The AI Playbook

The step-by-step guide to taking advantage of AI in your business

In partnership with

BARCLAYS

MMC Ventures

MMC Ventures is a research-led venture capital firm that has backed over 60 early-stage, high-growth technology companies since 2000.

MMC's dedicated research team provides the Firm with a deep and differentiated understanding of emerging technologies and sector dynamics to identify attractive investment opportunities. MMC's research team also supports portfolio companies through the life of MMC's investment.

MMC helps to catalyse the growth of enterprise software and consumer internet companies that have the potential to disrupt large markets. The Firm has one of the largest software-as-a-service (SaaS) portfolios in Europe, with recent exits including CloudSense, Invenias and NewVoiceMedia. MMC's dynamic consumer portfolio includes Bloom & Wild, Gousto and Interactive Investor.

MMC Ventures Research
David Kelnar – Partner & Head of Research
Asen Kostadinov, CFA – Research Manager

Explore MMC's cutting-edge research at mmcventures.com, MMC Writes (www.medium.com/mmc-writes) and @MMC_Ventures on twitter.

www.mmc.vc
@MMC_Ventures

Barclays UK Ventures

Barclays UK (BUK) Ventures is a specialist business unit within Barclays with an independent mandate to deliver new customer experiences at pace and scale – driving growth for communities, businesses and Barclays. BUK Ventures identifies, incubates and scales transformational new business lines and business models both within and outside of Barclays through organic build-out, commercial partnerships and venture investments.

BUK Ventures comprises a strong team of intrapreneurs, including Eagle Labs, who work to create thriving communities with the aim of connecting businesses of all sizes to the networks they need to succeed.

www.home.barclays

Explore our companion report The State of AI 2019: Divergence to understand AI today, what's to come and how to take advantage.

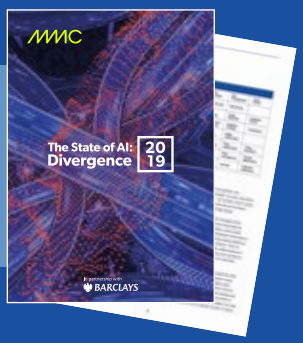

Contents

Contents

We'd value your feedback so we can improve the Playbook.
Get in touch at insights@mmcventures.com

Introduction

Your blueprint for AI

Artificial Intelligence (AI) is today's most important enabling technology. Leading startups, scale-ups and enterprises are using AI today to reimagine consumer experiences and business processes – unlocking revenue growth and cost savings at the expense of their competitors. How can you take advantage?

Embracing AI can seem daunting. How do I develop an AI strategy? Can I afford AI? Do I need an in-house team? Which algorithms should I explore?

We're excited to offer a blueprint for success. Our AI Playbook is an accessible, step-by-step guide to taking advantage of AI. We explain, without jargon, best practice in six core competencies for AI: strategy; people; data; development; production; and regulation & ethics. In a hurry? Every chapter includes a one-page summary and checklist of actions.

Whether you: are a founder, executive, information worker or developer; work at a startup, scale-up or enterprise; and have a budget of £500 or £5m – this Playbook will catalyse your journey.

Our research has been enriched by 400 discussions with Europe's leading AI startups, scale-ups and practitioners. Our special thanks to Dr. Janet Bastiman, MMC Ventures' AI Venture Partner, for her invaluable expertise.

At MMC Ventures we invest in, and support, the UK's most promising entrepreneurs. If you're an early stage company, get in touch to see how we can accelerate your journey.

David Kelnar
Partner & Head of Research
MMC Ventures

Email: david.kelnar@mmcventures.com
Twitter: @davidkelnar

Leading startups, scale-ups and enterprises are using AI today to reimagine consumer experiences and business processes.

Our AI Playbook is an accessible, step-by-step guide to taking advantage of AI. We explain, without jargon, best practice in six core competencies.

If you're an early stage company, get in touch to see how we can accelerate your journey.

Summary

Chapter 1: Strategy

- Recognise AI's potential for value creation. While you should not add AI to your initiatives for the sake of doing so, you risk losing competitive advantage if you fail to explore what AI can offer.

- Identify appropriate problems for AI to solve. AI is particularly effective at: assignment (identifying what something is, or the extent to which items are connected); grouping (determining correlations and subsets in data); generation (creating images or text based on inputs) and forecasting (predicting changes in time series data). All businesses will have challenges where the above apply and, therefore, where AI can be fruitful.

- Prioritise projects according to value and viability. Ensure you have a clear, concise specification of the problem and desired outcomes. Assessing viability includes considering whether your training data is balanced (free from bias), exhaustive (captures all relevant variables), diverse (captures rare situations) and is of sufficient volume.

- Timescales for creating AI are less certain than for traditional software development – and typically extend non-linearly with desired accuracy. Timescales vary according to the problem type, subject domain and data availability. Frequently, a prototype with limited accuracy can be developed within three months.

- Align your budget with your goals and deployment strategy. The budget an AI initiative requires will depend on multiple factors including the complexity of the problem, the availability and quality of training data, and the deployment strategy you select.

- AI deployment strategies include: calling third party Application Programming Interfaces (APIs); using managed AI services from third parties; building a small in-house AI team; and building an extensive in-house AI team. A large, in-house team is a multi-million-pound annual investment. Many companies develop a proof-of-concept using their existing development teams, and third-party APIs or paid services. Then, they create a budget proposal and begin with a small, in-house AI team.

- Seek sponsorship from senior executives. Support from management will be important for new AI initiatives to succeed. To build support, educate senior management regarding the benefits of AI while setting realistic expectations regarding timescales and results.

- Anticipate and mitigate cultural concerns about AI. To some, AI will be unfamiliar. Others will see their workflows change. Many people may be concerned about the impact of AI on job security. Frequently, AI will enhance an individual's role by offering 'augmented intelligence'. Address concerns proactively by highlighting the ways in which AI will support individuals' goals and enable team members to redirect their time to engaging aspects of their roles.

- Expect non-traditional security considerations. Protect against malicious activity via thorough system testing and exception handling.

- When your first project is underway, anticipate the longer-term aspects of your AI strategy. Consider: maintenance; data (budget to retrain your system as data evolves and increases); evolving algorithms (new techniques will offer better results in the future); scaling (extending useful AI systems to additional business units and geographies); innovation (a roadmap for new AI initiatives); and regulation (a strategy to comply with new legislation as it emerges).

To engage effectively with AI, separate AI myths from reality

Myth	Reality
"AI is a distant dream."	While general, human-level artificial intelligence will not be available for many years, there are many applications for AI that are viable today and offer companies cost savings and revenue growth.
"We don't have the budget to implement AI."	While a large, in-house AI team will require extensive investment, third parties offer access to AI services (via API) for as little as several hundred pounds. Further, as AI democratises, growing libraries of pre-trained models offer results at low cost. If you have a software engineering team, you can validate benefit from AI at minimal cost.
"AI is dominated by the big technology companies. There's no point in my company trying to compete."	While companies including Amazon, Google, IBM and Microsoft have developed extensive AI services, they lack the strategic desire, data advantage or domain expertise to tackle the many sector- or function-specific applications for AI. Today, a rich ecosystem of startups, scale-ups and corporates are deploying AI for competitive advantage.
"We can't use AI because our business requires explainable processes."	There are several ways to explain what is occurring inside an AI system (see Chapter 6). Some AI is directly explainable. With deep learning systems, where explainability is a challenge, it is possible to explain how input variables influence output.
"I can throw AI at my data and it will offer efficiencies."	AI is a tool that requires a structured problem and appropriate data to be effective.

Source: MMC Ventures

Chapter 2: People

- In AI, job titles vary and can be difficult to interpret. We describe characteristics and salaries for six key roles: Data/Machine Learning Engineer; Data scientist; Machine Learning Researcher; Head of Data; Head of Research/ AI; and Chief Scientist/Chief Science Officer. For each, individuals' capabilities vary across competencies in research, engineering, production and strategy.

- The composition of your team should depend upon the problem being solved and your approach to doing so. It is advisable, however, to avoid hiring solo talent. Begin with a small team, and ensure you have a robust AI strategy in place before expanding your AI personnel.

- We suggest team structures, first hires and next steps for six scenarios: "I want insights into internal data"; "I want to implement third party AI APIs"; "I want to outsource AI development"; I want to create bespoke AI models"; "I want to use a combination of bespoke and third party AI"; and "I have an idea that's cutting edge."

- Recruiters, conferences and universities are primary sources of talent. Traditional recruitment agents find it difficult to screen AI candidates, so engage with specialist recruiters. Conferences and meetups are powerful vehicles for talent acquisition; be active in the AI community, attend and speak at conferences, and grow your network to discover capable candidates. Engage with universities; post on their job boards, establish partnerships and pay for projects to engage students who may seek future opportunities with you.

