

Why Humans Better At Innovations Than AI?

Human Agency & Central Nervous System Matter

Jay Jiyuan Wu, Ph.D.

March 10, 2024

Executive Summary

It makes little sense, and is terribly boring, to compare the intelligence of a human sample with that of generative AI chatbots. The latter possess pseudo-intelligence as it is created, trained and empowered by humans with massive training data and billions of model parameters — one domain at a time. When it comes to existing knowledge in special domains or for special tasks, comparing human sample with AI chatbotx is like pitting a team of average college students against a team of Olympians. Humans are better than AI when it comes to innovations because we learn things differently: putting understanding before memorizing, and comprehension before specializing. Humans also learn through questioning, interactions, and practising in three dimensional world. The chatbots model parameters act like the number of pixels for a high resolution digital images. What is strong for humans is weak for AI, while what is strong for AI is weak for humans. Use of AI tools is strongly encouraged.

Right after publishing the white paper on launching a P&C insurance revolution, I find this interesting research report claiming that AI has outperformed human subjects in three standardized tests of creative potential.

This makes me wonder whether my statement in the white paper is accurate that says, “One thing we humans can still maintain an edge over GAI is the forward-looking innovations and creation.” Note “GAI” refers to the current era of Generative AI.

In conclusion, my words remain valid and accurate. But let us delve into the specifics of evidence and reasoning in this post.

1 Creative Potential or Unfair Comparison?

Here are the crucial details from the research report by the University of Arkansas on March 1, 2024, regarding the three standardized tests for creative potential on 151 human subjects:

“Alternative Use Task, which asks participants to come up with creative uses for everyday objects like a rope or a fork; the Consequences Task, which invites participants to imagine possible outcomes of hypothetical situations, like ‘what if humans no longer needed sleep?’; and the Divergent Associations Task, which asks participants to generate 10 nouns that are as semantically distant as possible.... the ability to generate a unique solution to a question that does not have one expected solution.”

1.1 A Discouraged Practice

The study authors have done something intuitive and seemingly sensible: finding out which perform better between a small sample of human subjects and the GAI chatbot.

Such studies are tempting, at a time when it seems everyone in the society is talking about large language models and GAI. And yet comparing responses from a single human sample to GAI chatbot answers is generally discouraged, not that it will not produce result, but the comparison will be boring, as the results are highly predictable and heavily favoring the AI — otherwise we would have witnessed a failure of AI training.

Let us consider the size of people answering the questions first. Going back to the question used in the U of Arkansas study, “What is the best way to avoid talking about politics with my parents?” This is an open ended question that demands open answers. The quality of open answers is related to the size of people thinking of the same question, directly or indirectly. Other things equal, the quality of the best answer obtained from 151 people at one time is unlikely to match the best answer chosen from a pool of 1.51 million people (or more), accumulated over decades or even centuries.

Size of the people is just one thing, selective quality is another. The comparison is similar to a 400-meter relayed competition between a random team of college students and a team of Olympians. In this scenario, both sides have the same number of players but it is the quality that sets them apart, way apart.

Another divergent associations task of generating 10 nouns that are as semantically distant as possible, is also in favor of AI because it has access to millions or more pieces of relevant information, far more than the live human sample participants this time in the study.

The problem is not in the comparison itself, but understanding what we are comparing of. We must ask ourselves one question first: Where are the answers of chatbot from?

1.2 Understanding Chatbot Answers

I have been going back and forth with Gemini on finding the best answer for this. When I first asked Gemini how it would summarize generative AI chatbot answers, Gemini offers a whole bunch of features like data-driven, natural language generation, contextual awareness, open ended, varied quality and not sentient (meaning the answers don't have their own thoughts or feelings).

This summary is interesting but not quite what I wanted. More specifically, I am not really interested in chatbot answers per se, but how they are linked with human knowledge and intelligence.

Switching to an inquiry of comparing chatbot answers and human knowledge and intelligence did not produce the result I have in mind, either. This time Gemini offers the following answer:

“AI chatbots excel at speed, scalability, and data-driven accuracy for routine tasks. Humans provide a broader understanding, creativity, and emotional intelligence for complex situations or interactions requiring empathy.”

