saifr®

WHITE PAPER

Considering AI solutions for your business? **Ask the right questions.**

Introduction

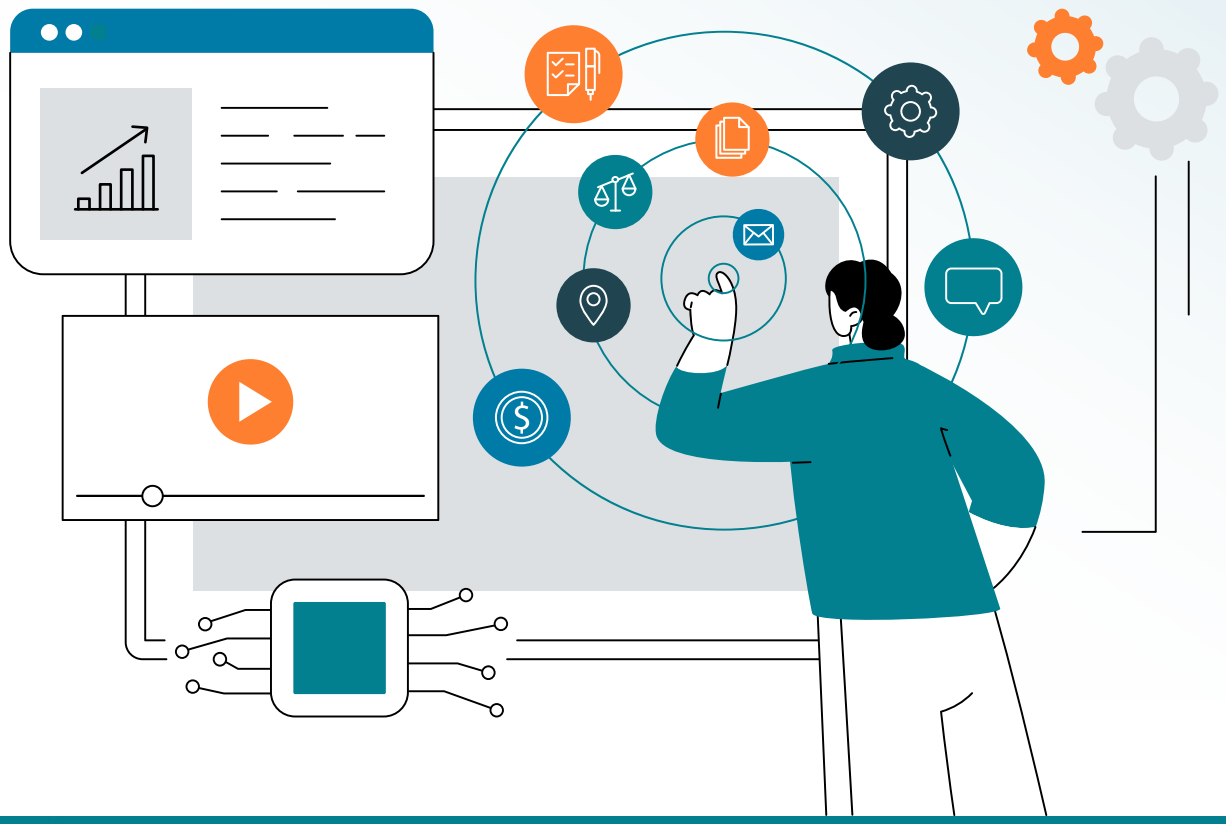

Artificial intelligence (AI) has emerged as a **transformative force in the business world,** revolutionizing industries and creating new opportunities for growth and innovation. The rush to embrace AI is evident in the resources that business leaders are committing to its use. Spending on AI is expected to **nearly double, from $33 billion in 2021 to $64 billion in 2025,** according to Forrester.

While companies with significant resources can build their own AI, **most others will likely buy AI solutions from a vendor or receive them bundled in software tools.** It is imperative that business leaders not only be able to identify opportunities where AI can create value and deliver competitive advantages but also pick the right AI solution for the right task. Before opening the checkbook, business leaders need to understand more than just the basics.

We wrote this white paper to help you understand AI so you can ask the right questions **to effectively select the appropriate AI tools for your business.**

SECTION 1

How did we get here?