- Diversity delivers economic value and competitive advantage. Review the culture in your company, AI team and hiring practices to ensure diversity, representation and inclusion.

- An effective job description should emphasise projects (the nature of the engagements on which the successful candidate will work), skills and impact. Most data scientists seek work that will 'make a difference'. To attract talent, demonstrate how the successful candidate's work will do so.

- When hiring, prioritise adaptable problem-solvers. In addition to having role-specific and technical skills, a strong AI candidate will: understand available tools to enable rapid research and development; appreciate when to release an imperfect solution and when to wait; and communicate and collaborate well.

- Optimise every stage of your recruitment funnel. We provide best practices for: CV screening; phone screening; technical testing; face-to-face interviews and post-interview follow-up.

- AI talent is in short supply. Challenge, culture and company are key for retention. In addition to an attractive financial package, consider: offering flexible working hours; offering challenging problems and minimising drudgery through automation; creating a culture in which diverse ideas are shared; avoiding 'lone workers'; ensuring your AI team receives recognition for its work; and supporting team members' publishing and presentation of work.

Most data scientists seek work that will 'make a difference'. To attract talent, demonstrate how the successful candidate's work will do so.

Chapter 3: Data

- For effective AI, develop a data strategy. A data strategy spans: data acquisition & processing; quality; context; storage; provisioning; and management & security. Define your data strategy at the outset of your AI initiative.

- Accelerate data acquisition by using multiple sources. Developers draw on several sources including: free resources (such as dataset aggregators); partnerships with third parties (companies, universities, data providers and government departments); and new, proprietary data.

- A high-quality data set has appropriate characteristics to address your business challenge, minimises bias and offers training data labelled with a high degree of accuracy. Develop a balanced data set - if you possess significantly more samples of one type of output than another, your system will exhibit bias.

- Primary forms of bias are: unwarranted correlations (between inputs and output classifications); erroneous assumptions which cause relationships to be missed ('underfitting'); and modelling noise instead of valid outputs ('overfitting'). Adjust for overfitting and underfitting by using different data volumes and model structures. Remove unwarranted correlations through testing.

- Ensure that the results of your internal testing will be maintained when applied to real-world data. Test early, and frequently, on real-world data.

- Managing 'dirty data' is data scientists' most significant challenge (Kaggle). Smaller volumes of relevant, well-labelled data will typically enable better model accuracy than large volumes of poor-quality data. To label data effectively: consider developing a supporting system to accelerate data labelling and improve accuracy; draw on existing AI and data techniques; and seek data labelled by multiple individuals to mitigate mislabelling.

- Understand the data you use. Ensure you capture the human knowledge regarding how your data was gathered, so you can make downstream decisions regarding its use. Capture data provenance (where your data originated and how it was collected). Define your variables (differentiate between raw data, merged data, labels and inferences). Understand the systems and mappings through which your data pass to retain detail.

- Store and structure data optimally to support your objectives. Storage options include basic file-based, relational, NoSQL or a combination. When selecting storage plan for growth in data volume, updates, resilience and recoverability.

- One in three data scientists report that access to data is a primary inhibitor of productivity (Kaggle). Develop a provisioning strategy that: ensures data is accessible across your organisation when needed; contains safeguards to protect your company against accidents; optimises system input/output; and maintains data freshness.

- Implement robust data management and security procedures consistent with local and global regulations. Personal data is protected by UK and EU law and you must store it securely. Draw on principles of appropriate storage, transmission and minimum required access.

The six components of an effective data strategy

Source: MMC Ventures

Chapter 4: Development

- There are many ways your company can engage with AI. Use third party AI APIs; outsource; use a managed service; build an in-house team; or adopt a 'hybrid' approach combining an in-house team with third party resources.

- Third party AI APIs fulfil specific functions to a moderate or high standard at low cost. Most solve problems in the domains of vision and language. Numerous APIs are available from Amazon, Google, IBM, Microsoft and also other smaller companies. Features vary; we provide a summary. APIs offer immediate results without upfront investment, at the expense of configurability and differentiation. Use an API if you seek a solution to a generic problem for which an API is available. APIs are unsuitable if you seek solutions to narrow, domain-specific problems, wish to configure your AI, or seek long-term differentiation through AI.

- Managed services enable you to upload your data, configure and train models using a simple interface, and refine the results. Managed services abstract away much of the difficulty of developing AI and enable you to develop a custom solution rapidly. Managed services offer greater flexibility and control than APIs, but less flexibility than an in-house team, and also require you to transfer data to a third party and may create dependencies.

- If a third-party solution is unavailable and an in-house team is too expensive, you can outsource your AI development. Whether outsourcing is appropriate will depend upon your domain, expertise, required time to value and data sensitivity. If outsourcing, specify desired frameworks and standards, who will provide training data, costs, timescales and deployment considerations. Outsource if you require trusted expertise quickly and a cheaper option than permanent employees. Avoid outsourcing if your data permissions prohibit it, you require domain or sector knowledge that an outsourcer lacks, or you wish to build knowledge within your own company.

- An in-house AI team offers maximum control, capability and competitive differentiation – at a price. A small in-house team will cost at least £250,000 to £500,000 per year. A large team requires a multi-million-pound annual investment. To develop an in-house team your company must also: attract, manage and retain AI talent; select

development frameworks and techniques; gather and cleanse data; learn how to productise AI into real-world systems; and comply with regulatory and ethical standards. Build an in-house team if you have a problem that cannot be solved with existing solutions, seek differentiation in the market, or seek to maintain control over your data.

- A 'hybrid' approach is ideal for many companies. Plan for an in-house team that will address your requirements to a high standard over time, but use third party APIs to solve an initial, simpler version of your challenge. A hybrid approach can be attractive if you seek rapid initial results, wish to limit spend until a business case is proven and want greater differentiation and resilience over time.

- To develop AI yourself you have choices to make regarding your AI 'technology stack'. The stack comprises six layers: hardware; operating systems; programming languages; libraries; frameworks; and abstractions. Not all problems require the full stack.

- Ensure your team has hardware with graphical processing units (GPUs) that support NVIDIA's CUDA libraries. Laptops with high performance graphics cards offer flexibility. For greater power, desktop machines with powerful GPUs are preferable. To train large models, use dedicated servers. Cloud-based servers offered by Amazon, Google or Microsoft are suitable for most early stage companies.

- Apply AI techniques suited to your problem domain. For assignment problems consider: Support Vector Classification; Naïve Bayes; K-Nearest Neighbour Classification; Convolutional Neural Networks; Support Vector Regression; or 'Lasso' techniques. We describe each and explain their advantages and limitations. For grouping problems, explore: Meanshift Clustering; K-Means; and Gaussian Mixture Models. For generation, consider: Probabilistic Prediction; Variational Auto-Encoders; and Generative Adversarial Networks (GANs).

With one network, GANs generate output from random noise; a second network serves as a discriminator

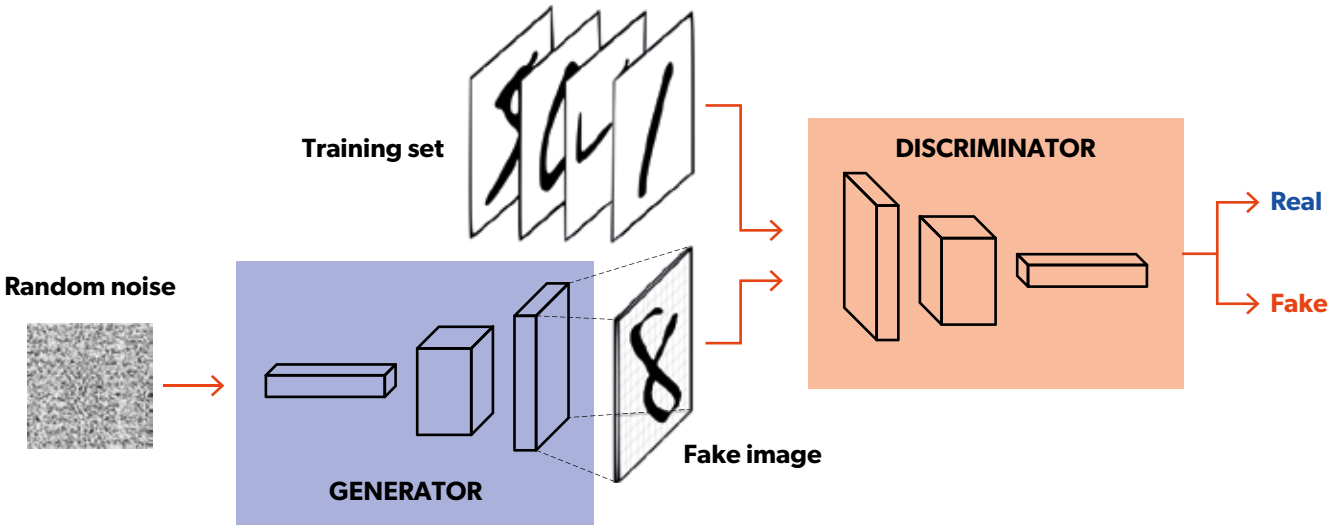

Source: https://medium.freecodecamp.org/an-intuitive-introduction-to-generative-adversarial-networks-gans-7a2264a81394

