ECKUITY CAPITAL



APPLES, Oranges AND AI

November 2024

APPLES, ORANGES, AND AI.

At our recently concluded AGM, one of the key points that generated a lot of curiosity was how we compare seemingly disparate industries and benchmark them alongside each other to make investment decisions. The bottom line is that – money is fungible. If I have \$1 dollar to invest in either a medical device technology or in AI improving surgical outcomes, how do I compare? Both markets have very different drivers, from the size of the addressable market, pipeline, reimbursement challenges, to pricing, competitive landscape, and eventual exits.

In a nutshell, we analyze risks and opportunities for each individual company, and then apply our proprietary "normalized MOIC" methodology which provides us unbiased insights and helps make objective investment decisions. For a detailed analysis of our process, please reach out to one of our team members and we will be happy to elaborate.

The analogy of apples and oranges is taking on a new meaning in the world of Al. If you had a chance to review our last note, we had promised to do a deeper dive into Al... and now we are keeping our promise. But when I last thought about writing a piece on Al, I had a few different themes in mind. With the passing of such a long time (around 90 days which looks like a lifetime in



the era of AI), I had to throw out half of my ideas. Not because they are irrelevant, but there are more use cases, more ideas, and more businesses that have replaced the initial ones.

19 DAYS VS. 4 YEARS

If you look at a time scale, these two numbers are obviously apples and oranges. You can't possibly have something that takes 4 years be done in 19 days or vice versa. Well, until now. As many of you would have read, Elon set up 100,000 H200 Blackwell GPUs in 19 days. Jensen claimed that anyone else would have taken the standard 4

years to execute. While this is a singular event given Elon's understanding of engineering, networking, and large-scale systems, the fact that we are able witness the world's fastest, easily the fastest indeed, supercomputer getting installed in under 3 weeks shows what's to come in the field of AI in the future.





to plan an installation of this magnitude, and then to fully equip it with liquid cooling and power to support the cutting-edge components would take another one year. The liquid cooling consumes massive amounts of power and just securing permitting for that is a mammoth task.

100% VS. 5%

Let's carry on the discussion about the power consumption. Once again, 100% vs 5% is like apples and oranges when looking at power consumption. Engineers from BitEnergy AI, a firm specializing in AI inference technology, have developed a means of artificial intelligence processing that replaces floating-point multiplication (FPM) with integer addition.

The new method, called Linear-Complexity Multiplication (L-Mul), comes close to the results of FPM while using the simpler algorithm. But despite that, it's still able to maintain the high accuracy and precision that FPM is known for. This method reduces the power consumption of AI systems, potentially up to 95%, making it a crucial development for our AI future.

Since this is a new process, popular and readily available hardware on the market, like Nvidia's upcoming Blackwell GPUs, aren't designed to handle this algorithm. So, even if BitEnergy AI's algorithm is confirmed to perform at the same level as FPM, we still need systems that could handle it. This might give a few AI companies pause, especially after they just invested millions, or even billions, of dollars in AI hardware. Nevertheless, the massive 95% reduction in power consumption would probably make the biggest tech companies jump ship, especially if AI chip makers build application-specific integrated circuits (ASICs) that will take advantage of the algorithm.

Power is now the primary constraint on AI development, with all data center GPUs sold last year alone consuming more power than one million homes in a year. Even Google put its climate target in the backseat because of AI's power demands, with its greenhouse gas emissions increasing by 48% from 2019, instead of declining year-on-year, as expected. The company's former CEO even suggested opening the floodgates for power production by dropping climate goals and using more



advanced AI to solve the global warming problem.

But if AI processing can be more power efficient, then it seems that we can still get advanced AI technologies without sacrificing the planet. Aside from that, this 95% drop in energy use would

also reduce the burden that these massive data centers put on the national grid, reducing the need to build more energy plants to power our future quickly.

While most of us are amazed by the additional power that new AI chips bring to every generation, true advancement only comes when these processors are more powerful and more efficient. So, if L-Mul works as advertised, then humanity could have its AI cake and eat it, too. (Source: Jowi Morales).

INTUITION VS. LOGIC

The eternal question of whether AI will be able to bridge this gap between apples and oranges. Intuition is until now squarely in the human domain, while logic is quickly being reshaped by AI. Let's come back to this after a quick historical event.

As bodies pressed against steel barriers, 97 people lost their lives in a crush of football fans. The 1989 Hillsborough disaster, Britain's deadliest sports tragedy, sparked a rush to judgment: football hooligans, the scourge of English sports in the 1980s, were blamed. It took decades of legal battles to reveal that police mismanagement, not fan behavior, was the true cause.