At this point I used a counter-argument as the inquiry: “But AI Chatbot answers are summarizing what we humans have learned currently and in the past.”

This time I finally received the answer I was looking for:

“You're absolutely right! AI Chatbot answers are heavily influenced by the vast amount of information we humans have accumulated over time.”

It then goes on by saying that AI Chatbots are products of human knowledge; the training data are created by humans; the learning algorithms that power AI chatbots are designed by humans; humans are constantly evaluating, refining, and improving these systems.

In its usual thoroughness, Gemini says AI Chatbots can go beyond simple mimicry but do pattern recognition that even humans may miss; AI models can

analyze massive amounts of data statistically; Chatbots can access and process information much faster than humans.

This is true and we can confirm that using the previous question: AI chatbots may not be trained with the exact question of “What is the best way to avoid talking about politics with my parents?” But encountering the words like “parents,” “politics,” “best way” and “avoid” is enough for it to generate a sensible answer or answers.

1.3 The Problem of Under-using Chatbots

Gemini ends its answer by saying “AI as a Tool, Not a Replacement.” That I completely agree but will add a quick comment here: The bigger problem with AI chatbots is *underuse*, not *overuse* (with the concern of them replacing human intelligence). We need far more people to use a chatbot out of personal curiosity and exploration, rather than use it only when there’s an assignment or last-minute task waiting.

Think about it: Never before have we had such powerful tools at our fingertips! They literally let us learn from the giants in an instant. In addition, they make it easy and convenient for everyone to check out the facts, find suggestions for almost anything ranging from writing to strategy.

Consider all these, if you’re not using Gemini (or your favorite chatbot) at least five times a day, you’re missing out on a valuable tool. I have been asking Gemini and Perplexity questions everyday for at least 20 times. For example, I asked about the underwear brand “Fruit of the Loom,” and find that the word “Loom” does not mean a basket, as someone once told me.

1.4 AI Chatbots Are Trained One Domain at a Time

My previous draft of this post says, “Any such comparison must be done carefully or cautiously, as they are fundamentally juxtaposing the mind power of a sample of human participants, at one point of time, with the collective intelligence of humanity, accumulated over a much longer time.”

But that is inaccurate as we are not comparing a human sample with the general “collective intelligence of humanity,” only collective intelligence for a specific task in a specific domain of knowledge.

This fact is not obvious that AI chatbot is a collection of knowledge points from numerous specific domains. This is because when you open a GAI chatbot, you can ask anything and everything, making it easy to believe you were talking to a single “person” who knows everything.

Nonetheless, AI model is trained *one domain at a time*. Gemini says so: “In most cases, AI models are trained for a specific domain or task. This is because the patterns and relationships the model needs to learn are specific to that domain... an image recognition model trained on pictures of cats would perform poorly if asked to identify different types of flowers.”

“Even in AI models that can learn across multiple domains, the training might still involve focusing on specific sub-domains within the larger dataset.”

1.5 Inequality in General Human-AI Comparisons

The best way to show what are being compared is to use math symbols. In the following inequality, “H” & “AI” denote responses from a human sample and a AI chatbot, respectively, with the human sample size (x) presumably smaller than the size of AI training data (Z); drawing at one time (t) versus accumulated over a long time (T), and finally for knowledge or intelligence in a specific domain

(s) versus general intelligence (G).

The first inequality compares the two *as they are*, meaning humans retain their general intelligence but AI has only specialized expertise. The study at University of Arkansas falls to this category.

The two superscripts “G” and “s” denote general knowledge and special knowledge, respectively:

$$(H)_{x,t}^G \neq (AI)_{Z>x,T>t}^s \quad (1)$$

Formula (1) applies to the scenario that is on most people’s mind when they talk about “humans vs AI” comparisons, as the inequality holds for three parameters: Sample size (x vs Z , Z larger than x), time spans (t vs T , T larger than t) and knowledge scope (s vs G , G larger than s). Note the only advantage of human sample is in their general intelligence. What is humans’ strength is the weakness of AI and what is humans’ weakness is AI’s strength.