Chapter 5: Production

- An unused AI system delivers no value. Develop a production process that smoothly transitions AI systems you have in development to live use.
- AI production follows a conventional development process and requires you to undertake research, develop a prototype and create a minimum viable product (MVP). Once in production, undertake cycles of ideation, research, development and quality assurance.
- Effective R&D requires rapid iteration. Initially, optimise for speed over quality. Releasing an early model into production for feedback is preferable to waiting until a research model is perfect.
- During the R&D phase, solicit feedback about prototypes from beyond the AI and production teams to minimise expensive redevelopment later.
- When moving from MVP to production, select an appropriate hosting environment. On-premise hosting is suitable for those with highly sensitive data and existing on-premise hardware, but is rarely preferred by early stage companies given high upfront costs, unpredictable activity levels and required security expertise. Hosting your own hardware in a data centre offers control and value over the long term. Upfront costs can be high, however, and managing a data centre can prove a distraction for young companies. Cloud hosting, which offers low upfront costs and high levels of flexibility, is well suited to many early stage companies – although annual costs can be double that of a self-managed data centre and cloud hosting may be unsuitable for highly sensitive data. Consider the physical location in which your cloud servers are hosted. Different countries have varying rules regarding data and you may be required to keep your data within its area of origin.
- Proving that AI systems are effective differs from the typical software quality assurance (QA) process. Test your AI system at multiple stages – during training, validation and continuously through its life. Efficiency is critical; automate testing to as great an extent as possible.

The AI production pipeline is similar to a normal development practice

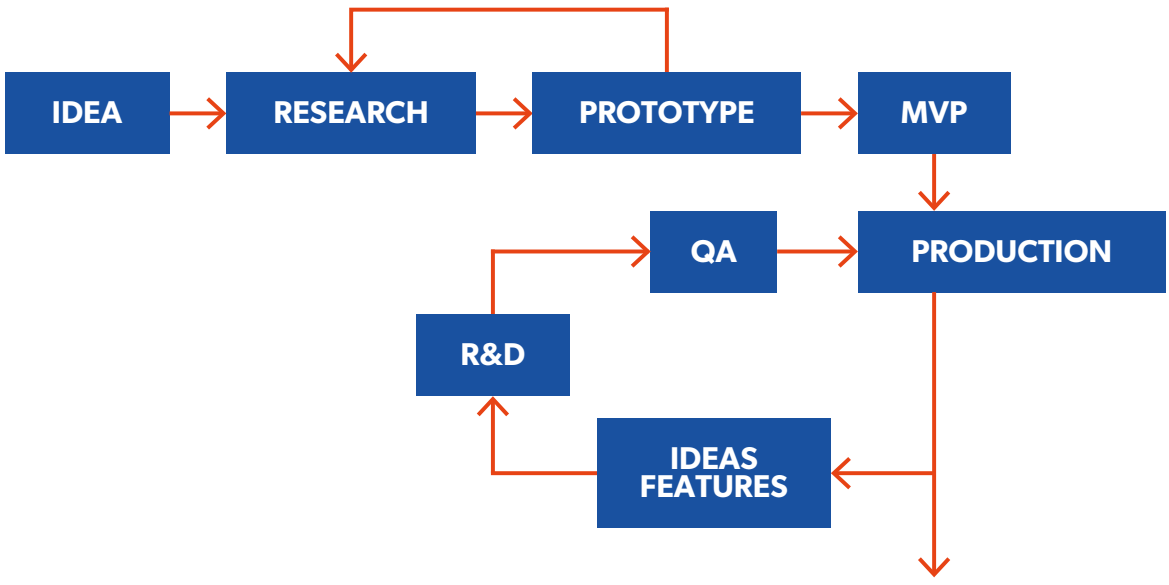

Source: MMC Ventures

- Understand the three common measures of 'accuracy' in AI – recall, precision and accuracy – and monitor all three to capture performance. Balancing precision and recall is challenging. Whether you elect to minimise false positives or false negatives should depend upon the nature of your sector and the problem you are solving.

- An effective maintenance programme will sustain your AI's intelligence. Beyond the maintenance you would typically perform on a software system, you should verify and update your AI system on an ongoing basis. AI technology is developing at pace. Invest in continual improvement to ensure your system avoids obsolescence.

Understand the three common measures of 'accuracy' in AI – recall, precision and accuracy – and monitor all three to capture performance.

Chapter 6: Regulation & Ethics

- As consideration of data privacy grows, and with the General Data Protection Regulation (GDPR) in force across the European Union (EU), it is vital to ensure you are using data appropriately. The GDPR applies to all companies processing the personal data of people in the EU, regardless of a company's location.

- Companies that are 'controllers' or 'processors' of personal information are accountable for their handling of individuals' personal information. Demonstrate compliance with GDPR data handling requirements and the principles of protection, fairness and transparency.

- Minimise the personal data you require, to reduce regulatory risk, and pseudonymise all personal data through anonymisation, encryption or tokenisation.

- In addition to standardising data handling requirements and penalties for misuse, the GDPR introduced considerations that can impact AI systems specifically. Verify that automated systems meet GDPR stipulations. Article 22 of the GDPR prohibits legal effects that result solely from automated processing being undertaken without an individual's explicit consent, when consent is required. Several legislative terms are subject to interpretation at this time. It may be prudent to make your system advisory only, and include a human check, if you are developing a system that could materially impact an individual's life.

- 'Explainability' – explaining how the outputs of your AI system are derived – is growing in importance. Convention 108 of the Council of Europe, adopted into UK and EU law in May 2018, provides individuals with the right to obtain knowledge of the reasoning underlying data processing systems applied to them. Explainability can be challenging in relation to deep learning systems. Explore varying approaches to explainability including Inferred Explanation, Feature Extrapolation and Key Variable Analysis. Each offers trade-offs regarding difficulty, speed and explanatory power.

- Develop a framework for ethical data use to avoid reputational and financial costs. The ALGOCARE framework, developed by the Durham Police Constabulary in partnership with academics, highlights issues you should consider when managing data. It incorporates: the nature of system output (Advisory); whether data is gathered

lawfully (Lawful); whether you understand the meaning of the data you use (Granularity); who owns the intellectual property (IP) associated with the data (Ownership); whether the outcomes of your system need to be available for individuals to challenge (Challenge); how your system is tested (Accuracy); whether ethical considerations are deliberated and stated (Responsible); and whether your model has been explained accessibly to as great an extent as possible (Explainable).

Companies that are 'controllers' or 'processors' of personal information are accountable for their handling of individuals' personal information. Demonstrate compliance with GDPR data handling requirements and the principles of protection, fairness and transparency.

How to select an approach to explainability

Use this approach if you:	Avoid this approach if you:
Inferred Explanation	
– Seek a high-level overview of your AI system – Believe correlation offers sufficient explainability	– Require detail regarding how variables lead to decisions
Feature Extraction	
– Require detail from within the network – Have a network type (e.g. images) where abstractions can be mapped onto input data	– Have limited time – Require precise impact of input variables, not general features – Are not using an assignment–based or generative AI network
Key Variable Analysis	
– Require detail about the importance of variables – Seek to prevent unwanted bias in your variables	– Have limited time – Seek to the publish your results – Wish to offer a layperson's guide to your model

Source: MMC Ventures

Chapter 1

Strategy

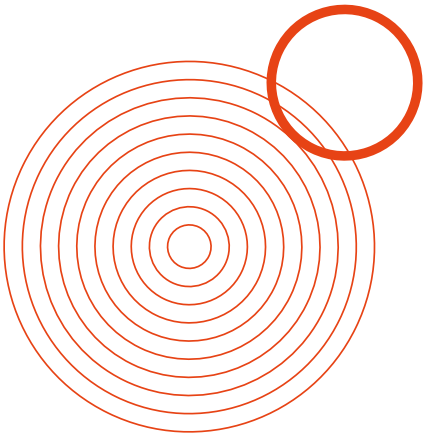

Summary

- Recognise AI's potential for value creation. While you should not add AI to your initiatives for the sake of doing so, you risk losing competitive advantage if you fail to explore what AI can offer.

