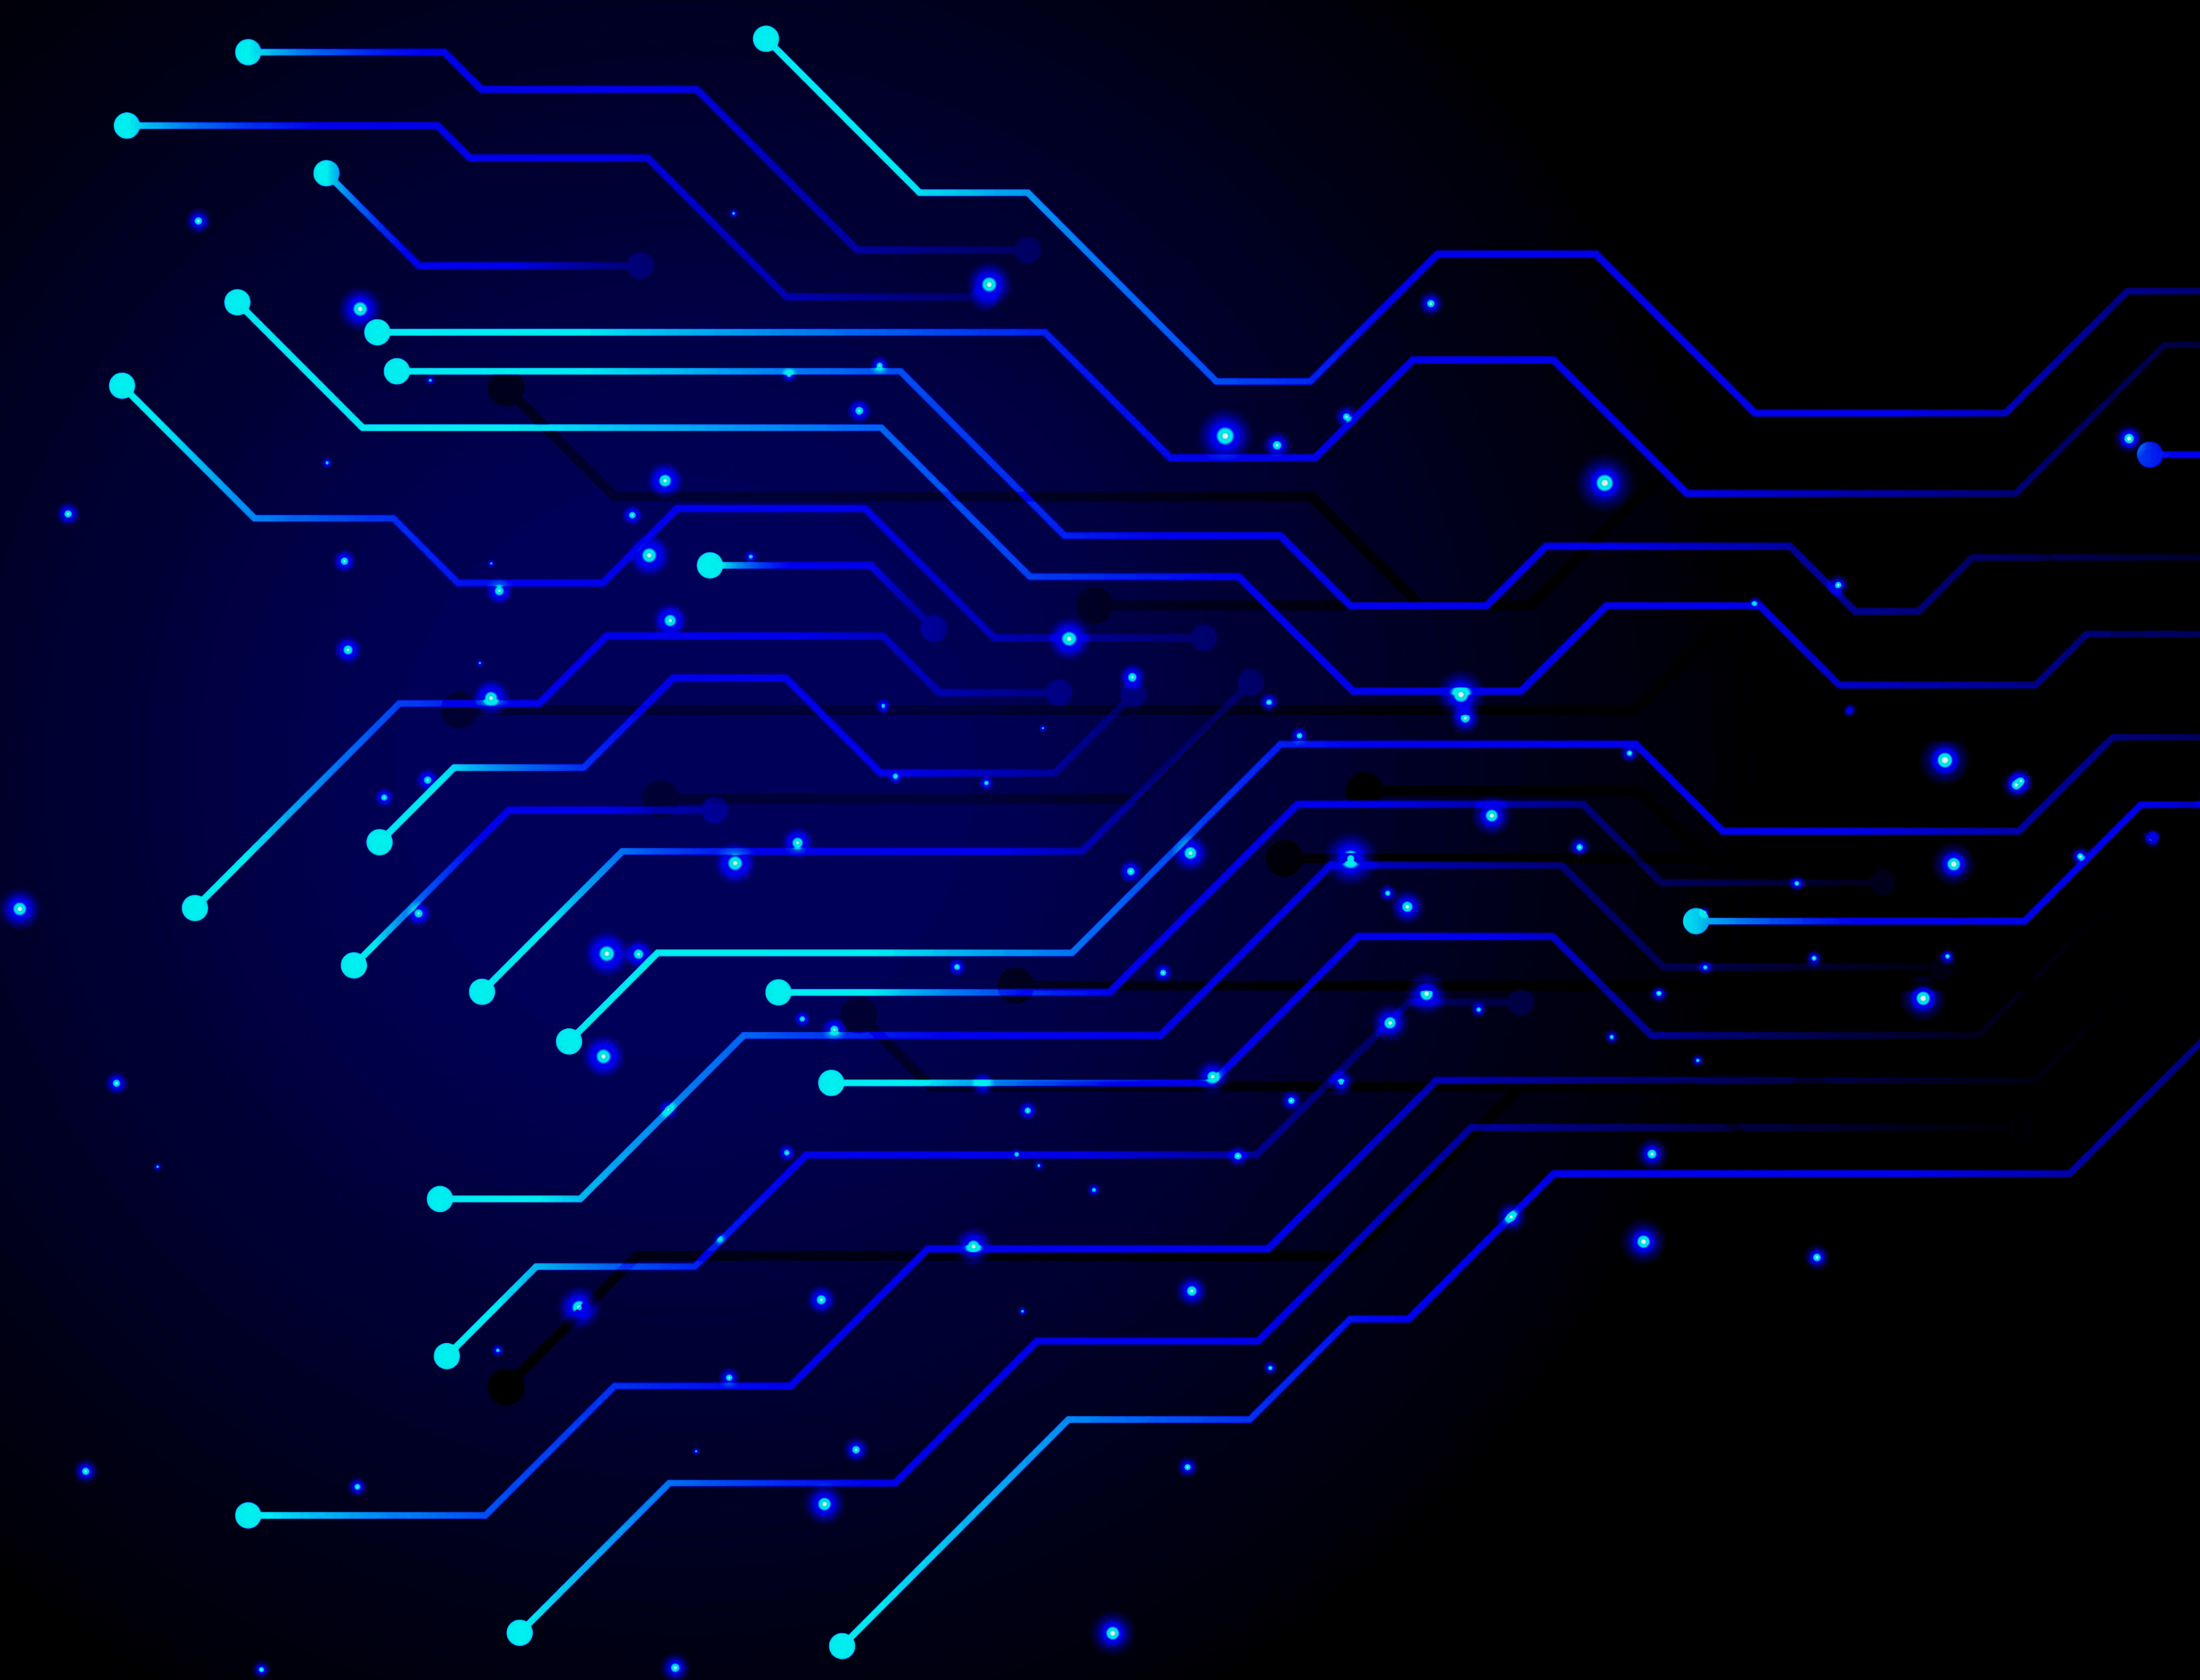

Artificial Intelligence

Adoption, Accessibility and Management

Authored By
Abhro Pal, Shobhit Gupta, Alexis Goodrich, Prof. Shalendra Jain

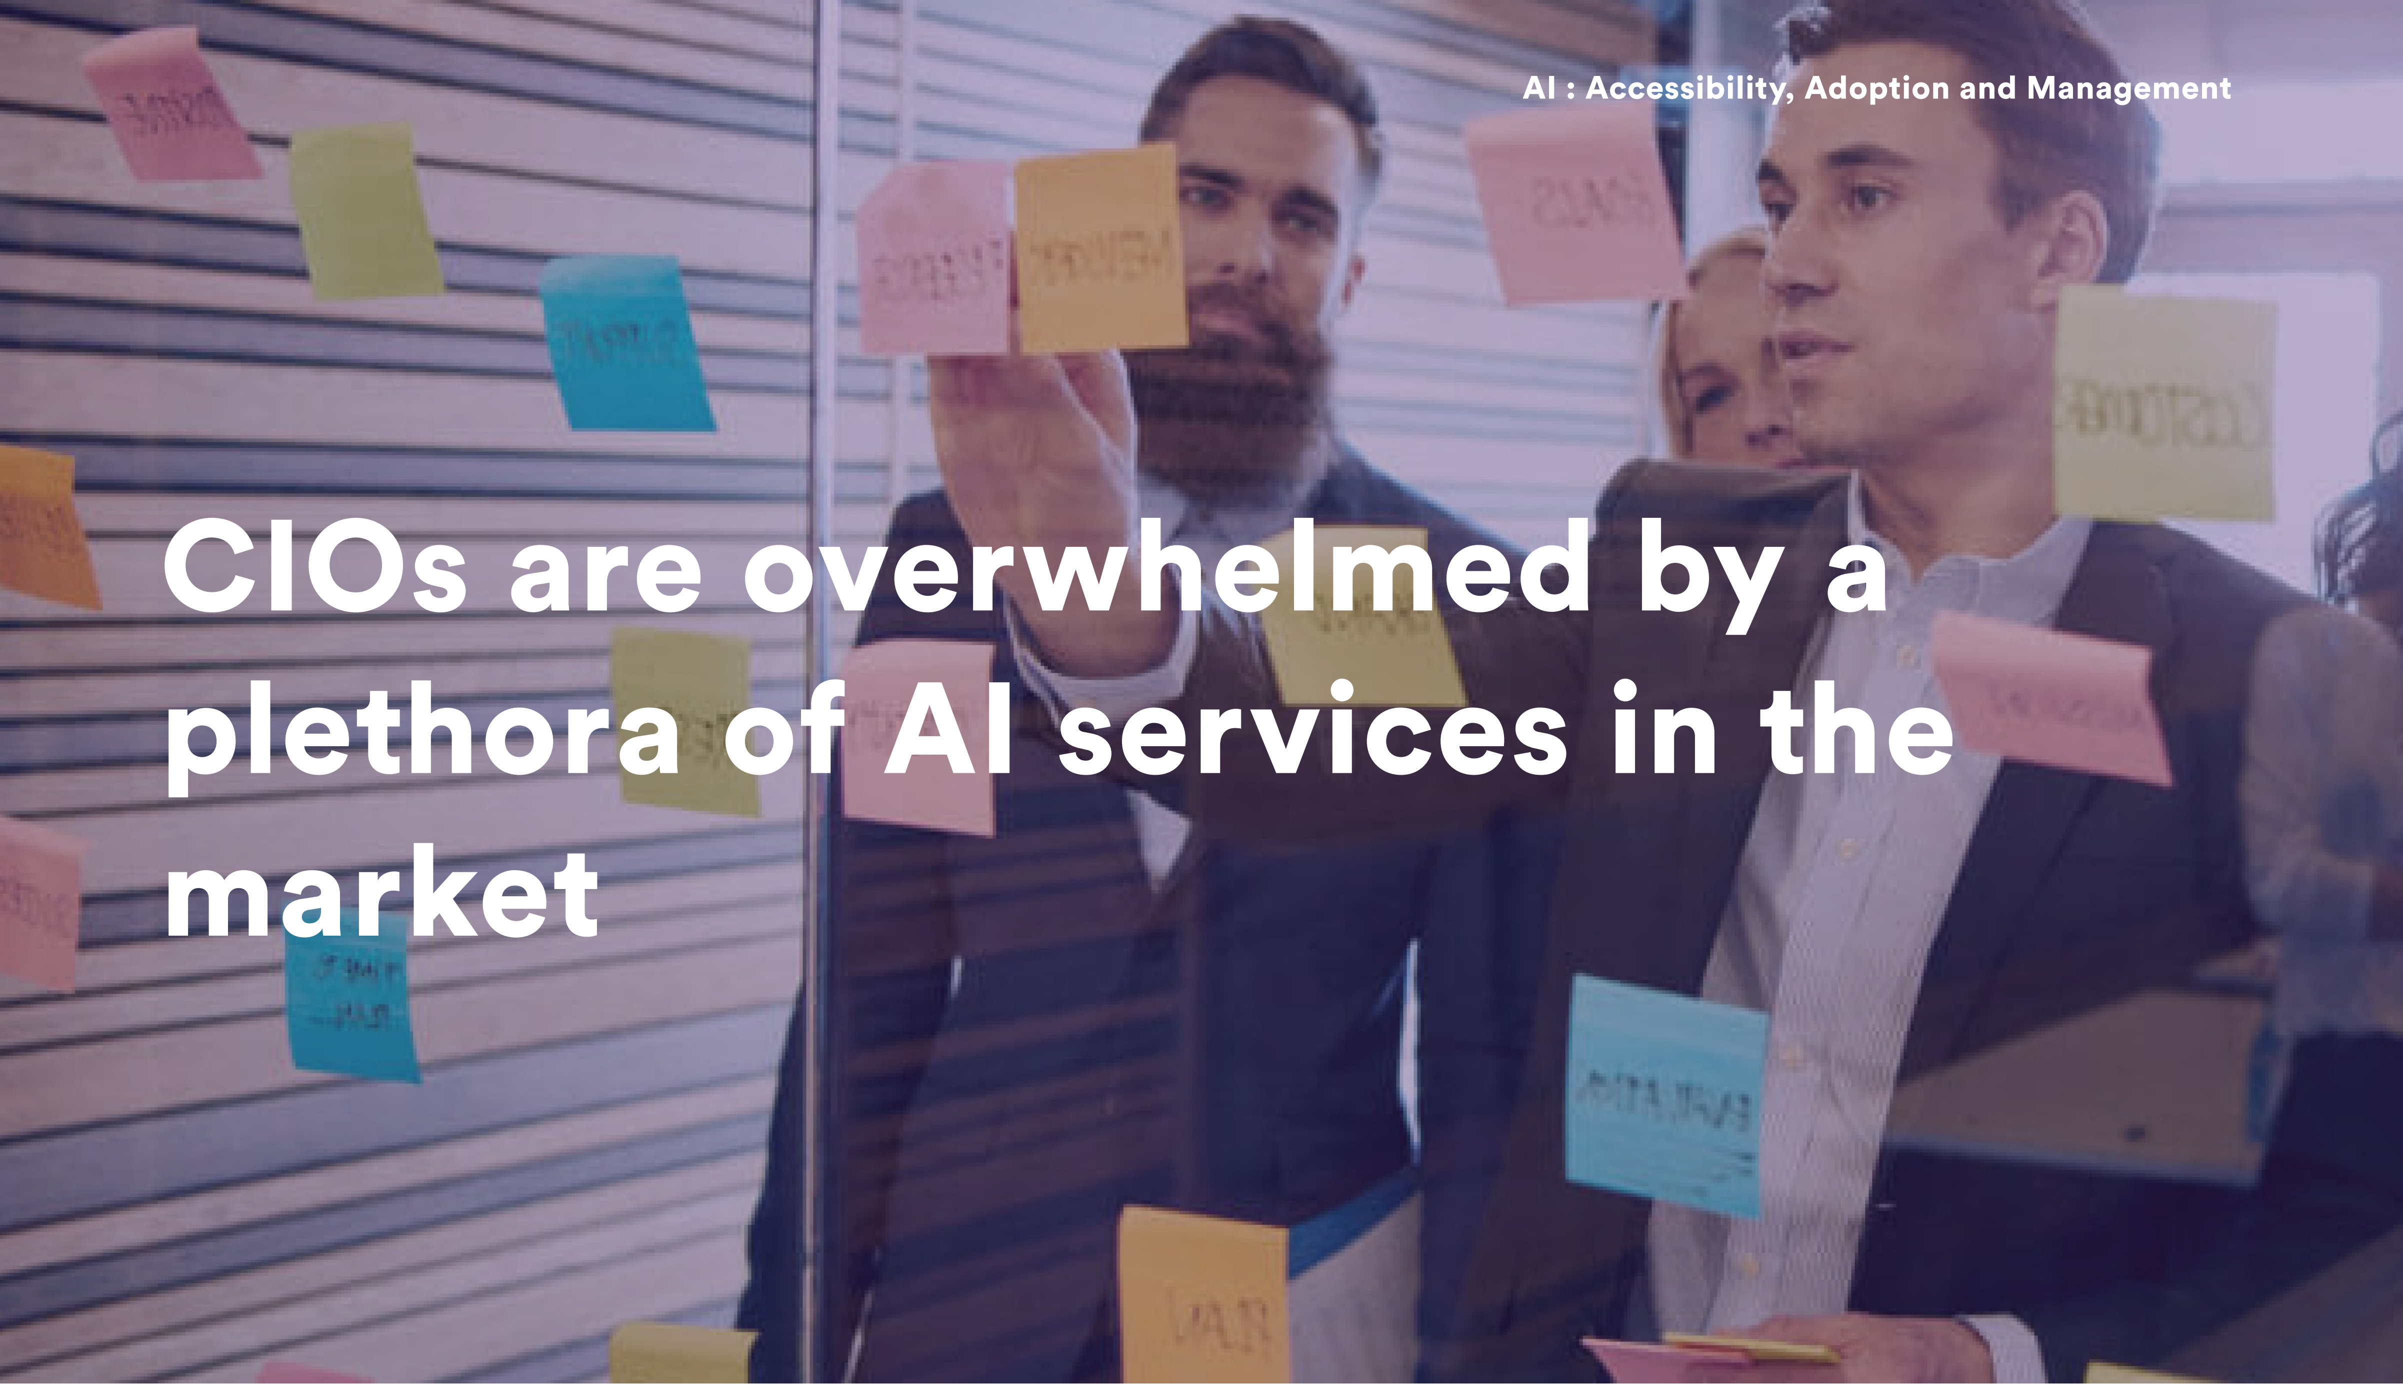

CIOs are overwhelmed by a plethora of AI services in the market

Leading Technology companies are increasing their AI services and products at breathtaking speed

AWS, Azure, and GCP combined have 300+ services in the market. While the technology companies are experimenting to move the speed of digital transformation from a horse cart to a Ferrari, business customers are racing to define where they want to go with it. In other words, there is a chance of technology myopia which can only be addressed if the customers have a clearer understanding of the current and potential offerings and their business value.

