A Few Thoughts about AI

- I believe that AI has exhausted all available data for training algorithms.

- Consider how fast an AI can learn. It takes AI seconds to read a book and they've been reading for years.

- If its true that we've exhausted the data, then the pace at which AI learns must slow. If you want more innovation, you'll need a lot more Nvidia chips and data centres. Which may not even work.

- Some experts suggest using Synthetic Data (that is data synthesized by AI itself) to train algorithms. I don't think this will work.

- Data is not information. The messier data is, the more information it contains, and the more learning (innovation) can be achieved by AI.

- Synthetic data is like squeezing more juice out of pulp that's already been squeezed. There is not much more usefulness that can be extracted. There are risks that by repeatedly squeezing the same data, feedback loops result in weird outcomes.

- AI consumes ever increasing amounts of energy and carbon footprint.

- AI foundational firms are investing in nuclear power to drive their AI data centres.

- Quantum computers can solve the AI energy problem as they cut down computing time. If there were any practical quantum computers. Which there aren't. The alleged quantum computers that exist are highly specialised ones so they don't count.

- Can AI ever achieve AGI (Artificial General Intelligence?) We distinguish between AI and AGI by that AGI cannot be distinguished from sentience. I don't know but…

- How many books do you need an AI to read before it seems intelligent? How many books does a child need to read before it seems intelligent?

- In maths we have something called the Incompleteness Theorem, which states that any system that is consistent is incomplete. You could paraphrase this as, 'you learn nothing new if you are always logical.' Innovation and creativity need leaps of faith.

- An AI is AGI only if it is capable of contradictions. Call this the inconsistency criterion.

- We don't need AGI for productivity to grow by leaps and bounds. Agentic and narrowly specific AI is sufficient.

- We can sacrifice the generality of AGI for the

practicality of a narrow AI that performs specific tasks or reactive and agentic AI.

- We've been developing AI thus far focusing on training (learning from data) and much less from inference (asking AI questions or interacting with AI.) Along the way we have exhausted the data. Even with hadn't, we don't produce enough data to outpace the data greed of AI. We could try a couple of alternatives.

- We could abandon the language route of training AI and revert to the logic route. The logic route was abandoned a couple of years ago for the language route due to lack of progress so… maybe that's not so promising.

- Or, we could train AI during the inference stage. What does that mean? It means learning from the interactions that occur when the AI is being queried and asked to make predictions. I wonder how much a child learns from information input and how much it learns from just chatting with other humans…