This is why Gemini tells us that, “The future likely lies in collaboration. AI chatbots can handle the speed and information retrieval aspects, while humans can focus on tasks that require deeper understanding, empathy, and creative problem-solving. Together, they can form a powerful team for various tasks.”

The inequality tells us not to be surprised when a human sample performs not on par with the AI model, or an AI model performs not on par with a human sample. In general, humans, even with just a small sample of humans, will do things better than AI when it comes to leveraging their general intelligence, including creating “new to the world” innovations.

1.6 Inequality In Special Domain Human-AI Comparisons

However, when it comes to intelligence in a specific domain or specific task, human samples in general will fail almost surely in all comparisons with AI, where the notion “*a.s.*” indicates an event will occur with a probability of 1, except a minuscule set of outcomes where it doesn’t:

$$P((H)_{x,t}^s < (AI)_{Z>x,T>t}^s) \text{ a.s.} = 1 \quad (2)$$

Formula (2) applies to special scenarios such as Google’s AlphaGo defeating the human GO champion in 2016, or an AI chess program defeating its human counterpart.

One perceivable reason we say “almost surely” rather than “surely” is that sometimes the AI training is not up to the latest, in which case the time leg will show up as AI performing inferior to humans in an infinitely long sequence of trials.

2 There Are Exceptions

Time delay is not the only scenario for AI model to perform inferior to humans. Sometimes one human can beat many or all others, even AI.

Think of the highly talented Indian mathematician Srinivasa Ramanujan (1887-1920) as told in this Wikipedia article, who “independently compiled nearly 3,900 results (mostly identities and equations)” during his short life of 32 years. It is entirely possible that Ramanujan alone would produce a better answer to a math problem than all the other people in the world, despite that he never even had a college degree. I recently watched the British biographical drama film of 2015, “The Man Who Knew Infinity,” and was totally awed.

Or think of Thomas Edison in the 2017 American historical drama film “The Current War,” inspired by the 19th century competition between Thomas Edison and George Westinghouse (and importantly, Nikola Tesla) over which electric power delivery system would be used in the United States.

Roughly 19 minutes into the movie, showing the scene when Edison, played by Benedict Cumberbatch, was interviewing Tesla, played by a baby faced Nicolas Hulot, and inviting Tesla to work for Edison Electricity, Edison claimed he and his team had made a small innovation every 10 days and a big one every 6 months.

If that was true, Edison would beat all other humans on the face of the earth back then and in an AI search Edison would stand out to be the best qualified person in human innovations.

But keep this in mind: even Edison, with all his brilliance, wouldn’t stand a chance beating an average chatbot today in a test of electricity knowledge — despite that a chatbot may not possess true understanding compared to human experts like Edison. This exemplifies the power of collective human progress, built on the advancements made by pioneers like Edison, Tesla, and Westinghouse. Innovation is a timed game, where time can change everything. To the extent AI models will be updated to catch up with the new innovations and new knowledge, human pioneers’ advantage will not be kept for too long, as their findings and contributions will be absorbed in the next AI model.

This blog cites the witty words from the co-founder of Coalesce, “If there were LLMs in the 1700s, and we asked ChatGPT back then whether the earth is round or flat and ChatGPT said it was flat, that would be because that’s what we fed it to believe as the truth. What we give and share with an LLM and how we train it will influence the output.”

In a way, we can think of chatbot as the Greek legendary hero Antaeus, the son of Poseidon (god of the sea and earthquakes) and Gaia (the earth goddess).

Antaeus was known for his immense strength that came directly from the earth. As long as he remained in contact with the ground, or his mother Gaia, he is invincible and his strength would be continually renewed.

Of course, his constant dependence on the earth was also his biggest weakness. Another Greek mythology hero Hercules finally figured out a way to defeat Antaeus by lifting him off the ground and killed Antaeus in the end.

Chatbots are like Antaeus and we humans are like Gaia.

3 Human Advantage in Central Nervous System

This section digs deeper into how human and AI learn differently.

