



H₂O.ai

8 Key Considerations for AI in the Enterprise

An AI Strategy Can Deliver Value from Your Data

The 8 key considerations when implementing an Enterprise AI Strategy



AI is unlocking new potential for every enterprise. Organizations are using AI and machine learning technology to inform business decisions, predict potential issues, and provide more efficient, customized customer experiences. The results can enable a competitive edge for the business.

AI empowers data teams to scale and deliver trusted, production-ready models in an easier, faster, more cost-effective way than traditional machine learning approaches. In this eBook, we'll explore eight key considerations you should keep top of mind as you develop an enterprise AI strategy for your organization.

Consideration #1:

Ask the right questions

A successful AI strategy begins with identifying the questions you want to answer for your business. Consider your needs, such as preventing customer churn, understanding customer dynamics, streamlining your supply chain, or determining the next best offer for consumers.

As inspiration, here are some examples of how AI and machine learning are being used in organizations worldwide today to address common business challenges:

- Financial institutions are using AI to evaluate credit risk more personally, providing the right applicants with the right level of credit.
- Retailers are using AI to enhance customer service with personalization and recommendations.
- Healthcare and life sciences organizations are using AI to help detect disease more accurately.
- Insurance companies are using AI to personalize a consumer's policy, or in the case of a claim, determine whether or not it is fraudulent, thus saving money.



Consideration #2:

Apply AI to business-specific use cases

The next step is to drill in further to determine specific use cases to which you can apply AI within your business. The best place to begin searching for use cases is in your existing business. Review your current processes and customer experiences and look for areas that need improvement.

Use cases do not have to be elaborate or complicated. Simple questions can help kickstart the ideation process, such as:

- What business problem am I trying to solve?
- How could predictive data improve my business?
- What results do I hope to achieve with this project?



Engage team members from multiple departments to ensure you have a complete view of the identified issue. Your goal is to take a 360-degree approach to each question.

Once you've implemented your AI solution and early results start rolling in, you can use this data to refine your use cases for future development.

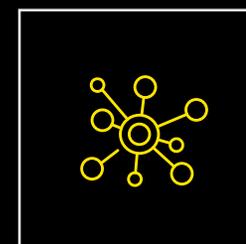
Consideration #3:

Where's your data?

Enterprise data is necessary for successful, pervasive machine learning. This rich data comes from a variety of sources, from customer transactions and social and marketing data to HR and finance information or even third-party data sources. In fact, addressing a single use case may require multiple data points drawing from different resources, including some that may not be immediately apparent.

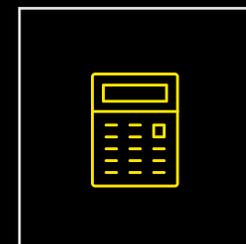


Consider a company that produces perishable products. The company must ensure a consistent temperature range at every step of the supply chain. Yet temperature data alone will not tell the complete story. In order to ensure a consistent temperature range, the company needs to examine a multitude of data points, including weather, shipping container options, and travel time.



To ensure you have a complete data set:

- Work with teams across your organization—from Marketing to IT—to identify a list of data repositories.
- Audit the list to determine data sources and any apparent issues or considerations.
- Identify gaps. Did you note anything missing? Take time to add new data sets to help solve your problem, as needed.



Consideration #4:

The fastest way to AI

AI provides insights—not just answers.

Using AI, you can generate insights to apply directly to your business. You may need someone to help you interpret the data, but you don't necessarily need a data scientist to analyze your AI. Look for team members who have an analytical mind or background, and assign them to your project.

You can also use an automatic machine learning (AutoML) solution to help turn data into actionable insights. AutoML platforms automate key aspects of the of the machine learning model, enabling data scientists, data engineers, mathematicians, statisticians, and others to build highly accurate, ready-to-deploy models, quickly. AutoML solutions can dramatically compress the model building cycle, enabling anyone in your organization—from the data scientist to the business owner—to get to answers faster.

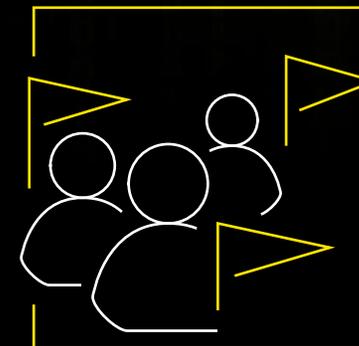


Consideration #5:

Data is a team effort

Developing the best possible solution to your business issue requires a team of diverse thinkers. In addition to your data scientist or analyst, your team should include a data engineer (usually in IT), a DevOps or Data Ops role, a business owner, and a sponsor.