- Identify appropriate problems for AI to solve. AI is particularly effective at: assignment (identifying what something is, or the extent to which items are connected); grouping (determining correlations and subsets in data); generation (creating images or text based on inputs) and forecasting (predicting changes in time series data). All businesses will have challenges where the above apply and, therefore, where AI can be fruitful.

- Prioritise projects according to value and viability. Ensure you have a clear, concise specification of the problem and desired outcomes. Assessing viability includes considering whether your training data is balanced (free from bias), exhaustive (captures all relevant variables), diverse (captures rare situations) and is of sufficient volume.

- Timescales for creating AI are less certain than for traditional software development – and typically extend non-linearly with desired accuracy. Timescales vary according to the problem type, subject domain and data availability. Frequently, a prototype with limited accuracy can be developed within three months.

- Align your budget with your goals and deployment strategy. The budget an AI initiative requires will depend on multiple factors including the complexity of the problem, the availability and quality of training data, and the deployment strategy you select.

- AI deployment strategies include: calling third party AI APIs; using managed AI services from third parties; building a small in-house AI team; and building an extensive in-house AI team. A large, in-house team is a multi-million-pound annual investment. Many companies develop a proof-of-concept using their existing development teams, and third-party APIs or paid services. Then, they create a budget proposal and begin with a small, in-house AI team.

- Seek sponsorship from senior executives. Support from management will be important for new AI initiatives to succeed. To build support, educate senior management regarding the benefits of AI while setting realistic expectations regarding timescales and results.

- Anticipate and mitigate cultural concerns about AI. To some, AI will be unfamiliar. Others will see their workflows change. Many people may be concerned about the impact of AI on job security. Frequently, AI will enhance an individual's role by offering 'augmented intelligence'. Address concerns proactively by highlighting the ways in which AI will support individuals' goals and enable team members to redirect their time to engaging aspects of their roles.

- Expect non-traditional security considerations. Protect against malicious activity via thorough system testing and exception handling.

- When your first project is underway, anticipate the longer-term aspects of your AI strategy. Consider: maintenance; data (budget to retrain your system as data evolves and increases); evolving algorithms (new techniques will offer better results in the future); scaling (extending useful AI systems to additional business units and geographies); innovation (a roadmap for new AI initiatives); and regulation (a strategy to comply with new legislation as it emerges).

Strategy: The Checklist

Identify use cases

- ☐ Understand the categories of problem AI can address
- ☐ Seek ideas and advice from AI practitioners
- ☐ Create a list of potential AI initiatives offering business benefit

Prioritise initiatives

- ☐ Develop a clear statement of the business challenge and opportunity
- ☐ Define measures of success
- ☐ Review the suitability of available data

Understand timescales

- ☐ Appreciate the need to iterate AI systems
- ☐ Develop realistic goals regarding accuracy and timescales

Develop a budget

- ☐ Understand budget requirements for different AI development strategies
- ☐ Select an initial and long-term development strategy for creating AI systems

Build buy-in

- ☐ Define the return on investment of your AI strategy
- ☐ Develop a clear, detailed implementation plan
- ☐ Educate senior management regarding AI and establish realistic expectations

Mitigate cultural and security concerns

- ☐ Educate and involve your workforce to address concerns
- ☐ Plan for non-traditional security challenges

Develop a long-term strategy

- ☐ Set aside budget for maintenance, updates and re-training
- ☐ Review new technology and evolving data sets
- ☐ Develop plans to extend your AI to additional business units and to undertake new AI initiatives
- ☐ Develop a process to ensure compliance with evolving legislation

AI is a powerful tool. Before you invest time and money in the technology, you need a strategy to guide its use. Without an AI strategy, AI will become an additional cost that fails to deliver a return on investment. Below, we describe how to: identify appropriate use cases for AI; select your first AI initiative; explore deployment strategies; anticipate timescales; predict required budget; and establish the cultural buy-in necessary for success.

Recognise AI's potential for value creation

AI is a powerful set of techniques offering companies tangible cost savings and increased revenue. Further, adoption of AI is 'crossing the chasm', from innovators and early adopters to the early mainstream. While you should not attempt to add AI to your initiatives for the sake of doing so, and should be mindful of its limitations, you risk losing competitive advantage if you fail to explore what AI can offer. Approach AI based on its transformational potential.

To engage effectively with AI, separate AI myths from reality:

Fig. 1. Separate AI myths from reality

Myth	Reality
"AI is a distant dream."	While general, human-level artificial intelligence will not be available for many years, there are many applications for AI that are viable today and offer companies cost savings and revenue growth.
"We don't have the budget to implement AI."	While a large, in-house AI team will require extensive investment, third parties offer access to AI services (via API) for as little as several hundred pounds. Further, as AI democratises, growing libraries of pre-trained models offer results at low cost. If you have a software engineering team, you can validate benefit from AI at minimal cost.
"AI is dominated by the big technology companies. There's no point in my company trying to compete."	While companies including Amazon, Google, IBM and Microsoft have developed extensive AI services, they lack the strategic desire, data advantage or domain expertise to tackle the many sector- or function-specific applications for AI. Today, a rich ecosystem of startups, scale-ups and corporates are deploying AI for competitive advantage.
"We can't use AI because our business requires explainable processes."	There are several ways to explain what is occurring inside an AI system (see Chapter 6). Some AI is directly explainable. With deep learning systems, where explainability is a challenge, it is possible to explain how input variables influence output.
"I can throw AI at my data and it will offer efficiencies."	AI is a tool that requires a structured problem and appropriate data to be effective.

Source: MMC Ventures

Identify appropriate problems

AI can be effective at solving problems – but it is important to begin with a clear problem in mind. Broad considerations are insufficient. When creating a list of potential AI initiatives, develop a precise definition of a problem you wish to address. "Always focus on the problem you're using AI to solve" (Tim Sadler, Tessian). Do you have a problem whose solution will add value within the business or to customers? Can the problem be solved using AI? AI is particularly effective in four problem domains: assignment; grouping; generation and forecasting (Fig. 2).

"Always focus on the problem you're using AI to solve."
Tim Sadler, Tessian

Fig. 2. AI is highly effective at Assignment, Grouping, Generation and Forecasting

Problem Domain	Definition	Examples
Assignment	• Identify what something is ***(classification)***	• Understand the sentiment of text • Recognise logos in images • Make a medical diagnosis based on symptoms
	• Identify how connected items are ***(regression)***	• Quantify the relationship between a preservative and product shelf life • Evaluate how consumer income affects propensity for impulse purchasing • Predict the purchase price of a second-hand vehicle based upon its condition
Grouping	• Given data, determine correlations and subsets ***(clustering)***	• Identify social subgroups within a customer base for enhanced targeting • Evaluate factors that correlate with patient melanomas • Identify themes in customer feedback surveys
Generation	• Given an input, create an image or text ***(generation)***	• Create a chatbot for customer service • Translate customer conversations to a different language • Create photorealistic media for advertising
Forecasting	• Given time series data, predict future changes ***(sequencing)***	• Predict weekly sales to avoid the loss of perishable stock • Determine the probability of equipment failure to enable proactive replacement • Predict exchange rate fluctuations

Source: MMC Ventures

All businesses will have challenges of the types above – and therefore problems to which AI can be usefully applied. The table below provides examples of popular AI use cases.

"The applications of AI are endless."
Timo Boldt, Gousto

Fig. 3. AI is being fruitfully applied to a wide variety of use cases

Sector	Example use cases			
Asset Management	Investment strategy	Portfolio construction	Risk management	Client service
Healthcare	Diagnostics	Drug discovery	Patient monitoring	Surgical support
Insurance	Risk assessment	Claims processing	Fraud detection	Customer service
Law & Compliance	Case law review	Due diligence	Litigation strategy	Compliance
Manufacturing	Predictive maintenance	Asset performance optimisation	Utility optimisation	Supply chain optimisation
Retail	Customer segmentation	Content personalisation	Price optimisation	Churn prediction
Transport	Autonomous vehicles	Infrastructure optimisation	Fleet management	Control applications
Utilities	Supply management	Demand optimisation	Security	Customer experience

Source: MMC Ventures

There are many ways to identify and evaluate potential AI projects, including:

- Network: to familiarise yourself with AI and its use cases, engage with your professional network and AI communities on LinkedIn and Meetup.com (on Meetup.com, communities can establish informal gatherings and there is a thriving AI community). Many community events are free. Ask an attendee for a coffee and you will find a useful sounding board for your questions and ideas. Informal advice is valuable; you can discuss whether AI might be suited to your use cases, why, and how to turn your idea into an initiative. "Find someone who is already using AI and bounce your ideas off them. Work out if your idea is possible. Have that conversation before even thinking about a consultant." (Miguel Martinez, Chief Data Scientist, Signal).
- Conferences: seek inspiration, talk with experts and understand industry best practises through conferences. Conferences tend to be high-level executive briefings, sales pitches or presentations of academic research. If you are early in your AI journey, prioritise events with multiple tracks for less experienced practitioners, or a mixture of levels so you receive an overview. Useful conferences will provide access to companies with successful AI solutions, which you can talk to for advice and collaboration. The cost of conference attendance varies from several hundred pounds to several thousand. Familiarise yourself with sessions before you book to ensure a return on investment.

