
The Six Roles You Need on Your AI Team



By [Jenny Watkins](#) (INSEAD GEMBA '14), CEO, Tadpole Consulting

**Can you build new technological capacities during a global crisis?
Yes - if you make the right decisions regarding your people.**

Let's say you were already sold on AI: You've gotten yourself and your executive team trained on the basics of AI; you've studied use cases; prioritised the business pain points that could align with AI solutions; and allocated funding to launch a few pilots.

But the COVID-19 crisis has thrown your well-laid plans for a loop and now you're in firefighting mode. The saying "Don't let a crisis go to waste" rings especially true as companies put innovation opportunities on the back burner. Today, more than ever, customers have heightened expectations and companies must act fast to deliver. This is the time to reinvent offerings and bring customers value alternatives. Technology will speed up economic recovery, so companies must strengthen their data and digital capabilities during this period - or suffer in the long term.

So, if you're ready to start assembling your cast of AI characters, who do you need? What skillsets should you look for to build your enterprise's AI efforts?

Here are the six roles you need on your AI team:

1. The strategist

Contrary to popular belief, your first hire should not be a data scientist. It should be someone who is even harder to find than a data scientist: a data strategist. It is someone who understands how the business operates and how technology and data science work together. This person acts as the bridge between business lines and technology to drive pilot projects that deliver measurable ROI.

2. The data engineer

Remember the frequently cited adage, "Data is the new oil"? Well, your data engineer is your miner. This person is responsible for collecting data from a variety of sources, then preparing and transforming that input. In [our COMA framework](#), the engineer works on compiling, organising and (sometimes) manipulating data, which can be time-consuming. Thankfully, today, start-ups like Trifacta and Forge.AI offer solutions that clean, label and move data in a more efficient way.

When hiring a data engineer, look for strong coding and engineering skills like SQL, Python, C++ and Java.

3. The data modeller

Often sporting strong maths and statistics skills, the data modeller looks for data patterns to predict outcomes (remember: AI, at its heart, is about patterns and predictions). This person also builds and trains models to determine events like propensity to buy, likelihood to churn, etc.

Tech giants often swoop up PhDs in maths and statistics, so start-ups are offering off-the-shelf models as an alternative. Before going this route, see our advice below.

4. The data in production person

This person converts prototype code to production, setting up a cloud environment to deploy the models. Other tasks involve managing version

control, improving response times and building APIs. Look for cloud expertise and software engineering skills.

5. The infrastructure and scale builder

Think of this role as providing the wiring and plumbing for a house. This person builds databases to store data and facilitate access, as well as maintain security and privacy. Strong software engineering skills and proficiency in cloud technologies are a must.

6. The data analyst/visualiser

This role is tasked with evaluating the model's performance and business value post-production. The person can build out A/B or multi-variate testing, to measure the impact of different variables. The data visualiser creates dashboards and translates data into actionable business insights. Look for expertise in software like R or Tableau, although start-ups like [Qlik](#) and [ThoughtSpot](#) can enable layperson visualisation across the enterprise.

Common mistakes

A frequent mistake is to focus more resources on finding data modellers than data engineers, production persons and infrastructure builders. Since the latter's tasks happen behind the scenes, their efforts tend to not be as visible or understood by executives than that of modellers. Keep in mind that the engineer to scientist ratio can often be 3:1 or 5:1 depending on the architectural data complexity of a company and initiative.

While tasks may intersect and certain profiles can perform multiple roles (depending on skillsets and available tools), be careful with expecting data modellers to do data engineering roles. This can lead to technical "debt": Data modellers may not have enough experience writing applications that scale in production, which is exactly what engineers and production builders are trained to do.

The reverse can also happen when data engineers handle data modelling. Here, statisticians are specifically trained on model accuracy and relevancy (i.e. picking the right model for the job), bias elimination and versioning. These aspects are less of a core competency for engineers.

Ultimately, a full-stack engineer or scientist can develop end-to-end expertise. Indeed, this is very valuable. However, remember that for large

organisations with siloed data and production, one person is not enough. Teams still need to be designed based on the company's business objectives, AI initiatives and data architecture.

So, where do you find these people?

Option 1: The borrow route

If there's a hiring freeze due to COVID-19 or your enterprise can't afford the salaries commanded by top talent, you can outsource the job to consultants or firms. Companies like [Experfy](#) offer data teams for hire. Start-ups selling models abound as well, but when hiring outside vendors that sell their models, check for:

1. **Industry knowledge.** This is very important. Data nuances are critical and the lexicon used needs to be relevant. Without proper industry knowledge, performance can be questionable. You need to probe under the hood because an off-the-shelf model doesn't perform the same way in production. Industry experience is needed to know the difference
2. **Validation frequency.** Data get stale and business environment changes. Ask how often a company re-validates its models. Make sure it's at least annually.
3. **Data compliance.** Data collection privacy, ethics and regulation need to be adhered to or your enterprise could be on the hook for noncompliance.

Option 2: The build route

This is the ideal route if your goal is to transform your organisation into a data-first enterprise. It takes more time and resources in the short term, but it allows you to reap dividends in the long run. To build, you should:

1. **Upskill your current staff.** If you don't have the right skillset, consider hiring an outside organisation to train your existing teams or send your talent pool to data camp. For example, software engineers can be upskilled to become machine-learning engineers so that tasks are automated. Ensure you pick a training partner well versed in the right training skills. Online MOOCs like [Coursera](#) and [Udacity](#) can also be an option. Just be careful to marry theory with practical on-the-job training.

2. **Invest in the right tools.** AI tools have become more and more sophisticated. From visualisers like [Domo](#) to AI platforms like [H2O.AI](#) and [Databricks](#), these tools are lowering the investment costs in infrastructure and talent. They are also enabling “citizen data scientists”, i.e. people who don’t know how to code but who can bridge data with insights and action.
3. **Form your “AI Marines”.** With the right mix of talent and tools, you can create a centre of excellence. This team can rotate in and out of business units to identify quick wins and cross-train various business lines. This can allow you to infuse data science and AI into your organisation in a centralised way. In time, as each business unit builds its competencies, the AI Marines can be absorbed into the broader enterprise and need not exist solely as a stand-alone sniper team.

Your data strategist will be key in determining the right strategy and mix of tactics above to properly build out your internal capabilities, so find this person first.

Then what?

Don’t forget your supporting cast: Your executive committee, business unit sponsors and governance team all have a critical role to play. Your AI team can roll out initiatives but it’s up to the supporting cast to ensure they’re properly framed, structured and nested in the right business context, and then enabled to gain the most ROI. Remember, you can’t buy nor simply hire AI – the hard work happens in-house.

Also, if your enterprise is serious about data-first transformation, consider hiring a coach to facilitate the transition to new decision-making and working habits.

[Jenny Watkins](#) (INSEAD GEMBA '14) is the CEO of Tadpole Consulting.

Don’t miss our latest content. Download the free [INSEAD Knowledge app](#) today.

Follow INSEAD Knowledge on [Twitter](#) and [Facebook](#).

Find article at

<https://knowledge.insead.edu/entrepreneurship/six-roles-you-need-your-ai-team>

About the author(s)

Jenny Watkins Jenny Watkins (INSEAD GEMBA '14) is CEO of Tadpole Consulting.