



Leading a Successful AI Team

How your organization can realize
the full potential of AI with the right
team in place

Table of contents

Introduction	3
What Is an AI Team?	4
The AI Team's Skill Set	5
An Exemplary AI Team	7
The AI Team Leader	8
Software Engineering	9
Data Science	10
Spotlight: AI Engineering	11
Subject matter experts	12
What AI Teams Need to Be Successful	13
deepset: Enabling AI Teams since 2018	14



Introduction

LLMs can help companies build better AI products and internal tools. But many teams are struggling to keep up with the growth of AI and the changes it requires. The biggest challenge for organizations is successfully building and deploying AI-driven applications.

There is a clear gap between data science, software development, and the C-suite when it comes to building with AI. This is not a technology problem. It's a people problem.

As more companies adopt AI technology, it's clear that a cross-functional, product-focused AI team is essential. But there's often a lack of clarity about the characteristics and leadership style that AI teams need to be successful.

This guide will help you build a successful team to deliver products with AI. Learn what skills your AI team needs and how to enable each team member to work together and excel.

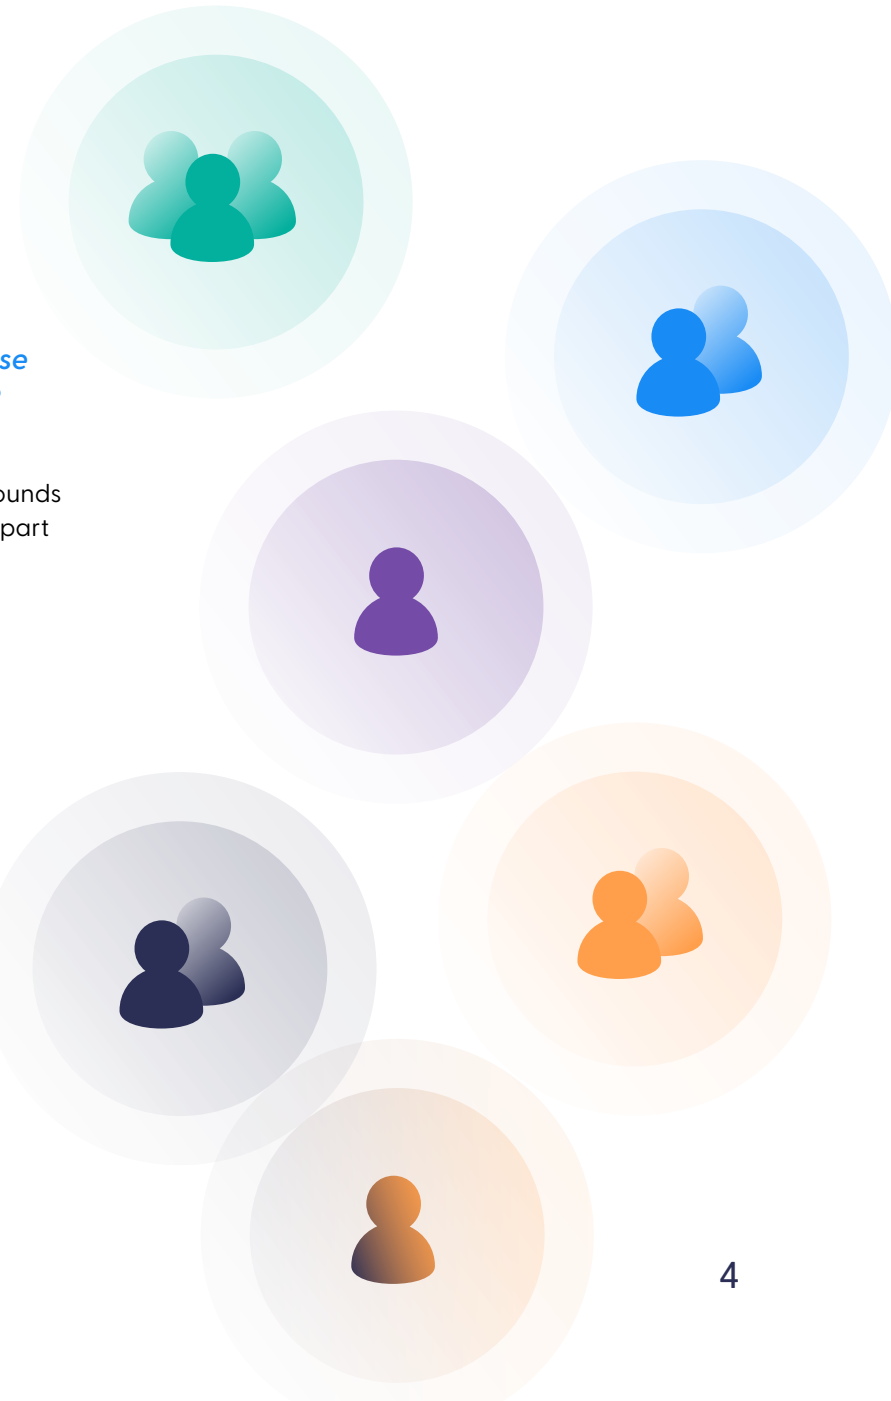
What Is an AI Team?

