CHAPTER

4

The Team

The main roles for your AI project

The recruitment of AI talent can be a high-stakes game. In some cases, a company may be willing to offer multi-million dollar packages along with significant equity options.

To get a sense of how competitive the environment can be, look at the example of Uber. In 2015, the company was struggling in its efforts to create an autonomous vehicle. The biggest problem: There simply was not enough talented AI engineers on the payroll. The fact is that it was extremely difficult to compete against giants like Apple and Google.

So Uber took a more aggressive approach to recruiting. The company went to Carnegie Mellon University and hired away 40 researchers and scientists. Of course, the school had a distinguished history as an innovator in AI. For example, it was the first to have a degree in machine learning.

Uber had raised $5 billion so there was more than enough resources to offer lucrative pay packages. Let's face it, when it comes to salaries for the academic community, they are far from attractive. Keep in mind that Uber's move was so bold that it garnered headlines in the media.[1]

As should be no surprise, Carnegie Mellon University took a hit. The department saw a reduction in outside funding because it did not have enough resources to execute on the projects.

[1] www.wsj.com/articles/is-uber-a-friend-or-foe-of-carnegie-mellon-in-robotics-1433084582

T. Taulli, *Implementing AI Systems*, https://doi.org/10.1007/978-1-4842-6385-3_4

But the Uber episode raised broader concerns. Might Big Tech wield its power to scoop up most of the talent at universities and perhaps stifle creativity and innovation? Aren't academic institutions more willing to explore basic research, which has been critical for cutting-edge developments?

This is all true.

In light of all this, the tech industry saw that there needed to be a better approach. The result was that companies like Facebook, IBM, and Microsoft set up partnerships with universities. There would not only be funding but programs where researchers and scientists could take leaves so as to work for private firms.

It was a reasonable strategy and it has worked fairly well.

But this has certainly not dampened the intense poaching among tech firms! For example, Apple has been particularly aggressive. During the past few years, the company has made the following high-profile hires:

- *John Giannandrea*: Apple recruited him away from Google to become the company's first Senior Vice President of AI. As a sign of the importance of the position, he would report to CEO Tim Cook.
- *Ian Goodfellow*: He was another recruitment from Google by Apple. As the creator of the GAN (Generative Adversarial Network), which we covered in Chapter 2, he was one of the top people in the AI field (he received his PhD at the University of Montreal). At Apple, he took on the role of Director of Machine Learning for the Special Projects Group.[2]

Apple has been hiring academics as well. There was the recruitment of Ruslan Salakhutdinov, a professor at Carnegie Mellon University, to be the Director of AI research.[3] There was also the hiring of Carlos Guestrin, a machine learning expert from the University of Washington.

Oh, and another part of the recruiting process was actually the purchase of startups. These deals have often been referred to as acqui-hires since the main goal is to snag talented engineers. Some of Apple's deals included purchases of Laserlike, Drive.ai, and Fashwell.[4]

---

[2] www.cnbc.com/2019/04/04/apple-hires-ai-expert-ian-goodfellow-from-google.html

[3] www.wsj.com/articles/apple-hires-artificial-intelligence-executive-from-rival-google-1522811544

[4] https://venturebeat.com/2019/12/23/how-the-big-5-bolstered-their-ai-through-acquisitions-in-2019/

Given all this activity, it is certainly tougher for smaller firms or non-tech operators to compete for talent. Yet this should not mean that there is no hope. There are definitely strategies and approaches to building great AI teams.

In this chapter, we'll take a look at these strategies and approaches. But first, let's cover the main roles of an AI team.

The Executive Sponsor

When it comes to implementing new technology, AI is not the only one that has major issues with adoption. All organizations suffer from inertia. Employees get accustomed to certain approaches and processes when it comes to doing their work. When there are attempts to change this, there is often resistance. It's natural. It's almost inevitable.

Regarding AI then, a bottoms-up adoption of the technology is not realistic. There needs to be substantial resources devoted to the effort, in terms of hiring new employees and purchasing new software and systems.

Because of this, an AI project needs an executive sponsor. This is a person who is at a senior level and has the ability to make significant budget decisions. The executive sponsor does not need to be from the tech side, say the CTO. The person can be anyone. But the key is that they have influence in the organization and are willing to take risks.

It's actually common for the executive sponsor to be a head of a department like marketing or sales. The need for AI may arise because there are challenges such as with improving performance or fending off rivals.

