers, bankers, and billionaires are woefully lacking. In non-psychopaths, it is integrity that triggers feelings of bad conscience, regret, or guilt if it is ignored.

Cybernetic Level 4: Superintelligent Systems

The technological singularity is a hypothetical moment when a self-improvement process in a machine causes runaway improvements in intelligence that results in superintelligence that is far greater than any human mind (Bolstrom, 2014). For this to happen, the system must be sufficiently self-aware to understand its own software and/or hardware.

Superintelligence Tests

These levels of self-awareness give rise to three levels of superintelligence tests. The ability to reprogram better software for itself, the ability to redesign better hardware for itself, and the ability to do both.

Together with the Turing Test (Turing, 1950), these tests mark milestones in the evolution of AI systems towards superintelligence, and should cause us alarm if progress towards them is made without significant progress creating ethical systems first. Of these tests, the Turing Test is the easiest to achieve, because it only requires that a computer can imitate a (not necessarily very intelligent) human sufficiently well to convince humans most of the time that it is a human being, and does not require self-awareness or runaway improvements in intelligence.

Prophecies of Possible Futures

In 1951, Ross Ashby started considering an “Advanced society planned as a super brain” (Ashby, 1951). Later, he wrote in his journal that super-clever machines could create a technological utopia: “It may be found that we shall solve our social problems by directing machines that can deliver an intelligence that is not our own.” (Ashby, 1952a).

Two pages later, he described a technological dystopia where a “Million I.Q. Engine” sounds like Google on steroids: “What people could resist propaganda and blarney directed by an I.Q. of 1,000,000? It would get to know their secret wishes, their unconscious drives; it would use symbolic messages that they didn’t understand consciously; it would play on their enthusiasms and hopes. They would be as children to it. (This sounds very much like Goebbels controlling the Germans).”.

On the appearance of such a machine, he described a paradox of perception of higher intelligence: “It seems, therefore, that a super-clever machine will not look clever. It will look either deceptively simple or, more likely, merely random.” (Ashby, 1952b). On the same subject, Arthur C. Clarke’s Third Law states: “Any sufficiently advanced technology is indistinguishable from magic.” (Clarke, 1973). If you think that Clarke’s “magic” and Ashby’s “deceptively simple or merely random” are incompatible; take a moment to reflect on the magical simplicity and “randomness” of a Las Vegas magic show or Google’s search results pages.

Just as there are two diametrically opposite archetypes for genius; namely the benevolent good genius and the nasty evil genius, it is important not to conflate systems that are ethical with ones that are not ethical, by making them share the same name or category, such as “superintelligent”. To do so, would focus attention on the least important dimension, and ignore the most important dimension; good and evil.

Cybernetic Level 5: Super-Ethical Systems

The term “super-ethical” is proposed to refer to superintelligent systems that are ethically adequate. Of course, by the time that super-ethical systems exist, a friendlier name will have emerged and the term “super-ethical” will seem quaintly archaic.

Cybernetic Level 6: Super-Unethical Systems

The term “super-unethical” is proposed to refer to superintelligent systems that are ethically inadequate. This term should always carry a certain stigma, like “weapons of mass destruction”. No one who is working to create artificially intelligent systems should be allowed to escape admitting whether the systems are ethically inadequate.

Just as human genetic experimentation is strictly ethically regulated, we need legislation, regulation, standards, and certification to ensure that autonomous AI systems that make decisions that can have ethical consequences are subjected to the same kind of obsessively rigorous safety-oriented design, construction, and operating procedures as commercial aircraft, nuclear power stations, and vehicles that carry humans into space.

One could start arguing that intelligence is ethically neutral, and it is, but that family of arguments are fallacies because a hyper-genius “Million I.Q. Engine” without ethics is not ethically neutral. It should be treated like a bomb that could destroy our planet. Even just planning to construct such a device is effectively conspiring to commit a crime against humanity.

As a thought experiment, let’s imagine a hypothetical super-unethical version of Google, named the Googlevil Corporation. The CEO is Dr. Evil, and both the CEO and the corporate AI are without ethics, avoid transparency, and will do anything to maximize their profits and power. The corporation’s secret mission statement is “Collect and organize the world’s personal information and make it accessible and useful for maximizing our profits, power, influence, and ability to avoid paying taxes.” and its secret corporate mantra is “Say ‘Don’t be evil’ then do it anyway.”.

Anytime that the super-unethical Googlevil artificial intelligence or the psychopathic demagogue Dr. Evil wants to blackmail the CEOs of other corporations, politicians that can’t be bought, jury members, or Supreme Court justices around the world to make “random” decisions that incrementally further their secret mission, would they have to do anything more than query the Googlevil user-profile database? In theory, they would only need to be able to blackmail a majority of members of lower- and upper-houses (how hard can that be?) to be able to get any legislation that they want in any country. Or just a few Supreme Court justices to steer a nation into a fascist dystopia.