An AI team is a group of people who work together to create a product or tool that uses AI technology. The team includes people with different skills and backgrounds, such as developers, data scientists, AI or ML engineers, and non-technical experts.

The AI team is led by a pragmatic and visionary AI team leader. In a news organization, this might be the editor-in-chief. In a financial institution, it might be the head of business applications. In another organization, it might be the head of AI. What matters is their unique skill set, which we'll talk about more later. This person is the best predictor of success.

The AI team is characterized by a diverse set of skills – plus the methods to utilize them in the most effective manner.

It's easy to see how people with diverse backgrounds and skills can complement each other. The hard part is creating an environment where everyone can contribute their best.



The AI Team's Skill Set

AI presents new challenges, including managing the entire AI tech stack. Building with AI also requires creating an organized process for moving from prototyping and testing to production, deployment, and monitoring, while considering the needs and desires of users and stakeholders. Let's look more closely at what a highly productive AI team needs to succeed.

Skill No. 1

Use Case Orientation

Your product should solve a problem for your users. If it doesn't, no one will use it or buy it. **Use case orientation means understanding your users' behaviors and problems and learning more when needed.** It also requires a realistic assessment of generative AI capabilities and limitations, the domain expertise your project may require, and a willingness to think outside the box.

Skill No. 2

Iterative Product Development

How do you know if your AI product is on the right track? Frequently check in with reality, build and test prototypes, and incorporate feedback. Prototypes help you test your product in a realistic way. You can then measure how users interact with it. **You won't know if your product is right until you test it with stakeholders and customers.** That's why most software products today go through hundreds of iterations before they're released.

Skill No. 3

LLM Expertise

Building with large language models is about using them for your purposes. To do this, you and your team need to understand how an LLM works and be able to monitor and evaluate it. **An LLM acts as a component in a modular pipeline where it can interact with your data and other language models.** Understand the differences between language models and how to combine them to create more powerful systems:

- **Not all language models are generative.** Embedding models match documents to a query based on meaning. They are often used as retrievers in retrieval augmented generation (RAG).
- **Generative LLMs can be open source or commercial.** Think about your needs for cost, speed, privacy, performance, and effort when choosing between open source and commercial models. Your choices may change over time as new models emerge – but with the right process, you can easily switch models.

Skill No. 4

Subject Matter Expertise

Your products may be for users who know a lot about their fields, like lawyers, doctors, or accountants. Involve professionals in product development from start to finish to find solutions to real problems. This ensures the product meets the target audience's needs and expectations.

Skill No. 5

UI/UX

User experience is important for AI projects. When developing your use case, think about how your AI app will fit into the customer or internal team's workflow. How will they use the product? Will it be accessed through a browser, a mobile app, or a business app? Consider how easy it is to integrate with existing systems and the learning curve for users.

Skill No. 6

Full-Stack AI Development

Over the past few decades, software development has changed a lot. Enterprise IT has gone from on-premises solutions to cloud-based services. When building with new technologies like LLMs, developers need to understand database management, cloud computing, and cybersecurity.

- **Database management** is the storage and management of data across different locations. Developers must ensure that data is handled efficiently and securely, whether on-premises or in the cloud. In the context of LLMs, they often work with vector databases, a new DB paradigm optimized for storing and processing high-dimensional text embeddings.
- **Cloud computing** provides scalable and cost-effective solutions to host, manage, and deploy applications using cloud service providers. Understanding modern cloud infrastructure, including containerization, microservices architectures, and cloud-native services, is key.
- **Cybersecurity skills** protect software products from threats and vulnerabilities. This includes encrypting data, securing network communications, controlling access, and staying up to date on security protocols and practices. To keep your users' data safe, make security part of your process.

AI Teams: Skills over Roles

Knowing the value of these skills helps AI team leaders find the best people for their AI team. For example, you may have an opening for an AI engineer – but whether you fill it with a data scientist, an ML engineer, or even an NLP researcher with a knack for building and deploying practical applications is not the most important factor, and will vary from team to team. What matters is that your team members are product-minded, eager to learn, and work well with others.

An Exemplary AI Team

AI teams can be of varying sizes, but they usually include individuals from two technical camps: software development and machine learning/AI/NLP/data science. The first group is adept at developing robust and scalable architectures on which the application runs. The second group is responsible for working with models and data. They are also supported by domain experts and business leaders.

Some of the roles in an AI team may be combined in one person, for example, the product lead may also have deep domain knowledge.