Prioritise projects according to value and viability

Once you have ideas for AI projects, beyond assessing the relative value of each to your company, determine the most viable by addressing the following questions. As well as enabling you to choose a feasible project, the answers will help you define project parameters and objectives.

- **Problem:** Does the project fall within the definition of assignment, grouping, generation or forecasting? If you cannot clearly define the type of problem, it may be a viable undertaking but is unlikely to be an ideal first AI project for your company.
- **Definition:** Can you state the problem clearly and concisely? If not, you will lack a clear definition of the system's purpose and will struggle to select and employ appropriate AI techniques.
- **Outcomes:** Can you define the levels of accuracy and speed the system must achieve to be successful? Avoid initiatives that lack these measures. If converting an existing manual process, do you know the accuracy and speed of the current workflow? If you are undertaking a new initiative for your company, define what will be deemed a successful outcome.
- **Data:** Do you have sufficient data to train and test a system? Without adequate, high quality data to train your system your initiative will fail. If you are choosing between a range of otherwise viable projects, select the engagement supported by the greatest quantity of high-quality data.

Without adequate, high quality data to train your system, your initiative will fail.

It can be challenging to assess data suitability. Typically, data must be:

- **Representative:** Data you use to train your AI model should reflect the data you will feed your system in its live environment. If the data differs significantly, results will be poor even if the accuracy of your system during training is high.
- **Diverse:** Even rare situations should be captured in available training data. Without diverse data, your system may not generalise effectively. Overall accuracy may be high, but your model will fail (misclassify, wrongly correlate or poorly predict) in less frequent situations.
- **Balanced:** A biased data set produces a biased system. Does your data have inherent bias? For example, are you analysing CVs for suitability to a role and most candidates are of the same gender? Liaise with individuals in your organisation who understand your data and can advise on its inherent bias.
- **Exhaustive:** All relevant variables must be included in the available data. For assignment and grouping problems, missing variables will lead to oversimplified results (unwarranted correlations). In other problem domains, you may be unable to derive utility from your system.
- **Sufficient:** While a smaller volume of high-quality data is preferable to extensive, poor-quality data, the volume of data you can acquire must be sufficient to train your algorithm well. For assignment problems, useful results frequently begin to emerge after approximately 1,000 examples for each output label. Some problems require more or fewer. For forecasting problems, you may require data spanning at least double the duration of the periodicity of the item forecasted. For grouping and generation challenges, typically output improves with data volume but again 1,000 examples are frequently a minimum. Typically, the more complex the challenge, the more data points you will require.

In Chapter 3, we explain how to develop a full data strategy to support your AI initiatives.

Timescales will extend non-linearly with accuracy

Timescales for AI initiatives can be less certain than for traditional software development. AI systems cannot predictably be developed once, tested and then deployed. Typically, multiple cycles of training are required to identify a suitable combination of data, network architecture and 'hyperparameters' (the variables that define how a system learns). These dynamics will vary according to domain, the nature of the problem and the data available. Accordingly, it can be challenging to predict or automate AI initiatives unless they are similar to projects you have previously undertaken.

While timescales will vary according to the problem you are addressing, the resources you have committed and the buy-in you have achieved, you can frequently develop a prototype within three months. It may take days to develop a first version of a system that offers 50% accuracy, weeks to improve the system to 80% accuracy, months to achieve 95% and much longer for additional incremental improvements (Fig. 4). For straightforward problems, expect a similar progression but over shorter timescales. For particularly challenging problems, which require extensive data to describe the problem or new techniques to solve it, this timeline may extend significantly. "Solving really hard problems using AI takes time and depth. It follows a different curve" (Fabio Kuhn, Vortexa).

Fig. 4. Timescales typically increase non-linearly with desired accuracy

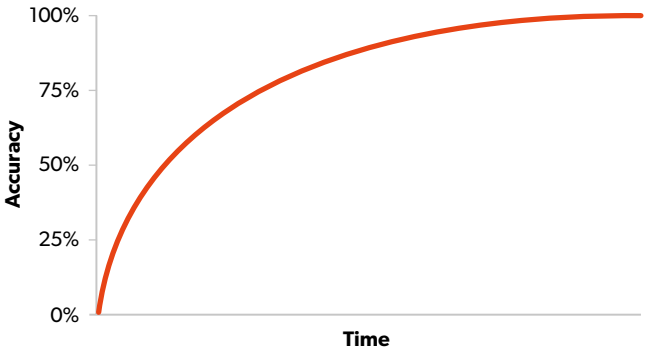

Source: MMC Ventures

"Solving really hard problems using AI takes time and depth. It follows a different curve. Endurance is key."

Fabio Kuhn, Vortexa

Align your budget with your goals and deployment strategy

The budget you require for your AI initiatives will depend upon multiple factors including:

- the nature, complexity and domain-specificity of the projects you undertake;
- available, and preferred, development strategies (use of third-party services versus an in-house development team);
- availability, quality and consistency of relevant data;
- a well-considered starting point;
- regulatory and ethical considerations to be addressed.

Costs will vary according to the development strategy you select.

Some challenges can be addressed with a readily-available third-party application programming interface (API). Others may be solved with a single pass of data through an existing, public domain network architecture. Others still will require extensive research and multiple iterations of training and adjustment to meet success conditions. Costs will vary according to the development strategy you select. The following strategies offer progressively greater functionality and uniqueness in return for increased spend:

- **Third-party APIs:** If another company has already solved your business problem, and you need only call the counterparty's service via an API to receive a result, prices can start as low as several hundred pounds. Using third-party APIs is the fastest way to deploy AI in your company and requires minimal time from your existing development team.
- **Bespoke third-party services:** To obviate the need for your own AI team, you can engage third-parties to develop and train your AI models. You will need to gather and prepare your own data and have a broad overview of the process of creating models. You are unlikely to require a budget of more than a few thousand pounds for training and running costs, plus the cost of an individual – ideally a data expert already in your business – with an understanding of AI to manage the process.

- **A small, in-house team:** A dedicated in-house AI team is likely to cost at least £250,000 to £500,000 per year, even for a small team. Whether you seek to repurpose publicly-available models, or solve unique problems, you will need to pay for: two to four individuals; the hardware they require to train and run their models and potentially extra hires for productionising the resulting system.
- **A large, in-house team:** An extensive team, recruited to solve problems at the edge of research, will require a multi-million-pound investment in personnel and hardware. This investment may yield a unique AI offering. It should only be considered as a first step, however, if your challenge cannot be solved with existing AI techniques and solutions, if you have access to unique data, and if you face significant restrictions on your ability to pass data to third parties.

We describe, in detail, the advantages and disadvantages of different development strategies in Chapter 4 (Development).

You may wish to develop a proof-of-concept, using your existing development team and third-party APIs or paid services, before creating a budgetary proposal. Most companies then start with the small, dedicated AI team.

Seek sponsorship from senior executives

Support from senior management in your organisation will be important for new AI initiatives to succeed. Your company may have a Board that strongly favours adopting AI; that sees AI as over-hyped and irrelevant; or has a healthy scepticism and seeks validation of benefits before assigning extensive resources. To build support within your company, define the focus of your first AI initiative and then set realistic goals. Your system will not, and need not, offer 100% accuracy. If it can save effort, even if results require human verification, you can deliver increased efficiency.

You can then present to senior management a plan that includes:

- a statement of the problem your AI will solve;
- a summary of outputs and benefits for the company;
- details of the nature and volume of data required;
- a viable approach with realistic timescales.

Leaders may be reluctant to invest in technology they do not understand. To achieve buy-in, it may be necessary to educate senior management regarding the benefits of AI while setting realistic expectations regarding timescales and results.

Anticipate and mitigate cultural concerns

When deploying AI, anticipate the potential for cultural resistance. For many in your team, AI will be unfamiliar. Some employees will see their workflows change. Many employees are concerned about the impact of AI on their job security.

Frequently, AI will enhance an individual's role by delivering what is termed 'Augmented Intelligence'. AI can bring new capabilities to an employee's workflow or free a human operator from repetitive, lower value tasks so he or she can focus on higher value aspects of their role.