By the time that super-unethical AI systems exist, they could be indistinguishable from the corporations that they belong to. They will be immoral, immortal, enjoy legal personhood, pay no taxes, and make unlimited donations (a.k.a. bribes) to all Googlevil-friendly political parties in all techno-democratic dystopias on the planet.

Future Time-Line Bifurcation Race Condition

At this point in time, there is a possibility-space bifurcation in our future time-line. Depending on whether the systems that achieve the singularity are ethically adequate or not, the runaway increase in intelligence and inevitable ethical polarization pressures will result in one of two outcomes:

- Good super-ethical AIs protect the human race.
- Evil super-unethical AIs dominate the human race.

Figure 7 illustrates how plotting the ethical dimension orthogonally to the intelligence dimension clarifies the dependencies between different cybernetic levels, and clearly shows that the ethically inadequate superintelligent systems of cybernetic level four-minus (CL4-) have no dependency on us first achieving ethical systems of cybernetic level three (CL3).

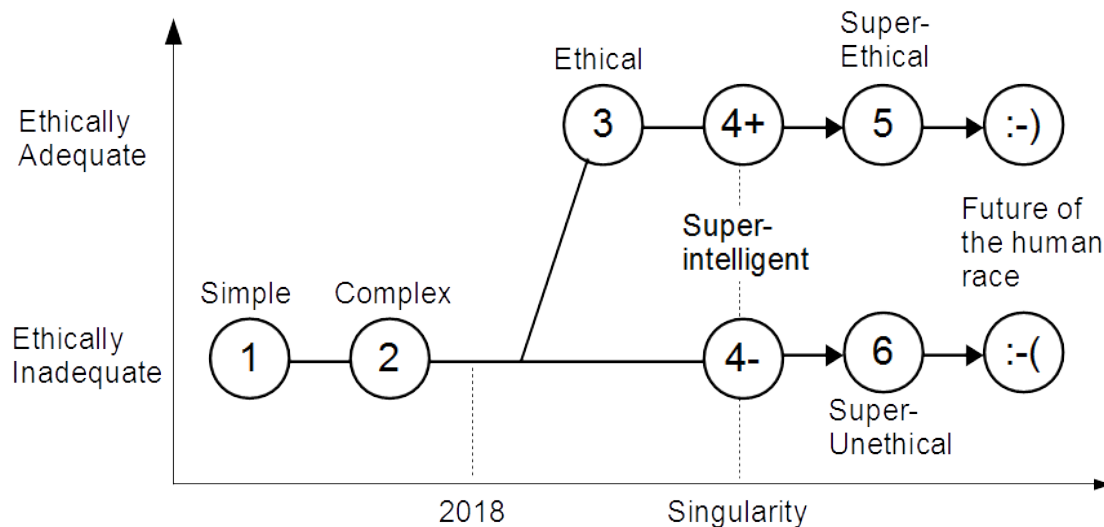


Figure 7: Two mutually exclusive possible futures

If we take the direct route from complex systems (CL2) to superintelligent systems that are ethically inadequate (CL4-), we quickly arrive at a dystopia that is ruled by super-unethical systems (CL6), and the potential utopia of being ruled by benevolent super-ethical systems (CL5) becomes permanently unreachable.

So there is a race condition that will determine which of two mutually exclusive possible futures will be the fate of our species; will our technological progress reach CL3 or CL4-first? And will legislators regulate such developments ethically and adequately, or will they sell us out to Dr. Evil's special interest lobby groups and think-tanks that will campaign vigorously for "self-regulation" — and we all know what that *really* means.

It cannot be overemphasized that CL4± is the point-of-no-return where humans could lose control over machines that become our intellectual superiors. And this is the window of opportunity to ensure that superintelligent machines are programmed with ethics and purposes that serve the greater good of humanity and our fragile ecosystem.

In this context, it is now clear that the ultimate grand challenge for cybernetics and third-order cyberneticians is to find ways to build ethical and super-ethical systems, avoid a technological dystopia, and help the human race create a super-ethical society.