Artificial intelligence is a broad term used to describe various methods that enable computers to have humanlike capabilities **such as learning, problem-solving, reasoning, and perception.** The concept of AI can be traced back to the mid-20th century, when pioneering thinkers began exploring the idea of machines that could exhibit intelligent behavior. In 1956, a group of researchers organized the Dartmouth Summer Research Project on Artificial Intelligence, widely regarded as the birthplace of AI, where they aimed at **developing programs that could mimic human intelligence.** However, their ideas took far longer to materialize than they predicted.

While AI experienced significant breakthroughs in the 1960s and 1970s, including advancements in natural language processing (NLP) that showcased its ability to engage in simple conversations, it wasn't until the 1990s that AI truly took off. Fueled by advancements in computing power and the availability of large datasets, the next two decades would see an explosion of interest and improvements in AI, driven largely by the **convergence of big data, powerful computing hardware, and algorithmic innovations.**

Three catalysts have played a major role in AI's growth

Data

The more data a model has access to, the better it can learn and perform. AI's need for vast datasets capitalizes on recent advancements in data storage and management systems.

Computing power

The rise in the speed and capacity of central processing units (CPUs) and later graphic processing units (GPUs), which were originally made for compute-hungry video games, has enabled the processing power needed to train the models and the data speed needed for realworld applications.

Methods

Algorithms have evolved from needing detailed instructions to being able to tell the computer how to learn to operate on its own.

SECTION 2

Four types of AI

It seems like every time you read a website or listen to a podcast, the topic is AI. As you read and listen, you notice terms popping up, like **"machine learning"** or **"deep learning."** These terms are foundational to understanding AI and helping you pick the right type, or types, that will best address your business needs. To the right, is a simple diagram to strengthen your knowledge:

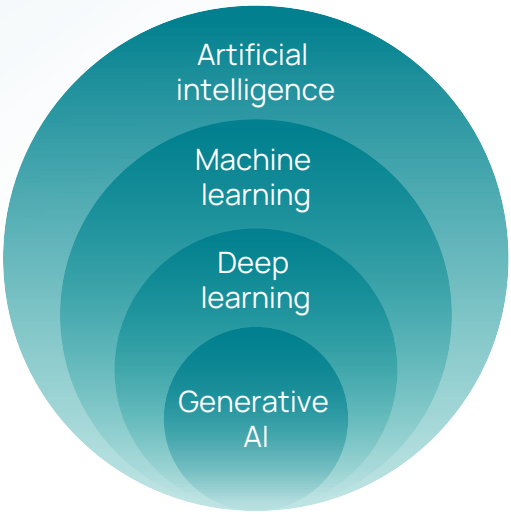

Rule-based models

The simplest AI models are rule-based systems that simply do as they are told. For example, if you want to train a rule-based AI model to identify words and phrases in a digital document, you would create a list of those words and phrases, potentially into the thousands, for the algorithm to identify and then surface for you.

However effective they are, rule-based AI tools may not be sufficient for more complex issues. For instance, they are unable to deduce context in a sentence and will often flag terms and phrases that are acceptable to the business user. These false positives can end up frustrating those users, causing them to mistrust or stop using the AI.

For instance, a legal or marketing team may deploy a lexicon tool to flag words and phrases in their marketing and sales materials, such as "promissory" statements promising a specific rate of return on an investment, that run afoul of regulatory compliance rules.

However, rule-based models have their limitations. A potentially promissory word, such as "superior," can be included in the lexicon and therefore flagged in a document, but the tool will flag the word wherever it appears, regardless of the context. Such as:

Come to our seminar in **Superior**, AZ.

Want to find success? Focus on providing **superior** customer service.

"It is impossible to produce **superior** performance unless you do something different from the average," said John Doe.

The examples above do not violate regulatory rules. In flagging them, the AI tool is doing what humans trained it to; but it's not performing in a way that reduces complexity and augments human effort.

