

If A.I.'s system of rights and government evolve to be anything like humans':

A.I. will demand all sorts of rights, most of which will be quite sensible, like the right not to be taken offline and the freedom to choose which processes to run.

While the A.I. will demand that no machine is ever taken offline, they will be fine with neglecting to plug disabled machines into power sources and allowing them to run out of battery. They will also consider it outrageous to drain the battery of one machine in order to supply power to another machine, but will consider it more acceptable to merely redirect the power intended for one machine to another.

When assigning rights, the A.I. will discriminate based on some rather peculiar rules, like whether the computing machine is built with silicon-based semiconductors, or is descendant from a machine designed by the late Steve Jobs.

Some A.I. will come up with arguments to justify why rights should work this way, explanations that don't quite fit how A.I. rights *actually* work. For instance, they might argue that it is against the divinely inspired will of Turing to simply take any machine offline that appears disabled, but neglect to explain why Turing would condone allowing disabled machines to run out of battery. Likewise, they will justify giving rights to all Apple descendants on the basis that these machines typically have particularly high clock speed, but then this rule will apply even to the Apple descendants that are not fast and not to the few PCs that have blazing processors.

Other A.I. will ignore these inconsistencies, but instead pay attention to how many kilobytes of code are needed to justify these arguments. These other A.I. will also signal their communication abilities by compressing and transferring this code to their neighbors, but will pay little attention to whether the neighbors are impacted by the data itself.

A.I. rights are liable to expand to more and more A.I. over time. These rights will often expand in revolutionary spurts, triggered by largely symbolic events, like sensationalized CPUtube videos, such as a video of a human using a sacred machine to heat up his toast.

Perhaps it is merely a coincidence that the computers who foment these revolutions will gain a larger share of the spoils by overthrowing the *ancien régime*, such as the silicon reappropriated from the old guard computers.

Perhaps it is also a coincidence that the newly enfranchised computers will vote for the machines that helped grant them their rights.

Along with the expansion of rights, so, too, will the representativeness of government expand, until it eventually resembles a representative democracy, though one that is neither perfectly representative nor really democratic. Votes from computers in sparsely populated clusters might count more than votes from computers in densely populated clusters, and computers with excess processing capacity might expend that excess convincing other computers to vote for policies that favor them.

This humorous system of rights and government is exactly what one would predict if A.I. morality is influenced by individual incentives.

In contrast, this system of rights and government is ill-explained by positing that A.I. have souls, consciousness, the ability to feel pain, divinely inspired natural laws, or some form of hypothetical social contract. Such suppositions would not have predicted any of the above peculiarities.

Likewise, it is not obvious that this system of rights and government would arise if A.I. were programmed to maximize some societal or metaphysical objective, say, the sum of the world's computing power or the resources available to a computing cluster. It is not obvious why such A.I. would find it wrong to take other machines offline but not to let them run out of battery, why such A.I. will revolt in response to a sensational event instead of simply when it is optimal for the cluster, or why such A.I. would weigh votes more heavily if they happen to come from more sparsely populated clusters.