3.1 The Costs & Benefits of AI Model

I find this interesting blog from stackoverflow that considers the costs and benefits of the existing large language models (LLMs). It begins by the following statements:

“Almost every week we hear news about the amazing performance and ever improving capabilities of large language models (LLMs) when it comes to creating human-like code and text. But alongside those, we see breathtaking dollar amounts ascribed to the cost of training those LLMs—reports and speculations regularly quote numbers in the tens and hundreds of millions. Future models may eventually crack the billion dollar mark. If you want a lot of advanced chips to train AI or plan to build your own hardware, rumors are now flying that trillions will be required.”

Burning truckloads of cash is nothing new for epoch-changing innovations, although the current GAI or LLMs’ efficiency leaves much to be desired. The blog tells us three things standing out to increase costs: “performance (defined as

performance against various benchmarks) scales smoothly on a power-law basis with **the size of the dataset, number of model parameters, and total compute** used to train the model.” (Emphasis added.)

To get a sense how many parameters the current LLMs need: “Microsoft researchers went bigger to create a general purpose language model. Their newer model, phi-2, has 2.7B parameters, well under what some of the state-of-the-art LLMs have but still double phi-1’s count.”

2.7 billion parameters in the model and still considered efficient?! I ask myself, and then Gemini, why so many parameters? This is way more than a human child would need to learn and to grow in a complex world.

3.2 Why AI Models Need So Many Parameters

Gemini explains why so many parameters are needed. I will quote it with my comments in parenthesis:

“1. Capture Complexities: The real world is intricate and nuanced, with diverse relationships and patterns. LLMs strive to understand and represent this complexity through learning from massive amounts of data. This data encompasses various formats like text, code, images, and audio, each with its own complexities.

(My Comments: A good analogy of a large number of model parameters is the number of pixels in digital images: The more pixels a picture has, the higher quality or resolution it has.

Some complex scenarios do call for many parameters. I was filling IRS Form 1120 the other day and found Gemini surprisingly adept at answering even the most specific questions. Its performance exceeded my expectations, perhaps even surpassing the capabilities of an average CPA. It must have come from intensive training inquiries with many parameters to fully capture the intricacy and

nuance.)

2. Model Relationships: Parameters act as adjustable variables within the AI model. By adjusting these parameters during training, the model learns to identify and represent the relationships between different pieces of information within the data. These relationships can be simple, like identifying the subject and verb in a sentence, or complex, like understanding the context and sentiment of a conversation.

(My Comments: A good example for a AI model to capture real world complicated relationships is Google’s AlphaGo, which defeated the international Go champion, Lee Sedol, in a five-game match held in March 2016. It was a historic moment, marking a significant milestone in the development of artificial intelligence (AI).

In order to beat humans, AlphaGo must have numerous parameters. I ask Gemini to explain why so, using AlphaGo as an example. The following is what Gemini says,

“Imagine AlphaGo as a highly trained Go player. The more games (data) it studies, the more it refines its understanding of the game’s subtleties and strategies. However, unlike humans who can develop intuition through experience, AlphaGo relies on the vast number of parameters within its neural network to represent this accumulated knowledge. These parameters act as a complex web of connections that encode the strategic nuances learned from the data.”

Gemini also tells us what make GO a hard game to simulate:

- **Massive State Space:** Unlike chess with a limited number of possible board configurations, Go has a much larger state space due to the larger board (19x19 grid) and the freedom of placing stones anywhere. This makes it difficult to pre-program every possible move and outcome.

- **Intuition and Long-Term Planning:** Go requires strategic thinking, considering not just immediate moves but also their impact on future board positions several moves ahead. This kind of long-term planning is challenging to achieve with traditional programming methods.)

3. Enhance Generalizability: With more parameters, an AI model can learn and adapt to a broader range of situations and contexts. This allows it to generalize its knowledge beyond the specific training data and potentially respond more effectively to unseen or novel situations.

(**My Comments:** Generalizability works best when learning or identifying patterns with nuances and subtleties from the training data and then applies to new or previously unseen data or situations.)