This initial blame is what Nobel laureate Daniel Kahneman, in his 2011 book Thinking, Fast and Slow, called "System 1 (fast)" thinking, which is intuitive. "Intuition is thinking that you know without knowing why you know," he wrote.

But the truth was revealed, thanks to families who demanded answers. Years of legal battles uncovered evidence of police mismanagement. This methodical analysis – Kahneman's "System 2 (slow)" thinking – finally exonerated the fans.

The medical profession, drug discovery, and operating room theaters all demand both swift intuition and careful reasoning. Now, AI is evolving to emulate this duality, potentially reshaping the medical field.

This is where neuro-symbolic AI comes into play – a hybrid approach that blends the strengths of neural networks (intuition) with the precision of symbolic AI (logic).



Neuro-Symbolic AI: Merging Intuition and Logic

Neural networks learn by analyzing patterns in vast amounts of data, like neurons in the human brain, underpinning Al systems we use daily, such as ChatGPT and Google's Gemini.

This data-driven processing aligns with Kahneman's "thinking fast" — rapid, intuitive thinking. While neural networks excel at finding patterns and making quick decisions, they can sometimes lead to errors, referred to as "hallucinations" in the AI world, due to biases or insufficient data. (Source: forbes.com).

In contrast to the intuitive, pattern-based approach of neural networks, symbolic AI operates on logic and rules ("thinking slow"). This deliberate, methodical processing is essential in domains demanding strict adherence to predefined rules and procedures, much like the careful analysis needed to uncover the truth at Hillsborough.

IBM's Watson exemplifies this, famously defeating chess grandmasters Garry Kasparov and Vladimir Kramnik. Chess, with its intricate rules and vast possible moves, necessitates a strategic, logic-driven approach — precisely the strength of symbolic Al.

Similarly, tax preparation software like TurboTax and H&R Block rely heavily on symbolic AI to navigate the intricate web of legal regulations and ensure accurate calculations. This meticulous, rule-based approach ensures each step is executed according to established guidelines.

ECKUITY'S NEXT-GENERATION PATIENT CENTRIC CONTINUUM OF CARE.

OpenAl's of model is not technically neuro-symbolic Al but rather a neural network designed to "think" longer before responding. It uses "chain-of-thought" prompting to

break down problems into steps, much like a human would. While it may appear to think, ol is not conscious or sentient. It's executing complex algorithms to produce this human-like reasoning, resulting in stronger problem-solving abilities.

Eckuity is using of modeling to build tomorrow's PATIENT 360° care continuum. We believe our approach is holistic, caters to major



stakeholders including the payors, providers, patients, large hospital chains, and big pharma and we firmly believe that any solution that neglects even one of these 5 is stakeholders will be unable to garner universal adoption. Additionally, ol showcases elements of agentic AI, where systems can act independently to achieve goals. This means that instead of just responding to prompts, AI agents can set objectives, plan steps and act to achieve them.

By breaking down problems systematically, of mimics human thought processes, considering strategies and recognizing mistakes. This ultimately leads to a more sophisticated ability to analyze information and solve complex problems.

In essence, of learns to reason by example, one step at a time. It serves as a bridge between Kahneman's concepts of thinking fast and thinking slow, aiming to deliver better reasoning with fewer mistakes. This approach paves the way for more advanced systems like AlphaGeometry that truly merge neural and symbolic approaches.

While OpenAI has garnered widespread attention with its new hybrid AI system, Google DeepMind is also making significant strides in this field with AlphaGeometry, announced in early 2024 in Nature. In tests, AlphaGeometry solved 83% of International Mathematical Olympiad geometry problems, matching ol's performance and nearly reaching that of human gold medalists.

Unlike ol, which is a neural network employing extended reasoning, AlphaGeometry combines a neural network with a symbolic reasoning engine, creating a true neurosymbolic model. Its application may be more specialized, but this approach represents a critical step toward AI models that can reason and think more like humans, capable of both intuition and deliberate analysis.

EKCUITY'S AI STRATEGY AND FOCUS AREAS

We believe AI will be a horizontal across all three of our vertical investment buckets: MedTech, HealthTech, and Biotech services.

Biotech Services: Biotech startups have made significant progress in using Alpowered platforms for drug discovery with drugs in late-stage clinical trials but have not shown them to be superior to the traditional way. This demonstrated that it is possible to take Al-generated drugs to the clinic, but approvals or exits are still an ongoing challenge.

However, Al-driven platforms for de novo drug design are still promising. These platforms use generative models to create novel molecular structures with desired

properties, significantly reducing the time required for lead optimization. Nextgeneration products focus on complex biologics that potentially enable the treatment of chronic diseases and difficult cancers. The success of these startups is often linked to their ability to form strategic partnerships with established pharmaceutical companies, access to high-quality datasets, and large cash reserves with strong investors to move multiple drugs forward in the clinic.

