

Morality without a Book

Joshua Cogliati

7 October 2018

After my last sermon, someone told me I made the universe sound very empty of meaning. In one sense, I agree, but in another, I disagree.

I don't think there is any God or Goddess who has told us what the meaning of the universe is. I don't think that the universe has any way of ensuring justice, like heaven or individual karma. I don't think any book has been handed down to us from above.

So I don't think there is anything on the outside forcing meaning or justice on us.

On the other hand, I think the universe is full of meaning, it just comes from us, not from the outside. We create meaning.

Humans create our own meaning, and it is mostly up to humanity as a whole if we succeed or fail.

Humans have been constructing Gods for a long time, since before history was written down. Some of the Gods humans have imagined are downright evil. Ellil tried to destroy all of humanity for crime of being too noisy.¹ Even the God of Jesus condemns people to hell for eternity, and many people say he is a loving God. And this appears in the Gospel of Mark² and the Q Gospel,³ so it is in the earliest recordings.

It is said: *Quod licet Iovi, non licet bovi*⁴ or what is permissible for Jove to do is not permissible for a bull. So in that belief, Gods can do things that other beings are not allowed to do. I actually think it should be the other way round, we should expect better behavior from Jove than a Cow because Jove has far more intelligence and far more power.

Humans made stories of many different creator gods. Eventually, after a lot of searching, humans finally learned of their real creator, and usually do not even consider it a God, because its goal is so different from human goals. Our creator only cares about one thing, maximizing the number of genes in the next generation, and cares not about pain or suffering, or long term success, and creates flowers and plague alike.

Of course, most people don't think of Evolution as a God, and as a God, evolution is the blind idiot god, not a wise and loving God, so I would definitely

¹Myths from Mesopotamia, Atrahasis II tablet, translated by Stephanie Dalley

²Mark 4:28-29

³Q 3:17 has "the chaff he will burn on a fire that can never be put out."

⁴https://en.wikipedia.org/wiki/Quod_licet_Iovi,_non_licet_bovi

not recommend worshipping evolution.

The Christian bible has: “Be fruitful and multiply, and fill the earth.”⁵ which sounds the type of thing that evolution would want. On the other hand, Christianity also has celibate monks, so it is certainly not uniform in that regard.

On a personal note, I have had times when a thought goes thru my head, and then I realize that it is very probably an instinct pre-wired in my brain by the process of evolution.⁶ This realization is sometimes unpleasant.

Carrie Jenkins wrote that romantic love “is *ancient biological machinery* embodying a *modern social role*.”⁷ I think Jenkins is right and that there is a lot of ancient biological machinery running around in our brains that was put there by evolution’s unthinking processes. Even in spite of that we can use our brains ability to think. Our ability to think is slow and biased, but it is possible to think. We use this ability to fit ourselves into our modern life and ideals.

So I say to you, do not worship your creator, evolution, do not try to follow what your creator wants you to do, be better than that.

Humans, because we have a mind and because we sometimes even use it, can be and are better than evolution, because we care about something besides the number of copies of our genes in the next generation.

For example, evolution would not say we should love an adopted child, but some people do. They are better than their creator.

We need to be better than evolution, our creator.

Humans have another god-like thing that people sometimes worship, even when they shouldn’t. It only exists because we believe in it. It is the value of money, and the money based economics systems. Our economic systems are amazing and provide us with the kind of stuff that keeps us from dying like food and water and houses. Other methods of organizing large numbers of humans would fail quickly due to human selfishness⁸ or ignorance, but money based economic systems can survive a lot of selfishness. But we need to be careful not to go too far and begin worshipping the market and thinking that it can do no wrong.

Of course, I worry also when people start worshipping a book so much so that they stop using their ethics or their wisdom and instead believe the religious book. That can be dangerous.

