

RESEARCH & DEVELOPMENT

How to Set Up an AI R&D Lab

by Foteini Agrafioti

NOVEMBER 20, 2018



PAPER BOAT CREATIVE/GETTY IMAGES

The moment a hyped-up new technology garners mainstream attention, many businesses will scramble to incorporate it into their enterprise. The majority of these trends will splutter and die out by Q4. Artificial intelligence (AI) is unlikely to be one of them.

AI is a transformative series of tools that can accelerate productivity, drive insight, and open up unexplored revenue streams. It's poised to revolutionize the way we do business and everyone in a leadership role should be thinking about it.

But few organizations are set up to do AI properly.

There's a common misconception that rebranding as an AI company is as simple as having data, infrastructure, and off-the-shelf data and analytics. The reality is that AI is complex, high-risk, expensive, and often requires significant business transformation. Most importantly, though, it demands an ultra-specialized talent pool that, according to latest reports, currently stands at only 22,000 PhD-level experts worldwide – a remarkably small pool.

When you consider that the annual market value predictions for AI techniques range between \$3.5 trillion and \$5.8 trillion, it's clear why the battle to recruit from this scarce group has become the industry's defining challenge. Right now, the biggest companies in the world are scooping up these experts to populate their teams because they understand that the key to building a robust, successful AI practice is to find and retain the right talent.

So, when the Royal Bank of Canada (RBC) approached me three years ago to help them grow their AI capabilities, this is what I advised: To go beyond data science and do real AI, you need to hire the right people, embrace research, and adapt your culture.

At the moment, AI is more of an open frontier than an industry-friendly space. Its applications are new enough that actual practitioners of AI don't yet exist at scale. This makes finding, retaining, and nurturing talent the field's most pressing challenge. If you want this capability in your organization, you have to hire the people with the perfect balance of data intuition and state-of-the-art knowledge. These people are almost all academics.

It's impossible to overstate the importance of this expertise level. Machine learning models run on mathematics whose subtlety requires a deep understanding of data domains. Simple tasks can be daunting to those without the right experience. Mistakes in the machine learning space are extremely common, but when applied to business they can have real-life consequences too - up to and including life-and-death.

Here's a textbook example of how easily these mistakes can occur. In 1973, the University of California at Berkeley compiled their graduate school admissions figures and discovered what appeared to be a significant bias against women applicants. The numbers were right there on the

page: 44% of male applicants were admitted to graduate programs, versus 35% of female applicants.

Fearing a lawsuit, school officials shipped the data to their statistics department for a closer look. A team led by Peter Bickel, now the school's professor emeritus of statistics, was able to decode the figures by parsing individual departments. When analyzed this way, the data didn't provide evidence of bias; the issue was that women were applying to harder, more competitive programs with lower admission rates than their male counterparts. In fact, women had slightly higher admission rates than men per department. (This doesn't mean there was no gender bias at play; it simply means the evidence that was being used to allege that bias - lower admissions rates for women - was not itself an indication of bias.)

Researchers will be familiar with this phenomenon, known as Simpson's paradox. Machine learning challenges have significantly increased in complexity since then and it takes years of training and experience to develop a well-honed intuition that can sniff these problems out. In the case of Berkeley, the administration had access to an in-house team of top-flight statisticians. But it's easy to see how inexperienced analysis could have tied the university up in a costly, lengthy legal battle. Instead, the university emerged with a more nuanced understanding of its admissions process.

In a knowledge-based economy, research becomes the means of production. This recognition should put an end to any misconception that as an AI-enabled business you can "get away" with not conducting in-house research. What leaders tend to miss here is that the scientific progress we've made in AI does not automatically render the technology ready for any environment. Each business carries its own unique challenges and requirements, from proprietary data types to operational constraints and compliance requirements, which may require additional customization and scientific progress.

This makes AI an extremely high-risk, high-return pursuit. To pursue research may, in this moment, seem like a novel and bold act. But it should be the norm. There is a vast competitive advantage that comes from owning these solutions for your industry. The companies that invest in research that adapts machine learning to their industry will generate extremely valuable intellectual property (IP).

