



Beyond the Hype.

A practical guide for
AI integration in business.



**“We are part of
the blossoming
era of data and AI”**

Table of contents.

01

Understanding Artificial Intelligence.

AI is a rapidly evolving field which will impact many aspects of our lives.

02

Generative AI.

The next frontier for AI - when AI doesn't just predict, but creates.

03

When not to use AI.

AI Missteps: When and Why to Abstain.

04

Identifying the best AI Use Cases.

“When the complexity of countless parameters strains human decision-making, artificial intelligence offers a guiding hand.”

05

The indispensable Role of Data Strategy.


“Strategy involves plans to achieve well defined objectives”

06

Architecting AI for Success: Building a Robust Data Infrastructure.

Driving value from your data strategy starts with the proper infrastructure.

The key takeaways.




Transformative Potential of AI

AI is reshaping businesses, offering transformative solutions and opening avenues for innovation.




Rising Momentum

AI adoption has seen over a 200% increase in the past five years, with companies harnessing its potential outpacing competitors.



Strategic Importance

AI and analytics are pivotal for the future strategies of most companies. It's not just a technological advancement but an existential necessity for businesses.



Human-AI Interaction

Successful AI integration requires understanding its fit within existing workflows, necessitating training and ensuring smooth communication between AI systems and humans.



Hypothesis Testing & Simulations

AI can be used to validate business hypotheses and run intricate simulations, aiding in decision-making processes.



Robust Infrastructure

Transitioning from AI testing to deployment requires a robust data infrastructure to handle and analyse vast amounts of data efficiently.

Introduction.

The rise of Artificial Intelligence (AI) is changing the way businesses operate and innovate.

In this whitepaper, we'll explore the facets of AI, from its transformative potential to practical use cases, and the essential role of data. Our goal? To equip businesses with a clear roadmap for AI integration and success.

The momentum behind AI is undeniable. A recent McKinsey report highlighted over a **200% increase in AI adoption over the past five years**¹. But it's not just about adoption rates; the real story is in the results. A select group of companies, those truly harnessing AI's potential, are outpacing their competitors, achieving higher financial returns and making bolder AI investments². When executed strategically, these increased investments in AI can lead to exponential gains. This trend is here to stay. As noted by a recent Gartner report, a significant **79% of corporate strategists deem AI and analytics as essential to their success in the strategy** of the upcoming years³.

For businesses, the adoption of AI and machine learning is not merely a strategic advantage; it is an existential necessity.

Companies that leverage AI have an edge in identifying trends, forecasting outcomes, optimising operations, and creating personalised customer experiences. They can turn data - an underutilised asset in many organisations - into a profit center.

Consider a global retailer using machine learning algorithms to optimise its supply chain, predict demand, automate pricing, and minimise waste. For instance, a real estate firm specialising in mall rentals, utilising AI to analyze consumer behaviour data, predict retail trends, and optimise tenant mix and rental pricing strategy. These are not edge cases; they are quickly becoming the norm.

Investing in AI is not about chasing the latest technology hype; it's about future-proofing organisations in an increasingly data-driven world. As a business leader, the question is not whether to invest in AI, but how quickly you can embrace it to stay ahead. But AI is not a cure-all. It's a tool, albeit a powerful one. Its potential is immense, but the path to using it successfully is filled with challenges. It calls for a rethinking of business models, processes, and culture.

Our goal? To equip businesses with a clear roadmap for AI integration and success.

Conclusion.

Consider a global retailer using machine learning algorithms to optimise its supply chain, predict demand, automate pricing, and minimise waste. Or on the other hand, for instance, a real estate firm specialising in mall rentals, utilising AI to analyse consumer behaviour data, predict retail trends, and optimise tenant mix and rental pricing strategy.

These are not edge cases; they are quickly becoming the norm.

Investing in AI is not about chasing the latest technology hype; it's about future-proofing organisations in an increasingly data-driven world. As a business leader, the question is not whether to invest in AI, but how quickly you can embrace it to stay ahead. **But AI is not a cure-all.** It's a tool, albeit a powerful one. Its potential is immense, but the path to using it successfully is filled with challenges. It calls for a rethinking of business models, processes, and culture.


Here at  **latitude**, our mission is to boost Europe's leading organisations by infusing AI into the **core of their being**. Our aim is not just to **implement AI solutions**, but to **inspire a transformational shift** in the way people think and act using the value of data & AI.

SECTION 01

Understanding Artificial Intelligence.

AI is a rapidly evolving field which will impact many aspects of our lives.

However, AI can be a complex and confusing topic, even for experts. This chapter provides an overview of AI, its applications and with a focus on the two main types of AI: **Narrow AI** and **Artificial General Intelligence (AGI)**.



Narrow AI.

Narrow AI is designed to solve a single task, like identifying spam emails or recommending music. It holds most of the current business value of AI and is where AI is most often applied in everyday business processes. Despite its Narrow focus, the impact and advancements of this type of AI can be significant, seen in systems like Spotify's Discover Weekly and Instagram's feed algorithm.



Artificial General Intelligence.

AGI (or Strong AI) is a more ambitious goal. It is the goal of many AI researchers to create an AI that can learn and perform any task that a human can autonomously. AGI would be able to do things like ethical reasoning, engage in meaningful conversations or potentially solve complex global challenges like climate change. It would also be able to learn new things and adapt to new situations.

The Difference Between Narrow AI and Artificial General Intelligence (AGI)

The key difference between Narrow AI and AGI is that Narrow AI is limited to a single task, while AGI is not. This means that AGI is more powerful and flexible than Narrow AI. However, AGI is also more complex and significantly more difficult to create. As of 2023, we are currently not able to create a full AGI. Different models are typically placed on a spectrum between these two fields. The new and popular generative AI models, such as **OpenAI's ChatGPT** or **MidJourney's visual model**, represent a significant step towards AGI.

The table below summarises the key differences between Narrow AI and AGI:

Feature	Narrow AI	AGI
Purpose	Solves a single specific task	Can learn and perform any task that a human can
Limitations	Can only perform the task it was designed for	Not limited to any specific task
Examples	Spam filters, self-driving cars, chess-playing computers	No real life examples yet

However, while many are drawn towards generative AI and the road towards AGI is luring, it's crucial not to lose sight of the power of traditional, Narrow AI. A recent McKinsey report suggests the added business value from new **generative AI use cases makes up just 35% of all potential Data, AI, and advanced analytics applications**⁴. This implies the importance and continued relevance of Narrow AI in the business world.

Subfields of Narrow AI and their use cases.

The continued development of Narrow AI is essential for businesses that want to stay ahead of the competition. By understanding the different subfields of Narrow AI and their potential business value, businesses can start to identify how they can use this technology to improve their operations and achieve their goals. The following matrix outlines some of the key subfields of Narrow AI and their potential business value. As can be seen, Narrow AI can be used to solve a wide range of problems, from optimising production lines to forecasting demand.

Field	Explanation	Examples/Use Cases	Business value
Operations research	Through classical mathematical optimisation to maximise or minimise a specific job	Optimising production lines, scheduling deliveries, and managing inventory	Improved decision-making, resource optimisation, and job scheduling
Machine learning	Allows computers to learn without being explicitly programmed	Forecasting demand, pricing products, recommending products, and automating customer service	Forecasting, pricing, predictive analytics, and automation
Deep learning	Uses artificial neural networks to learn from data	Processing text, recognising images, and understanding speech	Natural language processing, image recognition, and speech recognition
Natural Language Processing (NLP)	Deals with the interaction between computers and human (natural) languages. Quite often seen as precursor to ChatGPT	Analysing customer feedback, translating languages, and creating chatbots	Text analysis, machine translation, and chatbots
Reinforcement Learning	Allows agents to learn from trial and error	Allows agents to learn from training agents to play games, control robots, and trade financial assets. and error	Game playing, robotics, and financial trading

Conclusion.

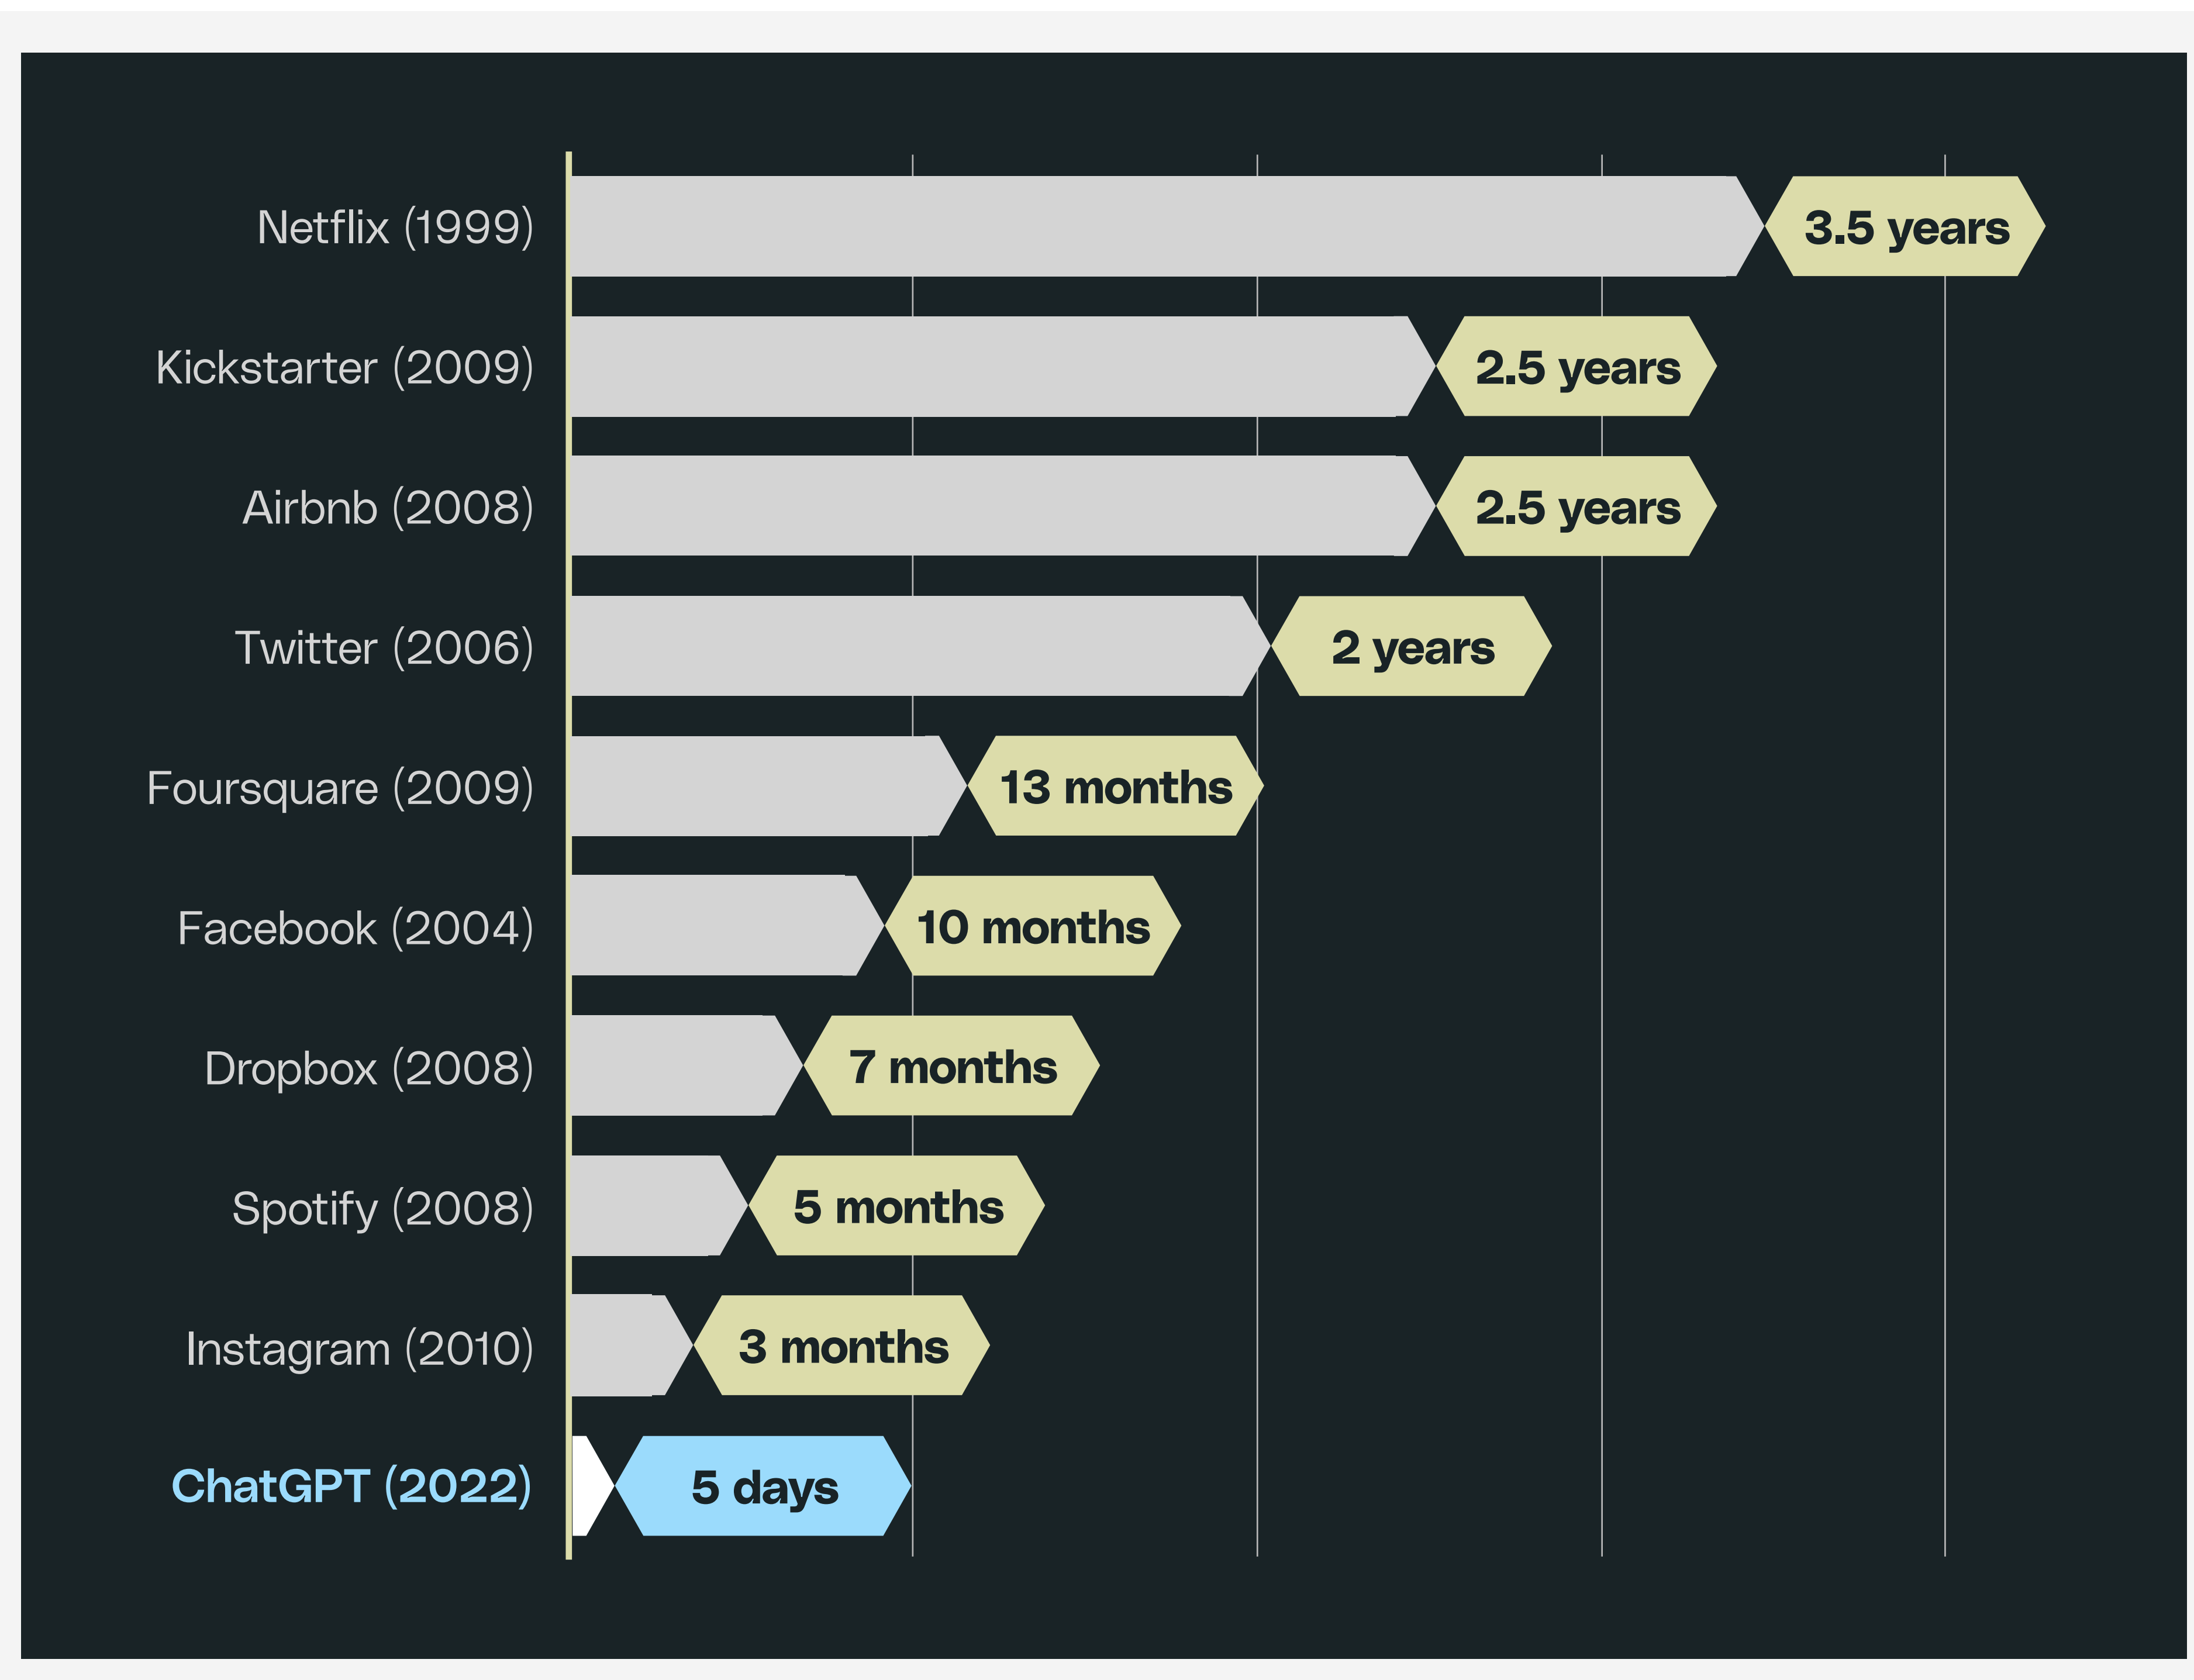
Generative AI, in its many forms, has quickly become a common phenomenon, strengthening capacities from content writing with OpenAI's ChatGPT, enhancing code development with GitHub's CoPilot and Code Interpreter to producing stunning designs through MidJourney. Amidst the rush, however, **it's important to not overlook the broader foundations of "classic" AI and Machine Learning**, which continue to hold significant business value at the moment.