HealthTech: We have seen AI emerge in the care coordination and navigation space—a sector differentiated from care search, as navigation also includes more advanced services such as benefits management, claims assistance, and point solution referral. Other established care navigation platforms have been using AI on the back end to support human care coordinators and more accurately predict benefits utilization.

Over time, consumers are likely to become more comfortable interacting with AI chatbots for their care navigation needs; however, we still see an important role for human care coordinators leveraging AI on the enterprise side, as many patients may still prefer human interaction to understand their complex care and benefit needs.

There is a growing recognition that popular GenAI tools are being used for mentalhealth-related purposes without direct clinical supervision. This runs the gamut from using LLMs to combat loneliness to asking GenAI questions about how to deal with difficult personal situations. In contrast, formalized, dedicated mental health chatbots allow users to engage with an AI agent intended for mental health care that can ask probing questions and provide guidance on cognitive behavioral therapy exercises.

On the payer side, rising administrative costs and shrinking margins are pushing many to consider AI workflow enhancements. However, there are material risks at play. UnitedHealthcare and Cigna both currently face lawsuits over their use of AI to make claims approval/denial decisions. Another challenge is the disconnect between payment models for payer tech and services, which are traditionally priced on a per-member-per-month basis, and the volume-based costs incurred by GenAI technology companies. As a result, buyers may question the long-term viability of the solutions being pitched to them. Finally, many payers are worried about second order implementation costs for AI solutions, including cloud service upgrades and data integration. We expect payer AI adoption to proceed, but cautiously. (Source: PitchBook).

Some of the most potentially transformative work is being done at a foundational level by companies such as Hippocratic AI and HOPPR, which are building healthcare-specific LLMs, and by companies such as Evidium, Glass Health, Kahun, and Xyla, which are creating knowledge graphs and referenced ontologies that allow LLMs to be grounded in reliable, explainable medical literature. However, it will take some time to see whether this work will bear fruit in the form of faster and more advanced AI adoption in healthcare IT.

MedTech: Recent advancements in AI may enable the earlier detection of major diseases such as cancer, Alzheimer's, and Parkinson's, and better screening for these conditions could have a meaningful impact on patient outcomes and care costs. Leading university researchers have published compelling studies on how AI can detect neurodegenerative diseases such as Alzheimer's and Parkinson's years ahead of current diagnostic timelines. (Source: PitchBook and "The AI Revolution That Could Slow Parkinson in Its Tracks," NeurologyLive, Neal K. Shah, April 19, 2024; "How AI Can Help Spot Early Risk Factors for Alzheimer's Disease," University of California San Francisco, Victoria Colliver, February 21, 2024).

And in the emerging category of whole-body MRI screening, startups such as Prenuvo and Ezra Health are using AI to check for several hundred distinct diseases and health conditions. ML has also powered significant technological advancements in cancer screening and precision oncology through blood test liquid biopsies; over the past year, advanced AI has been implemented into established precision medicine tests, such as Guardant Health's Guardant360 diagnostic test. (Source: "Guardant Health Introduces Guardant Galaxy™ Suite of Advanced AI Analytics to Enhance Its Portfolio of Cancer Tests and Accelerate Biomarker Discovery," Guardant Health, January 31, 2023). While automation has been used in ECG monitoring for decades, recent advances in Al have enabled Al-based systems to more effectively analyze heart data, predict acute cardiac events, and improve data interpretation.



Beyond cardiac monitoring, AI is increasingly being used in tandem with wearables for remote patient monitoring in hospital-at-home, post-discharge, and chronic condition programs. The rising consumerization of healthcare also has relevance in AI patient monitoring, considering the ambitions of large consumer tech companies to provide clinical-grade health tracking capabilities. Glucose levels, blood pressure, heart rate, temperature, and oxygen levels are all major vital signs that could benefit from the use of AI to improve remote health monitoring. Over the longer term, we expect regulatory bodies such as the US Food and Drug Administration to approve more consumer applications for AI-based monitoring and diagnostics, and as this occurs, there will be greater opportunity for both startup companies and established market players to leverage AI for future growth opportunities. In conclusion, we believe we are at the starting point of the general AI race. The starting pistol has been fired but the difference is that this is a sprint and a marathon at the same time. One must have the speed to market and, in some cases, be the first to market, but at the same time, one has to have the stamina to sustain over the long haul. It will not be easy, there will be winners and losers, and we at Eckuity are looking to proactively shape the winners of this *sprintathon* !