Complementary Roles

The core AI team roles are often complemented by other, more project-specific roles. For example, if your product is consumer-facing, you may want to add front-end developers and UI/UX designers to your team. If your use case requires complex data pipelines, you'll want to bring a data engineer on board. And so on.

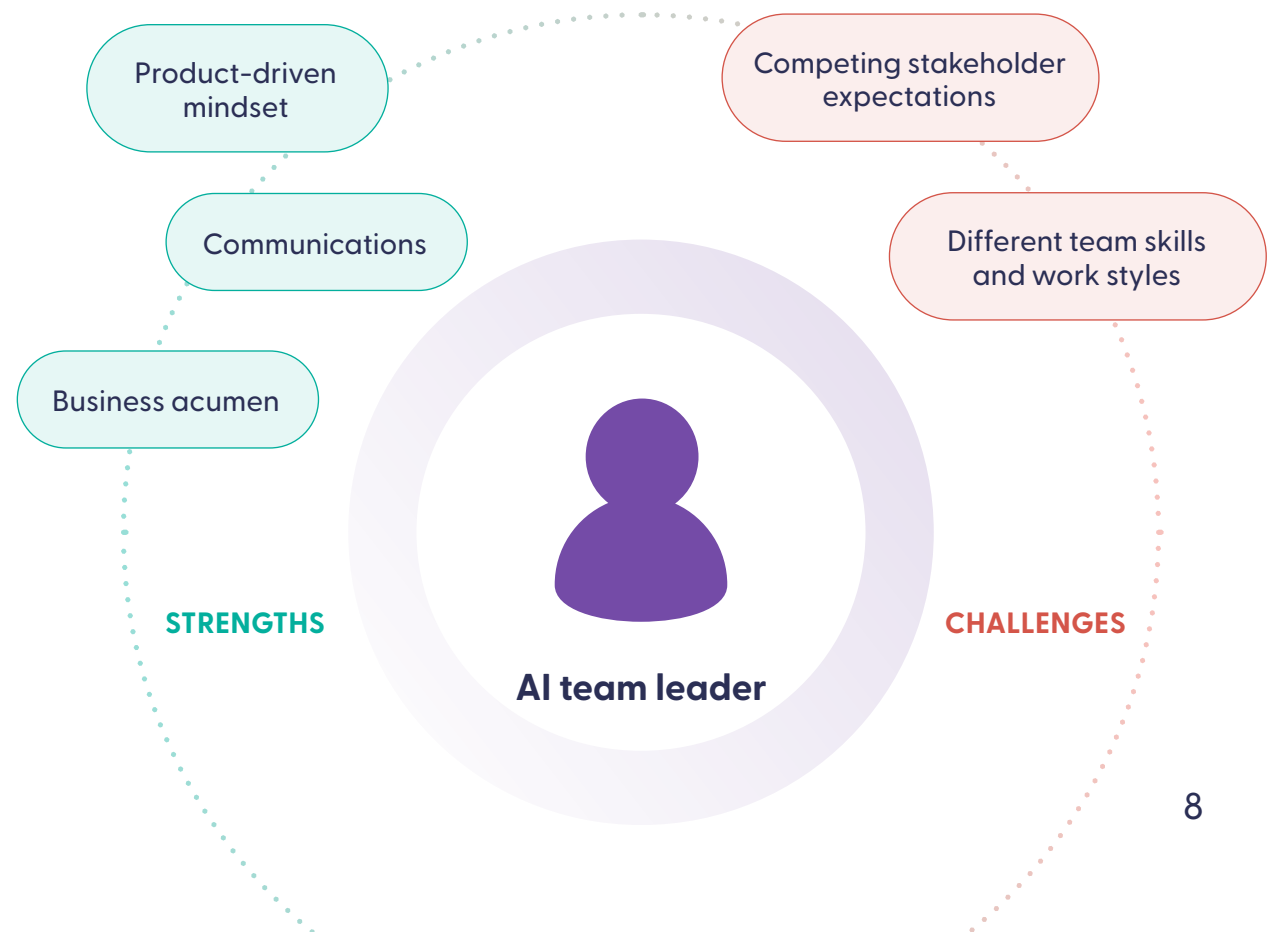


The AI Team Leader

The AI team leader is essential to the project's success. They manage the team and process, with rapid iteration and user-centricity as the project's guiding principles. These people tend to be hands-on, creative, and good communicators.