SECTION 02

Generative AI.

The next frontier for AI - when AI doesn't just predict, but creates.

From the moment ChatGPT was released to the public, it became clear in no time that the dawn of generative AI was a paradigm shift in the potential value AI can provide. This was shown by the fact that OpenAI's ChatGPT is currently the fastest technology ever to get to a million users (5 days vs. 10 months for Facebook).



Unprecedented Speed to Adoption.
Time to reach 1M users

What is different about generative AI than anything that has come before it? Where classical AI models were good at automating singular tasks that were generally considered simple to do to humans with a clear correct/wrong output, generative AI is focussed on creation. From generating your next high school essay with OpenAI's ChatGPT with its text-based outputs to creating visual outputs like new logo's and designs with MidJourney to the latest release of Meta and their AudioCraft models to generate audio-based outputs⁵. A blog of Andreessen Horowitz puts it perfectly:

But as most people are using these tools more and more in the way they are intended, we at Latitude are fascinated by non-standard use cases that can help businesses in novel ways which have not been possible before. In this chapter we would like to focus on two main topics. The first is looking at unconventional use cases and new business models being made possible by generative AI as well as how you can set up a dedicated generative AI model for your own company which can reason from all your internal data.

“Just like the microchip brought the marginal cost of compute to zero, and the Internet brought the marginal cost of distribution to zero, generative AI promises to bring the marginal cost of creation to zero⁶.”

Marc Andreessen



Generative AI in Content Creation.

Generative AI is being used to automate the content creation process across various mediums. Midjourney and Stable Diffusion are examples of tools that generate unique artwork and graphical designs based on a given prompt. In writing, tools like ChatGPT, Sudowrite and Verb.ai specialise in generating creative and engaging content. In video production, platforms like Descript and RunwayML are expanding from video editing to content generation.



Innovative Generative AI Use Cases Transforming Industries.

01

AI in Personalised Mental Health Care.

The prevalence of mental health issues, especially burnout, is rising, but generative AI could revolutionise how we handle these conditions. AI-powered therapy bots, such as [Woebot](#) and [Wysa](#), can provide personalised therapeutic interactions, essentially putting a **"therapist in one's pocket"**⁷. These bots can interact with users in real-time, providing immediate feedback, coping strategies, and emotional support, all tailored to the user's unique mental health profile.

02

AI-Driven Personalised Education⁸.

AI is transforming education by enabling personalized learning experiences⁹. Tools like [Quazel](#) and [Lingostar](#) use AI to facilitate language learning through real-time conversation and feedback. Other tools, like [Orchard](#), aid students in their assignments, providing personalised writing assistance and helping students overcome writer's block.

03

Generative AI in Personalised Gaming.

Generative AI has the potential to revolutionise gaming by creating personalized gameplay experiences. [AI Dungeon](#) and [Hidden Door](#) are early examples of games using AI to customize gameplay. In the future, we can imagine games where players can design custom avatars or even entire virtual worlds from scratch, offering a truly personalised and immersive gaming experience.

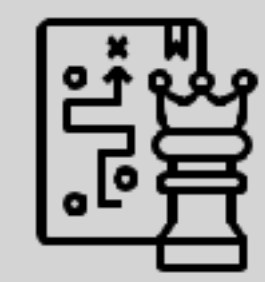
Training an own Gen AI model/Making Gen AI work for you.

One of the biggest challenges with generative AI is that it doesn't know the intricacies of your business and how you operate. Besides that, you also don't want to expose all your proprietary information to ChatGPT and risk the data falling in the wrong hands (which was one of the reasons Samsung banned access to ChatGPT for its personnel¹⁰). This combination of considerations can be resolved if there would just be a generative AI that would know all the details of your business without the risk of the data falling in the wrong hands. Moreover, what if by every prompt that the model processes the model becomes even better at suggesting and helping your business and colleagues be more productive?

A recent Harvard Business Review article outlined the main approaches currently available to **create your own proprietary generative AI model/LLM**¹¹. The three main approaches that are available.

Technique	Description	Example	Amount of data needed	Technical Difficulty	Potential Value
Design & Train a model from Scratch	Have a giant corpus of data, the technical expertise, time and investment to train a model internally.	BloombergGPT - A natural-language model based on 40 years' worth of proprietary financial data ¹² .	Very Large	High	€€€€
Fine-Tune an existing model	Take a pre-trained model and update the weights by training it on your proprietary data.	Elsevier has fine-tuned the weights of GPT3.5 to generate an assistant for researchers to get quick summaries of research topics they are unfamiliar with ¹³ .	Large	Medium	€€€
Prompt-Tune an existing model	Take a pre-trained model and provide it with all your data through prompts to get specific outputs. Unlike fine-tuning, weights of the model are not updated, but the model is "guided" using specific prompts.	Using GPT-3 to generate financial summaries by feeding it with specific prompts and a subset of proprietary financial data without changing its internal weights.	Moderate	Low	€€

But how to determine what to do for your use case? It's still early days in this new field, but with the experience that we have gained in more old-school Transfer Learning of classical neural networks, we have found that it becomes comes down to two main considerations:

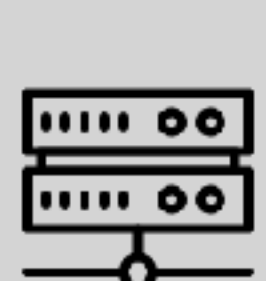


01

Capability Consideration.

First off, if you want to consider fine-tuning you will have to make sure that you have the ML engineering capabilities and infrastructure expertise to deploy a model. A well thought-out architecture for LLM applications is anything but trivial as demonstrated in this Andreessen-Horowitz blog “Emerging Architectures for LLM applications”¹⁴.

02



Data Consideration.

Do you have a vast dataset which is very unique and holds an untapped value which can be used in a novel way and open up a new business model or is it a smaller dataset that can be honed to automate and speed up internal processes?

There has been a big push for the fine-tuning approach, as was made clear when **OpenAI allowed for a fine-tuning approach on their platform allowing everyone with API keys to set up a fine-tuned GPT-3.5 model in late August 2023** with support for GPT-4 coming later this year¹⁵.

Conclusion.

Generative AI has brought in a new era of potential for artificial intelligence. From ChatGPT's rapid rise in popularity, it's evident that generative models have gone beyond traditional task automation to focus on creation, spanning text, visuals, and audio. This transformative technology offers tools for automated content creation, personalised mental health care, education, and gaming experiences. Yet, as businesses look to harness this power, they face challenges in ensuring the AI aligns with their unique needs. Whether designing a model from scratch, fine-tuning existing ones, or prompt-tuning to achieve specific outputs, it's crucial for organisations to consider their goals, data availability, and technical capabilities. As the world of generative AI evolves, companies must strike a balance between immediate value and long-term potential, ensuring their strategies remain adaptable and forward-thinking.

SECTION 03

When not to use AI.

AI Missteps: When and Why to Abstain.


In today's rapidly evolving digital landscape, AI has emerged as a powerful tool with transformative potential across numerous industries. From healthcare to finance, its capabilities have been praised, resulting in a surge of investments and integration efforts. However, amidst the race to adopt AI, it is essential for organisations to recognise scenarios where the implementation of AI may not be beneficial or even necessary. Unfortunately often still the real problem is not clarified enough and investments are made with limited or even negative returns.

Scenarios Where AI Implementation May Not Be Beneficial.



Simple Tasks

AI excels in solving complex problems. However, for simple tasks that don't require adaptive learning or analysis, traditional software or BI might suffice. Implementing AI for such tasks could increase costs without delivering any benefits.




Time Constraints

Training AI models can be time-consuming. If a solution is needed urgently, it might be more practical to use more straightforward methods or pre-existing solutions.



Data Scarcity

AI, especially deep learning models, require vast amounts of data to train. If an organisation lacks sufficient quality data, the AI model could produce unreliable outputs.



Fixed, Unchanging Environments

In static environments where conditions and requirements don't change, the adaptive nature of AI might be overkill. Traditional algorithms with fixed rules might be more efficient.



Highly Regulated Domains

In areas where decisions require clear interpretability or are subject to strict regulatory inspection, AI models, particularly black-box models, might not be appropriate.

Learn from the best

Two prime examples of when the application of AI did not return the expected results to say the least...

01

One vivid example where AI implementation has gone wrong was in late 2021, Zillow, a big online property site, decided to stop its "Zillow Offers" service and had to lay off nearly 2,000 of its workers. The main issue? Their AI system's predictions about home prices were off.

The "Zillow Offers" idea was simple: Zillow would buy homes based on a price predicted by their AI, do some renovations, and then try to sell at a profit. But, the AI's price guesses weren't always accurate. On average, it missed the mark by about 1.9%, and sometimes, especially for houses not even listed for sale, it was wrong by as much as 6.9%.

According to CNN, since the program began in 2018, Zillow bought 27,000 homes but only managed to sell 17,000 by September 2021. Unexpected challenges like the pandemic and a shortage of renovation workers didn't help either¹⁶.



In 2018, a significant bias issue arose in an AI recruitment tool developed by Amazon. The system was designed to sift through resumes and rank candidates according to their qualifications. However, it was discovered that the algorithm exhibited a notable bias against women.