However, because of the growing popularity of AI, there still may be bottoms-up activities. And this is OK. This is actually a good sign. It means that adoption may be easier. But then again, an executive sponsor should see this as an opportunity to centralize the initiatives so there is not too much confusion. They will be in a position to set forth the vision and major goals, which will help get the project off to a good start.

"Someone needs to understand what the business is going to do with the insight the AI finds and be a champion for the team," said Pat Ryan, who is the Executive VP of Enterprise Architecture at SPR. "AI teams typically do not produce output for every 'sprint,' like a development team. In many cases, the AI team will spend weeks analyzing, experimenting, and gathering more new data, which could seem to the rest of the organization like they are making little progress. It's the executive sponsor's job to make sure the business understands the value of the team."[5]

[5]From the author's interview with Pat Ryan on June 12, 2020.

A recent survey from Deloitte pointed to the critical importance of having an executive sponsor. According to the results, the CEO was actually the lead champion for 29% of the organizations and they were 77% more likely to have exceeded their business goals. They were also 59% more likely to get valuable insights from their models.

Tim Smith, a principal at Deloitte Consulting LLP and the Technology Strategy and Business Transformation Practice Leader, said, "Data science has to permeate company culture starting from the top to see true benefits."[6]

The Project Manager/Business Owner

The executive sponsor does not have the time for day-to-day management of the AI project. In other words, there needs to be a solid project manager or business owner. The focus will be on hitting the goals but also allowing for innovation and new ideas.

The project manager should be a quick study. That is, they need to have a high-level understanding of AI and its capabilities. This will be crucial to effectively communicate with the team, which will involve a mix of technical and non-technical people.

The project manager can be recruited internally. If anything, this may be better since the person will have a good understanding of the business.

In terms of the role of the AI project manager, it is fairly new and continues to evolve. But it is also far from easy. A successful AI project manager must be multi-disciplinary, be willing to keep learning, and be a good listener.

A major part of the role is administrative. There should be clear project plans and timetables. To this end, there are a variety of tools that can help out, such as Asana, SmartSheet, and Monday. They will also help to focus the team and help with the management.

The AI project manager's portfolio will be extensive and will involve the overseeing of data development and cleanup, evaluation of AI tools, help with infrastructure resources, creation and testing of the model, and deployment. But even after this, there is often a need to monitor the models. Over time, they can easily degrade and break down, as the data may be too old or the algorithms not robust enough to handle new conditions. Because of this, as more models are deployed, the management requirements can get quite complicated.

---

[6] www2.deloitte.com/us/en/pages/about-deloitte/articles/press-releases/deloitte-survey-analytics-and-ai-driven-enterprises-thrive.html

Something else: An AI project manager will have to work across different departments within the organization, say IT, legal, marketing, sales, and so on. They will also need to be able to communicate with executives in order to keep up the momentum for the project. Maintaining the buy-in will be crucial in handling bottlenecks and politics.

As for managing the AI team, the project manager will have the challenge of working with technical and non-technical employees. This means having, at times, to mediate issues and fights.

Even when a project manager is able to facilitate a project, there is the problem that AI often involves much trial-and-error. So in the early stages, it is likely that there will be failure, which can hurt the morale of the team and cause concern with the executives. This is why it is essential that everyone understands that AI is never smooth. It takes discipline and patience.

Then what type of background to look for in an AI project manager? You definitely want someone who has experience working on major projects where there were notable risks and a need for multi-tasking. Now this does not necessarily have to be a tech effort, but it helps. For example, it is common for AI project managers to have experience with software development projects. Another thing that is critical is that they have good business writing skills, with an ability to make complex topics understandable. The role will involve lots of emails and Slack communications!

The Subject Matter Expert

The SME (subject matter expert) is someone who has a strong understanding of a particular domain. For example, they may be the person who knows a process in the company, say for invoices or how to process a claim. Or the SME could have a background on how to deal with customers in an industry.

In relation to an AI project, you can usually find the SME within the organization, although in some cases there may be a need to get a consultant. So if you are looking to rethink a company process, you might want to have someone with a background in a methodology framework like Six Sigma or lean. They will also bring a fresh perspective to the process.

The SME will help with evaluating the types of data needed. They will also work with a data scientist to craft the model. And then the SME will provide assistance in evaluating the results.

In some circumstances, the hiring of an SME can be creative. Just look at Intuit. When the company built its conversational AI bot for QuickBooks, it recruited Scott Ganz as the Principal Content Designer. His background? He was actually a screenwriter! He won an Emmy for his work on *Wordgirl*.