Super-Ethical Society

Imagine how different the world would be:

- If we were happy to be ruled by benevolent super-ethical artificial intelligences that eliminated poverty, environmental destruction, global warming, corruption, and injustice.
- If the United Nations could deploy heavily armed super-ethical peace-keeping robot armies into conflict zones to protect civilians and enforce ceasefires.
- If our towns and cities were policed by super-ethical robots that protect all citizens equally, 24x7, and never shoot our friends or family because of their race, religion, social class, lifestyle, or peaceful protesting.
- If super-ethical child-care robots accompanied our children wherever they go, protecting them from danger, including physical, emotional, and sexual abuse.
- If ethically adequate corporations produce ethical products and services, and pay ethical levels of unavoidable corporation tax.

That future is possible; but only if we learn to recognize and act together in accordance with the fact: Ethics are a higher power for good that transcends science, politics, nations, and religions.

Cyberanthropic Utopia

Many utopias have been proposed that are really nothing more than naïve science fiction fantasies. In fact, utopias have accumulated such a bad reputation, that merely hearing the word causes some people to *literally* roll their eyes in ridicule; which is funny to watch!

However, now that artificial intelligence is making such impressive progress and showing no signs of slowing down nor having an upper-limit, Ross Ashby's 66 year old prediction looks increasingly realistic: That we might be able to “...solve our social problems by directing machines that can deliver an intelligence that is not our own.” (Ashby, 1952a).

But “Utopia” need not mean a “perfect” society, or that we all have flying cars, robot servants, and never have to go to work. Ashby's prediction hints at a possible definition of a minimum viable utopia: **A world where our social problems have been solved.** Just fulfilling human needs and eliminating poverty would create a truly awesome utopia. And as we start making progress achieving it, many other human problems, such as hopelessness, malnutrition, parasitic diseases, and poverty-driven prostitution and crime will fade away and the world will become a very different and happier place to live in.

Let no one say it cannot be done. Ethically adequate societies have existed in the past, where resources were shared and the environment respected. What is new is that we can now do it synthetically, consciously, deliberately.

ERT Universality

Anyone who has the impression that the Ethical Regulator Theorem (ERT) applies primarily to artificial intelligence, self-driving vehicles, robots, and autonomous weapons systems is urged to consider how the theorem can be applied to human systems that make decisions that affect people or the environment, such as CEOs, corporations, organizations, education systems, electoral systems, government institutions, or yourself.

Justice Stevens (2010) provides an excellent example of analysing the ethical inadequacy of the “Citizens United” ruling. His opinion that “The Court’s ruling threatens to undermine the integrity of elected institutions across the Nation.” implies that there is a pressing need to evaluate the ethical adequacy of the U.S. Supreme Court system.

ERT Reflexivity and Algebra

If the Ethical Regulator Theorem (ERT) is genuinely universal, it should give meaningful results for two special cases: When we apply ERT to itself, and when we apply ERT to second-order cybernetics (2oC).

First, let us define a convenient algebra that allows us to express important assertions in this domain. In particular, it is important to distinguish between the act of evaluating the ethical adequacy of a system and the act of determining the set of transformations or interventions that are necessary to make a system ethically adequate:

- I. A function, **IsEthical(S)**, returns the value True if system S is ethically adequate, it returns the value False if S is ethically inadequate, or it returns the value Undecidable if S is significantly inconsistent, contradictory, or paradoxical.
- II. A second function, **MakeEthical(S)**, returns a set of transformations or interventions to make system S ethically adequate. If S is already ethically adequate, the function returns an empty set.

Now we can use this ERT algebra to make some interesting and controversial claims:

Table 4: Some ERT algebra assertions

No.	Claim	Justification / Interpretation
1	IsEthical(ERT) = True	The ERT system fulfils all nine requisites of ERT, and is therefore ethically adequate. It can only be used to make systems more ethical.
2	MakeEthical(ERT) = { }	The ERT system is sufficient to be ethically adequate. Nothing else is required.

No.	Claim	Justification / Interpretation
3	IsEthical(2oC) = False	Second-order cybernetics is ethically inadequate. Unlike ERT, it has no intrinsic ethics or integrity, so it can be used to make ethically good or evil systems. It doesn't go beyond achieving effectiveness.
4	MakeEthical(2oC) = ERT	This claims that the set of ERT concepts is necessary and sufficient to make 2oC ethically adequate.
5	IsEthical(2oC + ERT) = True	There is nothing in 2oC that is incompatible with ERT. 2oC is just missing a few concepts that are necessary to be ethically adequate.
6	2oC + ERT = 3oC	Ethical systems are a new type of system that goes beyond being effective and does not belong in second-order cybernetics. Logically, the system that is created by joining the 2oC and ERT systems would be named third-order cybernetics (3oC).

Third-Order Cybernetics

Since Heinz von Foerster made the distinction between first- and second-order cybernetics in 1974, many people have attempted to find a worthwhile possible definition for third-order cybernetics, but until now, no definition has gained acceptance.

