

Fast, efficient, and explainable AI, meet Literal Labs

As part of our quick founder questions series – or QFQs – we spoke to Noel Hurley, CEO of Literal Labs about Tsetlin Machines, academic infrastructure, compute efficient AI and being on calls with VCs.

Temps de lecture : minute

17 December 2024

There had been a lot of research going around about the Tsetlin Machine approach to AI, but at that point in time, it had purely been academic. As the market moved forward, having fast, very energy efficient AI became, frankly, both more important and more valuable in the marketplace.

It had also become very clear that the energy levels and the computational requirements of neural networks were acting as blockers. My view was there needed to be an alternative approach to solve the computational complexity problem of neural networks, and our Tsetlin Machine ticked many of those boxes. That is what triggered commercialisation of the technology.

The academics in Newcastle University had been doing this collaborative research for about five years - Profs. Alex Yakovlev and Rishad Shafik. They had proven the technology's robustness in terms of base benchmarks and had built the necessary accelerator architecture. All the pieces came together at the right time.

Tell us about the business - what it is, what it aims to achieve, who you work with, how you reach customers, and so on?

What we're building is a new approach to energy efficient AI, which is based around a combination of propositional logic and Tsetlin Machines.

The technology delivers AI models that are fast, incredibly energy efficient, and reduced computational complexity for running said models. This also affects the cost point that you can achieve and their logic base, which makes them explainable.

We're focused on making this Tsetlin Machine technology available to the market. We're building training tools that allow our customers to take the data sets they already have internally and train Tsetlin Machines for their own applications using the tool chain we have.

By focusing on training tools we are able to build a business that will allow us to contribute to an open-source community and incorporate open-source contributions, which means it can benefit everyone - us, our customers, and other parts of the ecosystem.

The market we're focused on initially is the edge AI market. Edge AI is a space that is highly constrained in that at the edge of the network, you've got computation constraints and energy constraints, as many of these applications are battery based and cost constraints. We are an uber efficient technology, which is why we operate in a very tight knit space. Though we're not limited just to edge applications, it is an area where we're focusing our efforts on as we build the company.

How has the business evolved since its launch?

Literal Labs was founded in the summer of 2023 with very academic infrastructure from an engineering perspective.

As the business has evolved, we have started to build out the team. I brought in a new CTO in May, who was previously the Deep Learning Platform Lead at AstraZeneca, along with a Head of Product who comes from an IoT background.

We then focused on building a training pipeline that is scalable, which meant moving from academic environments with laptops and sporadic computing to a cloud-based, scalable platform that will allow us to build our products over the top.

We're in the phase of bringing academic technology into an industrialised product environment. This is where we've been spending our effort, and, frankly, it's where we're going to continue spending our time in the next 12-18 months. The ultimate aim is to put that product in the hands of our customers.

Tell us about the working culture at Literal Labs?

We are, and want to remain, an IP-heavy business. To do that, we're building a team of AI and tooling experts.

We have a deep vision of trying to develop AI that is good for all and has a light touch on the environment. That's key to the kind of business we want to build going forward. We want to be the number one company when it comes to energy and compute efficient AI.

We're constantly searching for excellence and elegance in design, under the banner of full stack thinkers. This means we're really thinking about the entire solution, always trying to put ourselves in the position of the customer, and truly understand everything from their perspective.

We then need to think carefully about how we operate, iterate, and converge. You have to experiment to find the best solution. You need to go round and round through that iterating process, which means that culturally, there shouldn't be a fear of failure or criticism of failure.

Once we've gone through this iteration phase and figured out what we want to achieve, then there is a convergence. Everyone gets together and we focus on delivery, delivery, delivery. It's not only good enough to have a good solution, your customers are going to expect you to deliver on time and on target.

Those are the four tenets of our culture - the search for excellence, full stack thinking, iterating, and converging.

How are you funded?

We had a SAFE note open when the company was set up, which is what we've been living off to date, and we're currently fundraising through venture capitalists.

What has been your biggest challenge so far and how have you overcome this?

I think there have been two. The first is on the engineering side, which is effectively, how you take a piece of technology developed within academia and turn it into an industrialised solution. This is a challenge we're continuing to work on.

The second challenge is of course that it's pretty hard to raise capital right now, especially for an early-stage business. It's hard to value a business like us without any track record. Though we have incredibly promising technology, I recognise that having spent a good portion of my career in corporate life, we are a risky prospect for investors. This is a perennial challenge for early-stage, academic spinouts.

How does Literal Labs answer an unmet need?

We're not talking about percentage improvements in computational requirements versus a neural network. We're talking about orders of magnitude sized improvements.

The premise is that with Edge AI, you can gain insight into what's happening within a production environment, within a machine. That insight will improve productivity, and that productivity will improve your profit line, and everyone wins.

But actually, because the computation requirements are so high, what's actually happened is that you can't afford to put the AI on the edge. You have to push it back into the cloud, which means that what you've done as a customer, is you now have a large cloud vendor as part of your value chain. We have reduced the complexity by orders of magnitude, so that AI can now, as promised, happen on the edge. And effectively, what gets transmitted back is the metadata from the decisions that those edge devices are making.

We believe that we will end up with AI solutions that cost orders of magnitude less, are easier to deploy, and will have an effect on where you can deploy them. You can deploy them across AI and many more applications now, because we've tackled this cost.

What's in store for the future?

For a company of our size, what we need to do now is close our funding and execute on our plan. In many ways, our technology is delivering the future. We know the advantages Tsetlin Machines will bring and we now need to get our training tools to market so that others can experience the benefits and access this game changing technology

What one piece of advice would you give others founders or future founders?

To raise capital, you have to kiss a lot of frogs before you find your Prince Charming, as it were. I confess I underestimated just how many VC meetings and calls you need to go through. It's ultimately a game of numbers. As a CEO, you need to dedicate most of your time to doing that, and in those early days, you're totally dependent on that funding.

The other thing I would say is to ensure you have a strong core team around you. For example, I brought in a CTO with a lot of industrial experience who buys into our vision and is developing that vision. I also brought in a head of product to focus on the defining our product and being that interface between potential customers and our engineering team

And finally, a more personal question! What's your daily routine and the rules you're living by at the moment?

As the CEO you are juggling a lot of balls, you are the "Chief Everything Officer"! Certainly, through this period, my day is dominated by VC calls. As the CEO, everyday you are effectively cycling between the financials, company management, fundraising, and developing the team and vision.

Noel Hurley is the CEO of *Literal Labs*.



MADDYNEWS UK

The newsletter you need for all the latest from the startup ecosystem

[SIGN UP](#)

Article by Noel Hurley