Address concerns proactively by highlighting the ways in which AI will support individuals' goals and workflows – and enable your team to redirect their time to the most engaging aspects of their roles. "We go through a change management program to educate the workforce. We explain that AI takes care of repetitive tasks so people can focus on bigger things" (Dmitry Aksenov, DigitalGenius).

"We go through a change management program to educate the workforce. We explain that AI takes care of repetitive tasks so people can focus on bigger things."

Dmitry Aksenov, DigitalGenius

Address non-traditional security considerations

AI systems can be attacked in non-traditional ways. If a classification or grouping system is given an input beyond the scope of the labels on which it has been trained, it may assign the closest label it has even if the label bears little relation to the input. Causes of confusion, more broadly, may be exploited. Malicious individuals have manipulated system inputs to obtain a particular result, or to disrupt the normal practise of AI systems (for example, by spraying obscure road markings to confuse autonomous vehicles).

Protect against malicious activity via thorough system testing and exception handling, undertaken from the perspective of an individual deliberately attempting to undermine or exploit your system.

A long-term strategy should incorporate evolution and extension

When your first project is underway, anticipate the longer- term aspects of your AI strategy. Then "obsess about capabilities to make your vision come true over five to ten years" (Timo Boldt, Gousto). Your long term AI strategy should consider:

- **Maintenance:** To maintain your system's intelligence, regularly test results against live data to ensure results continue to meet or exceed your acceptance criteria. Set aside budget for future updates and retraining and monitor for performance degradation. Chapter 5 provides a blueprint for maintaining AI systems effectively.
- **Data:** "Remember that AI is a capability, not a product. It's always improving" (David Benigson, Signal). Monitor changes in your data over time. As your business grows or changes focus, data fields (including time series data, languages and product characteristics) will evolve and expand. Retraining your system regularly should be a component of your long-term AI strategy. To develop a comprehensive data strategy for AI, see Chapter 3.

"Remember that AI is a capability, not a product. It's always improving."

David Benigson, Signal

- **Algorithms:** AI techniques are developing rapidly; what you create today may be less accurate and slower than systems you develop in 12 months' time using the same data. Ensure a member of your team understands advances being made in AI and can advise on when to apply them to your use cases.
- **Scaling:** A plan to leverage your existing AI systems by extending their deployment to additional business units and geographies.
- **New initiatives:** A roadmap of new use cases for AI within your organisation to deliver increased cost savings, greater revenue or both.
- **Legislation:** Developments in AI are being monitored by legislative authorities (see Chapter 6). Develop a strategy to comply with new legislation as it emerges.

"Plan for the long term and then obsess about capabilities to make your vision come true over five to ten years."

Timo Boldt, Gousto

Chapter 2

People

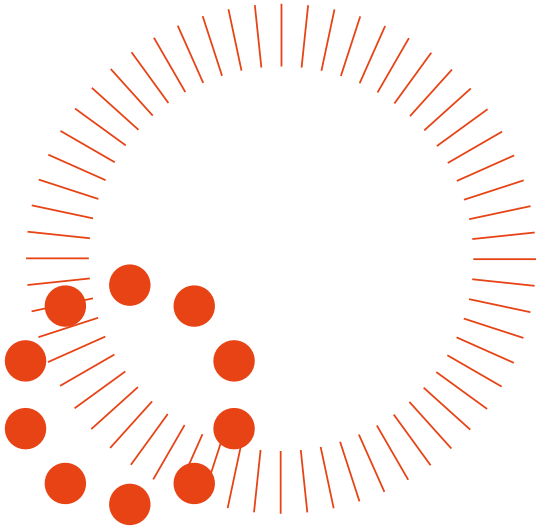

Summary

- In AI, job titles vary and can be difficult to interpret. We describe characteristics and salaries for six key roles: Data/Machine Learning Engineer; Data scientist; Machine Learning Researcher; Head of Data; Head of Research/ AI; and Chief Scientist/Chief Science Officer. For each, individuals' capabilities vary across competencies in research, engineering, production and strategy.

- The composition of your team should depend upon the problem being solved and your approach to doing so. It is advisable, however, to avoid hiring solo talent. Begin with a small team, and ensure you have a robust AI strategy in place before expanding your AI personnel.

- We suggest team structures, first hires and next steps for six scenarios: "I want insights into internal data"; "I want to implement third party AI APIs"; "I want to outsource AI development"; I want to create bespoke AI models"; "I want to use a combination of bespoke and third party AI"; and "I have an idea that's cutting edge."

- Recruiters, conferences and universities are primary sources of talent. Traditional recruitment agents find it difficult to screen AI candidates, so engage with specialist recruiters. Conferences and meetups are powerful vehicles for talent acquisition; be active in the AI community, attend and speak at conferences, and grow your network to discover capable candidates. Engage with universities; post on their job boards, establish partnerships and pay for projects to engage students who may seek future opportunities with you.

- Diversity delivers economic value and competitive advantage. Review the culture in your company, AI team and hiring practices to ensure diversity, representation and inclusion.

- An effective job description should emphasise projects (the nature of the engagements on which the successful candidate will work), skills and impact. Most data scientists seek work that will 'make a difference'. To attract talent, demonstrate how the successful candidate's work will do so.

- When hiring, prioritise adaptable problem-solvers. In addition to having role-specific and technical skills, a strong AI candidate will: understand available tools to enable rapid research and development; appreciate when to release an imperfect solution and when to wait; and communicate and collaborate well.

- Optimise every stage of your recruitment funnel. We provide best practices for: CV screening; phone screening; technical testing; face-to-face interviews and post-interview follow-up.

- AI talent is in short supply. Challenge, culture and company are key for retention. In addition to an attractive financial package, consider: offering flexible working hours; offering challenging problems and minimising drudgery through automation; creating a culture in which diverse ideas are shared; avoiding 'lone workers'; ensuring your AI team receives recognition for its work; and supporting team members' publishing and presentation of work.

People: The Checklist

Structure your team effectively

- ☐ Clarify the problem to be solved
- ☐ Identify a development strategy and associated hiring needs
- ☐ Understand the six core roles in AI teams
- ☐ Shape your team to reflect the competencies required
- ☐ Structure your team to avoid lone workers

Optimise your hiring process

- ☐ Understand the role, seniority and requirements for which you are hiring
- ☐ Develop a clear job description
- ☐ Leverage recruiters, conferences, meetups, universities and investors
- ☐ Embed best practices for screening, testing, interviews and follow-up
- ☐ Identify adaptable problem-solvers
- ☐ Recruit from diverse backgrounds

Invest in retention

- ☐ Maintain an inclusive and diverse culture
- ☐ Automate menial work
- ☐ Offer intellectually challenging problems
- ☐ Ensure the AI team receives recognition
- ☐ Support team members' publishing efforts

If you are building an in-house AI team, whether directly or via recruiters, it will be important to understand the roles you require and how to attract, deploy and maintain talent. Below, we provide a blueprint for structuring, building and retaining your AI team.

Hire for required competencies in engineering, production, research and strategy

Because AI is an emerging field, job titles vary and can be difficult to interpret. Further, people may describe themselves in different ways to market themselves for roles they want. There are at least six core roles in AI. In each, individuals' capabilities vary across competencies of research, engineering, production and strategy (Fig. 5).

Data Engineer/Machine Learning Engineer:

- understands data and can code AI models that are derivatives of systems already created
- focuses on engineering (creating code for applications and solutions to go live), not research
- create models but may lack finer understanding to push the boundaries of research.

Data Scientist:

- focuses on obtaining insight from data using scripts and mathematical techniques; manipulates data in a variety of programming languages for solutions to specific problems;
- typically has an academic background to PhD level
- stays current on contemporary research and be capable of implementing ideas from academic papers
- may lack wider development and AI skills, including understanding of the needs of live systems; typically produces reports, not applications.

Deep Learning Researcher/Machine Learning Researcher:

- focuses on research, not building business applications
- highly academic, typically with post-doctoral academic experience
- seeks to push the boundaries of technical solutions
- will have limited or no exposure of taking their work to the level of a live application.

Head of Data:

- understands the nuances of varying data sets and is sufficiently experienced to lead a team
- technically hands-on; works with her team to produce reports and applications
- may have data strategy responsibilities (responsibility for acquiring, managing and deriving value from data).

Head of Research/Head of AI:

- research-focused and with enough experience to to lead a team
- technically hands-on; may be sufficiently experienced to support the conversion of team output into live applications

Chief Scientist/Chief Science Officer/VP of AI:

- extensive experience in business as well as AI
- determine AI strategy and production pipelines; work with the Chief Technology Officer (CTO) to ensure the company's AI strategy can be executed
- typically report directly to the CEO
- experienced as a board level strategist.