The root of the problem lay in the training data used. The system had been trained on resumes submitted to Amazon over a 10-year period. Since tech is a male-dominated industry, most of the resumes came from men. This unintentionally taught the algorithm to favour male candidates. The AI downgraded resumes that included the word "women's", as in "women's chess club captain". Furthermore, it gave lower scores to graduates of two all-women colleges.

Amazon did attempt to neutralise the gender specific terms from the algorithm, but there was no assurance that the bias was completely removed, and so it was discontinued in 2018¹⁷.

02

Conclusion.

AI has reshaped many industries, leading to increased investments and widespread adoption. However, as seen with cases like Zillow, a hasty embrace without understanding can lead to challenges. Not every problem requires an AI solution. Organisations should evaluate the real-world applicability of AI against their specific needs, ensuring they choose the right tool for the job to achieve successful outcomes.

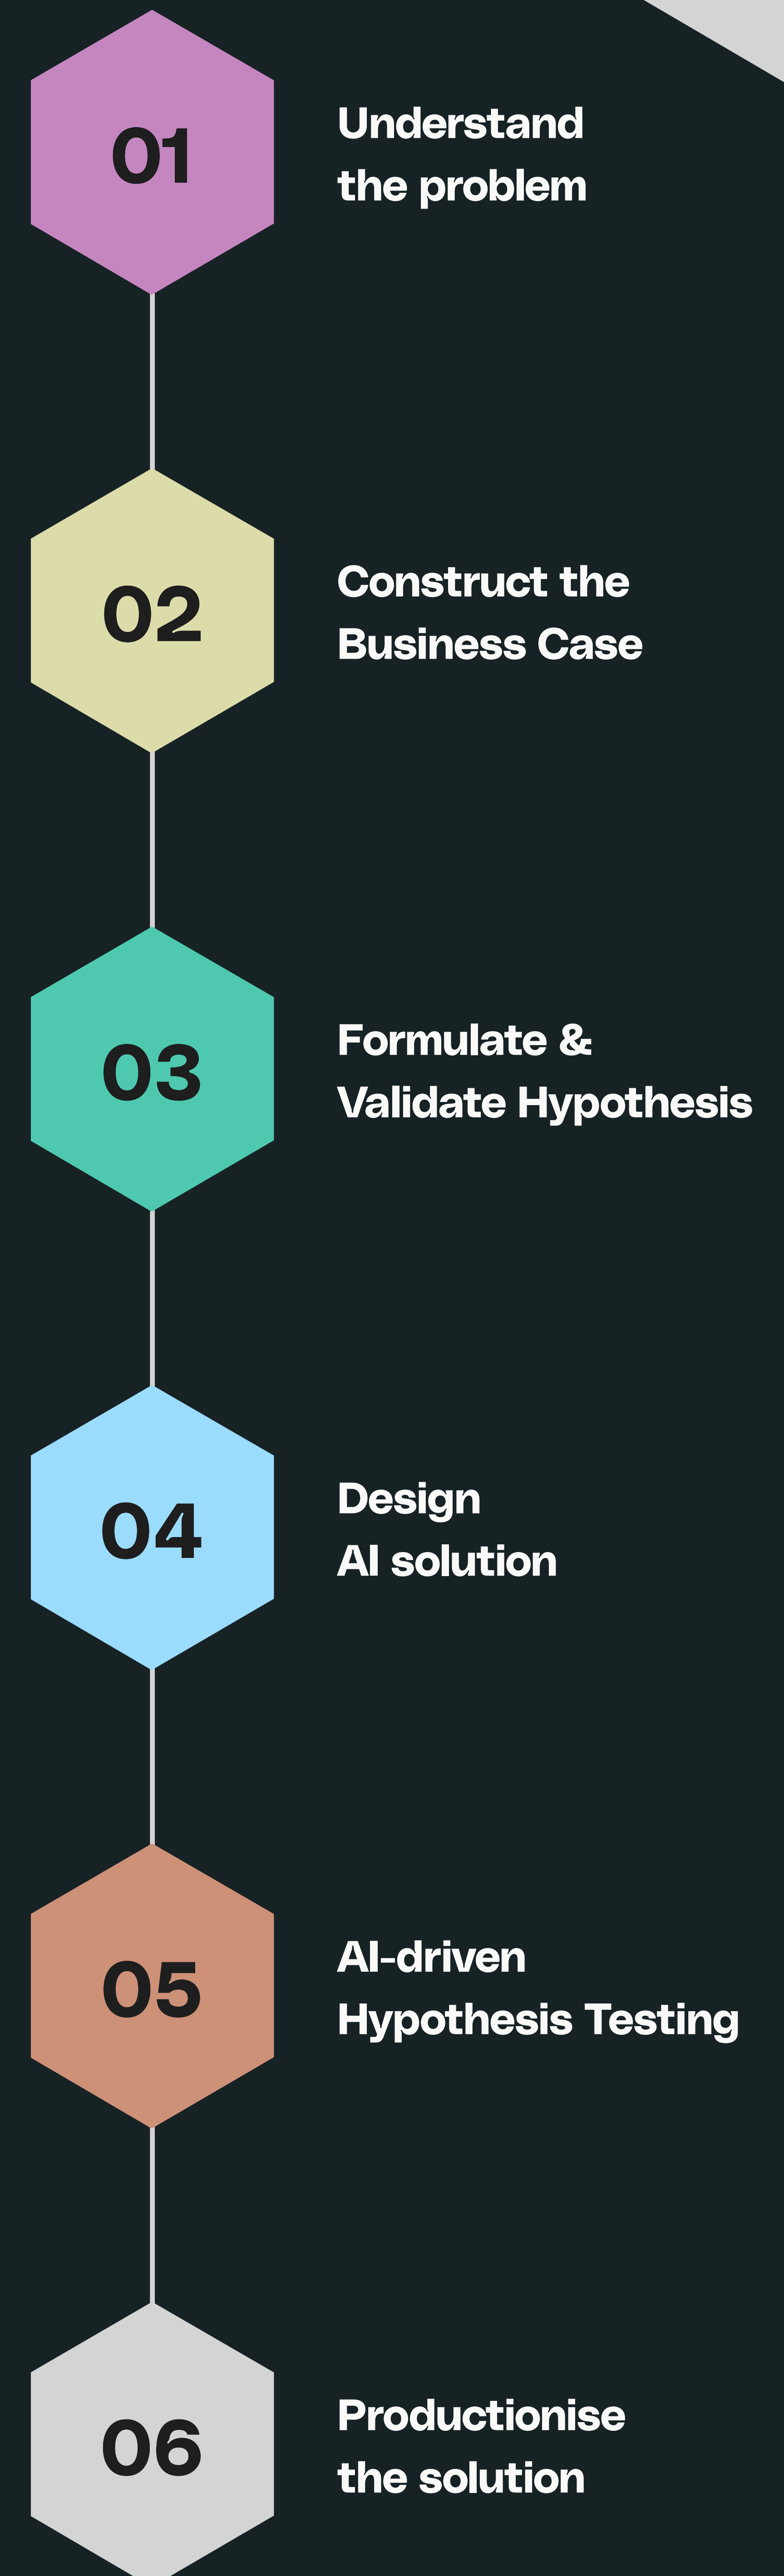
SECTION 04

Identifying the best AI Use Cases.

When the complexity of countless parameters strains human decision-making, artificial intelligence offers a guiding hand.

When in the process of identifying AI use cases, there is only one way and that is by **starting from the problem**. Often problems and solutions don't match in the world of data & AI, therefore we believe coming back to **Occam's razor is crucial** and it is important to work hard to create the simplest possible solution.

On the right hand side you will find six key steps on how to identify the best AI use cases:



01. Understand the problem.

The initial stage of the process is where the problem/goal is clarified. This means understanding the pain points, gathering requirements, and analysing existing solutions or methods that might be in place. It lays the groundwork for what you hope to achieve with AI.

By breaking down the overarching business problems into subproblems, it's possible to gain a more nuanced understanding and identify the best ways to tackle them. **At Latitude we have defined a tool to help you out.** [Download our problem statement canvas](#) to clearly define the problem before proceeding with the next steps. It often looks easier than it is and by carefully filling in the canvas it will make the next steps of the process much easier.

For example, consider the challenge of increasing sales of computer monitors in a particular region. This goal prompts several crucial questions:

4 crucial questions

01. What are the underlying factors affecting sales?
02. How can analytics be leveraged to enhance understanding and drive growth?
03. Is there a need for data-driven dashboards, or could AI and Machine Learning-based solutions be applied?
04. Who are stakeholders/decision makers to make this happen?

Six important considerations to see if AI is a fit for your problem.

- 01

Real-world Context

AI models, while powerful, often operate in ideal conditions during training phases. The real world, however, presents unique challenges, including dynamic changes, anomalies, and other unforeseen events that a model might not have encountered during training.
- 02

Infrastructure Considerations

Practical implementation may demand robust infrastructure, reliable connectivity, and other infrastructural support. Without these, even the best AI models might underperform or fail (see chapter 6 for more details on data architecture).
- 03

Human-AI Interaction

It's essential to understand how AI will fit into the current workflow. Will employees need training? How will the AI system communicate its results or decisions? Ensuring smooth integration with human operations is key for maximising the benefits of AI.
- 04

Data Assessment


Check the quantity and quality of available data. If your organisation lacks the data AI requires, or if data labeling is impractical, reconsider the decision.
- 05

Cost-Benefit Analysis

AI implementation is often resource-intensive. Evaluate the potential benefits against the costs of integration, maintenance, and potential risks. If the returns are not substantially higher than the investment, it might be better to explore other solutions.
- 06

Understand the Regulatory Landscape

Ensure that you are aware of all regulations that might affect AI's deployment in your domain. If compliance is difficult or ambiguous, it might be best to proceed with caution.



02. Construct the Business Case.

The business case is a critical component of any (AI) initiative, serving as the foundation upon which the project is built. It consolidates the rationale behind the proposed solution, detailing the expected benefits, costs, and potential risks. A well-constructed business case provides clarity, aligns stakeholders, and serves as a reference point throughout the project's lifecycle.

Objective Clarity: Begin by clearly articulating the problem you're aiming to solve with AI. This should be rooted in the understanding of the problem you've previously developed. For our computer monitor sales example, the objective might be: "Increase sales of computer monitors in region X by 20% over the next year using targeted marketing campaigns."

Expected Benefits: Quantify the potential benefits of the AI solution. This could include increased sales, reduced costs, enhanced customer experience, or improved operational efficiency. In our example, the expected benefit might be increased revenue from higher sales and improved brand loyalty due to more targeted marketing.

Cost Estimation: Calculate the expected costs of implementing the AI solution. This should at least contain:

- Development or purchase of the AI system.
- Infrastructure and integration expenses.
- Training and support costs.
- Ongoing maintenance and potential upgrade costs.
- Risk Assessment: Identify potential risks associated with the AI solution. This could range from technical challenges, like integration issues, to strategic risks, such as potential negative customer feedback. Each risk should be paired with a proposed mitigation strategy.

ROI Calculation: With the expected benefits and costs laid out, calculate the Return on Investment (ROI). This metric will be crucial to judge the project's viability. For our example, if the expected revenue increase from higher sales is significantly higher than the cost of the AI solution, the ROI would be positive.



03. Formulate & Validate Hypothesis.

Breaking down complex business challenges into smaller sub-problems is crucial for clarity. This method allows for a detailed analysis and tailored solutions. Take the goal of boosting computer monitor sales in a niche market. One approach might involve identifying the most effective age group for marketing. A hypothesis could suggest that focusing on the 35-40 age group would drive sales. Meanwhile, the null hypothesis would suggest no significant change in sales with this approach, and it's essential to remember that our aim is often to disprove this null hypothesis.

Having clear hypotheses provides a roadmap for systematic experimentation and validation. It's not just about making assumptions; it's about methodically testing strategies rooted in data. Using the sales example: the next step would involve validating whether targeting the 35-40 age group truly makes a difference. This means diving deep into past sales data to identify patterns related to age demographics and purchase behaviours.


Once validated, the real-world application of the hypothesis in the target market follows. This direct approach yields tangible results, either reinforcing or challenging our initial theory.

Through this cycle of forming a hypothesis, validation, and real-world testing, we move from assumptions to solid, actionable strategies, enabling smarter decision-making.

04. Design AI solution

Once we've confirmed our initial hypotheses, the next step is to determine the best tools for ongoing analysis. While traditional dashboards can provide a snapshot of results, for a deeper dive and predictive insights, AI and machine learning become invaluable. Taking the sales example: a dashboard might illustrate how targeting the 35-40 age group affects sales. But for nuanced insights, like pinpointing the optimal ad timings or refining campaign messaging, AI offers a more sophisticated lens.

For instance, when traditional analysis can't quite nail down the best time for campaign releases, an AI-driven model can step in, offering predictions on peak engagement hours. In terms of messaging, machine learning algorithms can sift through trending topics within our target age bracket, guiding us toward the most impactful language and themes. By integrating AI in this strategic manner, we refine our marketing approach, ensuring greater effectiveness and a stronger return on investment.



05. AI-driven Hypothesis Testing

Building on the insights gathered and the AI tools employed, the next progression is to iteratively test and optimise the hypothesis in real-world scenarios. The world of AI brings dynamic capabilities that allow for refined testing, ensuring that the strategies are not just theoretically sound but also practically effective.

In the context of our sales example: with a hypothesis validated and AI solutions in place, we can now run intricate simulations. These simulations might assess various marketing content crafted around our target age demographic, adjusting for factors like engagement times and keyword effectiveness. For instance, we might discover through AI analysis that certain keywords resonate more during specific times of the week.

The strength of AI-driven testing lies in its capacity for precision and adaptability. By integrating A/B testing within an AI framework, we can methodically evaluate different strategies in real-world scenarios.

For instance, using our sales example, two marketing strategies targeting the 35-40 age demographic can be tested concurrently. AI helps analyse the results, determining which strategy resonates more, and why.

As feedback from these tests is looped back into the system, the AI model refines its understanding, allowing for continuous optimisation of the strategy. This iterative process ensures our solutions remain agile, adapting to changing conditions and data. Ultimately, leveraging an AI-driven hypothesis testing framework translates to more accurate, data-informed decisions that consistently align with organisational goals.

