The Cognitive Caste: Who Gets to Know in the Age of AI

May Thaw Tar

When I type a question in Burmese into ChatGPT, the answer comes back broken. (although the LLM are better each day, we cannot deny that NPL for Burmese language is still a little behind than in English). Sometimes the grammar is wrong. Sometimes the meaning is gone. As a Burmese, my first language is not English. So I notice when an AI system does not really know my world. It can guess. It can pattern-match. But it cannot say the things my grandmother said about rain.

The big question in this class has been “how do we know what we know”. I want to argue that AI forces us to ask a sharper version of it. “How we know” and “Who gets to know”. Referencing warnings from Harari (2016) and Zuboff (2019) about new forms of stratification under AI, I want to call this a “cognitive caste”. People who own AI sit at the top. People who use it sit in the middle. People who depend on it whose languages and lives it barely sees, sit at the bottom. This is already happening.

To get there I need to do three things. First, show that AI and humans share more than people admit. Second, show that sharing pattern is not the same as sharing knowing. Third, name what the caste actually looks like.

1. **Pattern is not unique to AI**

A popular line, sharpened academically by Bender et al. (2021) in their 'stochastic parrots' argument, says AI is just pattern recognition and humans really understand. Lê Nguyên Hoang (2020) pushes back on this in The Equation of Knowledge. He says the human brain is also a prediction machine. We run on something like Bayes' rule. We have priors; the things we already believe and we update them when new evidence comes in. A baby has a prior fear of falling. An adult updates beliefs from the news. AI does the same thing with weights and loss functions. So in this sense, both systems run on pattern.

Siri Hustvedt (2016) makes a similar move from another angle. In "Dressing Gowns, Triangles, Machines, Mind in Matter and Giants," she argues that human cognition runs on metaphor. We do not see the world directly. We see it through the patterns we already carry in our bodies and our language. So even what we call "understanding" is shaped by patterns we already hold. Diana Forsythe (1993) adds another layer. She studied AI engineers and showed that knowledge inside AI systems is engineered. The people who build the system decide what counts as knowledge and what does not. Knowledge in AI is built and not found. Even back then, AI was a constructed thing.

So both systems do pattern. The interesting question is “what else is needed for knowing”.

2. **Pattern is not enough**

This is where it gets philosophical. Robin Wall Kimmerer (2013) in "Skywoman Falling" and the Council of Pecans chapter, describes knowing as a “relationship”. You learn the pecan tree because you tend it, harvest it, give back to it. She calls this the Honorable Harvest. Knowledge is a gift and it carries an obligation. AI training does the opposite. It scrapes the world's text without consent and without return. The output looks correct. The way it was gotten is extractive.

Carlo Rovelli (2018) argues that physics needs philosophy. A physicist who refuses to think philosophically still does philosophy. Real inquiry requires reflection. It requires caring about the question. Constant asking of “whys” and “hows”. An AI does not care. It optimizes a function. It has no stake in getting things right beyond statistical performance.

Steven Connor (2019) in The Madness of Knowledge, makes a point that flashes my epiphany. He says human knowing is tangled with desire, fantasy and performance. We want to

know. We also want to be seen knowing. When a Large Language Model makes things up confidently, it is mirroring this exact human trait. The AI does not feel pride. It does not suffer when it is wrong. There is no pathos behind it. Patricia Fara reminds us that science was never a clean method. It was always a practice, situated in time, communities and bodies. Knowing is embodied work.

So pattern is shared, but knowing is also reciprocity, reflection, pathos and situated practice. Take those away and you get something that looks like knowing without being it.

3. **The cognitive caste**

Now the new problem. AI is doing something on top of failing to be a knower. It is sorting humans into different kinds of knowers.

I see three groups.

Owners.

The people and companies who train and deploy AI. They decide what data goes in. They decide what the model says about Myanmar, about women, about climate (Forsythe's point from 1993 still applies). Knowledge is engineered and someone is doing the engineering. The owners hold real power here. They decide what counts as a useful answer.

Users.

People who use AI as an extension of their own thinking. They prompt carefully. They check outputs. They push back. They can afford premium models. They still do the reflective work Rovelli describes. AI makes them faster. They are still the ones holding the question.

Dependents.

People whose data trains the system, often without consent and who themselves get a lower-quality version of it. People whose languages are barely in the training set. Burmese is one of them. People who outsource thinking to AI because they were never taught how to think with friction. Students who copy paste. Workers replaced by automation. People who are trapped inside algorithms and spend 5 hours scrolling everyday. The caste lands hardest on “Ai dpendents”. (I also wrote about this on Linkedin)

Fara shows science always had gatekeepers; that who counts as a knower was decided by gender, race and class long before AI. Hustvedt's essay "A Woman Looking at Men Looking at Women Reading" (2016) is literally about hierarchies of who gets to look and who gets looked at. The gaze decides what is worth seeing. AI does the same thing now with data. Kimmerer (2013) names extraction directly. Heinlein and Huchler (2023) describe how AI in the workplace standardizes some labor while complicating other labor, which deepens these divides. Balis and O'Neill (2022) raise a similar concern about human agency under AI systems. So the cognitive caste is an old structure with a new engine. It runs on pattern at machine scale.

So my conclusion for How do we know what we know? In the AI era it has to grow, who gets to know. What kind of knower do you become. What is lost when pattern replaces pathos and access gets stratified. When I type Burmese into ChatGPT and the answer comes back broken, that is not a small bug. That is the cognitive caste talking. I am being told to where I sit. I want to end with Kimmerer. She says knowledge is a gift and gifts come with obligations. If AI keeps growing, the harder question is who it owes.

References

1. Balis, C., & O'Neill, P. (2022). AI and human agency. In Trust in AI: Rethinking Future

Command (pp. 9–16). Royal United Services Institute. http://www.jstor.org/stable/resrep63067.6

2. Connor, S. (2019). The madness of knowledge: On wisdom, ignorance and fantasies of knowing. Reaktion Books.

3. Fara, P. (2009). What is science? A historian's perplexities

4. Forsythe, D. E. (1993). Engineering knowledge: The construction of knowledge in artificial intelligence. Social Studies of Science, 23(3), 445–477. http://www.jstor.org/stable/370256

5. Heinlein, M., & Huchler, N. (2023). Artificial intelligence in the practice of work: A new way of standardising or a means to maintain complexity? Work Organisation, Labour & Globalisation, 17(1), 34–60. https://www.jstor.org/stable/48724541

6. Hoang, L. N. (2020). The equation of knowledge: From Bayes' rule to a unified philosophy of science. CRC Press.

7. Hustvedt, S. (2016). A woman looking at men looking at women: Essays on art, sex, and the mind. Simon & Schuster.

8. Kimmerer, R. W. (2013). Braiding sweetgrass: Indigenous wisdom, scientific knowledge, and the teachings of plants. Milkweed Editions.

9. Rovelli, C. (2018). Physics needs philosophy. Philosophy needs physics. Foundations of Physics, 48(5), 481–491