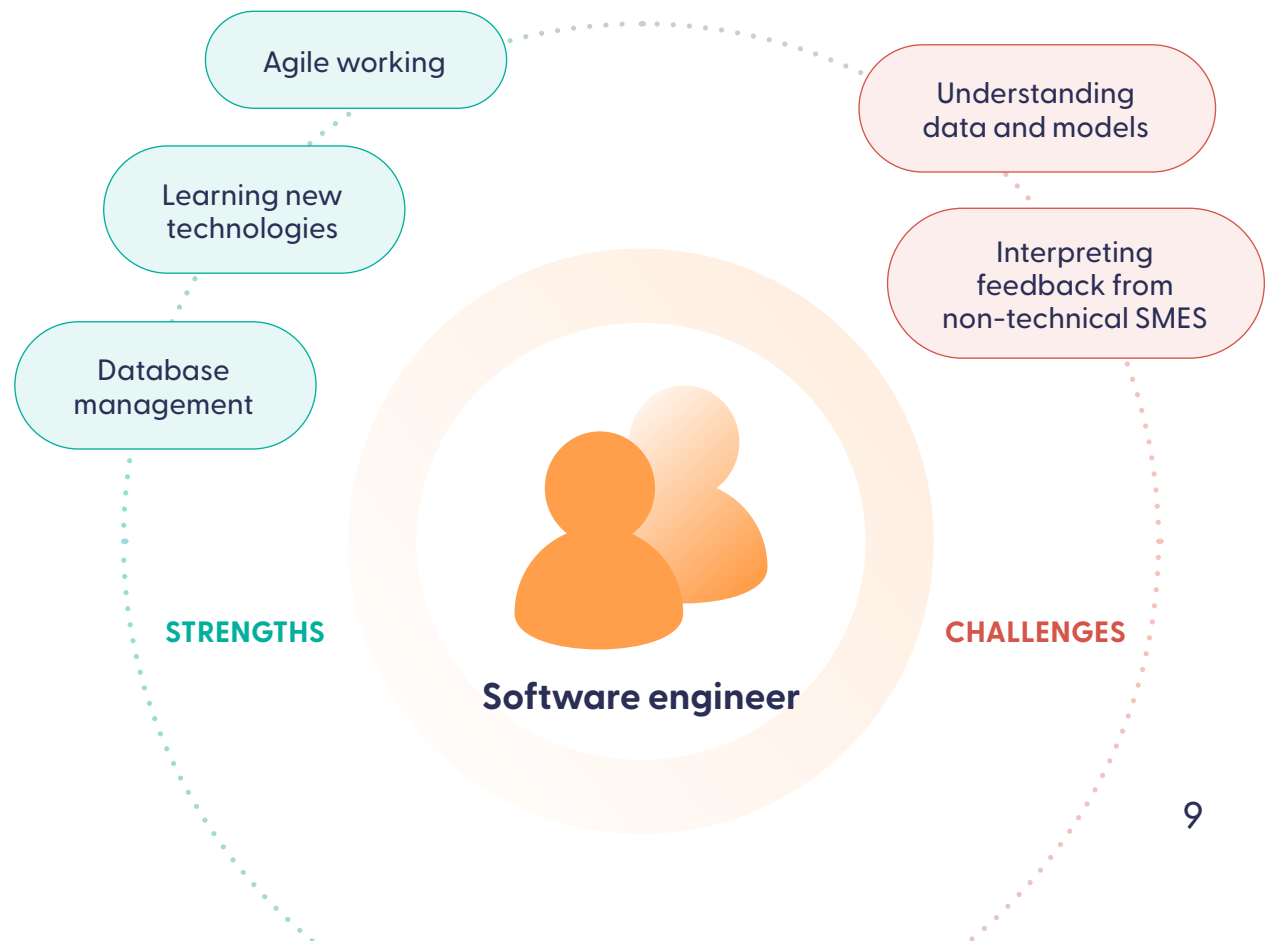
In practice, the role is relatively new and the specific skill set is rare, so it is often poorly defined by individual organizations. These organizations must recognize the importance of AI leaders in AI products. **The AI team leader sets the vision, explains how the product will address a specific use case, and presents a plan for bringing the application to market.** They understand what users need and how AI can help.

AI team leaders help different roles work together. They also connect the core team with key people, such as organizational leaders, subject matter experts, and end users. **They must communicate with the wider organization and users to ensure their AI application solves real-life challenges and delivers a return on investment.**