06. Productionise the solution

Moving from testing to full deployment means having the right tools and systems ready. A strong data infrastructure is key. For our sales example, when rolling out marketing strategies for the 35-40 age group, it's vital to have systems that can handle and analyse the data smoothly.

Investing in a good testing framework and solid data setup is essential. It's not just about launching the AI solution; it's about making sure it can keep improving with new data. With the right setup, we can continuously update our AI models, keeping strategies fresh and effective. In short, pairing a smart AI solution with a sturdy data backbone sets the stage for consistent growth and success.

Conclusion.

For businesses venturing into AI, the journey begins by clearly identifying the problem at hand. A systematic approach, grounded in understanding the problem, constructing a solid business case, formulating hypotheses, and leveraging AI tools, is crucial. By breaking down challenges, validating assumptions, and harnessing AI's predictive capabilities, companies can craft strategies that are both data-informed and actionable. Moreover, the significance of a robust data infrastructure cannot be overstated. When thoughtfully implemented and continuously refined, AI solutions can unlock unparalleled value, driving growth and operational excellence.



SECTION 05

The Indispensable Role of Data Strategy.

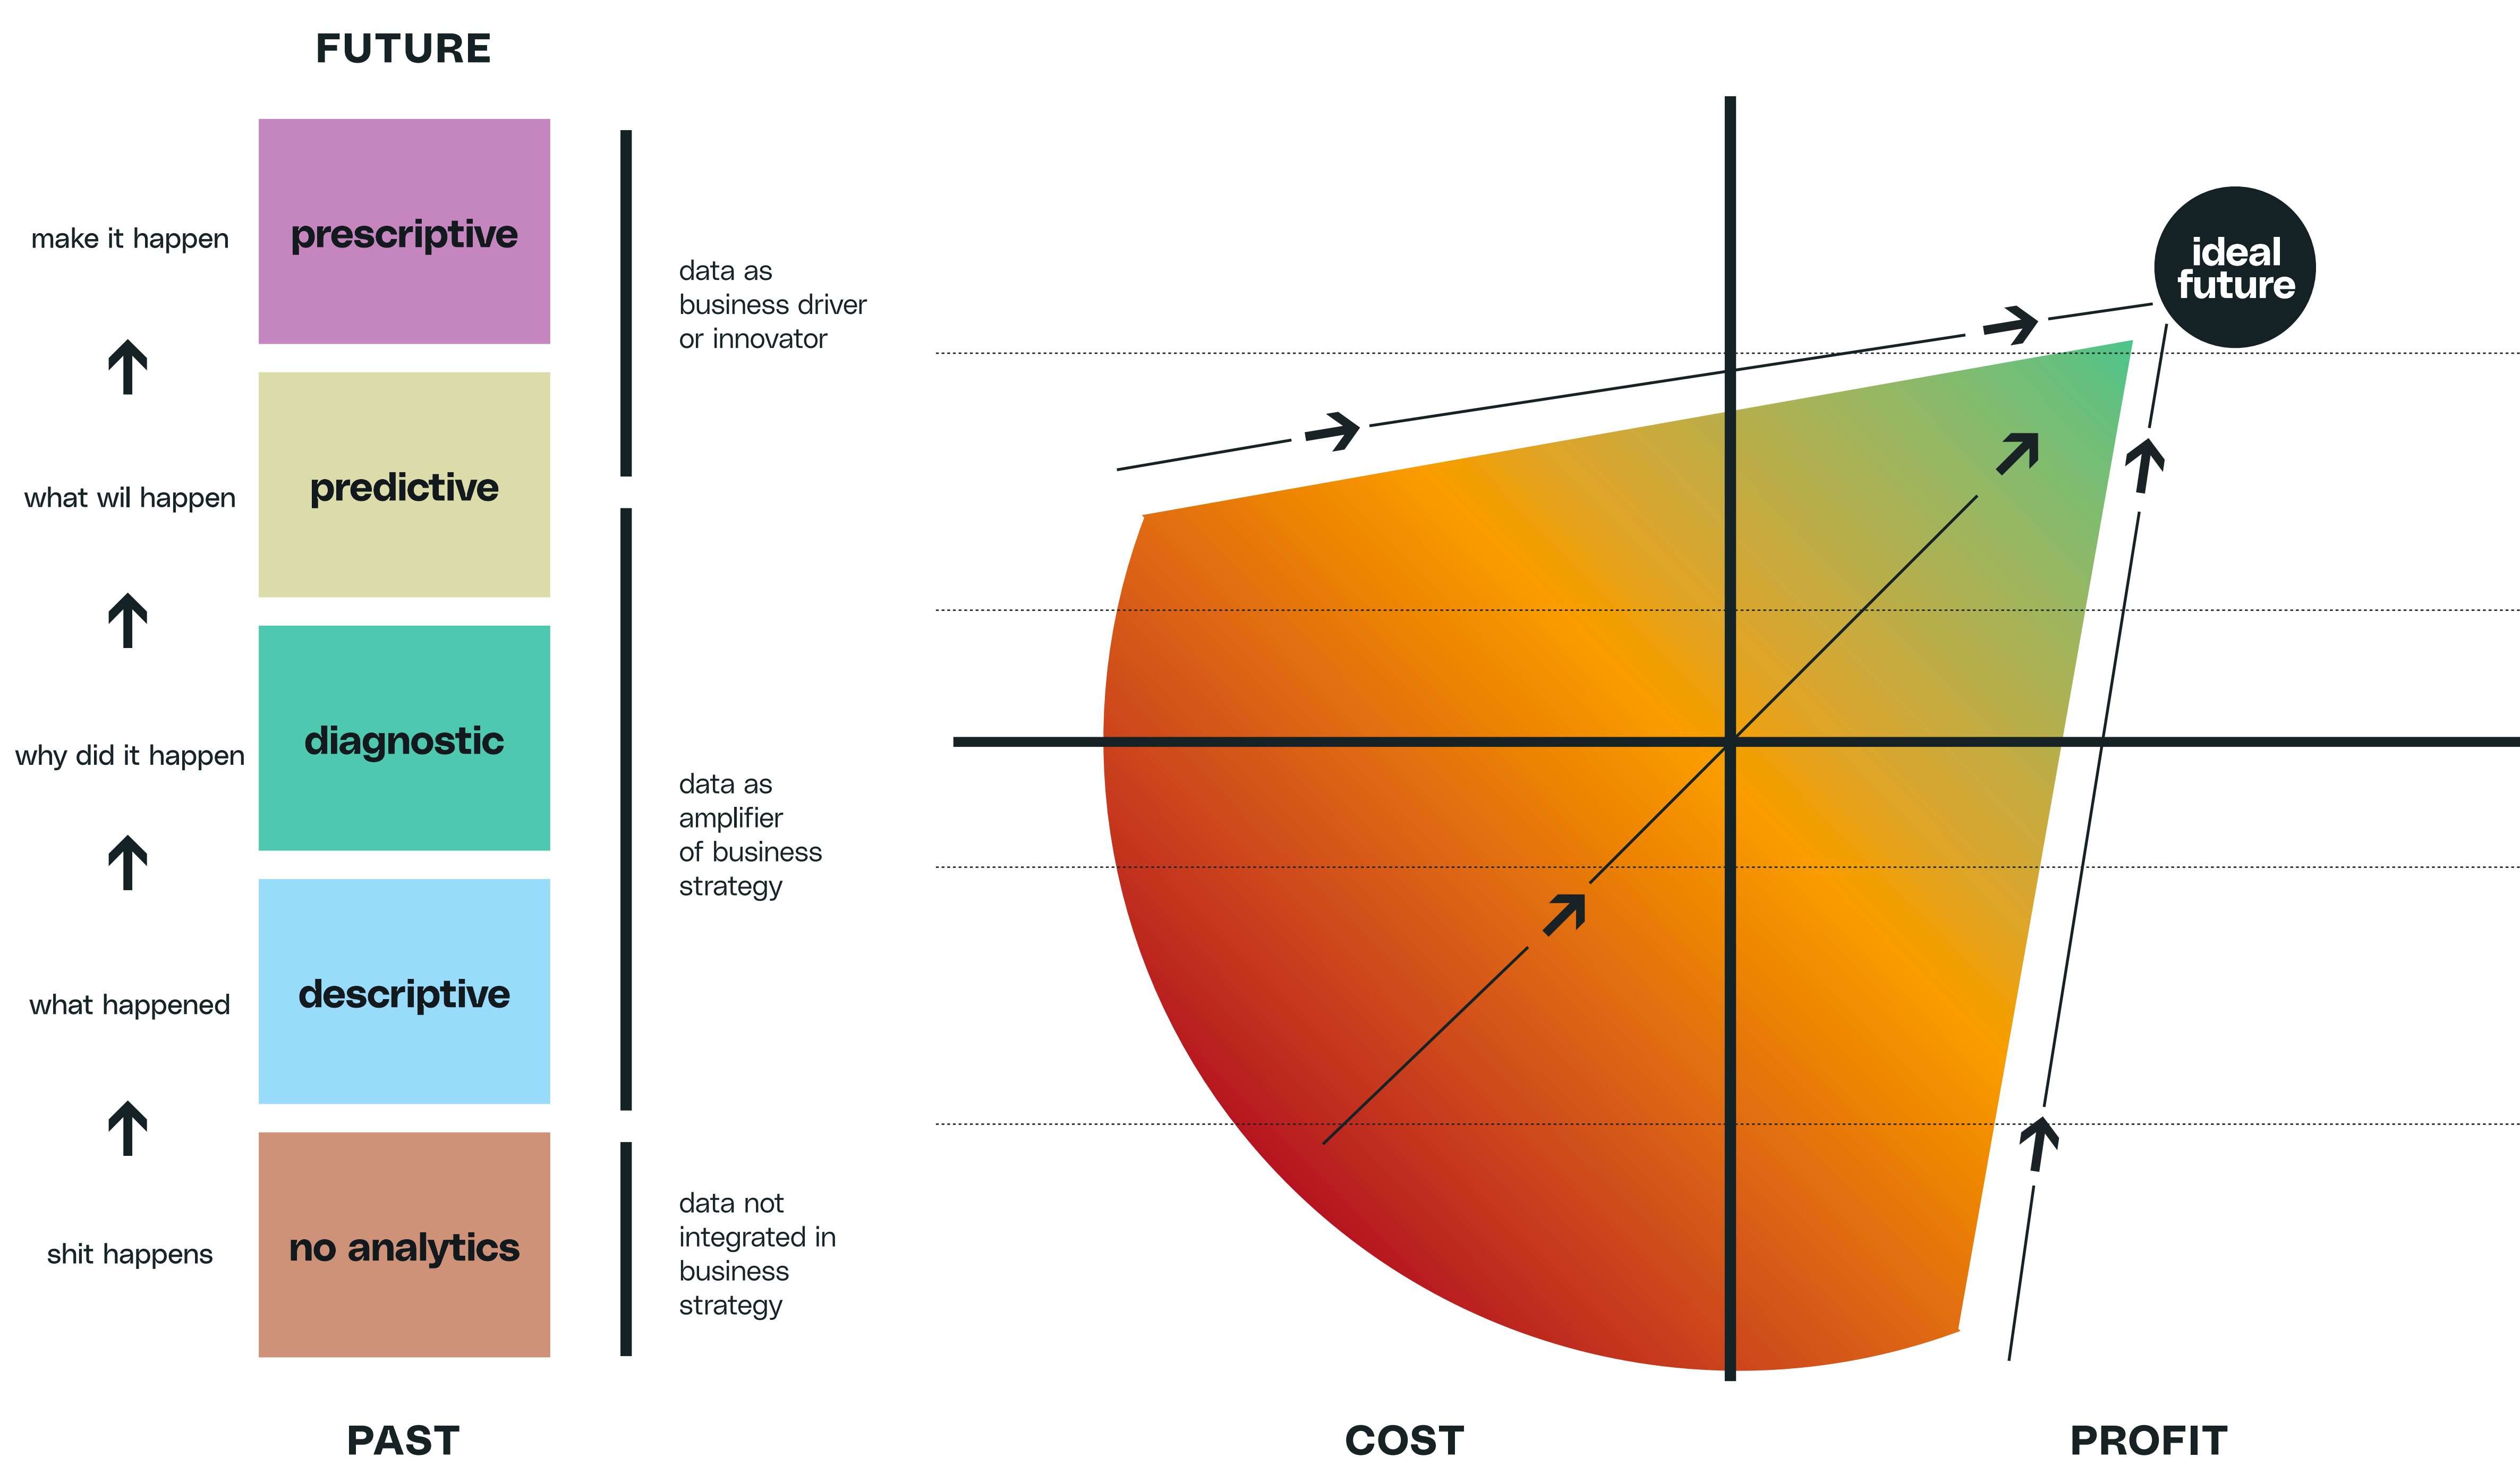
“Strategy involves plans to achieve well defined objectives”

When we talk about a data strategy in comparison to corporate strategy; the purpose remains the same, only the contents and execution differ. A data strategy should always be connected to or driving the defined objectives of the corporate strategy of the organisation. Similar to sailing, it’s essential to roughly know where you aspire to go before you embark on your journey.

Latitude’s Aspiration Model helps to drive value from data.

Terms such as 'maturity' are often used to describe the complexity level of how data analytics is implemented in an organisation. In human contexts, maturity means staying true to oneself and taking responsibility without drawing excessive comparisons or suggesting value differences. In many organisations, data is viewed as a cost center, receiving limited attention or having no clear goal. Without clear aspirations, it's challenging to determine a starting point or direction. So, the primary question becomes: “What do you aspire to achieve with data”

At Latitude, we perceive things from a unique angle. We believe that with the right approach and clear aspirations, **data can shift from being a cost center to a profit center.**



Depending on where your organisation is currently positioned along our aspirational model, the priorities and complexity of data use cases differ.

The processes underlying Latitude's Aspirational Model.

Below table provides an overview of the varying stages of data integration within organisations, illustrating the transformation from a nascent stage of data integration to a more advanced, proactive approach where data becomes a central driver of business innovation and strategy.

It categorises organisations into three distinct phases based on their maturity in data utilisation: those where data is not integrated into business strategy, those using data as an amplifier of business strategy, and finally, those leveraging data as a business driver or innovator.

Each phase is characterised by specific identifiers, actions, data cleanup approaches, percentages of low-quality data, people involved, processes, and technologies. It serves as a guide to understanding where organisations lie on the spectrum of data integration and what steps can be taken to advance to higher levels of data-driven business strategy.