Amongst various software services in the market, the hardest to understand today is Artificial Intelligence as a Service (AIaaS). Like Software as a Service (SaaS), AIaaS is a ‘cloud-first’ approach, with minimal internal support needed. For example, a website can integrate an AI chatbot functionality to provide customer service using a provider (e.g.Meya.ai) without writing a line of code. But unlike SaaS, AIaaS needs continuous feedback and data from the domain experts to operate.

Why are the CIOs Struggling?

Build or Buy dilemma – which solutions should be built in-house and which ones should be bought off the shelf and/or outsourced?AI services are new and complex to understand, thereby making it hard to navigate the dilemma. Market is currently overwhelmed with AI products and services. Additionally, there are AI startups providing solutions customized to industry needs.

An implementation strategy - Once the build vs buy decision has been made, CIOs need an adoption roadmap. The limited AI talent in the market, ambiguity to articulate the business value, and use of legacy software applications make it costly and time-consuming to adopt AI.

How can this report help?

We provide a framework that simplifies the build vs buy dilemma faced by a CIO as well as a case study to showcase the implementation strategy. Our conceptualization is based on interviews with 10 key stakeholders over zoom for 60 minutes each in AI startups, AI research, venture capital firms, academia, and technology giants (Madrona Venture Group, Phaidra, University of Washington, Microsoft, and Amazon AWS).

These stakeholders brought unique perspectives shaped by their backgrounds. Venture capitalists helped us gain insights on the long-term bets being made on AI.

42%

CIOs responded that AI will have strategic importance to them by as early as 2021.

AI will become one of the core competencies of all Fortune 500 companies in the next 3 to 5 years.