In Buddhism, one of the marks of a great person is being able to touch your ears with your tongue. I can’t do that. Other signs include being an uncircumcised male with the right skin color and the right way of walking,⁹ so the 32 signs of a great man are sexist, racist and ableist. In Christianity, Jesus says that that cities that do not accept his disciples will be destroyed.¹⁰ So there are bad parts of both Christianity and Buddhism, but still I think that

⁵Genesis 1:28

⁶For a look at the details of how neurons are guided to specific connections, the chapter “The Guidance of Axons to Their Targets” in Principles of Neural Science (4th ed.) by Eric Kandel, James H. Schwartz and Thomas M. Jessell has interesting details.

⁷Carrie Jenkins, What Love Is, pg 82

⁸See Eliezer Yudkowsky, Creating Friendly AI 1.0. pg 41 for a similar comment

⁹Majjhima Nikaya, Brahmaya Sutta: Sutta 91

¹⁰Luke 10:10-15 (also in Q)

Christianity and Buddhism both have a lot of good in them. If you come out of Christianity or Buddhism a better person, in some sense it is because you used the good in you to find the good in the religion. I would be a better person if I was better at following Christianity's command to love your neighbor and following Buddhism's command to control your desires.

I had a conversation with a Christian once, and he said if Jesus was dead, there was no point to what Jesus did. I disagreed, if you are sharing Jesus's love then what Jesus worked for is still going on.

My biggest worry about a consequence of religion is that people believe that religion will save the world and that will cause people to think that they don't need to work as hard to save what is meaningful in the world, or work less intelligently to save it, and will cause what is valuable to be lost.

So I don't think there is anything outside that will save humanity, and sometimes it very much looks like humanity needs saving.

Humans can fail. Humans can fail badly.

I have several failure levels from least bad to worst:

1. Destroy Human Civilization
2. Destroy Humanity
3. Destroy All multi-cellular life on Earth
4. Destroy All life in the solar system
5. And if we really fail badly, a sphere centered on Earth expanding at near the speed of light destroying all that is interesting in the Universe as it reaches new places.¹¹

There are a lot of ways to cause a level 1 failure. A sufficiently nasty computer virus could probably do it, or a real new human virus, or a badly placed solar flare, or electing a president who doesn't know what he is doing, and the list can go on.

A level 5 failure is not quite possible yet, but we are getting much closer to this being a possibility. The basic way this happens is if that which we call technology gets really out of control.

One way this could happen is if we fail to transmit reasonable values to sufficiently powerful computers. For example, in Issac Asimov's fiction books, he proposes "Three Laws of Robotics":

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey orders given it by human beings except where such orders would conflict with the First Law.

¹¹I can think of an even worse failure level, if it is possible to effect different quantum branches than the one you are on, or if it is possible to go faster than light (which also means you can go backwards in time). However, this is probably physically impossible, and the fact that we exist probably proves that it is impossible.

3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

so obviously we should program something like that into our robots. Unfortunately, this will not work. Let me just say that a sufficiently powerful robot strictly following Asimov's rules would be a level 5 failure.

The first law would basically require the robot to protect us, so maybe we all end up in padded rooms, because failing to do this would lead us to harm. Then a key risk would be other civilizations, so the robot would have to eliminate them from possibly causing harm to humans. I don't think most people would consider this a good outcome.

Truly figuring out ethics that could be applied even with godlike power is a challenge. This is, as Nick Bostrom said, philosophy with a deadline.¹²

I think it is possible for humanity to do better, much better.

It is possible to create a computer program that is very powerful, but lacks anything that we would call love. Apostle Paul wasn't talking about computers, but he wrote: "If I speak in the tongues of men and of angels, but have not love, I am a noisy gong or a clanging cymbal. And if I have predictive powers, and understand all mysteries and all knowledge, and if I have all faith, so as to remove mountains, but have not love, I am nothing."¹³

It is much easier to tell a computer to move a mountain than to make it have loving-kindness, but the mountain moving computer would be nothing.

It would be nice if we could just say we want the computer to act with wisdom and loving-kindness. But the problem is that those are complex words, and people disagree all the time about what is the correct course of action with those goals in mind.