	Data not integrated into Business Strategy	Data as an Amplifier of Business Strategy	Data as a Business Driver or Innovator
Identifiers	Many organisations are at this stage. Minimal analytics often lead to unforeseen costs or missed revenue opportunities.	Organisations recognise data's potential to reinforce business goals. Data is often viewed as a cost due to infrastructure and solution needs.	Data is seen as a profit generator. Potential for data-driven innovation in new business models or offers.
Key Action	Prioritise understanding the organisation's performance and its main drivers.	Develop strategies on how data can propel the business or reshape the market dynamics.	Focus on optimisation, infrastructure, processes, and team dynamics. Continuously evolve data practices for innovation and value generation.
Data Cleanup	Reactive, often after problems arise.	Proactively managed to ensure quality.	Fully automated processes to ensure data integrity.
% of Low-Quality Data	80%	40%	5%
People	Beginning stages of a data team, or a small existing data team.	Expanded data team with diverse roles, including analysts, engineers, and data scientists.	Data specialists in all business sectors with dedicated data strategy units.
Process	Ad-hoc data tasks and analysis without regular patterns.	Established data pipelines, feedback mechanisms, and consistent reporting.	Comprehensive data governance, seamless data integration across units, and integration of continuous improvement methods.
Technology	Emphasis on basic Business Intelligence (BI) and reporting. AI is explored with pioneering stakeholders.	AI initiatives target early adopters. Integrated analytics platforms support decision-making.	Business operations driven by AI, advanced analytics, and real-time data processing, irrespective of stakeholder perspectives.

How to start with Data Strategy?

After you've aligned the corporate strategy and defined the aspirations in terms of data. It's time to get started with the actual practical steps of your data strategy. This is boiled down to answering three questions:

1. What are our data & AI opportunities?



Offensive

Revenue generating: Leveraging data and AI for offensively **focuses on top-line growth.** This could involve using AI algorithms to personalise marketing efforts, tapping into big data to identify new market segments, or analyzing customer behaviour to optimise product offerings. By predicting what customers want or need before they even realise it, businesses can boost sales and expand into new territories.



Defensive

Cost saving: On the defensive side, data and AI can streamline operations and reduce expenses. For instance, predictive maintenance can forecast when a machine is likely to fail, thereby preventing costly downtime. Additionally, AI can automate routine tasks, reducing the need for manual labor and minimising human error, leading to significant cost savings.

2. Why should we chase these opportunities?



Business Case

Justification of your opportunities: Every strategic move should be backed by a solid business case. For data & AI, the justification lies in the potential ROI. By investing in AI-driven solutions, businesses can achieve faster growth, outpace competitors, and ensure long-term sustainability. **Data-driven insights can also uncover inefficiencies,** helping organisations refine their strategies and maximise profitability.



3. How do we implement these opportunities?



People

The success of any data & AI initiative depends on having the right talent. This involves hiring data scientists, data engineers, and data analysts. But it's not just about bringing in experts; it's also crucial to train existing employees to work in a data-driven environment.

A culture that values data and encourages continuous learning is key.



Process

Implementing data & AI requires a shift in traditional processes. This might involve adopting Agile methodologies, iterative testing, and regular feedback loops. Businesses should also establish clear data governance protocols, ensuring that data is handled, stored, and processed ethically and securely.



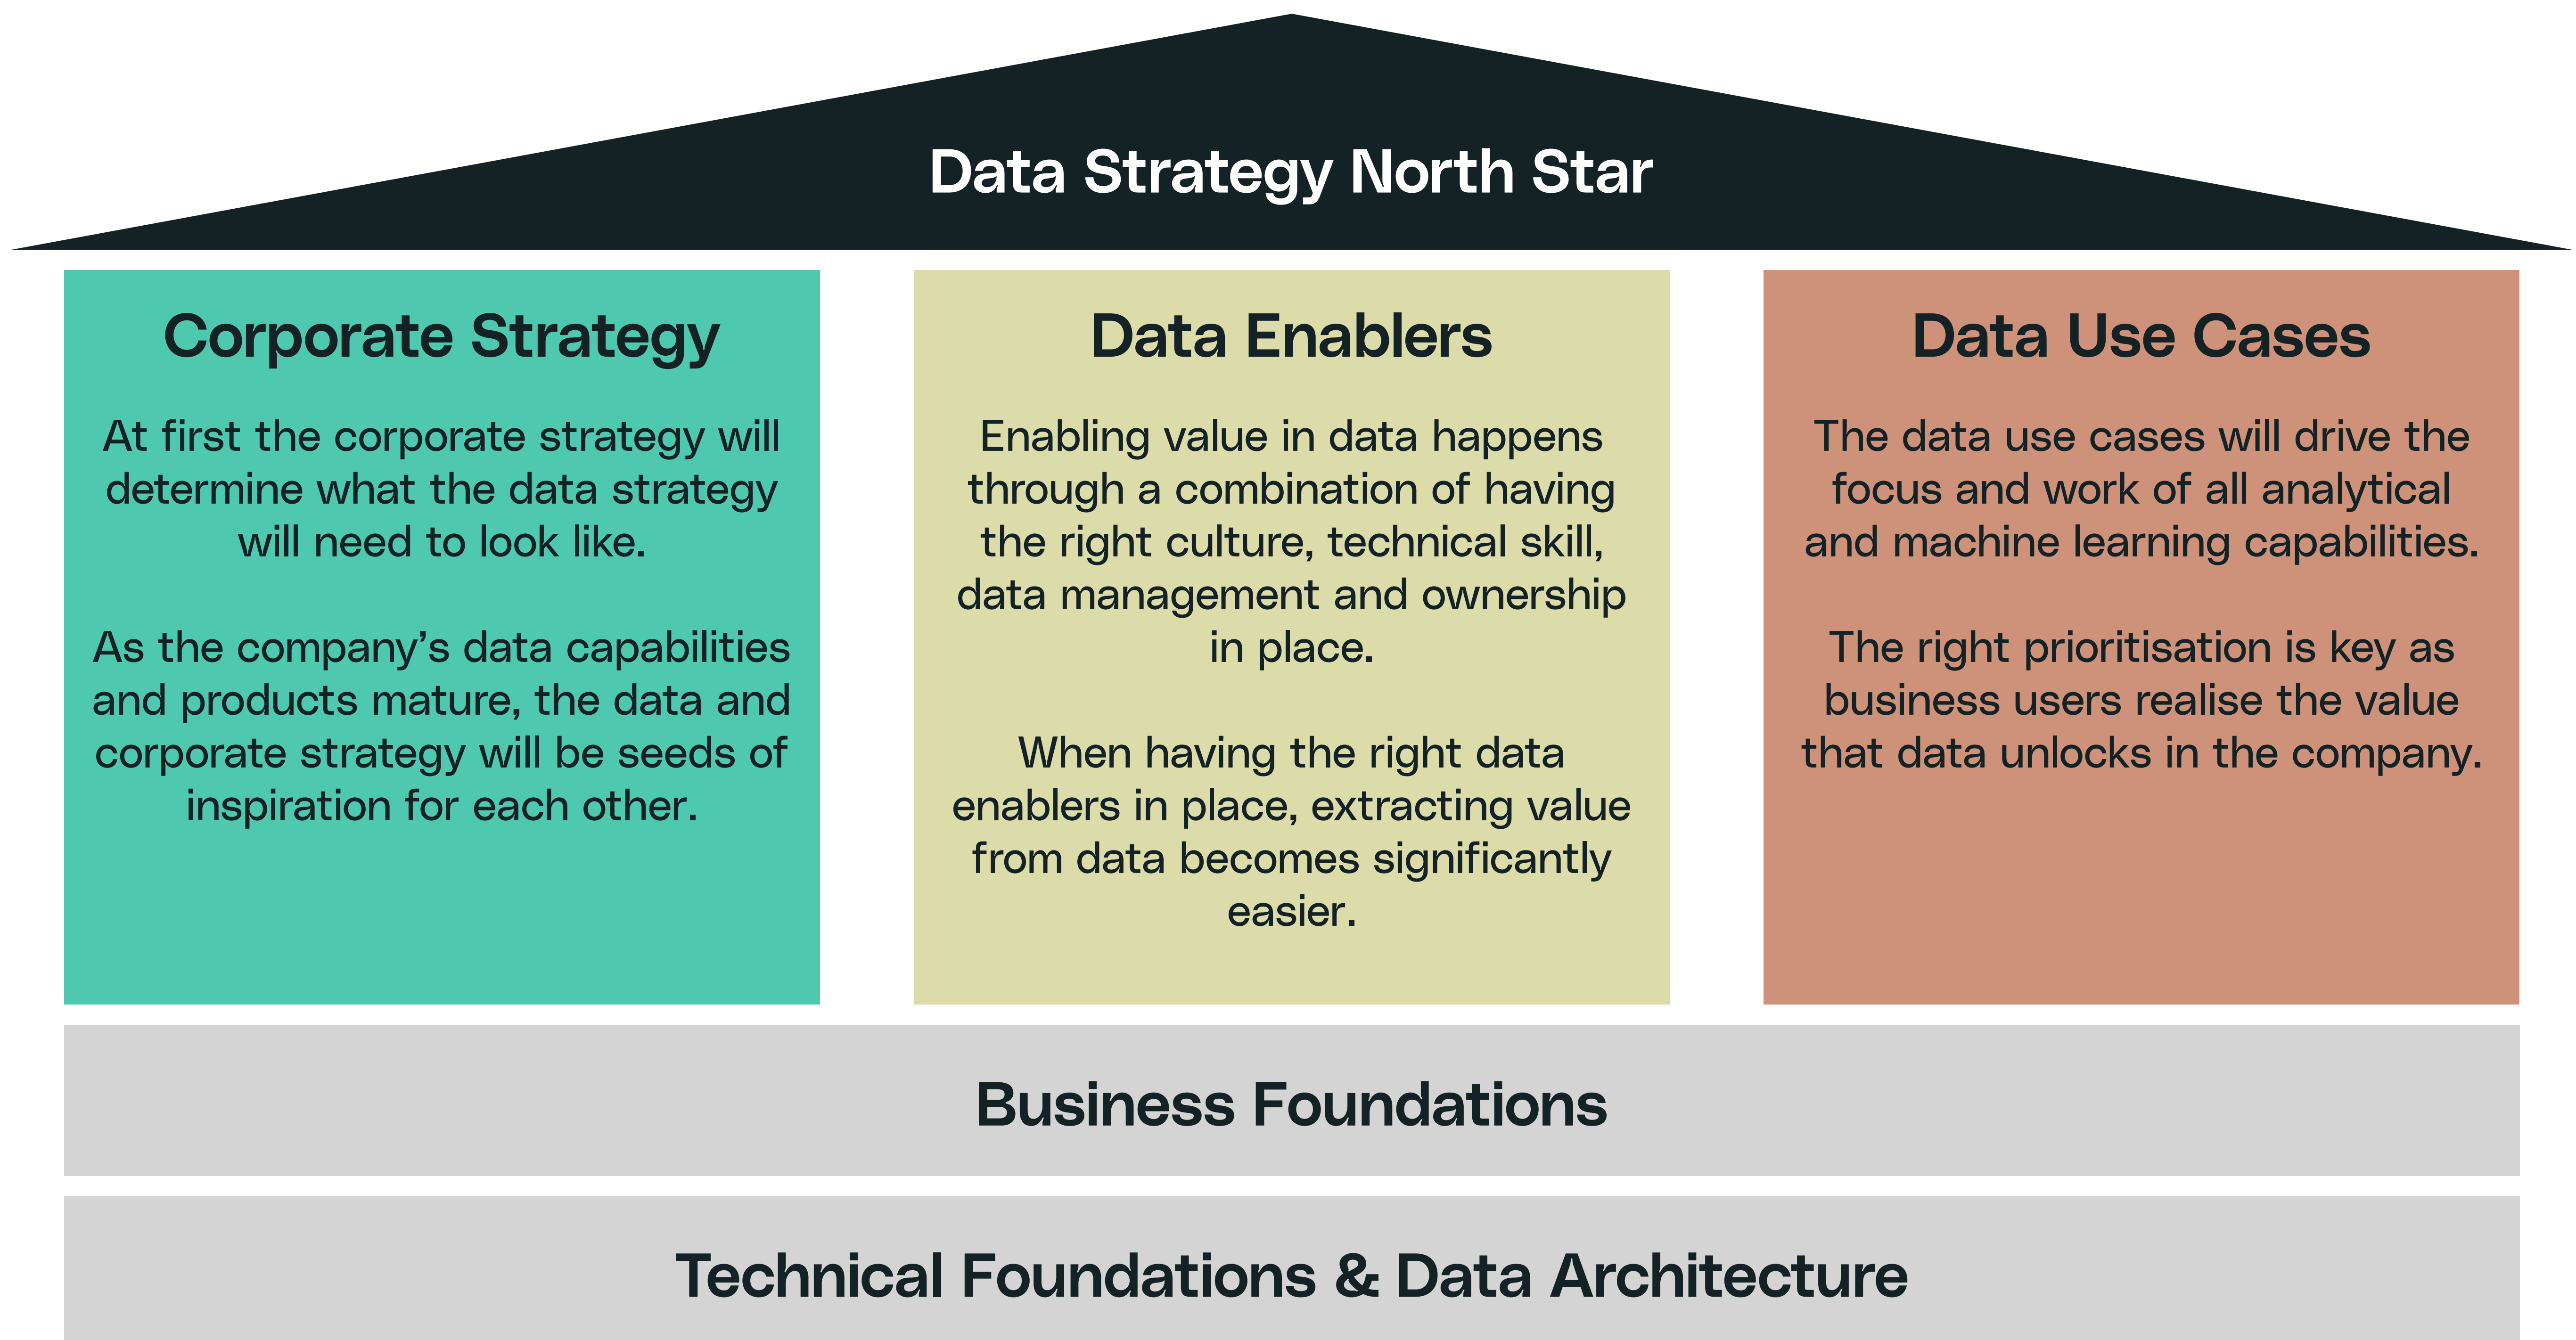
Technology

The backbone of any data & AI initiative is the technology stack. This includes investing in robust data storage solutions, such as data lakes or cloud platforms. AI algorithms require powerful computing resources, so organisations might need to invest in high-performance servers or cloud-based AI platforms.