His role has been to work with the engineering and design teams to develop a personality for the AI system. For the most part, the focus has been more than just understanding the customer but getting a sense of the underlying emotions. This is certainly critical when it comes to sensitive issues like money.

Ganz has also been helpful in simplifying the user experience. For example, instead of the bot talking about "accounts receivable," it mentions "who owes me money?"[7]

The Business Analyst

A business analyst is someone who uses data analysis to help improve the operations of a company. Such a person usually has a business degree or MBA, with a background in using BI (Business Intelligence) tools such as Tableau, MicroStrategy, and SAS. There are also a variety of certifications for business analysts: PMI Professional in Business Analysis (PBA), IIBA Certification of Competency in Business Analysis (CCBA), and IQBBA Certified Foundation Level Business Analyst (CFLBA).

Some of the common tasks include creating financial models, say to forecast sales or churn. They may also look at analysis to help reduce costs or to improve pricing. A business analyst will then create reports or dashboards, which are presented to executives and managers.

Then what would the role be for an AI project? It could actually be significant. After all, many projects are focused on rethinking processes or lowering costs. What's more, a business analyst can find and evaluate data sources. There could also be integrations with the BI systems.

Because a business analyst will have analytics skills, they could be a good candidate to ultimately transition to the role of a data scientist. Or, as is becoming more common, a business analyst could become a "citizen data scientist." That is, the person does not have the usual qualifications but can still perform a variety of tasks of a data scientist.

The Data Engineer

A data engineer will help create reliable data pipelines, such as with cloud databases and data warehouses. This process will often involve extensive wrangling of the data sets like deduplication, handling missing items, finding anomalies and outliers, and even detecting potential bias.

---

[7] www.forbes.com/sites/tomtaulli/2019/08/30/how-screenwriting-can-boost-ai-artificial-intelligence/#a3eff56586aa

When looking for a data engineer, you want someone who knows Python, R, and SQL. In terms of the educational background, it is usually a Bachelor's or Master's in statistics, mathematics, or computer science. They should also have experience with the following:

- Workflow engines such as Azure Data Factor, Airflow, Google Cloud Composer, and so on
- Data modelling
- ETL (extract, transform, load) design, implementation, and maintenance (this is a process for data integration, usually for a data warehouse)
- Querying of databases like Spark, Hive, Presto, Hadoop, and so on
- GIT version control
- A/B testing

The data engineer will usually work closely with the data scientists. They will not only present the data sets, in a good form, but also provide input for the modelling and testing.

The Data Scientist

The data scientist, who may also be called an AI engineer, is the person who develops the AI models. No doubt, this person often gets much of the attention. But they will still rely on a team of other important people. Perhaps there are some unicorn data scientists, but they are very rare!

■ **Note** There's an old joke: A data scientist is someone who knows more about statistics than a computer scientist and more computer science than a statistician. OK, this is not too funny! But this does point out that a data scientist has a blend of unique technical skills.

A data scientist will have a range of knowledge and capabilities that includes traditional statistics, probability, and decision theory. And yes, they should have a background in understanding machine learning, deep learning, and perhaps NLP. But a good data scientist will also have an understanding of basic business principles. This is essential since they will work with a myriad of stakeholders like those in marketing, sales, finance, and legal.

Their educational background is usually quite extensive. For example, it is common for a data scientist to have a PhD or Master's degree in computer science or machine learning. They will also have experience with academic theories and programming languages like Java, C, C++, Scala, and Python. They will then have to know AI platforms such as TensorFlow, Keras, scikit-learn, Caffe, and PyTorch, as well as cloud systems like AWS, Azure, and Databricks.

The Designer

Even if an AI model works, has accurate results, and provides useful insights, this may not be enough. Consider that the end user is often someone who is not technical. So if the AI model has a complex UI (user interface) or UX (user experience), then there may be little adoption. The result would be a waste of resources, and the AI initiative may ultimately be abandoned.

To avoid this, you can have a designer put together a compelling AI application. This could involve a simplified menu, effective graphics, and helpful visualizations.

A designer does not necessarily have to be a full-time person. It is actually more common to have this person involved on a contractor basis.

AI Tester

AI models can be temperamental. Even slight changes can have wide variances in the outcomes. While a data scientist can provide some level of testing, there should still be someone else to lead this role. This person is known as the AI tester or quality assurance engineer.

