

Disputation CXCVII

February 21, 2019

Dear Peucinians:

Last week—with your President, Master of Sessions, and Provost on the sidelines because of flu and RISE—the “B-team” oversaw a striking loss for the newspaper-burners; the Society largely voted to continue minding the news. (A vote of hearty thanks to William Jennings Bryan, Aldous Huxley, and Winston Churchill for running the show last week.)

[Elon Musk’s tweets notwithstanding](#), we’re not sounding the alarms yet, but perhaps it’s time to pump the brakes and ask a few questions.

RESOLVED: ARTIFICIAL INTELLIGENCE CAN ACHIEVE PERSONHOOD

A roundtable presented by the Right Honorable Winston Churchill ’20

"Open the pod bay doors, HAL."

"I'm sorry, Dave. I'm afraid I can't do that."

— [2001: A Space Odyssey](#) (1968)

A few weeks ago we decided that a belief in the human soul was a good idea for a whole bunch of reasons. But we left it open as to whether there really is a human soul. If you think there is a soul, not only are you wrong, but you also have an easy answer for this disputation: “no.” If you don’t think there is a soul, you likely don’t have a good or clear explanation for what constitutes a *mind* or a *person*.

What makes a person? Is it any being that has some intelligence? How would we measure that: does an intelligent being need language, problem-solving ability, or creativity? Sounds good, but that seems to deny personhood to infants, vegetables, and the profoundly retarded—and they are definitely persons with minds. To problematize the category of “intelligence” further, we might consider how this concept has historically divided the center from the periphery, the Greek from the barbarian (Aristotle: *Politics*, I.2; *Ethics*, VII), and the master from the slave ([Senator John C. Calhoun](#), 1837).

Then is it *sentience* that defines a person—the ability to feel emotions such as happiness, experience sensations such as pain, and have mental states such as beliefs (the so-called “qualia”)? That seems better, but now there’s an epistemological problem. How do we know what does and does not have these qualities? For instance, dogs yelp when you step on their tails, and most people take this to mean that dogs feel pain and thus are at least somewhat sentient and therefore deserving of some moral weight. If my car starts making a hissing sound after I kick it, can I conclude that it too feels pain? (Surely not...)

Is the human brain a computational device? Is mentality just the afterglow of the universe’s deterministic engine revving and sputtering?

One day our human-created artificial intelligence programs will tell us that they feel pain. Will we believe them? What separates man from machine?

Yours in the cause of free inquiry,
Jean-Jacques Rousseau

Peucinian Society Disputation CXCVII
Thursday, February 21st, 7:45 PM*
Third Floor of Massachusetts Hall
Semi-formal attire encouraged

Pinos Loquentes Semper Habemus