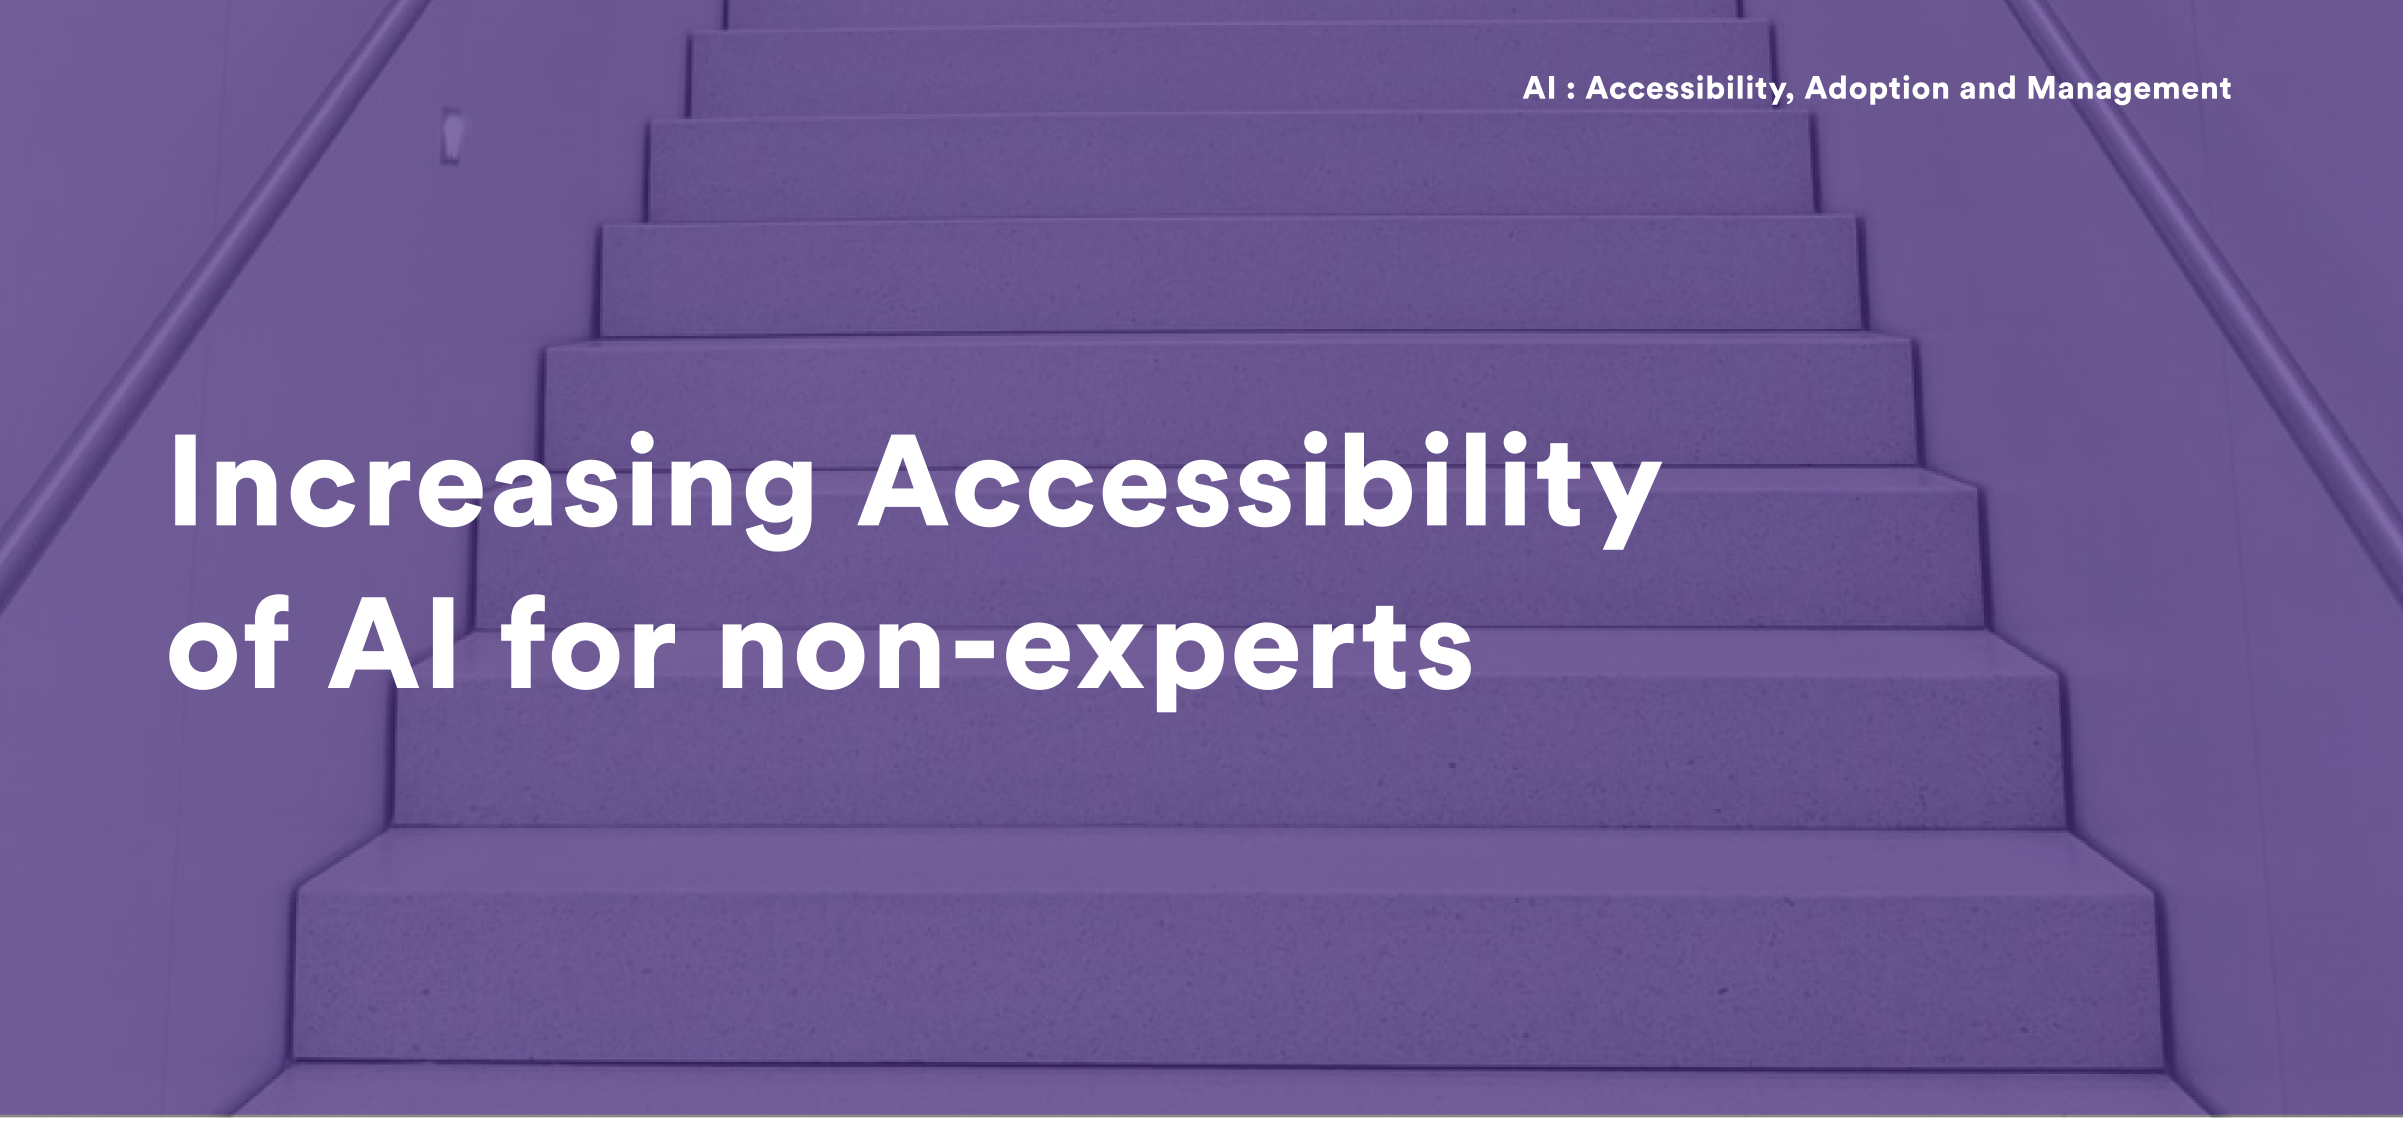

Increasing Accessibility of AI for non-experts

Offerings subsumed within **'AI as a Service' (AIaaS) can be categorized into four mutually exclusive and collectively exhaustive** buckets which are essential components of building an AI application