Additionally, tools for data visualisation, analytics, and machine learning (like PowerBI, TensorFlow, or Azure ML) can empower teams to extract meaningful insights from vast data sets.

“Every business, regardless of size, is now a data business, every business, therefore, needs a robust data strategy.”

How it all comes together in Latitude's Data Strategy Temple.



Latitude's Data Strategy Temple presents a methodical model for orchestrating a data strategy. The 'Data Strategy North Star' embodies the ultimate objective, setting the course for the journey ahead. Beneath, The trio of pillars—Corporate Strategy, Data Enablers, and Data Use Cases—translates strategic intent into actionable data insights. Grounded by Business and Technical Foundations, the temple encapsulates a comprehensive strategy, translating strategic intent into actionable data insights and demystifying the intricate choreography of a successful data strategy.

Conclusion.

Data strategy stands as a cornerstone in guiding businesses towards their objectives. **Just as a sailor needs a compass, organisations require a clear data strategy** that aligns with their overarching corporate goals. Recognising the potential of data, and shifting its perception from a mere cost center to a profit generator, can unlock significant value. Whether aiming for top-line growth through offensive strategies or focusing on cost-saving defensive tactics, the power of data cannot be understated. However, capitalising on this potential mandates a structured approach: understanding the business's current data maturity, aligning corporate strategies with data aspirations, and addressing the essential triad of people, process, and technology. In a world driven by information, a well-crafted data strategy isn't just beneficial—it's indispensable.

SECTION 06

Architecting AI for Success: Building a Robust Data Infrastructure.

Driving value from your data strategy starts with the proper infrastructure.

In today's rapidly evolving digital landscape, the significance of a robust data infrastructure cannot be understated. AI systems fundamentally rely on vast amounts of data to function, making the underlying architecture crucial for their success. The way data architecture is approached can either accelerate or slow down your AI and data strategies.

Depending on the position of data in your corporate strategy it differs where and what the next step in the evolution of your data architecture should be.



1. Definition of Data Architecture and its Role in Supporting AI and Data Strategy

Data architecture refers to the overarching structure, design, and principles that govern the management, movement, and integration of data within an organisation. It encompasses a broad range of tools, technologies, and practices aimed at ensuring the consistent, accurate, and efficient handling of data.

In the age of AI and digital transformation, data architecture is paramount. AI systems, in their essence, rely on large amounts of data to learn, reason, and predict. A robust data architecture ensures this data is available, accurate, and primed for AI processing. **In strategic terms, the data architecture you choose either accelerates or slows down your AI and data initiatives.**

Data Mesh: A Paradigm Shift

The Data Mesh concept offers a solution to the challenges posed by traditional architectures. Promoting a decentralised approach, it encourages individual domains within an organisation to treat their data as a product. This shift ensures better governance, maintenance, and utility of data. However, its adoption necessitates a change in organisational mindset and might not be suitable for all entities¹⁸.

For a deep dive in the concept, we highly recommend the “Data Mesh” book as written by Zhamak Dehghani - the creator of the Data Mesh concept¹⁹.



Zhamak Dehghani

2. Key Steps and Considerations in Setting Up an Effective Data Architecture.



Needs Assessment

Understand the current and future data needs of your organisation. What kind of data do you handle? What will your data needs look like in three years? These questions can guide architectural decisions.




Data Governance

Establish clear policies around data collection, usage, and sharing. Who has access? How is data quality maintained? A strong governance framework promotes trust and consistency.



Scalability

As your company grows, so too will your data needs. Your architecture must scale without causing service disruptions or significant reworks.




Integration and Interoperability

As businesses adopt various software solutions, data must flow seamlessly between them. Your architecture should promote easy integration and interoperability.



Security

Protecting your data is critical. Employ encryption, access controls, and regular audits to ensure your data remains in safe hands.



Future-proofing

Technology evolves rapidly. Opt for solutions that are adaptable and that can embrace future advancements, particularly in the AI sector.

3. The Impact of a Well-Planned Data Architecture on AI Implementation

A solid data architecture acts as **the backbone for successful AI implementations**. Here are concrete examples illustrating the benefits:

01

Real-time Decision Making

Companies like Uber and Lyft use data architectures that process millions of data points in real-time. This allows their AI algorithms to instantly match riders with drivers, predict ETAs, and optimise routes.

02

Enhanced Customer Experience

Netflix's recommendation engine thrives on its data architecture. By efficiently processing user viewing habits, ratings, and preferences, the AI system curates personalised content lists, improving user engagement and satisfaction.

03

Supply Chain Optimisation

Companies like Amazon and Walmart rely on sophisticated data architectures to predict product demand, optimise inventory levels, and streamline distribution. Their AI models process vast datasets to make accurate predictions, saving costs and improving service.

04

Fraud Detection

Financial institutions utilise AI models to detect fraudulent activities. With a robust data architecture, these models can analyse years of transaction data in moments, flagging suspicious activities with high accuracy.

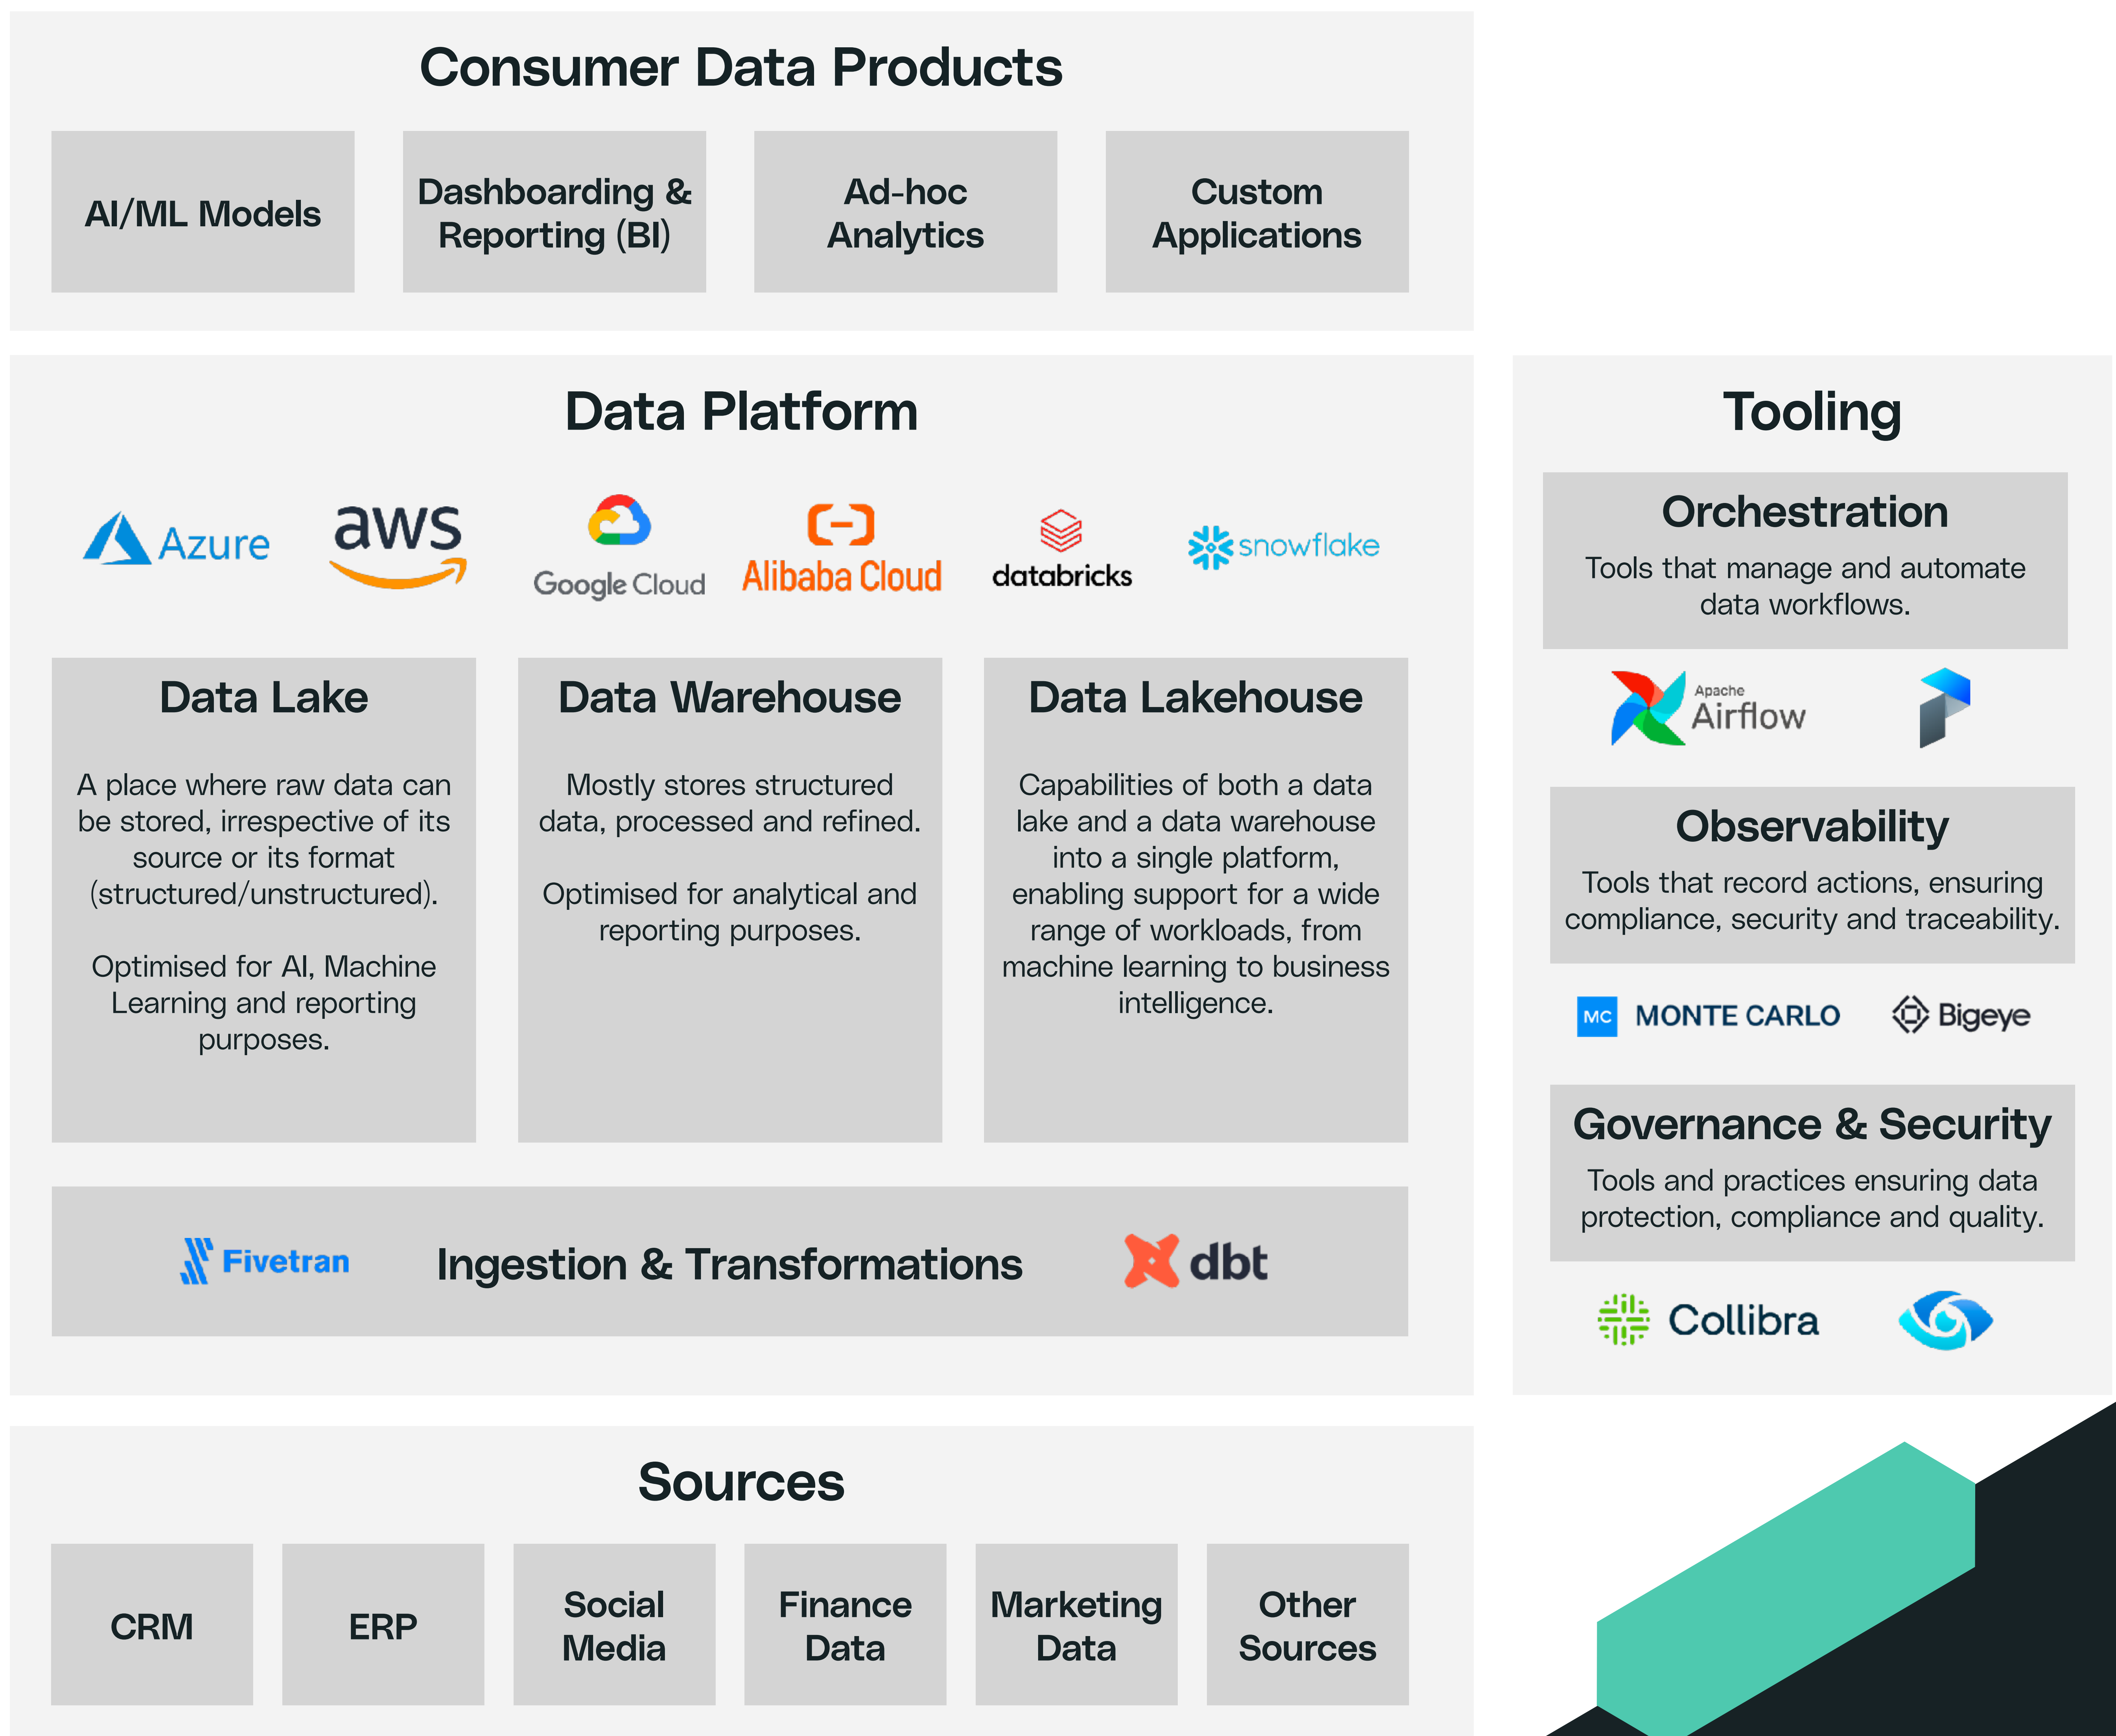
The Emergence of Data Platforms and Their Future Implications