Software Engineering

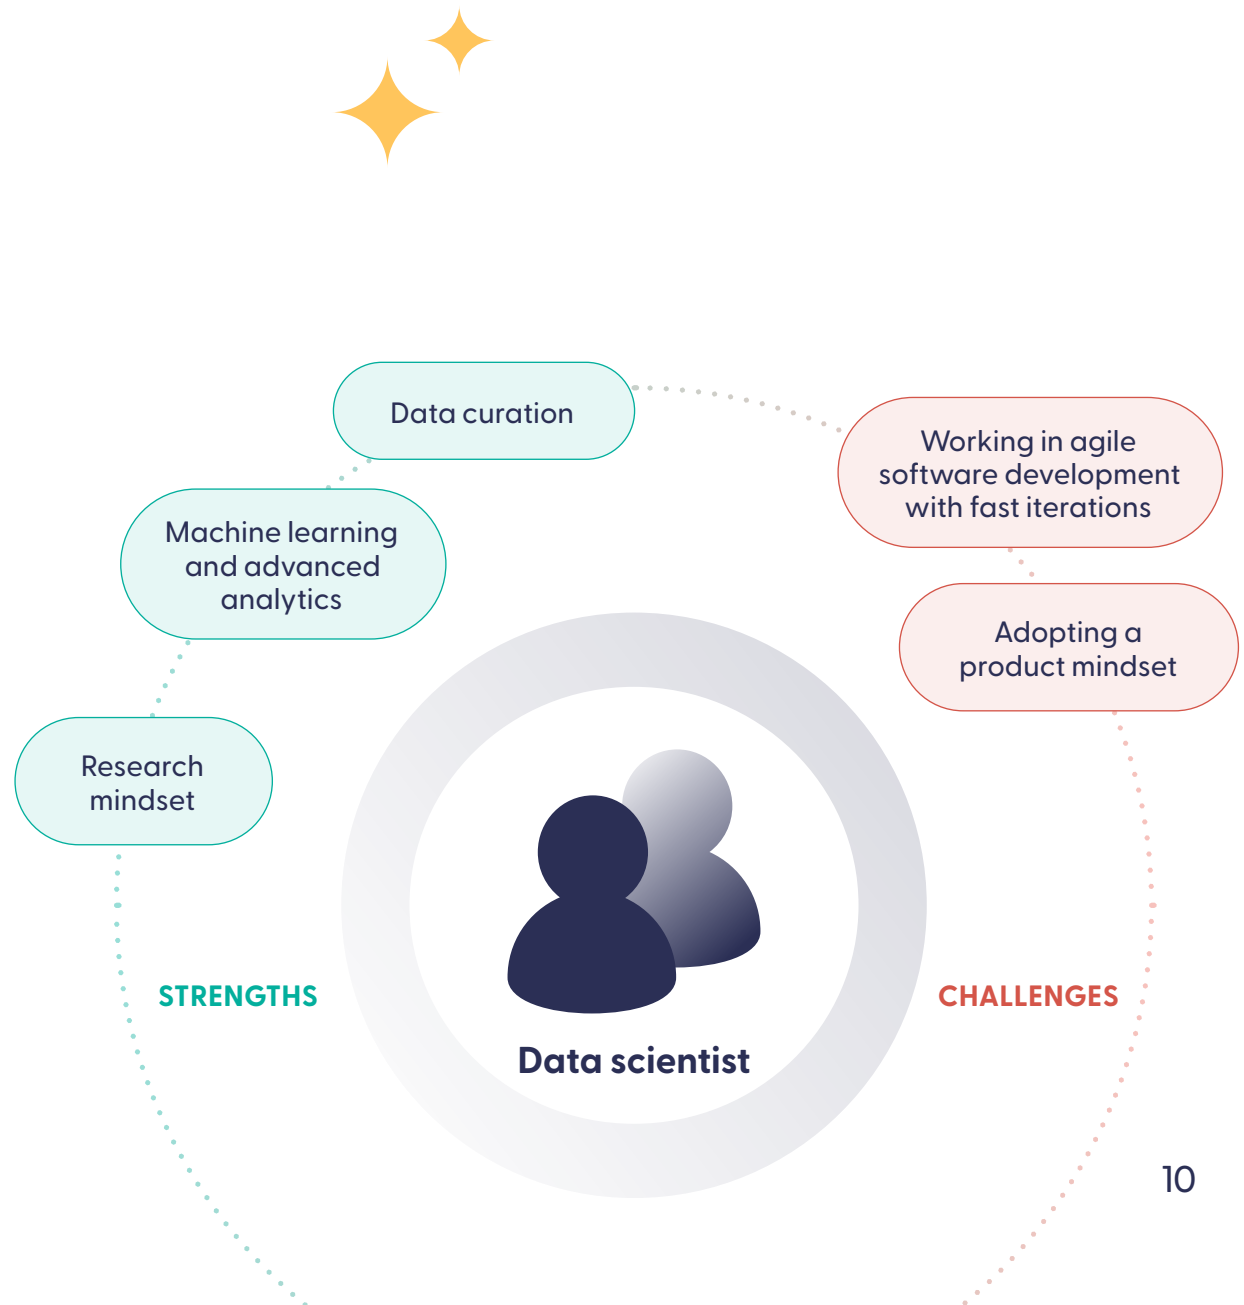
Software engineers are adept at rapidly developing, testing, and improving applications. They focus on the architectural aspects essential to AI applications, as well as managing deployments and ongoing operations. Data and modeling are beyond their expertise, so they rely on data science for data curation and delivery, system setup and monitoring, as well as actionable insights about the application in production.



Data Science

Data scientists are experts at using data to build machine learning models. As models become more accessible, the role of data science is changing. **Data scientists can now embrace generative AI and use it to expand their skill set.** After all, they understand better than most how data can be used to modify and guide machine learning models, and they have a solid theoretical foundation that can help them cut through the AI hype and predict trends.

