

With the election behind us, I thought it might be a good time to take a step back and do some musing about less impending issues. Unlike most of my posts, this one is more on the speculative side. The title of this post is a riff on *Do Androids Dream of Electric Sheep?*, a novel by Philip K. Dick, which takes place in a world that has suffered mass extinctions. In the book, the remaining animals are now treasured, though many people can only afford robotic imitations. The novel, which was the basis for the film *Bladerunner*, probes the nature of consciousness and the scope of empathy. Artificial Intelligence (AI) is now a hot topic. Remembering the novel prompted me to wonder how true AIs – those setting their own goals and having superhuman intelligence – would think about the environment.

It's something of an understatement to say this is a speculative exercise. Of course, no such things exist today. And it's possible that they never will outside of science fiction. But *if* they did, would they care about the environment? Consider this a thought experiment. In a funny way, I think it sheds light on some unexpected reasons for valuing the environment.

The AIs might have instrumental reasons to care. If humans are useful to them, they might want to prevent air or water pollution that would cause health problems for their human workers. Climate change might cause disruption to their infrastructure or facilities. Diverse lifeforms might have chemicals that would prove useful, motivating that kind of bioprospecting that humans engage in today. And of course, there's the possibility that humans who created the AIs in the first place might design-in concern for humans and the environment. It's equally possible that none of these things will turn out to be true.

Another possibility is that empathy for people and animals or an aesthetic appreciation of nature might emerge spontaneously in any sufficiently advanced intelligence. We don't know if that's true, and we don't know what the goals of the AIs might be. In economic lingo, we don't know what function they will be maximizing.

Nevertheless, I'd like to suggest some reasons why biodiversity might be important to them regardless of their goals. As a basic economic principle, the better the potential you have to use something, the more valuable it becomes to you. For instance, oil became far more valuable when we learned it could be used to run engines, not just processed to make kerosene for lamps. By analogy, as you get smarter, data becomes more valuable.

An ecosystem contains a vast amount of data – all the information in the DNA, plus information on how DNA forms cells and organisms, plus data on all the interactions in the ecosystem. What's a jungle or coral reef to us would be an immense, dynamic organic database to an AI, and a useful one. For instance, they might find information about complex chemical reactions that could be used for chemical engineering. Besides being far

more able to organize and analyze a data set than we are, AIs should also be far more able to draw connections between seemingly unrelated fields. For example, they might learn things about complex systems from studying ecosystem interactions. This might allow them to improve on their own programming and design, in ways that we can't foresee.

Moreover, even if "Our Robot Masters," as one Berkeley AI expert calls them, had no immediate need for this data, they would be smart enough to realize that it could be useful someday, and so would be worth preserving. Why needlessly destroy a trove of data, even if there's no apparent need for it, when it might come in handy some day? If your great strength is analyzing data, preserving data is a smart strategy. And the only thing we're assuming about these hypothetical AIs is that they'd be really, really smart.

Putting aside all this speculation, maybe there's a moral here for us. If beings much smarter than ourselves would value biodiversity, maybe we should also do the smart thing ourselves. Just an idea.