The rise of "data platforms" is reshaping the landscape of data infrastructure, with systems like Snowflake and Databricks potentially becoming foundational cornerstones. These platforms not only streamline current data operations but also hint at a transformative future for data applications and architectures. As the industry continues to evolve, recognising and harnessing the potential of these emerging platforms will be pivotal for any data-driven organisation²⁰.

The modern dataplatform.

A well-designed data architecture is crucial for unlocking the complete potential of data and AI, guaranteeing smooth integration, scalability, and adaptability. The Reference Data Architecture Diagram showcased below provides a schematic representation for organising, managing, and leveraging data in AI-centric solutions.


The tools marked with their respective icons represent prevalent practices, acknowledging that the list is not exhaustive and the sector is swiftly evolving. Nonetheless, these tools and platforms have demonstrated their worth and enjoy widespread recognition across various industries.






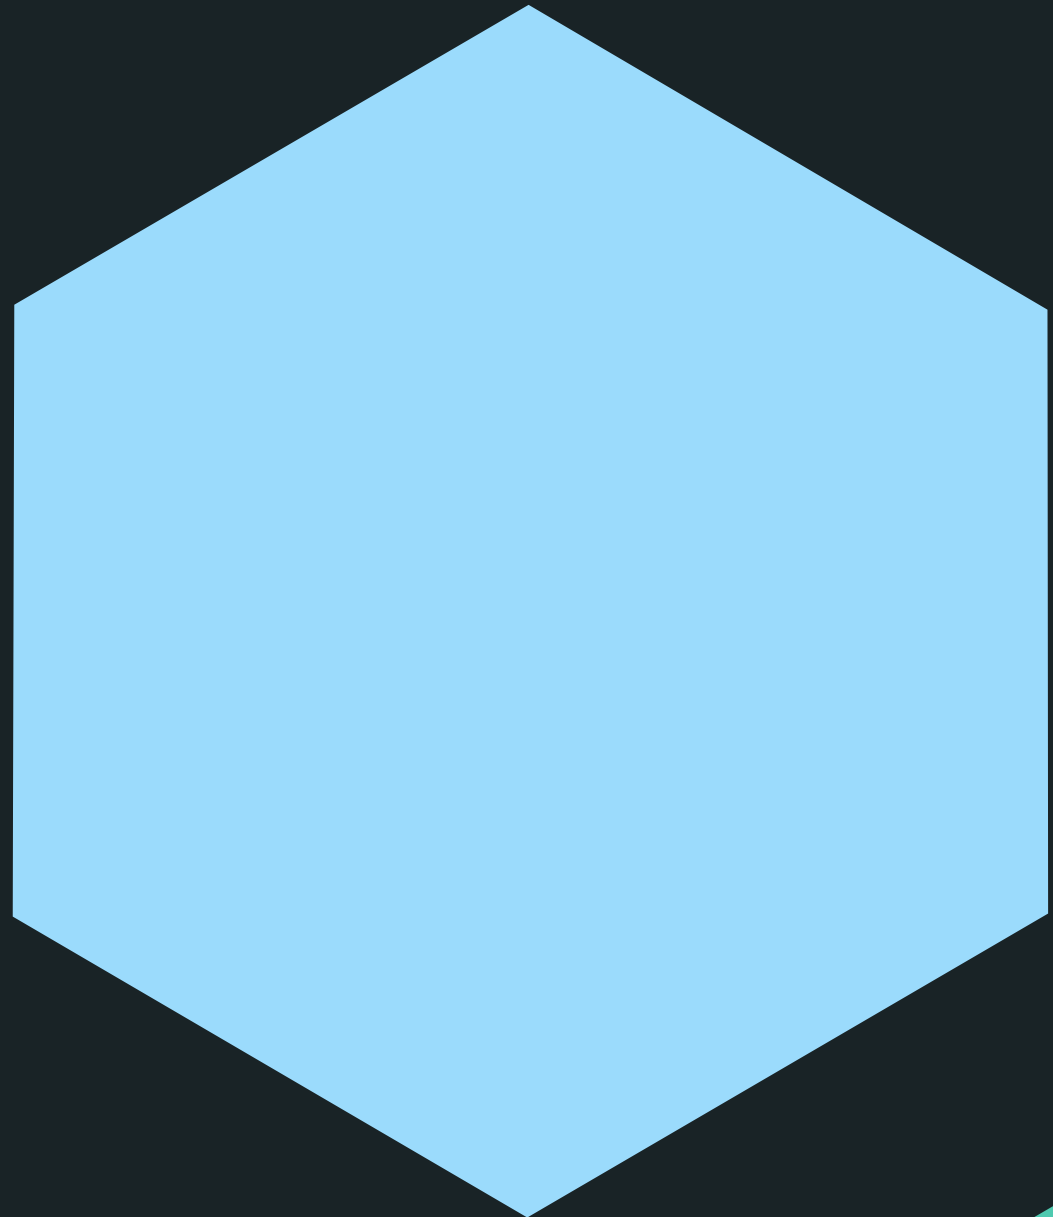
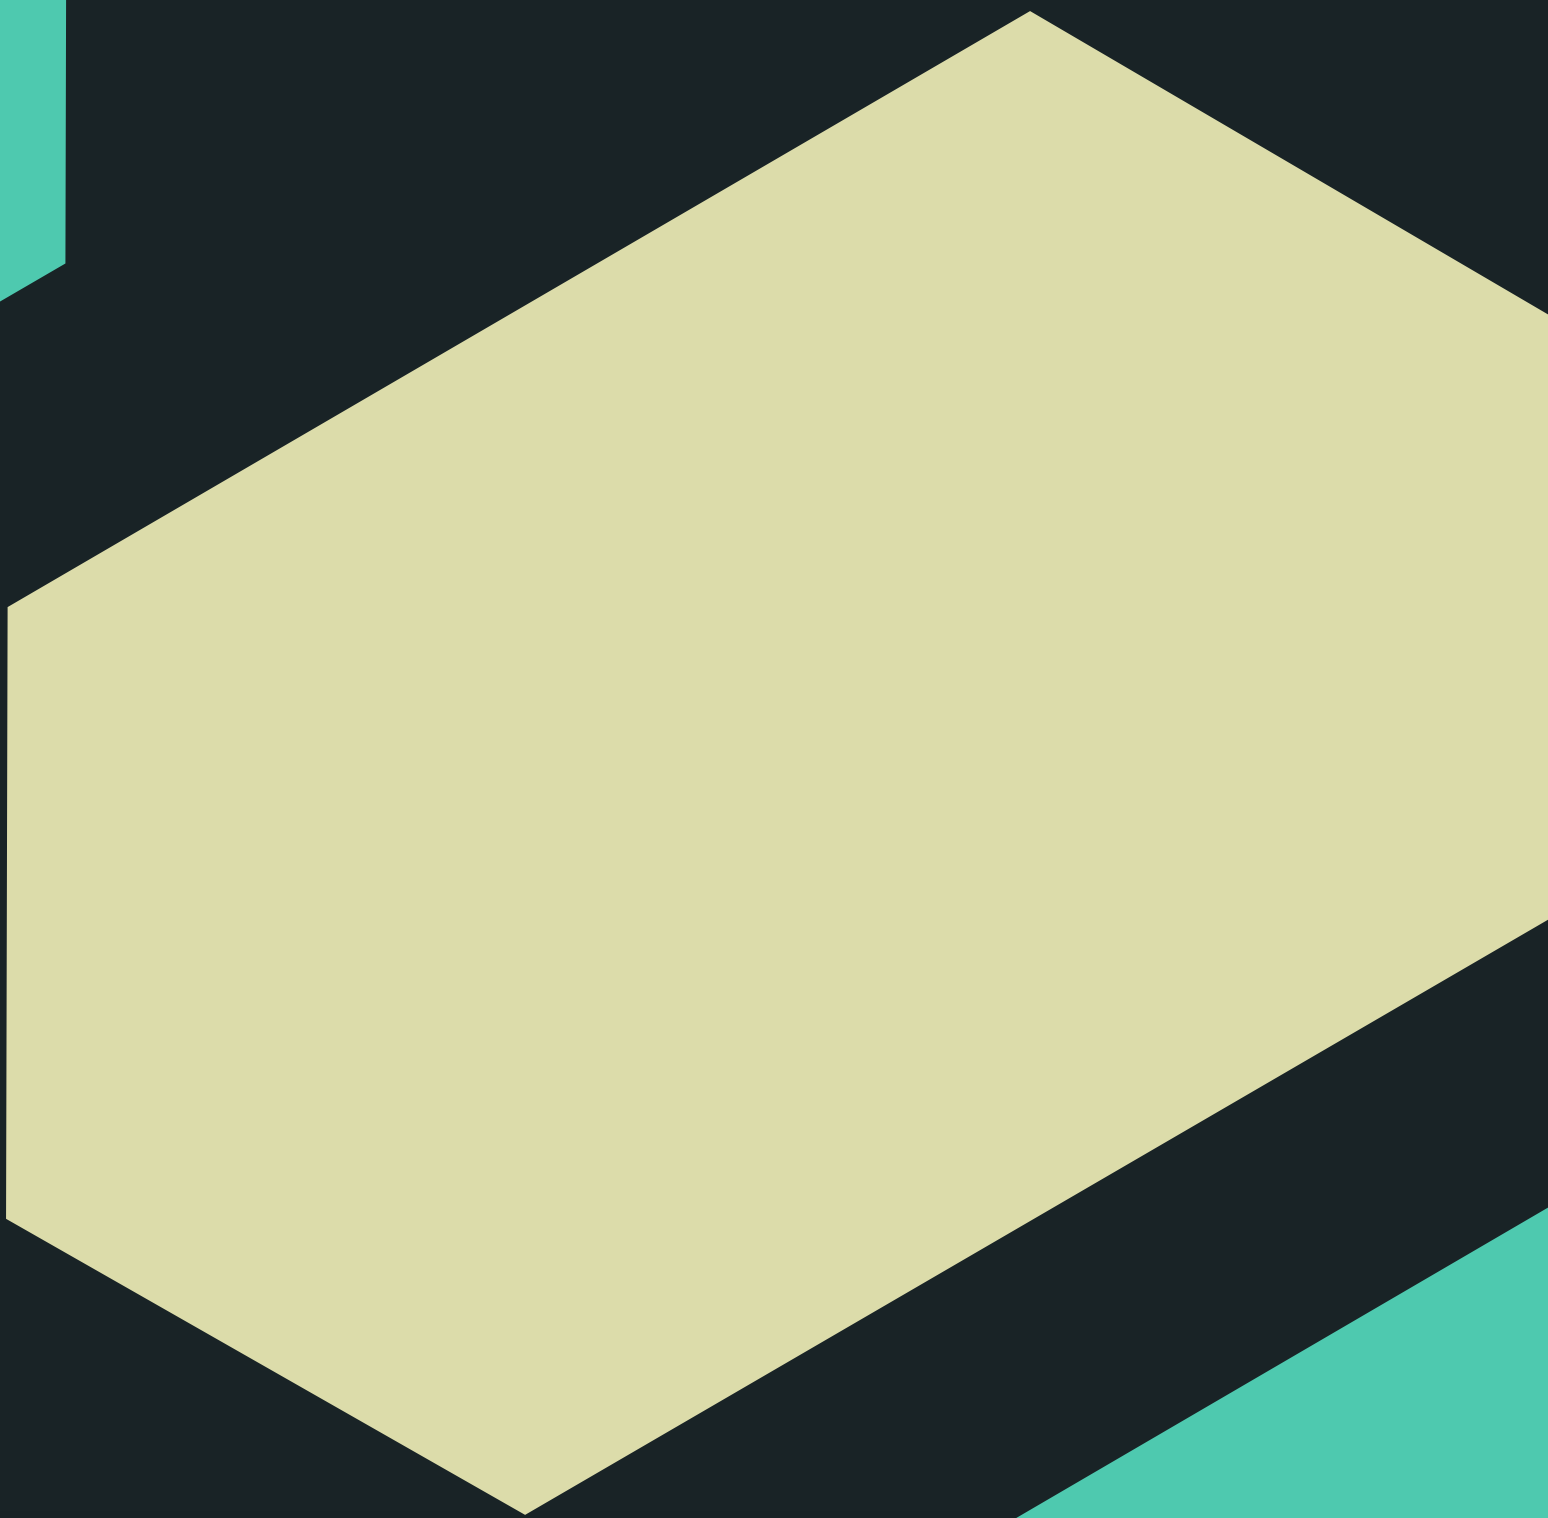
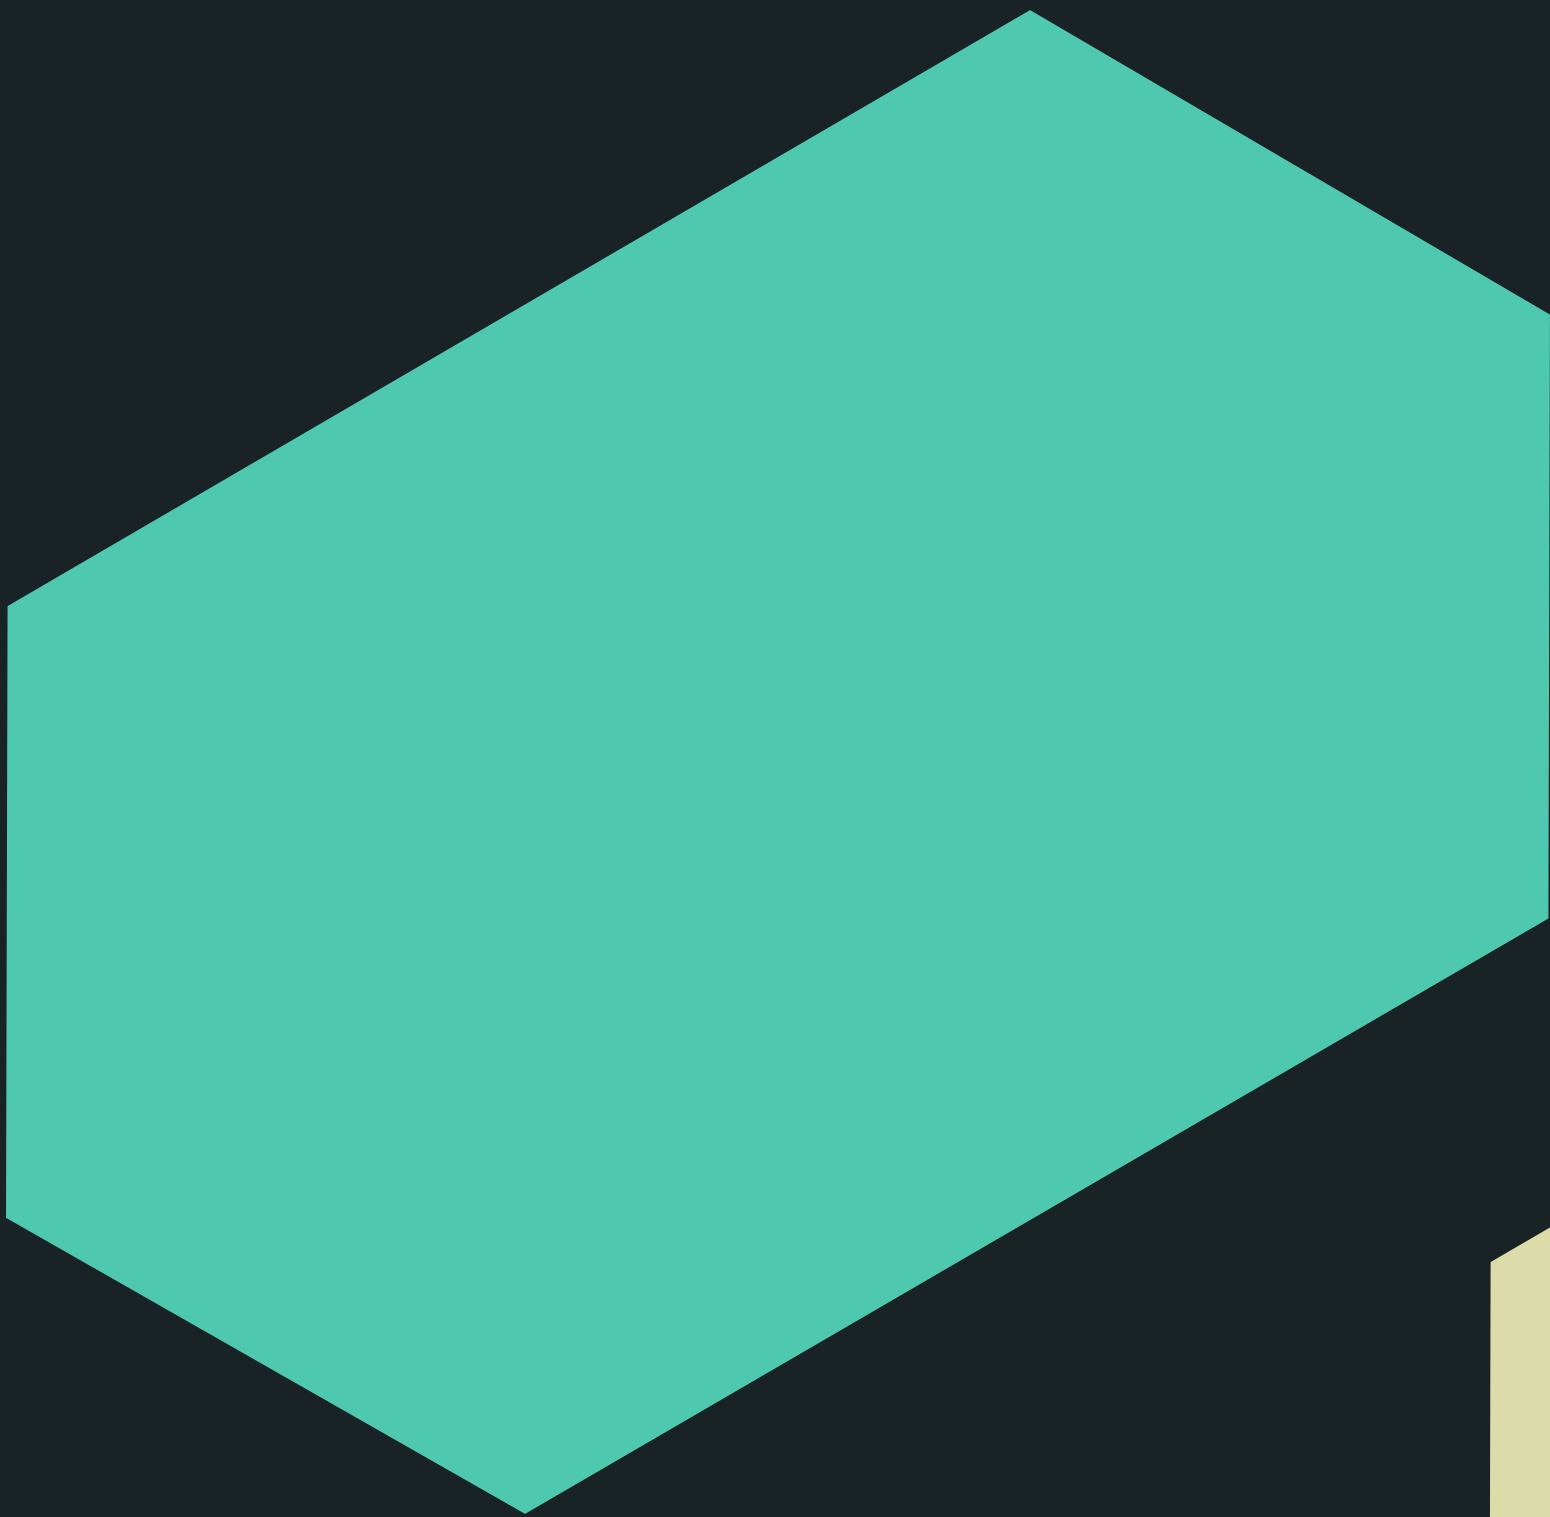