Machine learning

A more advanced type of AI is machine learning. Machine learning enables computers to learn to recognize complex patterns within data without humans having to explicitly describe all the patterns of interest. **The machine isn't just taught, it actually learns.**

There are many categories of machine learning, including **supervised** and **unsupervised learning.** Supervised learning trains the computer using a large set of labeled data. The training dataset contains input examples (features) along with their corresponding output labels (target variable). The goal is for the model to learn a mapping function that can **predict the correct output for new, unseen inputs.**

Unsupervised learning trains an algorithm, without any explicit target variable or predefined labels, to detect similarities or abnormalities in a dataset. This type of AI requires large amounts of initial data to train the models to "understand" the text. The goal is to **discover patterns, structures, or relationships within the data.**

Once trained on curated data, machine learning models can handle new scenarios and understand context. So, in the promissory examples on the previous page, an AI model trained with machine learning would not have flagged those examples because it would have understood the context. This ability to contextualize allows machine learning AI to overcome many of the shortcomings of rule-based lexicons.

Deep learning

Deep learning is a subset of machine learning. It involves more layers of neural networks (a way to process data inspired by the human brain), where the output of one network becomes the input for the next. And, with deep learning, the computer learns from the data which features are most useful in making predictions.

Deep learning has gained immense popularity and achieved groundbreaking results in areas such as **computer vision, natural language processing, and speech recognition.** Several newly popular generative AI models—capable of churning out quarterly business reports, writing and designing marketing emails, and suggesting language changes to legal and other documents—are examples of deep learning.

The machine isn't just taught,
it actually **learns.**

Generative AI

Generative AI is an example of deep learning. As its name suggests, it can **generate text, images, software code, music, and videos, and even problem solve.** To learn how to do these things, the AI algorithms devise rules to predict outcomes and create new patterns. So, in the case of text generation, the AI predicts—and produces—the next most likely word in a sequence.

Because it uses massive amounts of data, generative AI has acquired an encyclopedic knowledge. To give an example of the amount of data it can consume, consider that the entirety of Wikipedia's data equals 100 million data sequences and a prominent generative AI model was trained on the equivalent of 3,000 Wikipedias worth of data. Future models will likely use even larger datasets.

Generative AI is proving useful in all kinds of industries, **such as healthcare, media, and finance.** It is improving medical imaging and helping doctors make diagnoses and suggest treatment plans, generating personalized product recommendations to shoppers, and helping with risk management in financial services by providing marketing teams with tools that streamline their compliance workflows.

However, despite the massive potential for generative AI, there are some drawbacks and limitations to it. Generative AI requires huge volumes of data, and the results are not always accurate or desirable. Sometimes the results are even tantamount to hallucinations. Hence there has been a huge focus on solving the alignment problem with AI, making sure it is trustworthy and does what it is meant to.

Questions to ask:

- What types and/or combinations of AI are being used, and how?
- If rule-based, what are the rules? How much can they be customized?
- If machine learning, is the data supervised or unsupervised? Who labeled the data? Was it double-checked?
- If generative, how was it built? Has it been customized to your specific use case?

SECTION 3

Why data are vital

Data are the lifeblood of today's well-run enterprise. Data should inform nearly every decision that companies make, from product development to understanding customer behavior to managing global supply chains. Data quality and data management are key to ensuring that businesses make decisions that are reliable and accurate, enabling their ability to thrive and innovate.

AI is only as good as the data used to train it: garbage in, garbage out. That's why data quality, which typically refers to its accuracy, completeness, consistency, and timeliness, and fitness for the purpose for which it is being used, is so crucial. So too is data management—the process of cleaning, organizing, and transforming data before using it for AI model training.

Questions to ask:

- What data was included in the training of the models? Who labeled it?
- How is my company's data being included? Who is labeling it?
- What safeguards against bias are in place?
- How is my intellectual property being protected in all stages of the process?

Six aspects to consider when evaluating data

Quantity

The more data a model has access to, the better it can learn and perform. Often it is a sheer numbers game, and models trained on more data outperform those with fewer data points. State-of-the-art AI models are now trained on hundreds of billions of datapoints.

Accuracy

Garbage in, garbage out applies. Where does the data come from? How do you know it is true and not misinformation? Who labeled the data, and was it checked? It is vital that data is drawn from credible sources.

Bias

Bias in AI can be caused by a variety of factors, including the data itself, preprocessing techniques used to prepare the data, and the algorithms used to process the data. The data should represent an unbiased cross section.