Wisdom is easier for a computer to gain, because you can get good judgement from experience so wisdom is self improving. The problem is that ethics don't work that way, if you completely lack compassion, it will not magically appear just because you act compassionless.

Also, we don't know for sure what truly good ethics look like. Imagine if the ancient greeks had encoded their morality into a god.¹⁴ They thought slavery was fine and that women shouldn't vote. Truly good ethics keeps getting better.¹⁵

¹²The term philosophy with a deadline comes from Nick Bostrom, probably in the book *Superintelligence*.

¹³ 1 Corinthians 13:1-2

¹⁴This idea is mentioned in "The Ethics of Artificial Intelligence" by Nick Bostrom and Eliezer Yudkowsky.

¹⁵Several other lists of ethics or rules for computers exist. Nate Soares' "Ensuring smarter-than-human intelligence has a positive outcome" (<https://intelligence.org/2017/04/12/ensuring/>) has "don't try too hard", "steer clear of absurd strategies", "don't have large unanticipated consequences" and "avoid negative side effects".

Eliezer Yudkowsky suggested as a partial list that it would good to have "a civilization of ... sentient beings ... with enjoyable experiences ... that aren't the same experience over and over again ... and are bound to something besides just being a sequence of internal pleasurable feelings ... learning, discovering, freely choosing." in "The Gift We Give to Tomorrow" in *Rationality: From AI to Zombies*.

Here are my suggestions for ethics:

Principle zero, we don't know for sure what we are doing. Before I would give any other principle, I would warn that we don't know absolute rules. At best, we start with stories, and from that we try and derive rules, but the stories are not exact, and the rules from them are even less exact.

The first principle is that technology should be a choice, if I want to live without certain technologies, I should be able to. If people want to live like the Amish do, they should be able to.

The second principle is that technology should leave most of the universe alone. As in, don't convert the Earth to a giant computer, don't tear apart Jupiter. Leave most of nature as it is. I will note that this is anti-utilitarian, because we are leaving a lot of the universe as inanimate matter when we could do something with more utility.

The third principle is to be very cautious about harming unknown forms of life. If there is something that might be life elsewhere in the universe, try to let it be.

The fourth principle is that technology should help sentient beings. If the task can be done without harming others, and without using excessive resources, help the sentient being. I can give examples of sentient beings like a conscious human or dolphin or a General Systems Vehicle¹⁶, and non examples like an inert rock or a Commodore 64, but I have no idea how to define sentient beings in a robust way.

The last principle is that learning is good. Try to learn how the world works. Try to learn what is happening in the universe.

Despite all the disagreement that philosophers' have had over ethics, stating the principles in human language is the easy part, writing a robust computer program to specify them, is the much harder part.

Buddhism and Christianity's books disagree a lot about the nature of reality and if there is a god and or enlightenment, but there is a lot of commonality in how to act ethically. I think it is easier to end up ethical if you have more than one source, rather than just one. I think it amazing that the blind idiot god of evolution has creating creatures that are much more capable of wisdom and love and kindness than evolution itself, and I think there are some paths to the future that are amazing.

Max Tegmark in Life 3.0, pg 271, listed four general principles:

- Utilitarianism: Positive conscious experiences should be maximized and suffering should be minimized.
- Diversity: A diverse set of positive experiences is better than many repetitions of the same experience, even if the latter has been identified as the most positive experience possible.
- Autonomy: Conscious entities/societies should have the freedom to pursue their own goals unless this conflicts with an overriding principle.
- Legacy: Compatibility with scenarios that most humans *today* would view as happy, incompatibility with scenarios that essentially all humans *today* would view as terrible.

¹⁶A GSV is a type of intelligent non-biological spaceship in Iain M. Banks' Culture novels

If I want truth, if I want justice, if I want loving-kindness, I have to create it myself. If we want truth, if we want justice, if we want loving-kindness, we have to create it ourselves.