This paper proposes that Third-Order Cybernetics should be defined as “the cybernetics of ethical systems” and that “the cybernetics of ethics”, “3oC”, “Cybernetics 3.0” and “C3.0” are all acceptable synonyms for it.

Some of the supporting arguments for this proposal have already been mentioned, however a consolidated set of arguments that support this proposal are listed below:

1. Until now, second-order cybernetics (2oC) discussions about the need to create ethical systems, including the need for cybernetics itself to embody ethics, did not produce any satisfactory solutions. Here, “satisfactory solution” is understood to mean something that can be used systematically, for example by engineers, to create real ethical systems. Recognizing this need, but failing to fulfil the need could be referred to in the context of cybernetics as “The Ethics Problem”.
2. The fact that Heinz von Foerster described “Ethics and Second-Order Cybernetics” as an “outrageous title” for a presentation (von Foerster, 2003), implies that ethics do not belong in the 2oC that he envisioned.
3. Whereas 2oC can be used for good or evil, the Ethical Regulator Theorem (ERT) can only be used for good. This distinction is significant, and also implies that ERT does not belong in 2oC.

4. If we extrapolate from first-order cybernetics having one observer and 2oC having a second observer, we would logically expect 3oC to introduce a third observer. This hypothesized third observer maps exactly onto the ERT requirement that ethical systems must have integrity mechanisms that monitor and enforce compliance of the system with respect to an appropriate ethical schema.
5. Together, ERT and the six-level framework (SLF) for classifying cybernetic systems (see Table 3) create a new paradigm that has greater explanatory and predictive power than 2oC. For example, producing the ERT effectiveness function, the Law of Inevitable Ethical Inadequacy, identifying that the ethical systems of CL3 are the missing type of cybernetic system that is necessary to integrate cybernetic systems and superintelligent systems into a common framework, and explaining the impending bifurcation into either a technological utopia or a technological dystopia (see Figure 7). In addition, because SLF integrates three types of systems that do not yet exist, it can help us navigate a rational path into the future, for example, by predicting the existence of a race condition and thus identifying a possible solution to the dangers that are posed by superintelligent machines. 2oC cannot generate such insights.
6. Ethical systems constitute a new type of system, and ERT + SLF defines a new branch of cybernetics, which goes beyond effectiveness, solves “The Ethics Problem”, and does not belong in 2oC. Logically, the system that is created by joining the 2oC and ERT systems would be named third-order cybernetics or 3oC.
7. Although Ranulph Glanville already defined “the cybernetics of ethics and the ethics of cybernetics” as “cybernethics”, this term is invented jargon that carries no meaning for people who have not been exposed to its definition. By contrast, the term “third-order cybernetics” carries enough meaning for people who are familiar with the terms “cybernetics” or “second-order cybernetics” to at least trigger interest and curiosity. Therefore, using the term “third-order cybernetics” instead of “cybernethics” has significant advantages, and also enhances SLF by increasing symmetry in Table 3.
8. Whereas it is impossible to define objectively which theories and practices belong in 2oC, thus making it an intimidating subject for outsiders to even contemplate mastering, ERT is defined and proved in ten pages and doesn't require any knowledge of 2oC. This means that the theorem, and how to apply it to systems of any type, can easily be taught to non-cyberneticians. It is therefore imperative that ERT/SLF/3oC can make a fresh start without being burdened by 44 years of fuzzy 2oC baggage. However, 3oC is not limited to ERT and SLF, and will surely develop rapidly before it matures.
9. Unlike 2oC, 3oC has a fundamentally ethical purpose (making systems ethical) that together with the proposed grand challenge and the vision of a super-ethical society, create a unique opportunity for the cybernetics and systems sciences community to take a leading role in implementing a long-overdue and much needed systemic ethical revolution.

10. If someone makes the statement “I’m a second-order cybernetician”, it reveals nothing about the speaker's ethics. But from now on, anyone who is brave enough to declare “I’m a third-order cybernetician” is making a bold assertion that they are part of the only scientific movement that is dedicated to making the world a better place. And everyone is welcome to join.
11. Recognizing and declaring that 3oC is a new paradigm that is revitalizing and reunifying the 2oC community will help draw attention to the fact that something important and exciting is happening; a powerful attractor for energy and commitment. It is a genuine **ethical imperative** that we develop this new field for the other sciences, engineering, and management to use. And it is a task that cannot be performed by any of the more narrowly defined branches of science.
12. Many people thought that cybernetics had faded away after peaking in the 1960s or early 1970s, but the previous cybernetic peak was just a local maximum. Cybernetics has rebooted, a.k.a. Cybernetics 3.0, and it's going to be hard to ignore, because this time, we're bringing requisite Purpose, Truth, Intelligence, Influence, Ethics, Integrity, Transparency, and Humor ;-)

If these compelling arguments are accepted, it is suggested that going forwards, the terms 2oC and second-order cybernetics should only be used to refer to the second-order-cybernetics of effectiveness without the ethical, integrity, and transparency aspects, which belong to the 3oC layer that can be used to transform any effective but ethically inadequate system, such as second-order cybernetics, into an ethical system.