This team should think through the implications of using AI to streamline your processes and/or solve your business problems. For example, consider how your project may impact other teams internally. It's also important to let other internal teams know that the project is underway, and what you hope to learn. You may uncover intelligence or insights that are beneficial to other groups within the company—which can be used to build broader support of AI/ML initiatives across your business.



Consideration #6:

Time

Developing an AI model requires a steady drumbeat, and the cycle may be repeated several times before it's done. Throughout the project, the team is making decisions: which features to use, which algorithms to use, and how to tune them. Keep in mind that the decisions you make along the way may impact or complicate your schedule.

Model training can be time-intensive and repetitive: Once you have a model scored, you generate an automatic pipeline to deploy models. This can add time and require another skillset to the data science workflow, such as a developer.

While you can't invent more time, you can optimize to work more efficiently. Using an AutoML platform, you can automate key aspects of the data science workflow, significantly reduce time, and increase accuracy. Try using an AutoML platform to train models faster, and then generate an automatic pipeline to deploy models.



Consideration #7:

Trust in AI—but do so responsibly

As we've discussed, an AutoML platform can interpret or explain the model and its results. Machines are amazing learners. But they are only as smart as the data provided. For better results, ensure that a human is hands-on before and after the AI process to quality-check the results. After all, humans can make connections that a computer is unable to make.

Consider the following: an AI might approve a million-dollar loan based on an applicant's excellent credit history and income—but the bank's manager may not be comfortable making such a sizeable loan without additional vetting. In this case, a parameter could be set requiring all loans over a certain limit to be directed to a human loan officer for review and approval.

Also, ensure your team understands the parameters and knows how to interpret the output of the models. For example, be sure to consider and understand the regulatory or compliance constraints that you are operating within for your industry, such as HIPAA, fair lending, or similar.



Consideration #8:

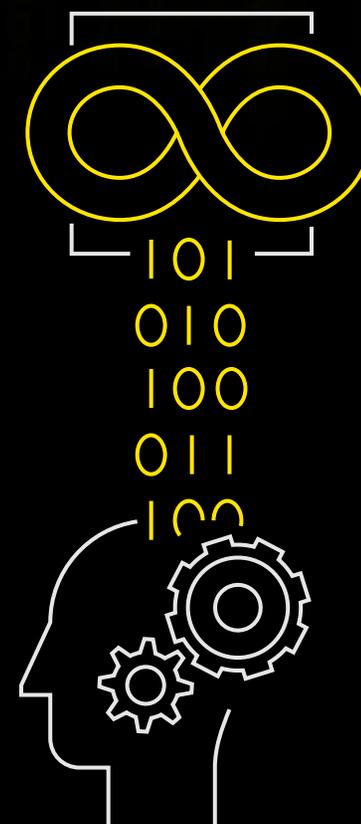
It's never over!

Your AI journey is just that—a journey, with constant iteration. Plan to rinse and repeat!

This means you must stay in charge and monitor your results. Your scored models can (and should) change based on the new data flowing in. Your team should be providing regular updates to key stakeholders throughout the process.

Some tips:

- Keep a human in the loop—don't leave it all to the machines.
- Use simple language to explain your models.
- Clearly state your results and how they apply (or can be applied) to the business.
- Back it up with repeatable models and documentation, so you can work more efficiently and train up your team.
- Stay alert and nimble—make changes as needed.



Get started with H2O Driverless AI

An industry leading and award winning automatic machine learning platform

AI in the enterprise can be a game changer and give you a competitive edge.

H2O Driverless AI is an automatic machine learning platform that helps scale your data science efforts. Driverless AI employs the techniques of expert data scientists in an easy-to-use platform to accomplish tasks in minutes that would normally take traditional techniques months.

With Driverless AI, everyone from data scientists to data engineers can develop trusted machine learning models. This next-generation automatic machine learning platform delivers unique and advanced functionality for algorithm selection, bring your own recipes, data visualization, automatic feature engineering, model interpretability, and low-latency deployment.

When you engage with H2O.ai on using and trialing H2O Driverless AI, we will give you help along the way, from teaching your data scientists and data engineers on how easy it is to use the platform, to having our customer data scientist help you with your use cases to get results.



Contact H2O.ai today to learn more