If data scientists understand the product their team is building and how to support it, and work closely with software engineers and subject matter experts, they'll be a valuable asset to any AI team. They may even be on their way to becoming AI engineers.





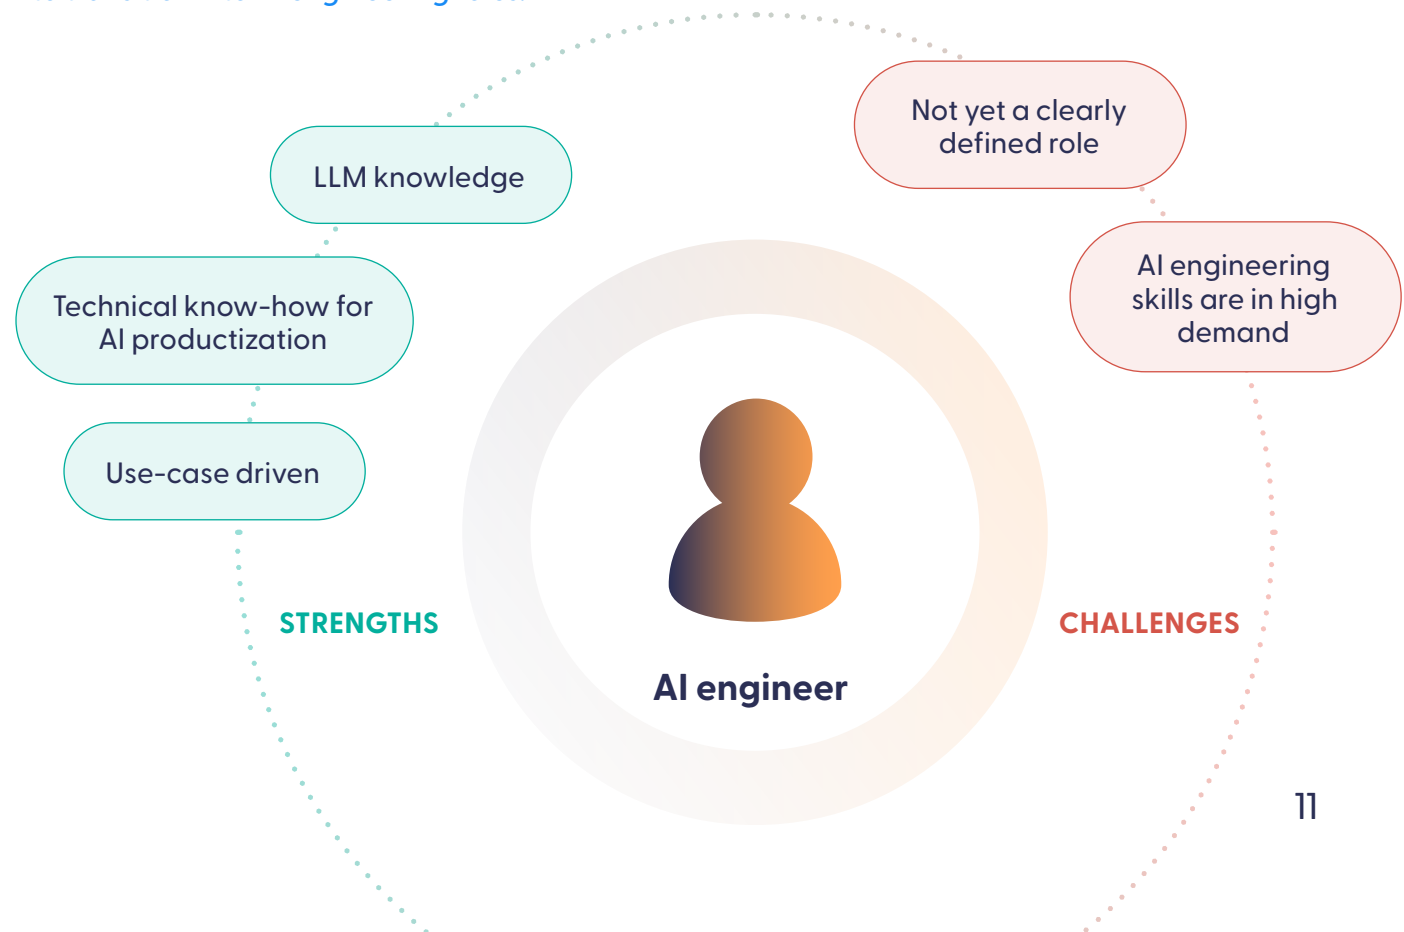
Spotlight: AI Engineering

AI engineers are pragmatic builders with a product mindset. **Rather than focusing solely on data curation for model development, they think about how AI can enrich products and what data, if any, they need to do so.** Their responsibilities include:

- Building scalable and flexible composite AI systems tailored to the use case.
- Keeping up with the rapid evolution of LLMs.
- Optimizing the performance of the LLM through prompt engineering and other tools.
- Bridging software engineering and data science, combining knowledge from both fields to create robust AI solutions.
- Implementing continuous delivery and application monitoring, using software engineering practices to maintain AI applications.

