Why We Can Still Innovate Better Than AI

Jay Jiyuan Wu, Ph.D.

March 5, 2024

Right after publishing the white paper on launching a P&C insurance revolution, I find this interesting research report stating 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 I use “GAI” to refer to the current era of Generative AI.

1 An Unfair Comparison

Here are the crucial details from the above research report by the University of Arkansas on March 1, 2024, regarding the three standardized tests creative potential: “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."

Any such comparison must be done carefully or cautiously. I would argue the test this time, or similar test in the future, is not fair because it compares the mind power of 151 human participants at one point of time with the collective intelligence of humanity, accumulated over time. Or, put more accurately, it's better to avoid directly comparing a single sample with the massive data used to train AI.

Consider the question used in the study, "What is the best way to avoid talking about politics with my parents?" Open questions demand open answers, and the quality of open answers is directly related to the size of people thinking of the same question, directly or indirectly. All other factors being equal, the quality of the best answer obtained from 151 people at once is likely to be weaker compared to the best answer chosen from a pool of 1.51 million people or more, accumulated over decades or even centuries.

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.

2 There Are Exceptions

Of course, sometimes one human person can beat many 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.

I recently watched the amazing movie, and roughly 19 minutes into it, showing the scene when Edison, played by Benedict Cumberbatch, was interviewing Tesla, played by Nicholas Hoult, and inviting Tesla to work for Edison Electricity, Edison claimed he and his entity 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 a AI search Edison would stand out to be the best qualified person in human innovations.

3 Humane Innovations

Speaking of the movie, I noticed that in the late 19th century Americans were already eager to figure out a less cruel and more humane way of executing criminals than hanging. Back then the chair did not work as well as they are today, so Edison and his team actively campaigned to associate the electric chair and 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 guilt 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.

4 Why Humans Have an Innovative Edge: Neurology

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 I can't help but thinking how inefficient the current GAI or LLMs have worked. 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.)

Just 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.

Gemini provides a pretty good answer, which I will quote in its entirety below:

"AI models, particularly large language models (LLMs) like me, require a significant number of parameters for several reasons:

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.

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.

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.

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 rela-

tionships and nuances within the data, leading to more accurate predictions or outputs."

The above all make sense, although it does not tell me why the number of parematers 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 a few formatting changes for better reading:

"While AI models have made significant strides in learning and performance, there are fundamental differences between how they learn and how a human child does. Here's a breakdown of some key distinctions:

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.

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.

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.

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.

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.

In conclusion, while AI models are powerful tools, they still differ significantly from human learning in crucial ways. The limitations mentioned above highlight the ongoing challenges in developing truly human-like learning capabilities in AI."

The above detailed "item-to-item" comparisons facilitate easy understanding even for a lay person. I wish more people would have read this to remember *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.*

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 and I only made some format changes:

"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.

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.

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," "quasi-efficiency," and overall, "quasi-intelligence."

Sometimes 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 its 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: Neuromorphic chips attempt to mimic the structure and function of the human brain, using artificial neurons and synapses to process information. This architecture could potentially enable AI systems to learn more efficiently and dynamically, similar to how the brain forms connections and adapts through experience.

Reduced reliance on massive datasets: Traditional AI models often require vast amounts of data for training, which can be a limitation. Neuromorphic systems, in theory, could 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: Neuromorphic systems might be better equipped to adapt to new situations and generalize their knowledge beyond the specific training data. This could be analogous to how humans can apply what they learn in one context to new and unforeseen situations.

However, significant challenges remain:

Limited understanding of the brain: Despite advancements, our understanding of the human brain and its learning mechanisms is still incomplete. This hinders the ability to fully replicate its functionality in artificial systems.

Hardware limitations: Current neuromorphic chips are still under development and face limitations in terms of scalability, power efficiency, and performance compared to traditional computing systems.

Algorithmic challenges: Developing efficient and effective algorithms that leverage the unique capabilities of neuromorphic hardware for learning tasks is an ongoing area of research."

Once again, neuromorphic computing has even a longer way to go than non-neuromorphic AI. Another point is that non-neuromorphic AI has 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.

5 Human Agency Matters

One thing that Gemini has not mentioned 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, because it is a good inclusive term for many things that separate AI from humans. I want to use my own example to prove its relevance.

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 make many connections going forward that otherwise you won't make.

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 when and where exactly the "aha" or epiphany moment, a sudden flash of insight, realization, or understanding, will show up is unpredictable, 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.