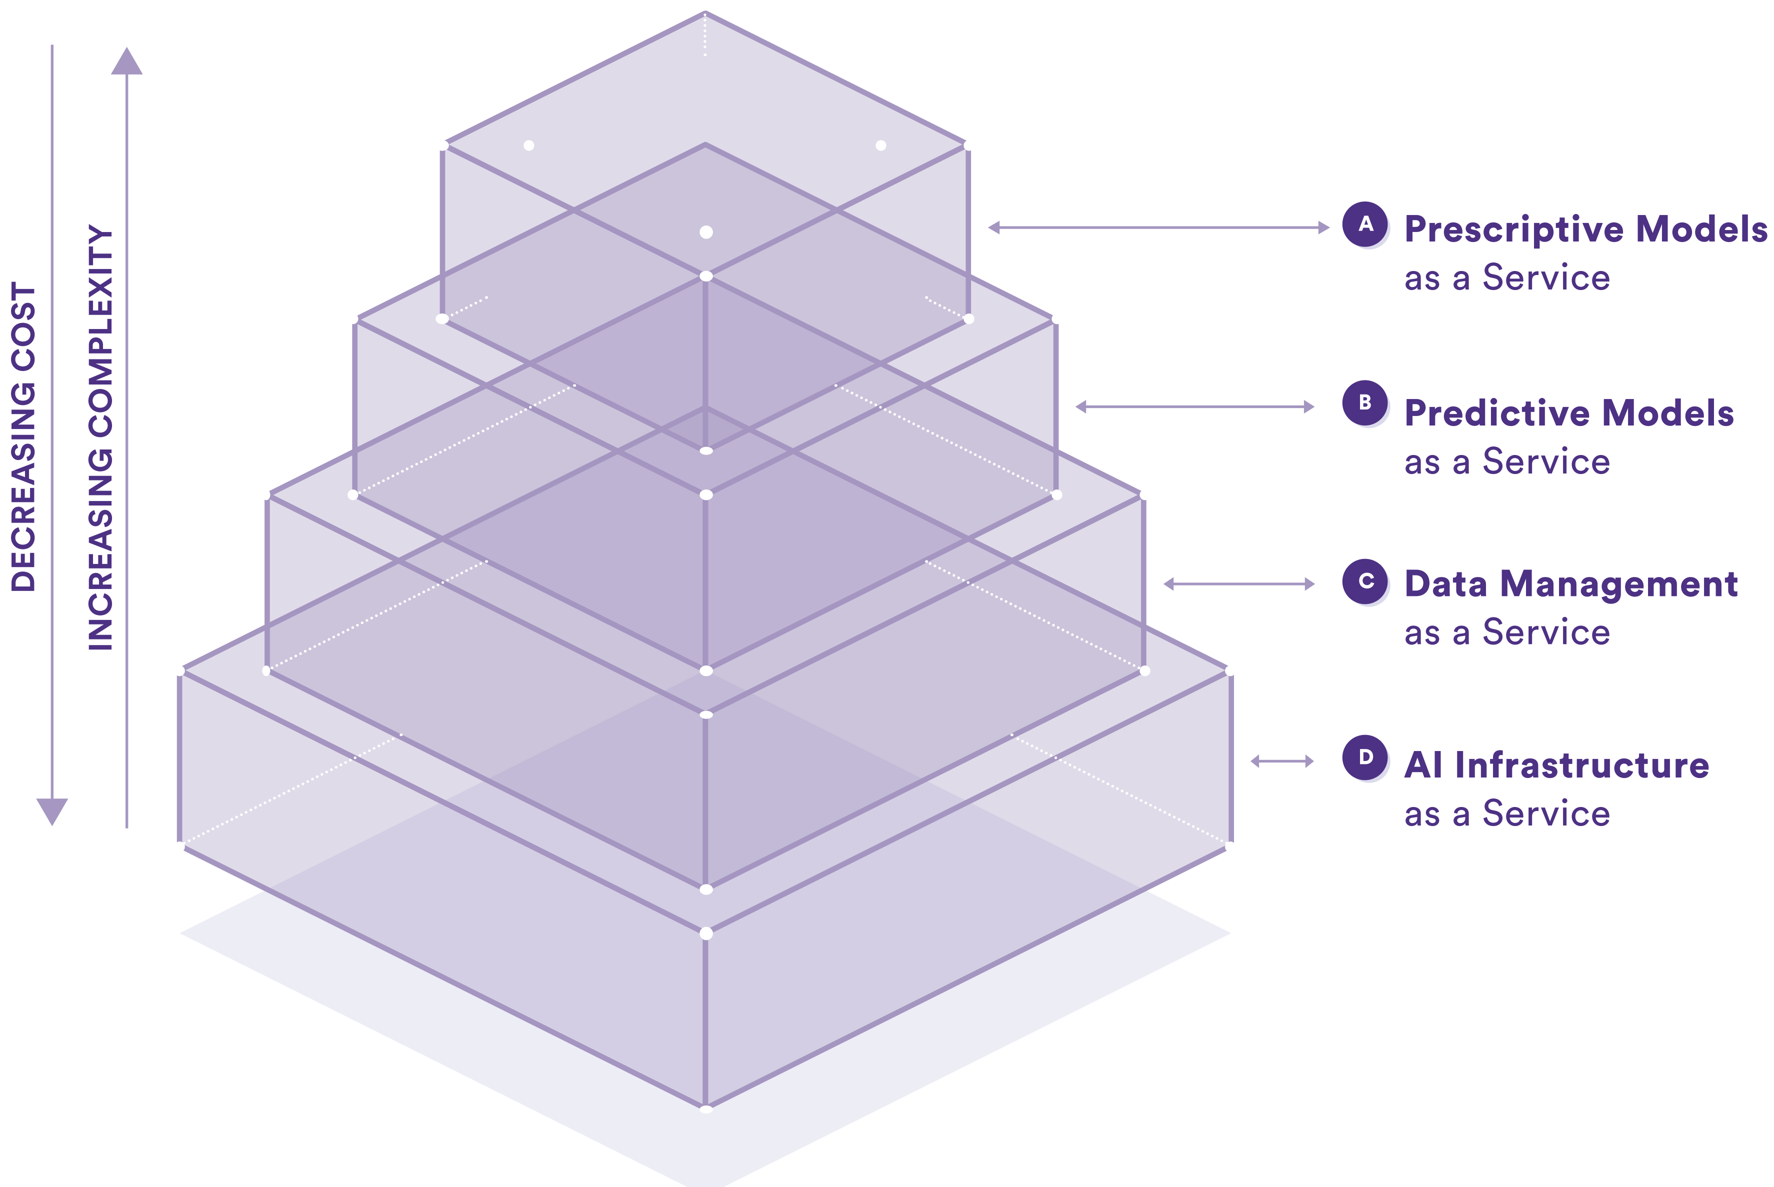

Fig 1: Pyramid Framework to analyze AI as a Service marketplace

What does each model mean?

AI Infrastructure as a Service (IaaS)

IaaS consists of cloud computing resources such as compute, network, storage and development platforms that are customized to build and deploy AI applications at scale. For example, Amazon Sagemaker is a fully-managed service that provides developers and data scientists a platform to build, train and deploy ML models on a pay-as-you-go model. While traditional software requires cost-effective storage and low-latency networks, AI applications need specialized hardware to train computationally-intensive machine learning models. IaaS provided by cloud providers removes the need of in-houses servers which are costly for building AI applications. The migration from on-prem servers to cloud services is a pre-requirement to enable cost-effective data management. For example, a Fortune 500 company using AWS for storing their customer data on S3/ DynamoDB (databases) can easily start using the AI personalization services such as Amazon Personalize and Amazon Pinpoint.

AI Data Management as a Service

Unlike traditional software, AI models need a feedback loop to improve performance over time. These models consume historical or real-time data to gain intelligence. Data management as a service is a set of tools used to acquire and prepare data before feeding them into AI models. It requires a series of engineering steps from acquiring, storing, normalizing, labeling, and feeding data into the models. For example, an AI application that can diagnose cancer has to be trained using tens and thousands of medical images. Data management as a service companies such as Cloudfactory and Scale.AI upload these images from medical imaging devices to the cloud over HIPAA compliant protocols removing personal identifiable information (anonymizing), converting the image into pixels (normalizing) that can be understood by machines, and labeling of the cancer causing nodes (a radiologist manually marking the image on cancerous nodules). Each of these steps in the data management process vary depending on the industry, data type and data size, and require humans in the loop.

Predictive Models as a Service are intelligent applications that can make predictions about future or unknown events based on patterns in historical or real-time data. A Predictive model uses machine learning and other statistical tools to forecast and provide a list of recommendations. An example of a predictive model is an AI recommendation engine that analyzes Netflix history to predict movies a customer might want to watch.

Prescriptive Models as a Service
Prescriptive models use machine learning to automate decision making and evaluate the best solution in complex environments. It uses the information provided by descriptive analytics and predictive analytics.The most valuable AI startups are the ones selling prescriptive models as a service. Prescriptive models rely on the feedback loop from subject matter experts to improve the efficacy of the decisions.

54%

of respondents in a Gartner survey identified lack of knowledge to be the biggest barrier in adoption AI

In our management journal, we attempt to remove the knowledge barrier of understanding AI products and services.

The first framework states that each and every AI product and service in the market can be categorized under one of the four mutually exclusive and collectively exhaustive buckets:

Infrastructure as a Service (hardware and platform services)
Data management as a Service (historical and real time)
Predictive modeling as a Service
Prescriptive modeling as a Service.

Showcasing Adoption of AI using the Pyramid Framework in a case study

Asterix Inc.

