How to Find Out if your AI project is a Business Case

Glenn Fitzgerald, Chief Data Officer, Platform Business, Europe at Fujitsu

What's it going to take to get AI more widely adopted in business? asks Glenn Fitzgerald, Chief Data Officer, Platform Business Europe at Fujitsu. Perhaps building use cases and testing on reliable infrastructure. Picture yourself taking an AI «test drive».

For all the talk that "AI is everywhere" these days, the reality is still that not many projects are making it out of proof of concepts (POCs). Even getting POCs off the ground can be tough.

Most people get the upside on offer with AI: the potential to automate processes that were just too hard to tackle without it; or the ability to inject much more meaning and nuance into customer and employee interactions. But there are high hurdles too.

Held back by infrastructure, cloud costs, and skills shortages
AI can be very compute-intensive. Few companies have the infrastructure available to handle it. Cloud alternatives exist but are prone to runaway costs that have frightened off many would-be AI trialists. The more data that needs to be trained, the more money the organization will need to invest. This can become cost-prohibitive very quickly, and performance can be impacted.

AI infrastructure – especially for AI training purposes - needs to be well-connected, with high-speed, secure access to the cloud and data brokers. Most legacy or existing cloud infrastructures cannot scale effectively to meet AI demands. Upgrading networks to handle the flood of data that's needed between you and the cloud is something that's regularly overlooked.

And the hundreds of open-source AI tools in libraries like Google's TensorFlow, Microsoft CNTK, and Core ML by Apple are great, but they all need to adapt to your organization and your use case. Few companies have the skills in place to do that.

New approaches to AI
Those high hurdles are now starting to come down, however. MLOps, for example, is a set of lifecycle management practices to deploy and maintain machine learning (ML) models in production reliably and efficiently. It is practiced between data scientists, DevOps, and ML engineers to transition an ML algorithm to production systems.

But that still leaves the question of how to test out an idea for the first time - in an organization yet to assemble the necessary infrastructure or skillsets.

At this practical level, there's now a new AI "test drive" facility based on PRIMERGY servers at Fujitsu in Europe that allows Data Scientists who want to try out AI but — before committing resources — want to find out if it's a viable use case and what (and how much) it's going to take to overcome the barriers.

The very hard hardware problem
The AI test drive overcomes one of the biggest stumbling blocks that data scientists often face – finding the hardware and running tests. AI could require massive amounts of computer power. However, most organizations don't have the resources to buy their own AI infrastructure – especially for testing purposes –, putting them at a disadvantage when running AI projects.

It's a chicken-and-egg challenge. Secure buy-in from stakeholders to invest in computing power and other resources requires a watertight business case that clearly demonstrates the potential return on investment. However, developing this business case is impossible without access to the computing power needed to determine the project's requirements accurately.

Fujitsu's AI Test Drive offers this opportunity - mainly free of charge. It provides the computing power, as well as the network capacity, open-source tool tweaking, and hand-holding that AI-inquisitive business users need to work out if they have a viable business case. By using the AI Test Drive, data scientists can remove the guesswork and create a solid, fully-costed argument for investors to realize the power of AI and drive their research further, faster.

Fujitsu will also advise if AI is the right approach in the first place — it often isn't. We estimate that in up to 60% of cases, conventional automation (using RPA, for example) offers a better, faster, lower-risk solution to the business challenges we see than AI.

GPUs were a temporary fix

In terms of technology choices, it's no longer a slam-dunk for Graphics Processing Units (GPUs). In the early days of AI, engineers quickly realized that GPUs, developed for gaming and other graphics-intensive applications, provided an off-the-shelf solution to the need for ultra-fast processing.

Since then, CPUs (Central Processing Units) and the software libraries that run on them have evolved to become much more capable of deep learning tasks. For example, through extensive software optimizations and the addition of dedicated AI hardware, , CPU-based systems have enjoyed improvements in deep learning performance. Additional tools simplify deep learning inference deployment for hundreds of pre-trained models and enables developers to quickly optimize and deploy the AI workloads with improved performance, from edge to cloud. Therefore, CPUs shine for many applications, such as high-definition-, 3D-, and non-image-based deep learning on language, text, and time-series data.

CPU-based systems are also generally simpler and more robust, for example, in edge environments. GPUs have higher power consumption and cooling requirements, while CPUs are available in various proven standard systems for easy deployment in the data center and on the edge.

That insight shows how Fujitsu's AI Test Drive has been built by working within a broad ecosystem, including industry-leading technology partners such as Intel®, SUSE®, NetApp®, and Juniper®. That ecosystem gives data scientists direct access to the technology they need to validate their theories and models and help them build a strong business case for investment.

Register here for a Fujitsu AI Test Drive: www.fujitsu.com/global/ai-test-drive

Glenn Fitzgerald
Chief Data Officer, Platform Business, Europe at Fujitsu

Glenn joined ICL in 1979 as an apprentice and has worked for the company, now Fujitsu, throughout a varied career, during which he has gained expertise in a wide range of IT fields, including manufacturing and production test, hardware, software, and firmware design, infrastructure implementation, project management, and business and ITC architecture.

In his role, Glenn is responsible for the development of the technical aspects of the Data-Driven Transformation Strategy within the Fujitsu Product Business, ensuring that the technologies utilized within that strategy have the necessary capacities to resolve customer business issues.

He is also developing the solution consulting capability within Product Business to assure the delivery of industry-leading solutions that support clients' varied businesses with the "art of the possible" at any point in time; a balance of technical feasibility, cost, timescales, risk, and flexibility.