One way to de-risk this pursuit is to pair fundamental research (pure science) and applied research practices together. The scientific pursuit is by definition open-ended: you can go on forever, pushing boundaries and conquering knowledge. (This is why universities will never go out of business.) Applied research, on the other hand, is designed to solve specific real-world problems. What applied researchers bring to the table is knowing when to stop researching and to focus on delivering a solution. Each practice serves to influence the other. It's a simultaneous push forward that is less likely to end up with no payoff.

For instance, one of the areas we're focused on at Borealis AI (the R&D arm of RBC) is natural language processing (NLP), the field of AI that can understand language. So far, NLP has proved most powerful for parsing and analyzing bodies of text in order to extract meaningful patterns. Our particular interest in this type of machine learning involves training NLP on news datasets to predict how global events can affect the trajectories of companies. The ultimate goal is to build software that can direct financial analysts in real-time toward relevant information within their respective industries.

It's important to note that the state-of-the-art in machine learning has not yet reached the level where it can solve every aspect of this problem. NLP algorithms perform best in constrained environments, such as question and answer systems, where a user can ask a computer questions from a finite list of possible queries. Since the computer knows what to anticipate in *language* form, the program can respond accordingly.

But when the goal is to understand something as dynamic as news and apply that information to the relationship between companies and world events, it requires the ability to contextualize, then track the evolution of very complex and time-dependent entities. This is where the marriage between fundamental and applied research comes into play. In our case, fundamental research aims to advance NLP to a place where it can independently perform high-level language-based reasoning and grasp complex relationships at the same level as humans do. Our applied researchers then ensure these solutions can become immediately applicable to financial services. This is how we build products while pushing the boundaries of science.

The last step to building an AI practice is to create the right environment. The world is an AI researcher's proverbial oyster right now. In an uncertain economy, they're among the rare few whose opportunities continue to become more fruitful and multiply. Offering potential talent a dynamic, comfortable, and unique workspace is a good start. You also need to be able to offer compelling datasets and interesting problems to work on. Computational power and a strong team to provide research mentorship are also required. But the ability to pursue curiosity-driven research is the real draw.

When you hire from academia, you're inviting a group of people into your organization who come from a very specific culture. They share values that are built on both the ethos of solving big, meaningful problems and having the ability to publish the results of their efforts. Researchers take pride in contributions they make, so these factors must be in place. This translates into reproducing some of the working conditions they've brought from academia such as while allowing the transparency of collaboration and open publication that serves to advance their community as a whole. Businesses working in closed environments need to reconsider that approach. If you are operating in this arena, it is up to you to prove your credentials and not the other way around.

Operating a business at the moment of this technological shift is a rare opportunity to seize upon an economic turning point. While we can't yet predict how it will re-shape the market, the prevalence with which AI is already embedded into our core technologies favors early adoption.

Unlike the last industrial revolution, however, investing in big machinery won't cut it. To truly have impact and remain relevant in the market, it's up to executive leaders to build the bridge between research and commercialization. Only in this collaborative vein will AI's true impact flourish.

Foteini Agrafioti is the Chief Science Officer at RBC and Head of Borealis AI. She is responsible for RBC's intellectual property portfolio in the fields of artificial intelligence and machine learning.

This article is about RESEARCH & DEVELOPMENT

 FOLLOW THIS TOPIC

Related Topics: ANALYTICS | TECHNOLOGY

 Loading...

 Loading...

Comments

Leave a Comment

POST

0 COMMENTS

 [JOIN THE CONVERSATION](#)

POSTING GUIDELINES

We hope the conversations that take place on HBR.org will be energetic, constructive, and thought-provoking. To comment, readers must sign in or register. And to ensure the quality of the discussion, our moderating team will review all comments and may edit them for clarity, length, and relevance. Comments that are overly promotional, mean-spirited, or off-topic may be deleted per the moderators' judgment. All postings become the property of Harvard Business Publishing.