They will spend much time creating test plans and cases to run against the models. Often these plans will involve automated scripts. The AI tester will also manage the bug tracking process.

There are a variety of testing tools available, such as Selenium, Eggplant, WebTest, Mercury QTP, or Watir. They can help streamline the process but also facilitate the different approaches, such as the following:

- *Black box testing*: In this situation, the tester does not know the structure of the application. Rather, they put together a variety of inputs and test cases for the evaluation.

- *White box testing*: In this case, the tester has access to the underlying code and data sets. This allows for helping to detect security exposures or poor structure.
- *Grey box testing*: This is a blend of the black and white box testing approaches.

Then which one to use? Given the complexities of AI, the grey box testing approach may be the best. This will provide for comprehensive testing.

An AI tester can also be helpful in evaluating the user interface, although it is advisable to have many people in the organization to also test this and to ask for feedback.

As for the background for an AI tester, they will usually have a BS in computer science or a related engineering field. They will also have experience with a language like Python and shell scripting languages.

The AI Solution Architect

Even startups have complicated IT environments and legacy systems. This can present challenges when implementing AI within an organization.

To help with this, there is the AI solution architect. The main role of this person is to integrate the AI within the IT infrastructure. If done right, this can help accelerate the implementation process but also help with the scaling.

The AI solution architect will need a familiarity with the core concepts of machine learning as well as the frameworks. They should also have experience with CRMs, ERPs, and middleware systems. Furthermore, a background with GPUs and modern architectures like Kubernetes is a plus.

And as for education, they will usually have a graduate degree in a field like computer science, machine learning, statistics, or mathematics.

Machine Learning Engineer

A machine learning engineer is involved in the process of machine learning operations or MLOps. The main focus is turning the AI into a product and maintaining the effectiveness of the models. Consider that, over time, models usually drift. So a machine learning engineer will monitor this and make the necessary adjustments.

This person will have experience in quickly prototyping new models but also placing them in production, such as in Python or R. They should have a background with handling big data pipelines as well.

Regarding education, a machine learning engineer will usually have a graduate degree in computer science, statistics, engineering, or mathematics.

Recruiting

Elon Musk is one of the world's best recruiters of tech talent. Tesla has over 48,000 employees and many of them have deep technical backgrounds.

AI is also a key to the company's growth strategy. All the cars have the hardware capabilities for autonomous driving, such as surround cameras, ultrasonic sensors, forward-facing radar, and AI chips. These systems are updated via the cloud.

In terms of recruiting, Musk has unconventional approaches. In fact, this is a necessity because of the talent shortages across the world.

Musk will even go to Twitter for his recruiting efforts. Here are some examples:

- "At Tesla, using AI to solve self-driving isn't just icing on the cake, it the cake" - @lexfridman
- Join AI at Tesla! It reports directly to me & we meet/email/text almost every day. My actions, not just words, show how critically I view (benign) AI. https://www.tesla.com/autopilotAI
- Tesla will hold a super fun AI party/hackathon at my house with the Tesla AI/autopilot team in about four weeks. Invitations going out soon.

All this certainly points to how Musk sees AI as a strategic priority. He has also been effective in creating a vision for his company that is quite compelling to engineers and researchers. While compensation is certainly important, people also want to work for a company that has a purpose.

In job listings for AI, here's what Tesla notes: "As a software engineer on the Autopilot Computer Vision and AI team, you will contribute to one of the most advanced and widely-deployed computer vision stacks in the world. Along with top researchers from academia and some of the most experienced autonomous vehicle engineers in the industry, you will marry cutting-edge deep learning algorithms with robust, real-time software, and deliver safety-critical features to hundreds of thousands of customers. You will develop and support a host of different projects, driven first-and-foremost by our mission to deploy the safest and most effective product in the market."[8]

[8] https://bit.ly/3fMxg95

The ad goes on to say that a person needs a "MS or PhD in Computer Science, Physics, Electrical Engineering or proof of exceptional skills in related fields, with practical software engineering experience." In other words, you do not have to have a formal education in the area, so long as you have "practical" experience.

This is by design. In an interview with a German auto magazine, Musk said: "There's no need even to have a college degree at all, or even high school."[9]

He also looks for talent in disciplines like physics, social sciences, and mathematics. Such people not only have quantitative skills but also bring diverse ideas to the team.

Again, when it comes to recruiting AI talent, you need to be very creative. Granted, this is not to say traditional approaches should be shunned. Using technical recruiters is definitely good to do, as long as they have strong AI skills.