Fig. 5. Roles vary across competencies of research, engineering, production and strategy

Role	Research	Engineering	Production	Strategy
Data Engineer	Low	High	High	Low
Data Scientist	Medium	Medium	Low	Low
Researcher	High	Low	Low	Medium
Head of Data	Medium	Medium	Medium	Medium
Head of AI	High	Medium	Medium	Medium
Chief Scientist	High	High	Medium	High

Source: MMC Ventures

Structure your team according to the problem and your approach

Salary range and job expectations will vary according to an individual's role (Fig. 6).

The composition of your AI team should depend on the problem being solved, your team's approach to doing so, and the level of integration required with your development team to support production. Bear in mind, however, the following principles:

- Do not hire solo AI talent, beyond an initial 'Head of…' role. AI professionals rely on collaboration for ideas and can feel isolated if they are a sole member of a larger team.
- Begin with a small team to validate the data and your team's ideas, regardless of the domain.
- Ensure you have a robust AI strategy in place (Chapter 1) before you expand your team.

Structure your team to avoid lone workers. AI professionals rely on collaboration.

Fig. 6. The expectations and costs of AI professionals differ

Job title	Salary range (£ thousands)	Job expectations
Data Engineer	45 - 90	• Directed problems • Full access to data • Create solutions that are deployed live
Machine Learning Engineer	60 - 90	• Challenging problems • Mix of APIs and models • Create solutions that are deployed live
Data Scientist	45 - 90	• Generate insights from data • Create models
Machine Learning Researcher	60 - 100+	• Find new methods • Solve complex AI problems • Create new models
Head of Data	80 - 120	• Own data strategy • Run the data team • Manage projects
Head of AI	80 - 120	• Define approaches • Run the AI team • Manage projects
Chief Scientist	110 - 180+	• Develop and deliver AI strategy • Board membership • Autonomy

Source: MMC Ventures

"I want insights into internal data"

- **Strategy:** Build an in-house data science team. Sensitive data will not leave your company and you can control the focus and outputs of your team.
  - » **First hire:** A Head of Data reporting to your CTO.
  - » **Next step:** Hire two or three data engineers or data scientists, or a combination, depending on the needs of the project.
  - » **Success factors:** These individuals will need time to understand your data and how it's gathered. Ensure they have access to all the data they need.

"I want to implement third-party AI APIs"

- **Strategy:** You will need individuals who understand your data and have the knowledge to implement and test third party APIs. If you have no budget to hire, an alternative approach is to find existing developers within your organisation who understand data well enough to manage the API integrations.
  - » **First hire:** Two machine learning engineers.
  - » **Next step:** For smaller projects, your engineers could report into the CTO or Head of Development. For larger projects, you may wish to hire a hands-on Head of AI as a team lead, to support the expanding team.
  - » **Success factors:** Review the available APIs and validate that they will address your use cases. Ensure you plan for changes to the APIs in future.

"I want to outsource AI development"

- **Strategy:** You need an individual who understands your project well to manage the outsourced relationship.
  - » **First hire:** A Head of AI, if you don't already have a suitable person within your team. A Head of AI can also enable you to bring the solution in-house, over time, if you choose to do so.
  - » **Next step:** Empower your Head of AI to manage project costs and timelines. Ensure you receive regular status reports for clarity on each delivery cycle.
  - » **Success factors:** Successful AI requires continuous feedback and iteration. Develop a good relationship with your AI provider and agree costs for updates up-front.

"I want to create bespoke AI models"

- **Strategy:** You are undertaking something unique with your AI solution and wish to keep your data and system knowledge in-house. This is the most common scenario for companies.
  - » **First hire:** A Head of AI, or Chief Scientist, reporting to the CTO.
  - » **Next step:** Let the Head of AI, or Chief Scientist, determine your strategy based on the problems you wish to address.
  - » **Success factors:** Allow budget for at least four further hires – more for larger projects. These hires will be a combination of Data Scientists and Machine Learning Engineers. You may need a Machine Learning researcher as part of this team – but ensure they are challenged enough.

"I'm going to use a combination of bespoke and third party AI"

- **Strategy:** You seek a fast start and third party APIs deliver enough for your minimum viable product. However, you want to develop bespoke AI in parallel, to deliver a unique value proposition.
  - » **First hire:** A Head of AI or Chief Scientist, plus at least two Data Engineers or Machine Learning Engineers to undertake the API work.
  - » **Next step:** Hire two or more Data Scientists.
  - » **Success factors:** Unless your hybrid approach involves examining research, avoid Deep Learning Researchers.

"I have an idea that's cutting edge"

- **Strategy:** Validate that your idea is feasible – as well as the problems and timelines associated with it.
  - » **First hire:** A Head of Research or Chief Scientist; two to three Deep Learning Researchers; plus potentially a Data Engineer to support them.
  - » **Next step:** If required, expand the team with Data Scientists to balance the team's skill-set.
  - » **Success factors:** Manage timelines closely and be prepared to assess when research isn't progressing to plan and alternative solutions should be explored. Maintain focus on the goal for the research and avoid research for its own stake.

Recruiters, conferences and universities are primary sources of talent

Unless you are a known company, advertising on your own website is unlikely to be effective. Alternative sources include:

1. **Recruiters:** Specialist recruiters exist for AI and data science. Poll your network to identify them. Unlike traditional development roles, recruitment agents find it challenging to screen AI candidates due to the research-intensive, data-specific problems they tackle. High quality AI recruiters have extensive networks of candidates, can identify candidates that fit your needs, and can save you more in the cost of your time than they charge in fees (typically 10%-45% of annual salary, plus bonuses).

2. **Conferences and Meetups:** Conferences and meetups are a powerful vehicle for talent acquisition. Be active in the AI community and grow your network. Every major city has an AI network you can join and there are conferences throughout Europe almost every week of the year. Many conferences offer job boards, in addition to which you can meet individuals at the event and undertake an initial screen. Even if timing does not align between you and potential candidates, making connections is valuable and you will become known as a potential employer.

   Consider speaking at events – it's easy to include a "we're hiring" closing slide and you may receive extensive interest. Some conferences require sponsorship for a speaking slot. This may be appropriate if you are discussing your general solution in a talk close to a sales pitch. You should never have to pay, however, if you have an advancement that will benefit the community. Submit your paper to an academic conference or one of the many high quality events that avoid sponsored talks and seek excellent speakers and and topics.

   Academic conferences – including NeurIPS, ICML, ICLR and AAAI – are busy events at which larger employers are highly active. These conferences are expensive and smaller companies can be overlooked. They are, however, valuable events if you seek exceptional researchers.

   Other conferences, including RE•WORK and M3, focus on the intersection of academic and business applications. These can be excellent environments in which to meet individuals who wish to move from academia to industry, as well as being stimulating environments for general conversation and ideation.

   Local meetups are even less formal and many offer events specifically for hiring. These are an excellent option if your team has not spoken before or if you wish to receive feedback on your approach. If there is not a meetup near you, start one.

3. **Universities:** If you are geographically close to a university with a strong AI department, or have alumni connections to one, the university may allow you to post on its job boards. While roles will be focused on students soon to graduate, alumni can also see the university's digital job boards and your role may be passed around networks of AI practitioners.

   You may also be able to work in partnership with a university by paying for projects. Typically, the head of a laboratory will accept a project for a fixed cost and the project will serve as a task for graduate students. Exercise care with this approach; the university may be motivated by gaining IP and publications, so ensure your agreement is appropriate. Students who have worked with you on a part-time basis in this manner are more likely to seek future opportunities with you. Similarly, there is a growing trend for Masters and PhD students to work part-time alongside their studies. If you can offer flexibility in this regard, you may attract exceptional candidates who are not seeking full-time work.

   Many universities offer excellent AI programmes and candidates. Examine their research pages and identify labs working on problems similar to yours. A small sub-set of universities with high quality AI programs includes:
   - **UK:** Bristol, Cambridge, Edinburgh, Imperial, Manchester, Oxford, Sheffield, Sussex, UCL.
   - **USA:** Carnegie Mellon, Harvard, MIT, Stanford, Yale.
   - **Canada:** Montreal, Toronto.
   - **Worldwide:** EPFL (Switzerland), Nanyang (Singapore), Politecnico de Milano (Italy), Technical University of Munich (Germany), Chinese University of Hong Kong.

4. **Investors:** If you have secured investment, leverage your investors. Ask them to introduce you to other AI-led companies in their portfolio so you can share ideas and recruitment opportunities. There may be excellent candidates who are no longer a fit for other companies – for example, due to a relocation – who would be ideal for yours.