4. Improve Accuracy and Performance: While not a guarantee, a higher number of parameters often leads to improved accuracy and performance on specific tasks. This is because the model has a greater capacity to learn complex relationships and nuances within the data, leading to more accurate predictions or outputs.”

(**My Comments:** This one is largely redundant after the other three points.)

3.3 How Human Child Learns

The above all make sense, although it still does not tell me why the number of parameters is so much larger than a human child needs to learn. This is the next question I asked and the Gemini answer is once again good, so I will quote it below with my own comments:

“Data Acquisition:

Humans: Children learn from a variety of multimodal experiences in their environment. They interact with the physical world, observe others, ask questions, and actively explore. This leads to a rich and diverse learning experience.

AI models: They learn solely from the data they are trained on, which is typically limited in scope compared to the complexities of the real world. This can limit their ability to generalize or adapt to new situations.

(My Comments: Well stated differences in four dimensions between human and AI learning. Humans learn in three dimensions and we use our eyes, nose, lips and ears, unlike AI that is only fed with two dimensional stuff (words, pictures, videos).

Humans integrate learning and doing and practicing. When a human baby was encouraged to taste orange juice for the first time, they can actually have a sip from the bottle, while the chatbots can only depend on humans to describe the taste.

Similarly, humans learn in groups by mutual observations, which among other things, preserve diversity because groups differ in attitudes and preferences. I was watching the 2004 German movie *Der Untergang* (*The Downfall* in English), a documentary-drama type telling the last 12 days in Hitler's life. Seeing the desperate war situations and hearing the bombing outside the underground bunker beneath the Reich Chancellery in Berlin, it's clear that what Hitler said or did was directly impacting everyone else. Likewise, how Hitler sensed how others were treating him also mattered.

Group observations also help reveal relationship intentions. In a group chatting, a heterosexual male would look straight into the eyes of the girl he is interested but rarely anyone else. This silent but real interaction is hard to show even in a video, unless the director purposely gives an exposure. In general, AI training will have a hard time to capture it.

Finally, humans learn by asking questions and by active exploration, other proactive ways of learning not doable by AI. How important is asking questions? In a way, AI can always find answers for you, but how much you get from AI

chatbots partly depends on how good your questions are. The best learning comes from asking smart questions that have not been asked before.

The other differences listed below are essentially the same as the four.)

Learning Process:

Humans: Children learn through a combination of supervised and unsupervised learning. They receive guidance and feedback from adults (supervised learning) but also actively explore and experiment on their own (unsupervised learning). This allows them to develop critical thinking, problem-solving skills, and learn intrinsic motivations.

AI models: They primarily rely on supervised learning, where the training data includes labeled examples of the desired outcome. This makes them efficient at specific tasks within the confines of the training data, but they lack the ability to learn independently or explore beyond those boundaries.

(**My Comments:** This difference between supervised and unsupervised learning is similar to “active exploration” listed above.)

Understanding and Reasoning:

Humans: Children develop a comprehensive understanding of the world through their experiences. They learn about objects, concepts, and relationships not just by memorizing facts, but also by building connections and understanding the underlying principles. This allows them to reason, make inferences, and apply their knowledge in new contexts.

AI models: They excel at pattern recognition and statistical analysis within the data they are trained on. However, they often lack the ability to truly understand the meaning or context of the information they process. This can limit their ability to reason creatively, make sound judgments, or adapt to unexpected situations.

(**My Comments:** Holistic comprehension only happens to humans but not

AI, at least not at this point. Gemini correctly points out the limits associated with specialized learning. In a way, we purposely train AI to focus on one line of business at a time, and sacrificing the capability of making horizontal connections across lines.)

Generalization and Adaptability:

Humans: Children are remarkably adaptable and capable of generalizing their learning to new situations. They can take the knowledge and skills they acquire in one context and apply them to new experiences, even if they haven't encountered them before.

AI models: While some AI models can achieve a degree of generalization, their ability is often limited to situations that are similar to the data they were trained on. They often struggle to adapt to completely new contexts or situations not explicitly included in their training data.

(**My Comments:** This dimension has been listed earlier.)