Headquarters: California USA

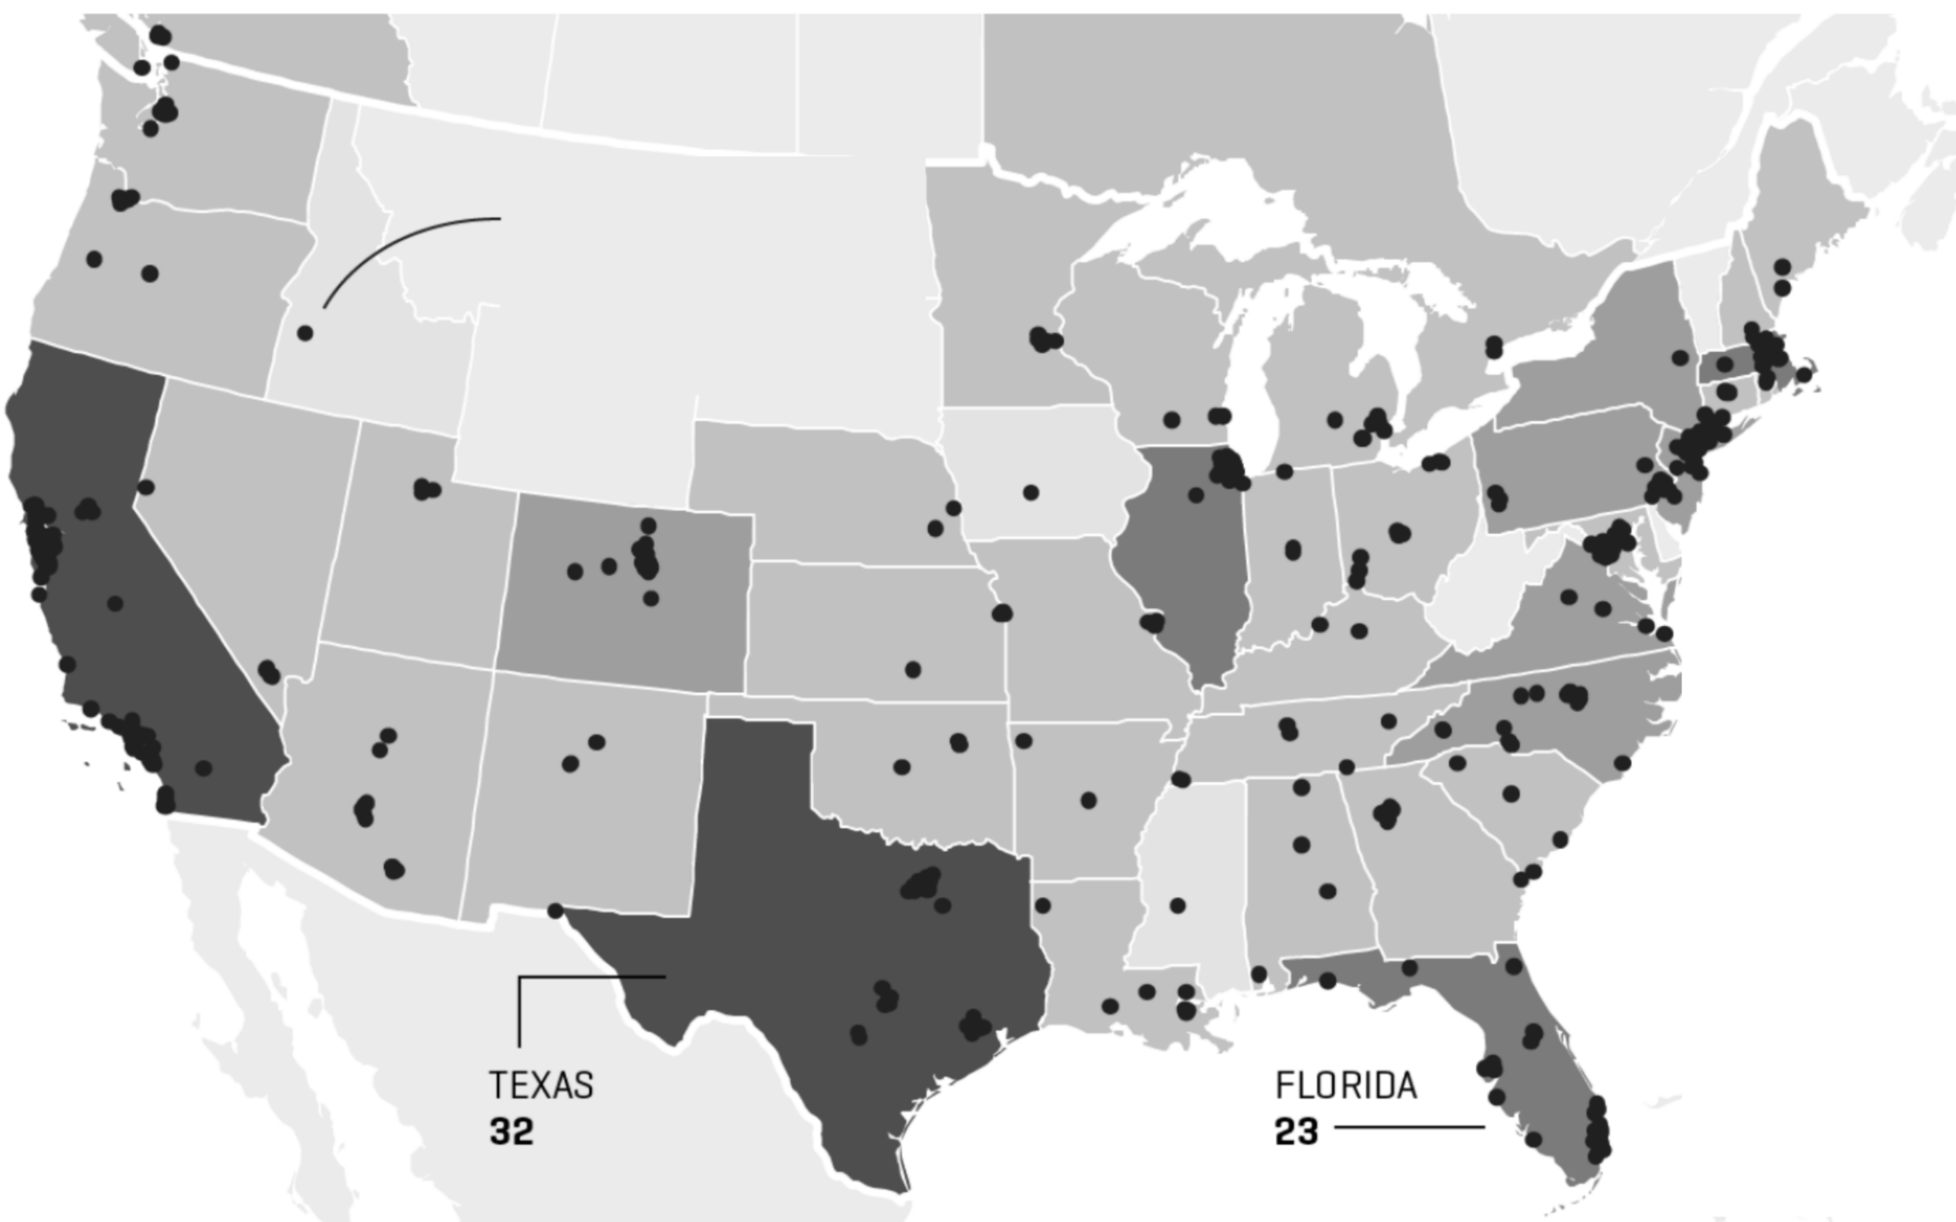

A Fortune 500 company named "Asterix" is a US-based quick service restaurant chain. It has 1000 stores across the country and needs to adopt digital platforms to reach out to its tech savvy customers.

How should Asterix adopt AI?

1

Data Migration

Customer data of 5M+ Asterix users such as names, email addresses, location, loyalty points, most recent orders, etc. are stored in Excel, third party applications such as SAP (used for their billing and management and inventory management), Salesforce (used for managing vendor and store relations) and data warehouses (Amazon Redshift to run business intelligence on inventory management). A solution architect at AWS advises Asterix to aggregate all their data from these disparate sources into a central repository such as Amazon S3, Amazon Elastic File System, Amazon FxS, etc. Given that Asterix does not have prior experience of migrating their data from different silos onto the cloud, they hire a global technology consulting company to move huge volumes of customer data securely and cost-effectively on the cloud. Hence, Asterix buys infrastructure as a service from AWS to start the AI adoption in consultation with a solution consulting firm.

2

Data Management

Once the data is readily available on the cloud, unstructured data needs to be converted into a standard structured format that can be fed into machine learning algorithms. Accenture helps in integrating AWS services such as AWS Glue to extract, transform, and load customer data that exists in different formats into a standard template called data cohorts. For example, a personal information cohort will only have data associated with the following fields: name, phone number, email addresses, and mailing address while a digital cohort will have data associated with a name, mobile application ID, email address most recently bought items in the app etc. After executing all the steps of cleaning, classifying, and preparing the data, Accenture pushes the data into a central data storage system called data lakes (a new S3 bucket). By taking the help of Accenture and AWS service to prepare the data for AI models, Asterix buys data management as a service.

Predictive Modeling

Asterix builds an internal team of software engineers with no background in AI to write code for ingesting data from different data cohorts into predictive AWS services such as Amazon Personalize and Amazon Forecast. The output of predictive services generates new cohorts on the basis of a specific user behavior called behavioral cohorts. The internal team of software engineers is required to define rules of segmentation in Amazon Personalize to customize the output of the model based on the business need. For example, the 5M+ customer data is now segmented into various logical buckets such as churn customers, new customers, loyal customers etc. on the basis of probabilistic similarities and dissimilarities between their customer journey. The decision of the CIO to buy predictive models as a service from AWS and customize them using an internal team of software engineers saved huge costs of hiring AI talent and accelerated the go-to-market of AI-driven marketing.

Prescriptive Modeling

Ultimately, the marketing team at Asterix can use the recommendation based on predictive models to reach out to different behavioral cohorts with different marketing communications to convert them into active customers. The marketing teams can further automate the marketing campaigns by using prescriptive models. Amazon Pinpoint can be used as an orchestrator to send an email, sms, notification to an end-customer on the basis of parameters such as responsiveness, privacy settings, time of the day and learnings from previous marketing communications to personalize the campaign. For example, a churn customer can be offered a discount coupon to attract them into making a purchase. The communication can be delivered to them during lunch time through the mobile application notification to maximize engagement and call-to-action. As the needs of Asterix grows in the future, instead of buying prescriptive models from AWS, the engineering team can hire a team of data scientists to build models for content and delivery personalization to maximize the click through rate. If such in-house prescriptive models indeed lead to higher sales due to continuous tweaking and improvement overtime, AI can become a business moat or competitive differentiation for Asterix.

The Build or Buy Decision Matrix

CIOs and business leaders face a decision, whether to hire a team of AI experts in-house to build AI ground-up, buy off the shelf AI from AIaaS providers, or adopt a hybrid approach that prioritizes future competitive advantages. In Figure 2, we attempt to provide a framework to approach this dilemma.

The decision to use AI as a Service is easy for commodity use cases such as customer service, human resources, and marketing as they do not possess any competitive advantage for the core business. But for data sensitive use cases which define the competitive edge of a company, the decision is harder. For example, AI to optimize delivery routes and timing for a logistics company (such as FedEx) based on both real-time and historic data leads to FedEx outperforming DHL in terms of reduced cost of doing business and faster delivery. Hence, FedEx will like to build AI in-house while keeping their customer service outsourced to AIaaS providers such as Meya.AI.

Similarly, Dow Chemicals will like to build AI for energy management in-house since Dow will optimize raw materials costs for chemical plants and build a competitive advantage over BASF. **The key decision metrics to solve the build vs buy decision is similar to the conventional outsourcing model: keep strategic initiatives in-house to ensure patent and IP protection, enhance speed and accuracy of execution, and avoid the high costs of customizing AI as a Service to suit their needs.**

We recommend that CIOs buy AI infrastructure and platform services from cloud providers instead of investing in on-premise AI infrastructure, largely due to costs.