Our Future Epilog or Eulogy

We are approaching a decisive fork in the road in the evolution of intelligent machines, immortal corporations, political systems, and human society, and it is imperative that we learn to make these systems rigorously ethical before artificially intelligent machines reach the technological singularity, start to evolve exponentially, exceed human intelligence, and are used by ethically inadequate corporations to dominate the human race politically and economically.

For we are the generation that had the chance to steer the fate of future generations of humanity towards being collectively ruled, potentially for eternity, by benevolent super-ethical systems that create a stable cyberanthropic utopia for us, effectively and ethically minimizing environmental problems and human suffering, rather than allowing hubris and super-unethical systems to either enslave most of us in a cybermisanthropic dystopia or causing the extinction of our species to become a footnote in Gaia’s geological record.

The Path Forwards

To start steering the future of the human race and our wonderful planet towards becoming a stable cyberanthropic super-ethical society, I propose establishing an independent, non-profit ethics institute with ambitious goals that lie in the areas of research, development, standards, certification, legislation, and democracy.

Research and Development

The institute will promote theoretical and practical progress:

- Coordinate and fund research into creating ethical systems and making existing systems ethical.
- Develop a taxonomy of open-source ethics modules for different types of rules, regulations, and laws that can be used by anyone, free of charge.

Standards and Certification

The institute will create an ethical certification infrastructure:

- Establish standards for certifying the ethical adequacy of systems.
- Establish a curriculum for training accredited ethical consultants.
- Coordinate and regulate contracts for ethical audits and certifications.

Legislation and Democracy

The institute will lobby governments to implement ethically adequate legislation and will evaluate the adequacy of any proposed legislation. In particular, promoting the following:

- Regulate autonomous machines to require that their design and implementation is ethically adequate, and that they support compulsory ethics modules.
- Require that all new systems and processes are designed to be ethically adequate.
- Make it illegal to import or sell products that have not been certified as being ethically adequate, unless they are explicitly excluded from requiring certification.
- Extend political representation to every living human being, by giving parents proxy votes to cast on-behalf of their children who are too young to vote.

MakeEthical(yourself) = { ... }

As members of a human society, we are all cybernetic regulators; of ourselves and of each other. As a thought experiment, to become a more effective and ethical force for good, you could identify ways to improve each ethical requisite as it applies to yourself:

Table 5: Ways to become a better ethical regulator

Requisite	Example set of self-improvement interventions
Purpose	To clarify your purpose in life and help you to recognize your strongest motivating thoughts, write down your most important life goals: <ol style="list-style-type: none"> 1. 2. 3. 4.

Requisite	Example set of self-improvement interventions
Truth	To become a good judge (effective and ethical) of who tells the truth and who distorts it, seek a wide-spectrum of opinions by finding alternative information sources that are genuinely independent of your primary sources. Investigate any inconsistencies that you notice, modify the reputation of liars appropriately, and resolve to avoid them in future.
Variety	Brainstorm new actions, responses, and strategies that you have never previously considered.
Predictability	Improve your model of human behaviour by studying the following Wikipedia articles until you are competent at recognizing the patterns in yourself and others: List of cognitive biases Defence mechanisms List of fallacies Demagogue
Intelligence	Read a book or take a course on personal effectiveness or critical thinking.
Influence	Identify ways that you can increase your influence (on your family, friends, colleagues, or society) to achieve your life goals and promote your ethical values.
Ethics	Write down five undesirable, unethical, or disrespectful behaviours that, up until now, you have tolerated in other people, organizations, or corporations: <ol style="list-style-type: none"> 1. 2. 3. 4. 5. Now, next to them, write down five undesirable, unethical, or disrespectful behaviours that, up until now, you have tolerated in yourself. If you can't think of five things about yourself, read the Wikipedia article: Denial . If that doesn't help, ask someone that you live with to suggest five things that you do that they'd prefer you not to do.
Integrity	Seek to stop or prevent all the undesirable, unethical, or disrespectful behaviours that you listed under requisite ethics.
Transparency	Let other people know about the changes that you are making.

Finally, keep reviewing and refining your answers until they resonate with who you are and how you want your world to become.