Continuous Learning:

Humans: Children continuously learn and adapt throughout their lives. They are constantly building upon their existing knowledge and skills, even after they reach adulthood. This ongoing learning process allows them to stay relevant and adapt to a changing world.

AI models: Once an AI model is trained, its learning process typically stops. It may require additional training data or retraining to update its knowledge or adapt to new information.

(**My Comments:** This is a good point unlisted before but in a way it is a part of active learning or exploration that only occurs to humans.)

In sum, the above detailed "item-to-item" comparisons facilitate easy understanding even for a layperson. Bear the followings in mind:

- Supervised vs Unsupervised learning;

- Continuous vs Interrupted learning;
- Adapted vs Limited learning;
- Generalized vs Specialized learning;
- Comprehensive vs Pattern recognition learning, and finally,
- Learning from the whole world vs learning only from the training data.

3.4 The Central Nervous System

While the above differences are interesting and important, I want to dig deeper into the neurological basis between AI and human child, which can explain why the latter does not need so many parameters but the former does.

The way I ask the question is indirect, merely mentioning the fact that, “Human child does not require so many parameters to learn.” Gemini gets that and its answer is good:

“It’s true that a human child doesn’t require the same number of ‘parameters’ as an AI model to learn. However, it’s important to understand that comparing learning processes between humans and AI models directly isn’t entirely accurate. Here’s why:

1. Different Learning Mechanisms:

Humans: We learn through biological processes involving the brain and nervous system. These involve complex interactions of neurons, hormones, and various neural pathways. This biological learning process is dynamic and adaptable, constantly evolving as we experience the world.

AI models: They learn through statistical analysis and pattern recognition within the data they are trained on. This involves adjusting parameters within

the model to improve its ability to identify patterns and relationships in the data. This is a fundamentally different process from human biological learning.

(**My Comments:** While we know much about how AI models are trained, we still know relatively little how learning occurs with our CNS or central nervous system. For example, how the Indian mathematician Ramanujan was able to come up with so many new math discoveries without ever getting a college degree is still a myth. We can have a brain anatomy of Einstein but that won't tell us much about how his brain actually functioned. To understand his genius, we'd need to see his Central Nervous System (CNS) working in a living person, not someone who has passed away.)

2. Focus on Parameters vs. Underlying Processes:

Parameters: While the number of parameters in an AI model is often used as a measure of its learning capacity, it's not a perfect analogy to the complexities of human learning. In humans, factors like brain structure, neural connections, and cognitive processes play a much more significant role in learning than a single numerical value.

(**My Comments:** Once again, we have too many unknowns in human learning. One thing seems to be sure: Humans have much stronger synergy efficiency so that they do not require so many "parameters" in reasoning and in decision making.)

3. Efficiency vs. Versatility:

AI models: They can be trained on massive datasets and achieve remarkable efficiency in specific tasks, often surpassing human performance in specific areas. However, this efficiency comes at the cost of limited understanding and adaptability.

Humans: We may not be as efficient in processing information as AI models, but our learning process allows for greater versatility, creativity, and the ability

to generalize our knowledge to new situations.

Therefore, it's not accurate to say that humans 'learn with fewer parameters.' Rather, we have a fundamentally different learning system that relies on complex biological processes and allows for a broader range of learning capabilities despite not requiring vast amounts of data or specific parameters."

I don't know about you but I feel relieved after reading the above: Depending on what dimension we look at, human brains and our neurological system still have an upper hand when compared with the GAI system. As of today, AI only possess what may be called "quasi-understanding" and "quasi-intelligence."

For specific domains of knowledge AI does things faster, better and cheaper than us, but that is only because we trained it to do so, and only because AI sacrifices continuous, holistic, proactive and associative learning to get where it is now. After all, learning is a biological process with humans but a mechanical process with AI.

Where AI is weak, humans are strong — as the way it should be as we humans created, trained and empowered AI.

AI still has a long way to go to catch us. And that is comforting.

But Gemini has more to tell us. From the above:

"It's important to remember that AI research is constantly evolving, and advancements are being made in areas like neuromorphic computing which aims to mimic the structure and function of the human brain.