Conclusion.

In the AI-focused business realm, the right data infrastructure is critical. This foundation supports efficient AI operations and adapts to future demands. It guarantees data safety, system integration, and paves the way for real-time analytics. With evolving tools like Snowflake and Databricks, businesses should actively refine their structures. In sum, a strong data strategy and infrastructure are vital for maximising AI's potential.

Concluding words.



In this whitepaper, we've illuminated the transformative potential of AI, emphasising its power not just as a standalone technology but as an integrated solution within a strategic framework. The key point is in understanding and addressing genuine business problems, underpinned by a robust data strategy and architecture. As organisations navigate the AI landscape, it's important to remember that successful AI adoption hinges on a blend of vision, solid data infrastructure, and a problem-centric approach, ensuring that every AI endeavour is both purposeful and impactful.



**“We boost Europe’s
leading organisations
by infusing AI into the
core of their being”**

About the Authors.

Gianmarco Disario

As one of the founders of Latitude, Gianmarco is a multifaceted data scientist/machine learning engineer. But Gianmarco's passion isn't just about numbers and algorithms. He approaches every problem with a strategic mindset and a love for sharing knowledge and building relationships

His experience is broad and diverse. Gianmarco has collaborated with global brands, including Adidas, where he helped them steer their pricing strategies through the Corona pandemic with the use of time series analysis, deep learning and causal inference. As the pioneering Machine Learning Engineer at Otrium, he was instrumental in setting up their Data Science and Machine Learning strategies. His work there was pivotal, from driving smart pricing engines to crafting strategies for company-wide KPIs.

His combination of technical prowess, strategic thinking, and a genuine desire to connect makes Gianmarco a linchpin in any data-driven initiative.



Jelle Stienen

Jelle Stienen - the other founder of Latitude - is an expert in Data & AI, leading the way at Latitude, dedicated to harnessing the power of business data. With hands-on experience working with major brands like T-Mobile, Adidas, and Polaroid, he's provided strategic insights, set up data-focused departments, and led innovative data science & engineering projects.

Starting Latitude, Jelle aimed to help organisations not just collect, but truly leverage their data, always with the end-user in mind. With him in charge, data transforms from mere numbers to a strategic asset, driving real growth and change.

With a history of partnering with industry giants and a knack for practical solutions, Jelle ensures Latitude delivers clarity and tangible outcomes.



Need Help?

Thinking about your next steps with AI? The landscape is vast, and the benefits are clear: turning dormant data into profit, optimising operations, and staying ahead. But diving into AI should be a journey of clarity, not complexity.

At Latitude, we prioritise simplicity and purity, ensuring a tailored approach that aligns with your unique needs. We provide insights and solutions that truly resonate with your strategic needs. Reach out and let's get started.

Get in touch with the authors by scanning the Whatsapp QR codes

Reference List.

1. McKinsey & Company. (2022, December 6). The state of AI in 2022--and a half decade in review | McKinsey. [www.mckinsey.com. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review](https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review)
2. Stanford University. (2023, April). The AI Index Report – Artificial Intelligence Index. [Aiindex.stanford.edu. https://aiindex.stanford.edu/report/](https://aiindex.stanford.edu/report/)
3. Gartner Survey Finds 79% of Corporate Strategists See AI and Analytics as Critical to Their Success Over the Next Two Years. (2023, July 5). Gartner. <https://www.gartner.com/en/newsroom/press-releases/2023-07-05-gartner-survey-finds-79-percent-of-corporate-strategists-see-ai-and-analytics-as-critical-to-their-success-over-the-next-two-years>
4. Economic potential of generative AI | McKinsey. (2023, June). [www.mckinsey.com. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier/](https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier/)
5. Introducing AudioCraft: A Generative AI Tool For Audio and Music. (2023, August 2). Meta. <https://about.fb.com/news/2023/08/audiocraft-generative-ai-for-music-and-audio/>
6. Wang, M. C., Sarah. (2023, August 3). The Economic Case for Generative AI and Foundation Models. Andreessen Horowitz. <https://a16z.com/the-economic-case-for-generative-ai-and-foundation-models/>
7. Khullar, D. (2023, February 27). Can A.I. Treat Mental Illness? The New Yorker. <https://www.newyorker.com/magazine/2023/03/06/can-ai-treat-mental-illness>
8. Alphonso, G. (2023, March 3). Council Post: Generative AI: Education In The Age Of Innovation. Forbes. <https://www.forbes.com/sites/forbestechcouncil/2023/03/03/generative-ai-education-in-the-age-of-innovation/>
9. Skates, A. L. (2023, February 8). Five Predictions for the Future of Learning in the Age of AI. Andreessen Horowitz. <https://a16z.com/the-future-of-learning-education-knowledge-in-the-age-of-ai/>
10. Gurman, M. (2023, May 2). Samsung Bans Generative AI Use by Staff After ChatGPT Data Leak. Bloomberg.com. <https://www.bloomberg.com/news/articles/2023-05-02/samsung-bans-chatgpt-and-other-generative-ai-use-by-staff-after-leak>
11. Davenport, T., & Alavi, M. (2023, July 6). How to Train Generative AI Using Your Company's Data. Harvard Business Review. <https://hbr.org/2023/07/how-to-train-generative-ai-using-your-companys-data>
12. Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). BloombergGPT: A Large Language Model for Finance. ArXiv:2303.17564 [Cs, Q-Fin]. <https://arxiv.org/abs/2303.17564>
13. Van Noorden, R. (2023). ChatGPT-like AIs are coming to major science search engines. Nature, 620(7973), 258–258. <https://doi.org/10.1038/d41586-023-02470-3>
14. Radovanovic, M. B., Rajko. (2023, June 20). Emerging Architectures for LLM Applications. Andreessen Horowitz. <https://a16z.com/emerging-architectures-for-llm-applications/>
15. GPT-3.5 Turbo fine-tuning and API updates. (2023, August 22). Openai.com. <https://openai.com/blog/gpt-3-5-turbo-fine-tuning-and-api-updates>
16. Business, R. M., CNN. (n.d.). Zillow's home-buying debacle shows how hard it is to use AI to value real estate. CNN. <https://edition.cnn.com/2021/11/09/tech/zillow-ibuying-home-zestimate/index.html>
17. Dastin, J. (2018, October 11). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
18. Disario, G. (2022, October 6). Is Data Mesh the solution for all? Medium. <https://medium.com/@gedisario/is-data-mesh-the-solution-for-all-1109dd002847>
19. Dehghani, Z. (2022). Data mesh : delivering data-driven value at Scale. O'reilly Media.
20. Casado, M. B., Jennifer Li, Martin. (2020, October 15). Emerging Architectures for Modern Data Infrastructure. Andreessen Horowitz. <https://a16z.com/emerging-architectures-for-modern-data-infrastructure/>