Cloud companies have the scale to offer a vertically integrated AI stack that includes machine learning libraries, custom AI hardware and cost-effective centralized repository to store all structured and unstructured data. As more and more developers use cloud-based AI development tools, building AI products on the cloud becomes faster, cheaper, and more efficient than building them in-house.

The cloud companies share the benefits derived from economies of scale with the customers in the form of high-end compute and new product offerings at reduced cost structures. Therefore, it makes sense for CIOs to channel their in-house resources towards building prescriptive and predictive models instead of maintaining internal servers.

	Now	Future
AI Infrastructure as a Service (IaaS)	Buy	Buy
Data Management as a Service (DMaaS)	Buy	Buy
Predictive Models	Buy	Build
Prescriptive Models	Build	Build

Fig 2: The Build or Buy Decision Matrix

Buy Build

We predict that businesses adopting AI will buy data management as a service from startups and cloud companies instead of building an in-house data management team for a simple reason: maintaining an in-house data management team is expensive, time consuming and offers no competitive advantage. Building the data management platform requires hiring human labelers, developing a taxonomy of data labeling, and prior experience in managing data labeling workflows. Additionally, third party service providers offer data management as a service at competitive prices due to the lower shared cost structure. For example, if a company wants to convert their unstructured data (images and videos) into structured data (objects labelled in an image) to be used in machine learning models, they can use a third party service for labeling such as Mechanical Turk, Ground Zero, and Amazon Transcribe offered by AWS. It costs the customer as low as 20 cents per label (identified and marked in an image by a human labeller). Assuming that a typical complex ML model (AI to detect cancer by analyzing MRI scans) requires 1 million labels to get trained, it leads to $200,000 in recurring costs.

However, if data sources start at a consumer level (Netflix, EA Sports, Tesla), acquiring new customer data in real-time can contribute to a competitive differentiation. The reason is that the technology required to acquire, store, and feed data into machine models real-time from devices (mobile phone, televisions, and laptops) introduces technical challenges (internet infrastructure, on device processor bandwidth, and cyber security).Therefore, B2C technology companies adopting AI will need to build an in-house team for data-management as a service since it is highly customized to their use case and will likely yield competitive advantages.

We recommend that CIOs buy the predictive models off the shelf from AI service providers for three reasons. First, predictive analytics use cases for established technologies do not have a competitive differentiation. Standard use-cases such as data analytics, facial recognition, translation, sentiment analysis of evolved technologies such as big data, computer vision and natural language processing are available either in open-source or offered as a service by cloud providers. For example, open source AI frameworks such as TensorFlow provide standard algorithms (the semi-trained models) that are available free of cost and have no business value. Consequently, there is no need to build established use cases in-house and reinvent the wheel. Second, pre-trained supervised models such as Speech Services/Amazon Transcribe which converts voice to text and Azure Cognitive Services/Amazon Rekognition/Google Vision which converts video to identifiable objects are available at competitive price points. It will be costly to build them in-house and hence CIOs should buy them from the AI service providers. Third, predictive analytics tools that have a wide customer base deliver high-quality recommendations due to economies of scale. As more and more customers use these analytics engines for forecasting, the more accurate the models get over time. For example, for a well-established use case such as customer analytics, a CIO of Kroger can buy predictive analytics models from AI services providers and run them over their customer data lake to identify the list of customers who are most likely to purchase a specific item. In addition, CIOs can hire a solution provider such as Deloitte and Accenture to customize the models to their unique needs for a fraction of the price of building in-house. For example, a healthcare provider can hire Deloitte to guide them through managing patient healthcare records in accordance with HIPAA protocols and anonymization before uploading it to the cloud where off-the-shelf AI-driven diagnostic tools can then derive insights.

In order to take predictive models to the next stage of AI by automating the action taken on the predictions, we need prescriptive models. Continuing the e-commerce example, a prescriptive model can trigger a personalized email and notification to the customer to close a sale based on the actions taken before. These models are hard to build even for top technology companies for the following reasons. First, prescriptive models require a strong feedback loop from industry experts who possess domain knowledge as they attempt to automate decision-making in business processes. Second, there is a limited supply of AI experts in the world who possess the skills to build complex models to automate decision-making. Third, articulating and solving prescriptive use cases require complex interplay of industry expertise, system engineering, software engineering, Finally, the data sets for industry-specific use cases are hard to procure for technology companies. **Thus, a partnership to build prescriptive models is a win-win scenario for technology companies and conventional businesses since the former provides the AI talent and tools at scale while the latter provides access to industry experts and proprietary data.** In a conversation with Jim Gao, CEO of Phaidra.ai, he predicted “In the long term, AI is going to become a core competency of any business and hence they need internal teams to evolve out of the partnerships with cloud companies and become self-sufficient.” At that stage, prescriptive models become their core focus that enables them to achieve competitive advantage. For example, Walgreens recently partnered with Microsoft to build AI expertise in healthcare in which Walgreens brings the domain expertise and Microsoft brings the technology and tools to build AI ground-up in the healthcare delivery space. In an interview, anonymizing this person, Product Marketing Leader at Medtronic said ”We partnered with IBM cloud to build our AI powered diabetes app - SugarIQ. While IBM brings in the tools to build the app from a technology standpoint, we provide healthcare experts and data to train the AI as per the needs of our patients. In working with IBM, we are transforming into the agile product methodology of product development which is a huge cultural shift to adopt AI.”

AI is simple to implement. CIOs and business leaders face a decision, whether to hire a team of AI experts in-house to build AI bottom-up, buy off the shelf AI from AIaaS providers, or a hybrid approach that prioritizes future competitive advantages. Our recommendation is that a CIO buy commodity parts of building AI such as AI Infrastructure as a Service and AI Data Management as a Service from the AI service provider while building and collaborating with AI service providers to nurture Predictive and Prescriptive modeling.

Management: How to execute a partnership with a AI as a Service provider?