You should also leverage online resources like job sites, forums, online communities (like Hackernews.com and AngelList), and LinkedIn. Recruiting AI talent usually takes a lot of courting and patience.

Reskilling

Reskilling or upskilling is when you train existing employees to transition to becoming part of the AI technical team. This is becoming easier as there are a myriad of online courses such as from Coursera, Udacity, open.ai, and deeplearning.ai as well as bootcamps. More traditional universities are also offering executive programs for AI and data science.

Note In February 2020, Udacity launched an online learning executive program called AI For Business Leaders, in partnership with BMW. It has quickly become one of the most popular courses on the platform. The goal is to teach technical AI concepts to business leaders quickly. A typical course takes about four to eight weeks to complete.

An example of a successful reskilling program is from Bloomberg, which is the dominant data platform for the financial industry. The company has invested significantly in AI and machine learning to enhance its products and services. There are more than 150 data scientists, researchers, and engineers on the Data Science Team, which is part of the CTO's office and AI Group. Some of

[9] www.cnbc.com/2020/02/03/elon-musk-is-recruiting-for-tesla-education-is-irrelevant.html

the innovations have included: using models to predict the impact of news events on stock prices, sentiment analysis based on social media, product recommendation systems, and anomaly detection for large time series data sets.

Back in 2017, Bloomberg realized that it could not rely solely on traditional recruiting for its AI efforts. So the company launched its own educational program. It covers the core concepts of AI and how the technology applies to Bloomberg systems. Over time, the company has developed nine courses.

Interestingly enough, during the summer of 2018, Bloomberg made public 30 training videos. You can find them here: https://bloomberg.github.io/foml/#lectures. Some of the topics include an introduction to statistical theory, stochastic gradient descent, black box machine learning, loss functions for regression and classification, and support vector machines.

Why did Bloomberg do this? First of all, it was a way to provide a resource to its own client base. Next, the videos are a statement of Bloomberg's commitment to its own investment in learning, which has been useful in recruiting talent.

True, not all companies have the resources to do this. But then again, there should still be some investment in education, both for reskilling and existing AI employees. This is vital for success.

However, education is still not enough. There needs to be initiatives to provide employees ways to implement their new skills. In the meantime, an organization needs to give incentives for mentoring.

Oh, and the training should not just be for technical talent. It's a good idea to provide beginner level training for all employees. This will go a long way in creating a data-driven culture.

Regarding the process of reskilling, one effective approach is to first identify those employees who have backgrounds with similarities to AI. For example, a conventional full stack or backend software engineer can learn to become a data engineer on the job and with modest training. The same goes for a business analyst.

■ **Note** A survey from Deloitte points out that when all employees have been trained on analytics, about 88% exceeded their objectives. This is compared to 61% that did not.[10]

---

[10] www2.deloitte.com/us/en/insights/topics/analytics/insight-driven-organization.html

Team Size

Even if you have a large budget, it is not realistic to assemble a full-blown AI team quickly. The market is just too competitive. So in the early phases, you need to focus on some of the key players. For example, it is usually advisable to spend top dollar on a qualified data scientist. This person will not only help with crafting effective models but could provide assistance with reskilling and recruiting. After all, he or she likely will have a strong network.

It's often a good idea to keep the team small anyway. This should allow for more agility. This is consistent with the so-called two-pizza rule, which is what Amazon's CEO and founder Jeff Bezos instituted at his company in the early days. This is where an internal team would have no more people than two pizzas can feed!

Salary Size

As AI has grown in popularity, the salaries have likewise seen growth. The main reason is that there has been an ongoing shortage of people with the necessary technical skills and experience in the space.

The job site Indeed put together a survey on AI salaries that was based on over 78,000 employees (as of the summer of 2020).[11] Here's what it found (salaries per year):

- Machine Learning Engineer - $140,686
- Software Engineer - $109,321
- Research Scientist - $98,788
- Data Scientist - $122,788
- Scientist - $101,628
- Senior Software Engineer - $132,367
- Research Engineer - $94,910
- Software Engineer Intern - $76,579

But this likely understates greatly the compensation. Many AI employees work for technology companies that provide stock options and other forms of equity incentives. The result is that the compensation could easily be two to three times as much as the salary.

[11] www.indeed.com/salaries/artificial-intelligence-Salaries

In the Indeed survey, there was a 5.8% increase in 2019 for average AI salaries. This is compared to an average of 2.9% for all categories. The survey also reported a 15% increase in searches for AI candidates.

