March 15th, 2024
Koko Xsu

Material Intelligence
In the 2010s, many researchers and entrepreneurs attempted to solve problems involving *material* matters using artificial *intelligence*. These projects varied from autonomous vehicles to decoding neuro signals to designing new proteins to low energy Bluetooth. The majority of these achieved only modest success, and as a result, a lot of our scientific thought has been shaped by that scar tissue - "lidar is better than pure vision", "EEG is too noisy", etc.

It's time to rethink Material Intelligence (MI).

What Changed
The explosion of AI in the last five years has reopened these problems. The point of end-to-end neural networks is to learn underlying representations without needing to explicitly hardcode them - just imagine trying to write a sentiment analysis function using only if-then statements. Claude Shannon's Information Theory tells us that with more entropy there exists more information. Modern AI is our tool to interpret that *information* from the *entropy* (software problem), whereas previous attempts focused on increasing the *channel capacity* (sensor problem). A few things have changed since the 2010s:

1. We have better abstract frameworks.
   It is now extremely clear to me how to solve any material problem using artificial intelligence. This is my proposed Material Intelligence Framework:
   a) Identify a problem involving material matters
   b) Collect a little data
   c) Put that data through some transformers architecture
   d) See if it works at least a little bit
   e) If so, collect a little more data and train the model again
   f) See if its performance scales superlinearly with data quantity (is the d^2/dx^2 positive?)
   g) If so, you're onto something and it's time to climb the scaling laws curve
2. We have better mechanical tools.
   We now have modern tools like batch norm, transformers, a better understanding of how to train embedding models, etc. This modern ML toolkit is still making its way around the scientific community.
3. Research coevolves with engineering.
   Scientists design problem-solving approaches using tools at their disposal, and as tools get better so do their approaches. Take robotics for example: modern robotics AI as a field really only emerged in the last five years. The first wave of robotics in the early 2000s led to the adoption of Kuka and Amazon Robotics, but those hardware were not designed with modern robotics AI in mind. It turns out that for precision manipulation tasks, having a camera under the robotic hand makes training the vision model so much easier. So now, the next generation of robotics hardware will be designed with software considerations like this in mind, and as our understanding of the software limitations

advances, we repeat this cycle and tweak the hardware. Because AI is so new, this research-engineering coevolution has yet to happen for most Material Intelligence problems.

Areas of Interest

There are endless material problems to be solved using modern AI. I'll list a few here:

- Neuroscience
  - The 2010s scientific movement shifting from EEG to EMG is a reflection of an industry shying away from the entropy problem toward solving the channel capacity problem. EEG was too noisy to extract useful information from because of its proximity to the brain, so it's easier to interpret EMG data on the wrist that has been pre-filtered by the motor cortex. Now that interpretation tools are better, I expect a comeback of brain-based approaches.
- Ultrasound
  - Hardware in ultrasound is in the early innings of exponential growth, and the applications that are possible are already mindblowing. Preliminary research has shown promising applications beyond detection, namely in cell stimulation. With the MI Framework, I expect lots of useful technology to emerge from ultrasound in the coming decade.
- Biology
  - Lots of ink has been spilled on this. Virtual cell, platform discovery, biomanufacturing, etc all have problems that can be explored using the MI Framework.
- Aerodynamics, Cryptography, Wearable Health, Quantum Materials, IoT…

Building Companies

So you've found a problem and took it through the MI Framework. Now what? The next steps to building a company around that problem are usually the following:

1. Build a commercial data collection device that ideally adds some immediate user value, leading to large-scale adoption and data collection.
2. Take the data and train better models.
3. Take what you learn from the model limitations and tweak the hardware. Iterate.
4. Eventually reach a final hardware product that runs the model at high enough performance that it becomes a magical product.

This process was pioneered and best demonstrated by Tesla.

Lingering Problems

However, the main problems people run into here are:

- How much data is enough?
  - It's hard to estimate the exact shape of the scaling laws curve for each problem.
- My hardware is useless with a bad model?
  - It's hard to find a commercial device that delivers user value without good AI.

If you have clever solutions to these problems, or if you're working on a problem under the MI Framework, reach out!