The ultimate goal of AI research isn't necessarily to replicate human learning perfectly, but rather to develop intelligent systems that can complement and collaborate with human capabilities."

The notion of neuromorphic computing gets my attention. I then ask the follow-up question: "Will neuromorphic computing in the future make AI learn more like humans?" Gemini responds by saying the following:

“Neuromorphic computing holds promise for developing AI systems that learn and process information in ways more similar to the human brain compared to traditional AI models. However, it’s crucial to understand that achieving true ‘human-like’ learning in AI is complex and involves overcoming several challenges.

Here’s how neuromorphic computing might influence AI learning:

- **Brain-inspired architecture** to mimic the structure and function of the human brain, using artificial neurons and synapses to process information.
- **Reduced reliance on massive datasets** to be more data-efficient due to their inherent parallelism and ability to learn from smaller datasets, potentially resembling how humans learn from diverse but limited experiences.
- **Improved adaptability and generalization**, analogous to how humans can apply what they learn in one context to new and unforeseen situations.

Of course, we have a long way to go as our understanding of the human brain and its learning mechanisms is still incomplete. We also have hardware limitations, as current neuromorphic chips face limitations in terms of scalability, power efficiency, and performance compared to traditional computing systems. Finally, developing efficient and effective algorithms that leverage the unique capabilities of neuromorphic hardware for learning tasks is an ongoing area of research.

Another challenge is that the non-neuromorphic AI models have already encountered many outcries, so it is only natural that more people will feel nervous about neuromorphic AI. The road will not be straight, even though it does hold more promise in reducing the amount of training data and model parameters.

4 Human Agency Matters

One thing that Gemini has not mentioned, or only mentioned implicitly, is the human agency effect, another advantage for human innovations. The authors of the Arkansas study correctly point out that, “AI, unlike humans, does not have agency,” which means AI lacks the capacity to act independently and make choices based on one’s own desires, intentions, and understanding of the world — unlike a normal human.

Human agency is a big deal and a good inclusive term for many things that separate AI from humans. A lot can be said about it.

4.1 Intrinsic Preferences Drive Innovations

One dimension of human agency is the intrinsic desires and preferences. Recall the movie “The current war” that shows in the late 19th century, Americans were already eager to find a less cruel and more humane way of executing criminals than hanging. Initially the electric chairs did not work as well as they are today, so Edison and his team actively campaigned to associate electrocution with AC (Alternating Current), referred to it as “Westinghousing.”

Compared with that, the ancient China would qualify as one of the most barbarian societies in the world, as the emperors had the liberty of beheading, skinning and lingchi anybody presumably guilty or disliked by the rulers: 砍头, 剥皮, 凌迟. The last way of execution, “Lingchi,” was especially brutal and does not have a corresponding English word that fully captures its meaning. Lingchi involved slowly cutting a human to death, one piece of flesh at a time, until only the skeleton remained.

Skinning was another inhumane method, where a person’s skin would be pierced and then sometimes filled with straw to create a gruesome, human-

shaped scarecrow as a deterrent.

I see a big cultural difference here: Americans could not stand for hanging (or any elongated suffering before death), and they asked, “Is there a more humane way to carry out capital punishment?”

Such a question was never asked by Chinese. Of course, you do not see beheading, skinning or lingchi today in China, but that has something to do with the higher killing efficiency of guns and rifles than swords and spears — more than the intrinsic humane considerations.

4.2 Mind Changes Drive Diversity

I watched a 2021 Japanese movie last night that was thought-provoking. It tells a real-life story in the American occupied Japan after WWII. A Japanese MD, who is also the associate professor in a Tokyo hospital, was involuntarily involved in an experiment to test whether, or for how long, humans can survive with just one lung — using the captured American B-29 pilots or POWs (prisoners of war) whose lungs were perfectly functioning.

This of course violated the ethic rules for medical professionals and so the MD was sentenced to death by the International Military Tribunal court — waiting to be hanged in the prison.