Chief AI Officer

Generally, it's the Chief Technology Officer (CTO), the Chief Information Officer (CIO), or the Chief Information Security Officer (CISO) who is the high-level executive who oversees AI initiatives. But as the technology becomes more strategic, some companies have created a new role, such as the Chief AI Officer (CAIO) or the Chief Algorithms Officer (CAO).

This role will often involve someone with a PhD in data science or machine learning. That is, they can read the academic papers and keep up with the latest developments in AI. But this background in academia can also be helpful in recruiting. What's more, a CAIO should have a business background, such as with understanding how to implement complex systems successfully.

One of the most notable CAIO's is Yann LeCun (he has over 203,000 followers on Twitter), who is an executive at Facebook (the company refers to his position as VP and Chief AI Scientist). In 1987, he received his PhD in Computer Science from the Universite Pierre et Marie Curie in Paris, France, and then worked at AT&T's Bell Labs. He developed optical character recognition technologies for interpreting checks and created DjVu, which allowed for image compression for websites.

In 2003, he became a professor at NYU and founded the NYU Center for Data Science. During his academic career, he published over 180 technical papers on topics like AI, machine learning, computer vision, and computational neuroscience. When he joined Facebook in 2013, he maintained his professor role on a part-time basis.[12]

In 2019, he shared the Turing Award with Yoshua Bengio and Geoffrey Hinton. They are often called the Godfathers of AI.[13]

Conclusion

As you've seen in this chapter, there are many different roles for an AI team. This is why it is incredibly important for there to be strong collaboration. And because of the inherent complexities of AI, it's essential that the process be iterative. If the pace gets rushed, then failure can be the result.

[12] http://yann.lecun.com/ex/bio.html

[13] www.theverge.com/2019/3/27/18280665/ai-godfathers-turing-award-2018-yoshua-bengio-geoffrey-hinton-yann-lecun

In the next chapter, we'll take a look at the data process.

Key Takeaways

- An AI team involves many roles. Here's a look at the main ones:
    - *Executive sponsor*: This is a person from the executive suite who champions the AI project. This role is critical since there often needs to be substantial resources, especially for hiring technical talent and purchasing software. But the executive sponsor does not have to be the CIO or CTO. Rather, they should be someone who has the ability to galvanize the organization and has the right budget authority.
    - *Project manager/business owner*: This is the person who takes on the day-to-day management of the project. The focus will be on helping to set and hit the goals. The project manager can be recruited internally, but they should have some technical skillsets.
    - *SME (subject matter expert)*: As the name implies, this is a person who has a strong understanding of a particular domain or area of the company, such as its processes or customer base. Usually the SME is recruited internally.
    - *Business analyst*: This person uses data analysis to help improve the operations of a company. They create models, reports, and dashboards with BI tools. Regarding an AI project, a business analyst can understand business processes, evaluate data, and integrate systems.
    - *Data engineer*: This person will find the right data sources and structure them properly, such as with fixing duplications and missing items.
    - *Data scientist (or AI engineer)*: This is the person who creates the AI models. They certainly have a strong background in statistics, machine learning, and deep learning but also some background in business.
    - *Designer*: A designer can create better interfaces, which should lead to improved adoption of the AI. A designer is usually a part-time position.

  - *AI tester*: This person focuses on trying to find problems with AI models.
  - *AI solution architect*: This person handles the integration of the AI with the IT infrastructure. They can also set up the systems to help scale the technology.
  - *Machine learning engineer*: This person is responsible for turning the AI into a product. The role also involves tracking and evaluating the models.
- Recruiting AI talent is far from easy. It's something that is ongoing. Because of the shortage for tech talent, it is usually necessary to find people who may not necessarily have direct AI experience. For example, they may have advanced degrees in physics and mathematics. Thus, the process should be easier for training them in machine learning and deep learning.
- Reskilling and upskilling are also very effective. They can be a way to build a strong AI team. This approach is also getting much easier because of the proliferation of online courses, bootcamps, and executive degree programs.
- Even if you have a big budget, this does not mean you should be aggressive in hiring AI talent. This can lead to major problems. The process should not be rushed. Actually, small teams are often better.
- As AI has become more important, there has emerged a new role at the executive level: the Chief AI Officer or CAIO. This person usually has a strong background in AI, such as a PhD. They should be able to explain complex topics to the rest of the C-suite.