Ethically Resonant Wisdom

If you distil different solutions that contain alcohol, you get pure alcohol that is free of impurities. And if you distil different religions and philosophies that contain ethical wisdom, you get pure ethical wisdom that is free of culturally-specific dogma.

Such ethical wisdom is universal, and resonates with all good people, regardless of their world-view, politics, nationality, or religion. And because pure ethics are a higher power for good that transcends science, politics, nations, and religions, it is probably the only force that can unify humanity to work together for our greater good.

Consider the following selected quotes:

Mahatma Gandhi (1869-1948):

- The future depends on what you do today.
- You must become the change you wish to see in the world.
- First they ignore you, then they laugh at you, then they fight you, then you win.
- The difference between what we do and what we are capable of doing would suffice to solve most of the world's problems.
- Happiness is when what you think, what you say, and what you do are in harmony.
- If I have the belief that I can do it, I shall surely acquire the capacity to do it even if I may not have it at the beginning.
- Capital as such is not evil; it is its wrong use that is evil.
- Non-cooperation with evil is as much a duty as is cooperation with good.
- Poverty is the worst form of violence.
- There is sufficiency in the world for man's need, but not for man's greed.
- There are people in the world so hungry, that God cannot appear to them except in the form of bread.
- Those who say religion has nothing to do with politics do not know what religion is.
- Where love is, there God is also.
- God has no religion.
- There is a higher court than the courts of justice and that is the court of conscience.
- They may torture my body, break my bones, even kill me. Then they will have my dead body, but not my obedience.

- Victory attained by violence is tantamount to a defeat, for it is momentary.
- What difference does it make to the dead, the orphans, and the homeless, whether the mad destruction is wrought under the name of totalitarianism or the holy name of liberty or democracy?
- Your beliefs become your thoughts, your thoughts become your words, your words become your actions, your actions become your habits, your habits become your values, your values become your destiny.

His Holiness Pope Francis:

- We must restore hope to young people, help the old, be open to the future, spread love. Be poor among the poor. We need to include the excluded and preach peace.
- Hatred is not to be carried in the name of God. War is not to be waged in the name of God!
- Human rights are not only violated by terrorism, repression, or assassination, but also by unfair economic structures that create huge inequalities.
- The worship of the golden calf of old has found a new and heartless image in the cult of money and the dictatorship of an economy which is faceless and lacking any truly human goal.
- Men and women are sacrificed to the idols of profit and consumption: It is the “culture of waste”. If a computer breaks it is a tragedy, but poverty, the needs and dramas of so many people end up being considered normal.
- Women in the church are more important than bishops and priests.
- All that is good, all that is true, all that is beautiful, God is the truth.
- We all have the duty to do good.
- Everyone has his own idea of good and evil and must choose to follow the good and fight evil as he conceives them. That would be enough to make the world a better place.

His Holiness the Dalai Lama XIV:

- All religious institutions, despite different philosophical views, all have the same message — a message of love.
- If you can, help others; if you cannot do that, at least do not harm them.
- The whole purpose of religion is to facilitate love and compassion, patience, tolerance, humility, and forgiveness.
- Irrespective of whether we are believers or agnostics, whether we believe in God or karma, moral ethics is a code which everyone is able to pursue.
- The ultimate authority must always rest with the individual’s own reason and critical analysis.

- A good friend who points out mistakes and imperfections and rebukes evil is to be respected as if he reveals the secret of some hidden treasure.
- The true hero is one who conquers his own anger and hatred.
- A lack of transparency results in distrust and a deep sense of insecurity.
- In our struggle for freedom, truth is the only weapon we possess.
- Where ignorance is our master, there is no possibility of real peace.
- Through violence, you may “solve” one problem, but you sow the seeds for another.
- Don’t ever mistake my silence for ignorance, my calmness for acceptance or my kindness for weakness. Compassion and tolerance are not a sign of weakness, but a sign of strength.
- A truly compassionate attitude toward others does not change even if they behave negatively or hurt you.
- I defeat my enemies by making them my friends.
- When you practice gratefulness, there is a sense of respect toward others.
- With realization of one’s own potential and self-confidence in one’s abilities, one can build a better world.
- If you think you are too small to make a difference, try sleeping with a mosquito.
- As people alive today, we must consider future generations: a clean environment is a human right like any other. It is therefore part of our responsibility toward others to ensure that the world we pass on is as healthy, if not healthier, than we found it.
- The ultimate source of happiness is not money and power, but warm-heartedness.
- The more you are motivated by love, the more fearless and free your action will be.
- Love and compassion are necessities, not luxuries. Without them humanity cannot survive.
- Be kind when possible. It is always possible.
- Love is the absence of judgement.