The tricky part was that the MD did protest, in vain, against the inhumane treatment of POWs, but was forced to participate as an assistant in the surgery conducted by the full professor. This is understandable considering the hierarchical structure of Japanese society influenced by Confucianism, where a full professor would hold coercive power over an associate professor.

Sadly, the full professor committed suicide before the trial and the MD in our story could not use him as a witness to defend himself. Fortunately, his wife was a strong and brave woman who decided to appeal on behalf of her husband. She

travelled to Tokyo several times to meet the Japanese interpreter at the trial, who then passed her appeal to the American prosecutor.

Things seemed hopeful, as the prosecutor was sympathetic to the case and agreed to reconsider — except that the MD must make a petition from inside the prison to trigger the appeal process.

Shockingly, the MD decided NOT to make the petition because he thought he deserved the death as he did help kill the American pilots. Instead of writing a petition, he wrote a last letter to his son and daughter, asking them to be nice with their mom and saying goodbye to them.

One reason for his change of mind is the personal observation of a military leader in the prison, who admitted more war crimes than he really committed, taking responsibilities for his subordinates. Another more direct reason is hearing the expert witness in the court by another MD, who claimed he would never assist in any human subject research surgeries hurting POWs, even if he would be killed by the Japanese military for refusing the order.

In the end, the MD changed his mind again after a tearful meeting with his son and daughter in prison, when his daughter begged him not going away. He ended up writing the petition and was re-sentenced to 10 year labor services.

I am glad that justice eventually served for the MD, but the moral of the story is that even a single human can have complicated and sometimes conflict feelings, thoughts and preferences, all of which can play a role in their decision making. This intrinsic diversity for a single person, plus the inherent diversity for a group or a society of humans, could be lost in a single answer from the chatbot.

On the other hand, I start to think that having numerous parameters in the underlying AI model is a good thing, because it allows the chatbot answers to be flexible and fluid, offering a significantly different version with seemingly trivial and subtle changes in the inquiries.

4.3 The “Hammer & Nail” Effects

The best lesson for me is to keep some topic, hypothesis or theoretic framework in mind, and it will work just like what people say, “When you have a hammer, everything looks like a nail.” This may have a negative connotation but another way to put it is that you need to have a prepared mind, which will inspire you to see connections and generate ideas you normally won’t.

One way agency works is through our determination and our subjective attention focus. Since late last year I have been thinking about P&C insurance revolution, and have not stopped writing and reading on it. One thing I find is that even when you are reading something seemingly totally irrelevant to insurance, ideas and thoughts will still come to you in a way somehow magically related to insurance.

This really happened to me: When I was reading the news of Super Bowl shooting in Kansas it suddenly came to me that this was a useful story for my idea of partitioned risks and shared policies.

Similarly, when I was visiting my relative’s house and saw a LEGO game there, it suddenly came to my mind that I can use LEGO game to illustrate the difference between legacy and new insurances. I came up with the hypothetical example in which a vendor is selling a preassembled, ready-to-play LEGO game, and compare that with the legacy insurance.

There are other examples that I don’t remember but my experience is that we can’t predict the next “aha” or epiphany moment (i.e., a sudden flash of insight, realization, or understanding), nor is it important. What is important is to have the topic or theme in mind, and somehow our brain will work out its magic.

4.4 Do Not Train Humans Like Machines

We have seen how humans learn things more efficient than AI training. The best thing to do is to emphasize (1) comprehension and understanding more than memorizing; (2) learning through practicing, questioning and exploring more than cramming; and (3) associations across domains more than having a subject imbalance.

But sometimes we may follow the wrong approach and force humans to learn like how AI models are trained. The biggest real life example is China, where parents voluntarily force their kids to learn like machines, with teachers controlling everything, emphasizing rote learning and practicing on questions with single right answers. Schools face so much cut-throat competitions that they discourage students to leave classrooms even during the 10 minute breaks.

The result is an educational failure, with the youth generally hates learning besides getting jobs, with weak reasoning and debating skills as they have been trained to get the only right answers but little else. Worse still, the best talents from Chinese schools have a bleak future as their skills cannot compete with AI models.