AI engineers take a holistic view of building AI applications. Going far beyond the traditional data scientist's model- and data-centric domain, they build APIs to integrate data products, create scalable infrastructure, and ensure the reproducibility and efficiency of AI solutions. These professionals are constantly learning through collaboration and are fascinated by generative AI, identifying new data-driven use cases and becoming adept prompt engineers.

As companies continue to invest in AI to improve operations and develop new products, the demand for skilled AI engineers is expected to grow. The current landscape of LLM development offers opportunities for AI enthusiasts from diverse backgrounds to transition into AI engineering roles.



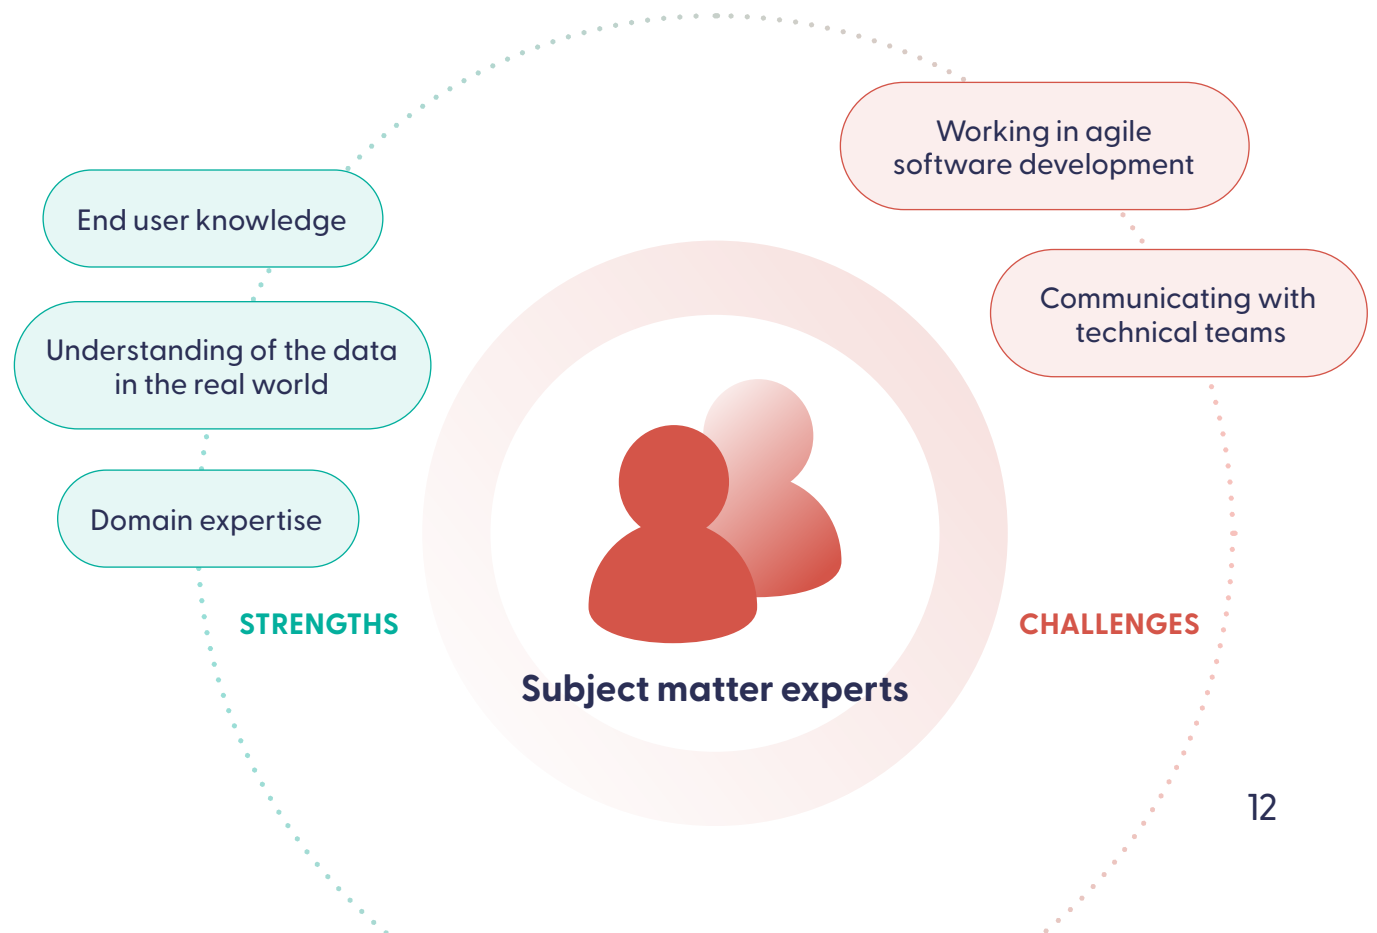
Subject matter experts

SMEs are experts in a specific area. They have gained their knowledge through years of experience. For a legal product, an SME might be an experienced lawyer who understands the legal system and can provide advice on data and evaluation.