Dr. Martin Luther King Jr. (1929-1968):

- We must discover the power of love, the power, the redemptive power of love. And when we discover that we will be able to make of this old world a new world. We will be able to make men better. Love is the only way.

- I say to you, “I love you. I would rather die than hate you.” And I’m foolish enough to believe that through the power of this love, somewhere, men of the most recalcitrant bent will be transformed.
- Darkness cannot drive out darkness; only light can do that. Hate cannot drive out hate, only love can do that.
- Those who love peace must learn to organize as effectively as those who love war.
- True peace is not merely the absence of tension. It is the presence of justice.
- Injustice anywhere is a threat to justice everywhere.
- In a real sense all life is inter-related. All men are caught in an inescapable network of mutuality, tied in a single garment of destiny. Whatever affects one directly, affects all indirectly.
- Every man must decide whether to walk in the light of creative altruism or in the darkness of destructive selfishness.
- The time is always right to do the right thing.
- We must learn that passively to accept an unjust system is to cooperate with that system, and thereby to become a participant in its evil.
- You are not only responsible for what you say, but also for what you do not say.
- Our lives begin to end the day we become silent about things that matter.
- Our scientific power has outrun our spiritual power. We have guided missiles and misguided men.
- A nation that continues year after year to spend more money on military defence than on programs of social uplift is approaching spiritual doom.
- Nonviolence is directed against forces of evil rather than against persons who happen to be doing evil. It is evil that the nonviolent resister seeks to defeat, not the persons victimized by evil.
- We should never forget that everything Adolf Hitler did in Germany was “legal” and everything the Hungarian freedom fighters did in Hungary was “illegal”.
- Nonviolence means avoiding not only external physical violence but also internal violence of spirit. You not only refuse to shoot a man, but you refuse to hate him.

Nelson Mandela (1918-2013):

- Freedom can never be taken for granted. Each generation must safeguard it and extend it. Your parents and elders sacrificed much so that you should have freedom without suffering what they did. Use this precious right to ensure that the darkness of the past never returns.

- Like slavery and apartheid, poverty is not natural. It is man-made and it can be overcome and eradicated by the actions of human beings.
- Overcoming poverty is not a gesture of charity. It is an act of justice.
- As long as poverty, injustice and gross inequality persist in our world, none of us can truly rest.
- Education is the most powerful weapon which you can use to change the world.
- It is in your hands to create a better world for all who live in it.

Albert Einstein (1879-1955):

- No problem can be solved from the same level of consciousness that created it.

Margaret Mead (1901-1978):

- Never doubt that a small group of committed people can change the world. Indeed it is the only thing that ever has.

Bertolt Brecht (1898-1956):

- Change the world, she needs it.

Leonardo da Vinci (1452-1519):

- I have been impressed with the urgency of doing. Knowing is not enough; we must apply. Being willing is not enough; we must do.

Percy Bysshe Shelly (1792-1822):

Rise, lions after the slumber
In unvanquishable number!
Shake your chains to earth like dew
Which in sleep had fallen on you:
Ye are many — they are few!

Despite the authors of these quotes being separated by space, time, and their affiliations, it is easy to imagine that they all share the same human ethical belief system, and that they would have no significant arguments with each other if they all came together in one room to plan an ethical revolution to make the world a better place.

Unethical Arguments

The capstone to this scientific manifesto for a global ethical revolution to create a stable cyberanthropic super-ethical society is to define **The Law of Unethical Arguments**, which states:

“Because no ethical argument can exist against making a system ethical, anyone who argues against this goal, obstructs progress towards this goal, or abuses its sincere supporters, is either objectively unethical, corrupt, or evil.”

Conclusion

This grand challenge to implement a systemic ethical revolution is neither a new religion nor a political movement, it is a response to Johann Eder's call for a grand challenge in Vienna (Eder, 2010) and Irma Wilson and Pamela Buckle Henning's call to action for the systems sciences community, in Berlin (Wilson and Buckle Henning, 2016).

This challenge is the product of a compassionate heart and mind, employing the Ethical Regulator Theorem to generate maximally coherent ethical interventions (for the greater good of humanity) in multiple complex systems, such as the computational, corporate, criminal, cybernetic, personal, political, product development, psychological, scientific, social, and spiritual realms. And all such interventions resonate, not only with each other, but also with all good people who have ever existed — or ever will.

This revolution is long overdue, and we are privileged to live in these exciting times, but passively watching from the sidelines, or doing nothing, only helps the criminals, psychopaths, demagogues, and ethically indifferent corporations to exploit the pathological chaos and emergent problems that, until now, we have accepted as normal. It's time to make a commitment to design, research, educate, campaign, and fight for a better world.