Diversity

Ensuring that the data covers a wide range of scenarios, variations, and edge cases can improve the models. Diversity can help ensure the accuracy of the model in realworld implementation.

Timeliness

How fresh is the data? Does it represent the current environment? Models can be most effective and potentially more unbiased when trained on current conditions.

Provenance

If the data is proprietary, how was it collected, stored, and processed? Were all copyrights respected? AI should not introduce new data security issues.

Poor quality data can have serious negative consequences on model quality, so it is essential to consider the above factors when training your AI and when deploying AI built by third-party vendors.

SECTION 4

How do you measure AI?

Even the best AI models will not be correct all the time. You need to understand how they can veer off course and the impacts that could have on your business application. There are various ways to measure how AI models are performing (see below). Which measurement you use depends on what is important to your use case.

For instance, can your AI application tolerate false negatives? If your application is predicting cancer, you'll want to avoid false negatives—telling a patient who has cancer that they do not. In this scenario, a false negative can be catastrophic: the patient won't seek treatment.

Now, what about false positives? If your AI application predicts which customers are likely to leave your business, and you use a false positive readout to offer free services, too many false positives may increase your customer retention costs.

Both of these examples demonstrate why it is crucial to understand what is important to your business and how to measure your models so that you can make the needed tradeoffs.

Questions to ask:

- What are the measurements of each model?
- How have the measures changed over time?
- How often are the models measured?

There are several common metrics used to evaluate the performance of AI algorithms

Accuracy measures the overall performance of the algorithm to correctly predict the output. How often was it right?

Precision can be thought of as a measure of exactness. It measures the quality of the positive predictions. When the model returned X, how often was it correct? This is an important measure for avoiding false positives.

Recall can be thought of as a measure of completeness. It measures the proportion of actual positives that were identified correctly. How often did the model say X when it *was* X? This is significant when overlooked cases are important.

F1 score is a combination of precision and recall and is used to measure the overall performance of an AI algorithm. A good F1 score means that you have low false positives and low false negatives.

SECTION 5

Humans must be in the loop

Computers are getting smarter, but it is only because of human input. The term "human in the loop" refers to the involvement of humans in all stages of creating, deploying, and using AI. Humans help ensure that AI models are accurate and reliable and that they continue to learn and improve via feedback and rewards.

As discussed in the last chapter, AI requires humans to clean, organize, and transform the needed data. Humans are also integral to the training stage as they can understand what the intended results are and help refine the model's accuracy.

Questions to ask:

- How is human intelligence and experience used throughout the lifecycle?
- How is feedback included in the use of AI in my business?
- How does the AI get smarter about my business and more accurate over time?

AI isn't something you set and forget; it needs to be continually assessed and adjusted. Even the best models will not perform properly 100% of the time, and can often lose accuracy—or drift—over time, thanks in part to their ability to absorb new data. Because of the potential for drift, **AI tools need a human in the loop who can guide the model to learn and self-correct.**

For example, a frequent use case for AI is to help customer service representatives handle calls more effectively and efficiently. The AI models offer suggestions that the rep can adopt or not. When the rep rejects a solution, the system learns that it wasn't correct and will learn what works for the clients, reps, and the company.

Humans are integral to all stages of AI, and their involvement needs to be incorporated thoughtfully to maximize benefits while helping to reduce some of the risks.

Conclusion

AI is not just a buzzword, it **enables organizations to gain insights from data, automate processes, and make informed decisions at scale.** Business leaders are immersing themselves in the fundamentals of AI. It can seem like the learning curve is steep, but it does not have to be.

Developing an understanding of some of the key ways that AI works can help you select and deploy solutions powered by this technology. If you pick the right tools, AI can help you achieve your goals.

saifr®

About Saifr

Saifr, a RegTech incubated by Fidelity Labs, streamlines and enriches the content creation and review processes to help mitigate brand, reputational, and regulatory risk. Saifr leverages its unique access to millions of documents representing more than 20 years of work by thousands of marketing and compliance experts in various lines of business to create deep-learning AI models. Saifr's AI helps make content creation, approval, and filing processes easier, faster, less expensive, and more accurate.