CIOs need partnership to build models

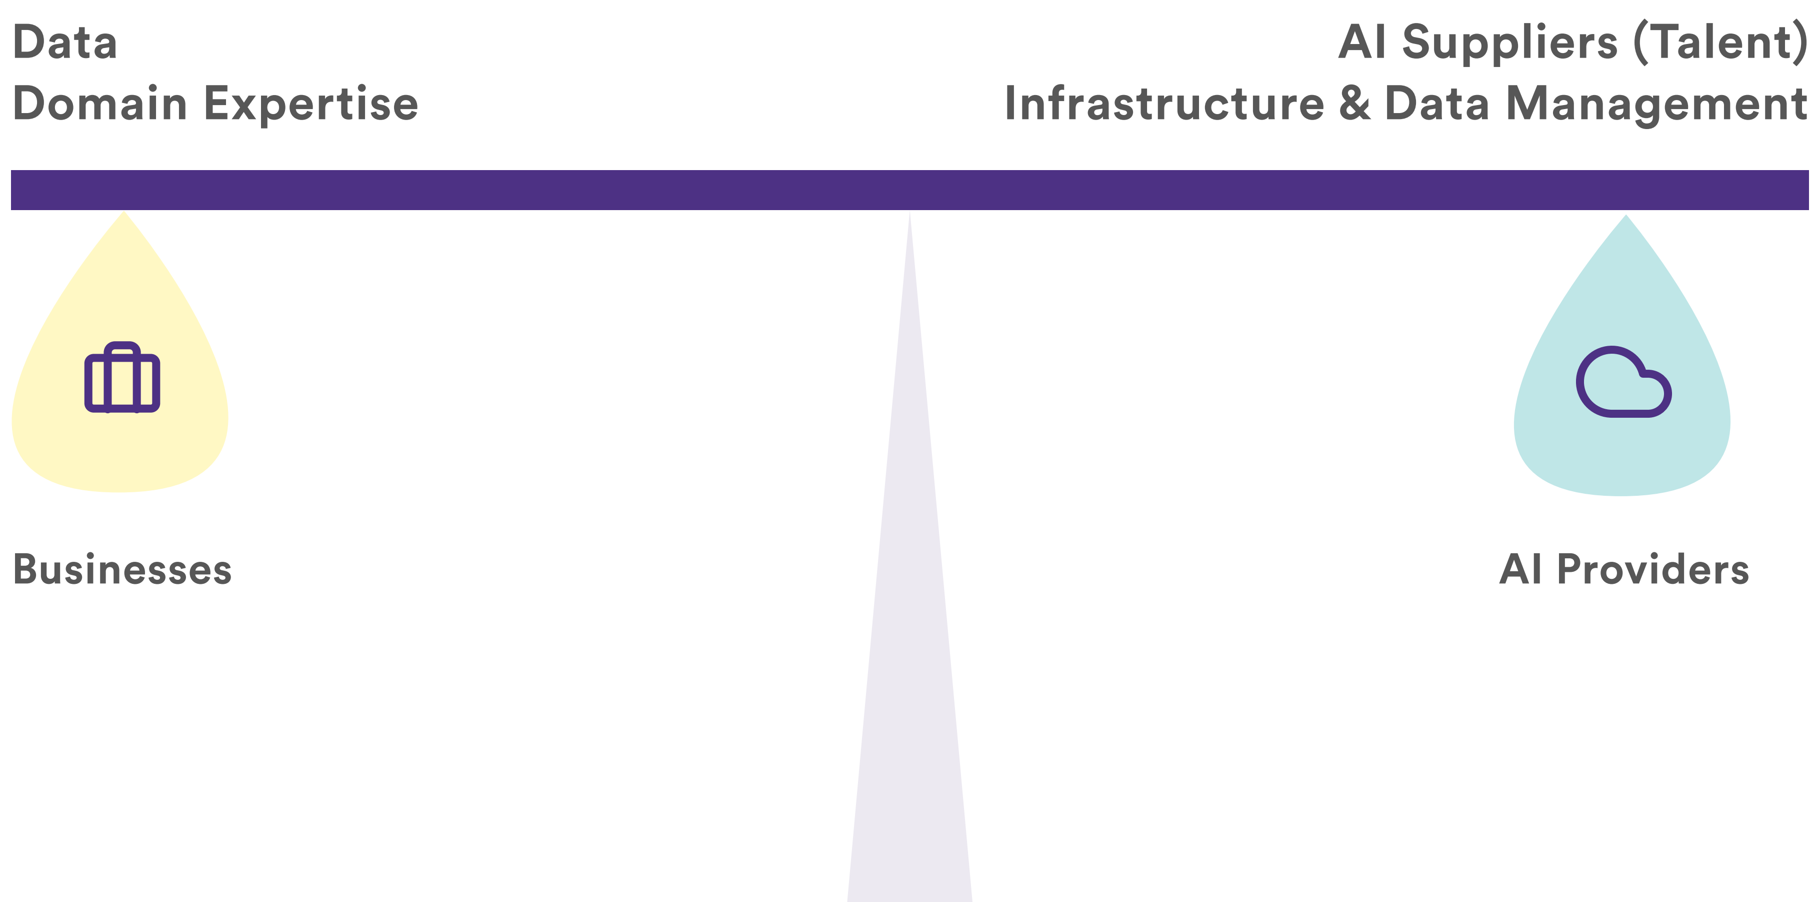

AI is simple to scale up. In order to execute on AI adoption, CIOs need to (a) build an in-house team of software developers to customize the AI services provided by vendors, (b) structure partnership with AI service providers where the buyer brings domain expertise and proprietary data while the supplier brings AI talent and technology creating a symbiotic relationship. (c) recognize that the power balance between the buyer and supplier is horizontal today

A step by step approach to AI Management

1

Leverage new generation of software architecture

As monolithic software becomes cumbersome to maintain, enterprises are shifting to microservices – a software architectural style that allows services to be loosely coupled to each other and be independently deployed. Key benefit of this Lego-like AI model building is that applications can be easily assembled and deployed at scale because an update to a small part of the application requires rebuilding and redeploying only a small number of services. We interviewed Steve Singh, former CEO of Concur and Docker. "CIOs typically only have about 20% of the annual budget for innovation. A significant portion of their budget, roughly 80%, is already spoken for in things like infrastructure costs, existing applications, both internally built and external. While there are inefficiencies throughout any budget, in terms of infrastructure, most companies are only using 15% of the compute infrastructure they operate. With the advent of nearly infinite, low cost compute from the major cloud providers, and the availability of technologies like containers [Docker & Kubernetes], microservices, machine learning, companies can dramatically reduce their long-term infrastructure costs. Those savings can be used to drive the innovation needed to compete."

2

Build your first AI Team

Instead of hiring costly AI engineers and data scientists, software engineers with minimal experience of building AI systems can help convert monolithic legacy business applications into microservices-driven next generation architecture. An investment in such a team and re-training employees on functional machine learning skills will help businesses partner with AI service providers in building prescriptive models. By reducing the time-to-market of an AI application from months to days, cloud-native microservices dramatically increase the pace of AI application development. These AI-infused, and API enabled microservices will sit on top of the public cloud. Cloud infrastructure makes these apps cheaper to build and maintain. Most SaaS products have a gross margin of 85%. If intelligence can increase the gross margin to 95%, CIOs will have 10 points to play with and invest in innovation projects. For example, a financial service provider can build prescriptive models for fraud detection and use it across different banking services such as loans, payments, personal banking, transactions etc", says Steve.

3

Structure the Partnership Agreement

An AI services vendor will be part of your AI adoption journey for 5-10 years. The existing investment in public cloud does not lead to a vendor lockdown since the AI applications are containerized (they can be built over a cloud provider X from data sets in a cloud provider Y, and finally deployed over a cloud provider Z through a multi-cloud strategy). Hence, open the request for proposal (RFPs) with all AI providers in the market. Once you receive proposal, decide on the basis of (a) bundled price offering of Infrastructure, Data Management, and Models as a Service, (b) existing customer base with similar use-case as per your needs, (c) non-compete agreements which prevents cloud providers from entering into partnerships with their competitors.

Thus, IT Procurement and partnership teams can structure the conversation with the vendor and negotiate on the following topics a) Corporate Data Governance - access to the proprietary data of the customer by the AI vendor for future use, b) Pricing - business model of payments to the vendor in terms of pay as you use model or sharing part of net-new revenues generated by AI, and c) Economic Moat - ownership of the IP generated in co-building AI tools due to your historical data, business use case application, and domain expert feedback. AI application agreements are different from SaaS application agreements since the former is constantly learning using client data.

4

Engage a third party solution provider

Technology solution provider such as Accenture, IBM Services, Infosys have been instrumental in helping Fortune 500 adopt AI especially in data sensitive industries such as Financial Services and Healthcare. If your company identifies with such a sector, hiring AI consulting partners could help to a) design a best in class system for your use case due to previous experience of the consulting companies, and b) save time and energy of your in-house AI resources by outsourcing non strategic workflows to third party solution architects.

5

Keep up with AI innovation and keep adding new data source to differentiate

The executive team needs to stay updated on new sources of data which can contribute to their data lake, strengthening the competitive edge due to AI. For example, health care providers have unique access to patient history which can be valuable for drug development. Thus, pharmaceutical companies should build data sharing agreements with partner providers for long term differentiation with respect to competitors. For example - Medtronic is a leader in providing medical monitoring tools for diabetes. It quickly understood that computer vision(new technology) can give insights into the food based consumption(new data sources) of diabetic patients. It acquired Nurtino, a big data, personalized nutrition platform that offers nutritional values and data through wearables and other technologies. In order to add new data sources to stay competent in AI, Medtronic recently acquired Klue (an Apple Watch app that detects your motion when you eat) to supplement SugarIQ.

6

Cultural shift in organization to implement AI at scale

The CEO of Alphabet, Sundar Pichai, announced at the Google I/O summit in 2017 -“Google will be become an AI first company”. Google is at the farthest end of the spectrum in adopting AI where they fully automate first and work backwards to see where they still need humans in the machine loop. On the conventional end of the spectrum, an organization augments its employees with predictive analytics tools to make them more productive and expand their capabilities. We urge our readers to find their place in the between these ends of AI adoption.

In the AI world, the biggest cultural shift in terms of project management is the adoption of agile methodology. AI applications have a continuous feedback loop that provides insights to companies to improve their apps in real-time. This increasing feedback will urge the company to process feedback at a rapid pace to keep improving their predictive and prescriptive models and maintain their competitive edge, drastically changing project planning from waterfall model spread across months/ years to an iterative product development model across weeks. For example, Ford would be creating a multi-year product roadmap in launching new versions of Mustang G20 . The company wide business units operates in the same timeframe to gather customer feedback, roll out promotions, manage operations, and initiate manufacturing. But in an AI dominated world of autonomous vehicles, everything changes. Every car on the road is learning each mile, gathering insights and sharing it centrally to Ford servers. Ford will be rolling out updates in a Mustang G20 every two weeks. Thus, each business unit needs to pivot towards a more dynamic and collaborative decision making framework to succeed in company wide adoption of AI.

Authored By

Abhro Pal

As a HPC (High Performance Computing) solutions architect for Siemens software, Mr. Pal partnered with customers in the automotive and aerospace industry to enable digital transformation. Most recently, he helped devise a positioning strategy for Microsoft Azure's HPC
and AI products/services on the Cloud, and is part of Azure as a Product Marketing Manager in the HPC and AI space.

Shobhit Gupta

Shobhit is a serial technology entrepreneur, having built high tech startups in the Internet of Things(IoT), Intelligent Mobile Applications, and Machine Learning Infrastructure. He has also consulted at Deloitte Strategy in the TMT practice and is joining Amazon Web Services to build products around Artificial Intelligence and Personalization. Prior Projects- Blockchain Infrastructure Commercialization, AI Ops, and AIaaS for Medical Imaging.

Alexis Goodrich

With seven years of experience in the start-up B2C space, Alexis has launched products, developed customer personas and founded a lead generation business. She has been a product manager at Vulcan, working on its wildlife conservation software using Machine Learning and Computer Vision, and has consulted for Intel and Microsoft on AI marketplaces.

Professor Shailendra Pratap Jain

Professor Jain is James D. Currie Professor of Marketing and International Business at the Foster School of Business, University of Washington, Seattle. His research interests include brand strategy, psychological segmentation, AI, COVID, consumer addictions, advertising, and
influence of emotions on consumption. His work has appeared in leading journals and he is anAssociate Editor with several of them (Journal of Marketing Research, Journal of Consumer Psychology, Journal of the Academy of Marketing Science, and International Journal of Research in Marketing).