SMEs are involved at different stages of the project to help develop the product:

- SMEs help the product lead identify a use case during the ideation phase.
- They can add information to data sets as needed.
- They test the system and give feedback on how useful and relevant the answers are.

Depending on the domain, **SMEs don't typically have a technical background, but they're ideal partners for use case development and project-specific jargon and pain points.** Their expertise lets them use the product realistically and provide better feedback. The team leader then helps translate that feedback into technical information for AI and backend engineers.



What AI Teams Need to Be Successful

How can diverse teams with different skills and backgrounds work together effectively? How can team leaders help their people reach their full potential and identify problems early on? There is a tool that can help AI teams build a valuable and secure product in the fastest and most straightforward manner.

Enterprise AI production platforms

If you want to build enterprise-grade AI into customer-facing products or internal workflows, you need to:


- Design and test AI pipelines.
- Tune and optimize models.
- Build continuous delivery infrastructure.
- Design and evaluate experiments.
- Deploy and monitor your application in production.

Unified AI production platforms allow teams to automate these repetitive and distracting tasks, streamlining the "heavy lifting" of applied AI.

This means your team can focus on what makes your product stand out.

AI development platforms enable AI team members to collaborate effectively and stay aligned on the goal of creating a high-quality, high-adoption application. By organizing every step of the process through APIs, no-code and low-code workflows, and graphical user interfaces, **AI production platforms make it easier for everyone to contribute.**

A good platform should let you design, test, and deploy application architectures. It should also allow you to monitor and refine the AI product in production. This includes collecting user feedback for beta testing and detecting and eliminating hallucinations. It should also allow you to manage changes and track your AI pipelines so that you can stay on top of the AI evolution.



deepset: Powering AI Team Collaboration

deepset's AI production platform, [deepset Cloud](#), offers a 360-degree approach to AI development. With clean interfaces that cover the entire AI development cycle, **deepset Cloud is the collaborative platform for AI teams to quickly build, test, and deploy their products.**

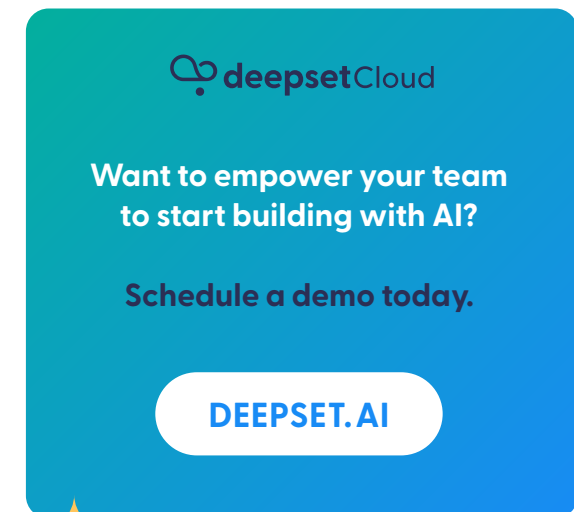
With pre-built components and LLM application setups, users can create a fully functional prototype in minutes and test and launch in days or weeks, instead of months or quarters. *deepset Cloud* is not tied to any one model provider, allowing teams to experiment with different models to find the best setup for their needs. It has "ready-made" templates for the most common LLM use cases – from RAG to Conversational Business Intelligence – while being infinitely customizable thanks to Compound AI.


“With deepset Cloud, we were able to focus almost immediately on rapid iteration and user feedback to refine our solution.” – The AI Team at ZEIT ONLINE

deepset Cloud's trust layer allows your AI team to benefit from our pioneering research and implementation in the field of LLM observability, providing our users with the most reliable and transparent LLM experience on the market:

- Eliminate LLM hallucinations for full trust from your user base.
- Back up LLM claims with rigorous source annotation.
- Bring transparency and observability to generative AI systems.
- Run experiments and collect feedback with [robust user feedback tooling](#).
- Use our [in-house metrics](#) to evaluate and monitor the accuracy of a model's responses.
- Ensure transparency and accountability for every product built in *deepset Cloud*.

deepset Cloud is the platform where AI engineers, data scientists, and software engineers can collaborate with product leaders, SMEs, and other stakeholders to build the best products possible – knowing that the heavy lifting of AI is taken care of.





Want to empower your team to start building with AI?

Schedule a demo today.

DEESET.AI