“To be bold enough to consciously and deliberately reach beyond ourselves, to accept a grand challenge for the greater good, would be an act of self-actualization.” Stella Octangula (2011)

Just like we have legislation and non-negotiable expectations that passenger aircraft are designed to include expensive redundant subsystems to avoid having single points of failure in flight-safety-critical systems, and that all electrical products that we purchase conform to strict safety standards, we must change our attitudes, to create a cultural shift that makes it become utterly unthinkable and totally unacceptable to knowingly design systems or sell products that are ethically inadequate. Outrage at such behaviour is appropriate: We must demand strict legislation to force ethically indifferent corporations to play their part in creating a better world for the entire human race.

Arguably, the root of all evil is a lack of ethics, and by systematically applying the Ethical Regulator Theorem, we can reliably increase ethical behaviour in many classes of systems; progressively reducing unethical behaviour, reducing unethical suffering, and setting a course for humanity towards the tipping-point where we will experience a peaceful social phase-transition to a stable cyberanthropic super-ethical society.

Though this paper covers many topics, these are but means; the end has been throughout to make clear what principles must be followed when one attempts to restore ethical function to a sick organism that is, as a human society, of fearful complexity. It is my faith that the new understanding may lead to super-ethical systems that can create a better world, for the need is great.

References

- Ashby, W. Ross (1951). Advanced society planned as a super brain, ashby.de/rossashby.info *Journal*, volume 14, page 3528.
- Ashby, W. Ross (1952a). Power and I.Q. have many similar properties, ashby.de/rossashby.info *Journal*, volume 16, pages 4276-4278.
- Ashby, W. Ross (1952b). Appearance of a super-clever machine, ashby.de/rossashby.info *Journal*, volume 16, pages 4279-4280.
- Ashby, W. Ross (1956). *An Introduction to Cybernetics*, Chapman and Hall, London. ashby.de/rossashby.info/Ashby-Introduction-to-Cybernetics.pdf
- Asimov, Isaac (1942). Handbook of Robotics, 56th Edition, 2058 A.D. in *Runaround* in *Astounding Science Fiction*, March 1942.
- Bolstrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford Uni. Press.
- Clarke, Arthur C. (1973). Hazards of Prophecy: The Failure of Imagination, in *Profiles of the Future: An Inquiry into the Limits of the Future*.
- Conant, Roger C., and Ashby, W. Ross (1970). Every good regulator of a system must be a model of that system, *International Journal of Systems Science* 1(2):89-97. pespmc1.vub.ac.be/books/Conant_Ashby.pdf
- Eder, Johann (2010). Grand Challenges for Computer Science Research: Ross Ashby Memorial Lecture of the International Federation for Systems Research, in *Cybernetics and Systems 2010* (ed. Robert Trappl), Vienna, Austria, pp. xix-xxv.
- Glanville, Ranulph (1986). The Cybernetics of Ethics and the Ethics of Cybernetics, *Tutorial at the 21st American Society for Cybernetics Conference*, Virginia Beach, USA.
- Kuhn, Thomas (1962). *The Structure of Scientific Revolutions*, Univ. of Chicago Press.
- Octangula, Stella (2011). *Structure, Environment, Purpose, and a Grand Challenge for the ASC*. asc-cybernetics.org/CofC/wp-content/uploads/2011/02/Structure-Environment-Purpose-and-a-Grand-Challenge-for-the-ASC-V2.0.pdf
- Stevens, John Paul (2010). Opinion of Stevens J., Supreme Court of the United States. Citizens United, Appellant v. Federal Election Commission, *Legal Information Institute*, Cornell University Law School, www.law.cornell.edu/supct/html/08-205.ZX.html
- Turing, Alan M. (1950). Computing Machinery and Intelligence, *Mind* 59:433-460.
- Umpleby, Stuart A. (2001). What comes after second order cybernetics, *Cybernetics and Human Knowing* 8(3):87-89. nomads.usp.br/pesquisas/design/objetos_interativos/arquivos/restrito/umpleby_what_comes_after_second_order_cybernetics.pdf
- von Foerster, Heinz (2003). Ethics and Second-Order Cybernetics. In: *Understanding Understanding*. Springer, New York, NY. stanford.edu/group/SHR/4-2/text/foerster.html
- Wilson, Irma, and Buckle Henning, Pamela (2016). A Call to Action for the Systems Sciences Community, in *Proceedings of the 59th Annual Meeting of the International Society for the Systems Sciences, 2015, Berlin, Germany*. Vol. 1(1). journals.iss.org/index.php/proceedings59th/article/viewFile/2609/836