5. **Job Boards:** Job boards can be effective at attracting applications. However, with a public posting on a generic job board you are likely to receive numerous applications that are poorer in quality or fit. The time, and therefore cost, required for an individual in your company to review them can be considerable. Specialist AI job boards, including those on Kaggle (www.kaggle.com/jobs) and StackOverflow (https://stackoverflow.com/jobs?q=AI) are typically searched only by people already part of these communities and typically offer higher-quality candidates.

Diversity delivers economic value and competitive advantage

Many problematic and embarrassing AI systems have been developed because the teams that created them lacked diversity. Without different perspectives on data and results, you may create systems that offer narrow appeal and broad offence. Review the culture in your company, AI team and hiring practices to ensure diversity, representation and inclusion.

Further, in a competitive market for AI talent, leadership in diversity will enable you to attract exceptional candidates who may otherwise have overlooked your position – or been minded to accept an alternative.

Job descriptions should emphasise projects, skills and impact

When hiring, ensure you understand the role, seniority and the minimum requirements for which you are hiring. If your team uses Python exclusively, for example, do not hire someone who only works in R or Matlab. Missing skills will impact your costs directly, as individuals take time to address gaps. Describe the projects on which the successful candidate will work. Do they relate to computer vision, language, prediction, generation or other? Use the industry-standard terms of classification, regression, generative AI, sequencing and clustering for easier comprehension. Describe the expectation for the role as well as the difficulty of the problem.

You will also need to sell your company and the domain. Most data scientists seek work that will 'make a difference'. To attract talent, frame your problem to demonstrate how the successful candidate's work will do so.

Fig. 7. Example job description

- **Title:** Deep Learning Researcher
- **Role:** Working on exclusive medical imagery, you will be solving classification and generative problems beyond the current state of the art.
- **Team structure:** As part of a dedicated AI team reporting to the Chief Scientist, you will work closely with Machine Learning Engineers who manage the AI production pipeline.
- **Skills:** Python 3.6, including numerical libraries; Tensorflow; Keras.

Source: MMC Ventures

When hiring, prioritise adaptable problem-solvers

Hiring exceptional talent is challenging. While it can be tempting to prioritise mathematics candidates with first class degrees, people with the largest number of academic papers, or those with the most appointments, consider your company's needs. Exclude individuals with poor communication or collaboration skills, and people who cannot adapt to the fast pace and fluid nature of industry. In addition to role-specific skills, a strong AI candidate will:

- have sufficient technical skills to solve AI problems
- understand available tools to enable rapid research and development processes
- appreciate when to release a solution, even if it is imperfect, and when to hold back a release
- communicate and collaborate well.

Seek adaptable, intelligent problem solvers – not individuals limited to following TensorFlow tutorials. Are you recruiting an individual to undertake research, or do you need

someone to obtain insights from data? Depending upon the problems to be solved, limited academic experience may be unproblematic if the individual possesses required skills. Individuals may have gained experience through alternative initiatives, including hackathons and competitions.

While the interview process for AI roles is similar to the process for other technical roles, there are differences. A conventional developer, for example, would undertake a technical test at a face-to-face meeting. AI candidates cannot demonstrate their ability to build an AI model at an interview given the time constraints. Involve existing members of your AI team throughout the process. The best candidates will complement existing ideas while bringing something new to your team.

Optimise each stage of the recruitment funnel

1. CV screen

- Maintain your specification for minimum skill-set. While no candidate will be perfect, identify your red lines and do not waste candidates' time, or yours, taking the wrong individuals forward.
- Prioritise candidates who have stayed over a year in past roles; it can take months for a new team member to understand the data specific to your business.
- Evaluate candidates' ability to contribute to your business above their academic experience, even for research posts. An individual with a decade of experience researching an obscure problem may not adapt to your natural language challenge.

2. Phone screen

- Identify candidates' passions and motivations; evaluate if they will be a good fit for the projects you have planned.
- Ask candidates about their contributions to previous projects and seek individuals who can explain their contributions clearly.
- Let the candidates ask you questions for half the available time; good candidates will want to know as much about the company, team and projects as you will about them.
- Progress only the candidates who can demonstrate their passions and exhibit good collaboration and communication skills.

3. Technical test

- A technical test (Fig. 8) will be expected but it is important to be reasonable. Do not set a task that requires more than four hours to complete; candidates' time is limited and you should not take advantage of them.
- Offer a problem representative of the work they would undertake in your team, with (subject to privacy constraints) real data. Ideally the problem should have a trivial solution, which will highlight individuals who do not consider data complexity.
- Technical tests should be specific to the candidate; individuals can upload information about your tests to websites, such as Glassdoor, which can give candidates an unfair advantage. Similarly, recruitment agents may prime their favoured candidates with the problem upfront. Kaggle.com can be an appropriate environment in which to run technical tests.
- For research-based roles, implementing code from an academic paper may be a suitable test.
- If a candidate offers a solution which demonstrates that the candidate thinks beyond the trivial, invite them for a face-to-face interview.

Fig. 8. Example technical test

"Write a script to identify and remove duplicates in the following data set."

The candidate is given a set of 50 frames from a video. Some are identical; some close; some have the same composition but different subjects; and some are unique.

- A strong candidate will understand that this is a data preparation problem and will consider the impact of this data on training or testing a model.
- A trivial solution, indicative of a lack of experience, would be to identify identical images.
- A better solution would group the images, identify a dominant image from each group as an output, and discard the rest.
- An excellent solution would understand that the object of the images may be important and provide a script in which the important characteristics could be selected.

Source: MMC Ventures

4. Face-to-face interview

- By this stage you should have a small number of exciting candidates.
- Discuss each candidate's technical test. Can they critique their own solutions? What would they do given more time? These questions will provide insight into how candidates think and plan their time.
- Add a thought experiment with extension challenges: how would the candidate solve a problem in sub-optimal circumstances? What if there are large gaps in available data, or if data quality varies? What if the business required a 50% improvement in predictive speed? Thought experiments will enable you to understand a candidate's creativity and how they will perform in a dynamic environment. If a candidate can only follow steps described in AI tutorials, their impact on your business will be limited. Similarly, exercise caution with candidates who express annoyance when faced with changing business requirements, or who advocate long timelines for any change. They may not have the skills and temperament to thrive in an early stage company.
- For research-led roles, or where you seek a candidate who will digest state-of-the-art academic papers, ask the candidate to bring and present a recent paper written by someone else. Evaluate whether the candidate can explain another person's concepts in simple terms. Invite an interested non-technical person to join this part of the interview and ask simple questions. Screen out candidates who cannot communicate the work to the wider business, or become frustrated with simple questions.
- Discuss, in as much detail as possible, upcoming projects. See how excited the candidate becomes and prioritise candidates that are eager and propose solutions.

5. Post-interview

- Provide immediate feedback and, if you still have other candidates to consider, manage expectations regarding a decision.
- AI practitioners are data-led and do not appreciate uncertainty. Come to a decision promptly and make an offer quickly.

Challenge, culture and company are key for retention

AI talent is in short supply. When you have attracted high quality professionals to your team, it is important to retain them. While an attractive financial package and benefits are necessary, large companies can – and will – offer higher salaries. Retain AI talent by catering to team members' other needs:

- Offer flexible working hours. AI models can take a long time to run; if they finish at night or during weekends, good candidates will want to re-engage and alter parameters.
- Offer challenging problems and minimise drudgery. Steps that can be automated should be. Ensure your team has support from other parts of the business to do so.
- Ensure your team has appropriate hardware. Crippling your team to save £1,000 is a false economy.
- Create a culture in which intellectual debate is encouraged and diverse ideas are shared. Advances in AI are the result of multiple scientific fields bringing different perspectives to the same problem. Individuals with different backgrounds and education see things differently; combining their ideas will present novel solutions. An environment in which all opinions can be voiced and debated will enable your AI team to solve problems faster and motivate your team.
- Regardless of level, ensure your AI team comprises more than one person. The 'lone AI worker' is a frequent challenge for early stage companies. The scientific nature of AI instils a need for collaboration and the testing of ideas. While it can be exciting for an AI team member to be the first in your company, and to develop your company's prototype, months of solo work on your company's AI initiative, even amidst collaboration with your broader team, can be intellectually isolating and drive attrition.
- Ensure your AI team receives recognition for its work. If individuals have worked for months to develop an effective AI model, it will be dispiriting for the team that adds the front-end to receive sole credit.
- Decide early your approach to intellectual property that projects will produce. Ensure your team understands whether there is a patent strategy, if team members may publish results, and if they can present them at conferences. Many AI practitioners have academic careers they wish to sustain. If your company can provide support to their efforts to publish and present, it will be deemed a benefit. Balance this, however, with an understanding that developing research to the standard of an academic paper may require further work that will not benefit your company.

Chapter 3

Data

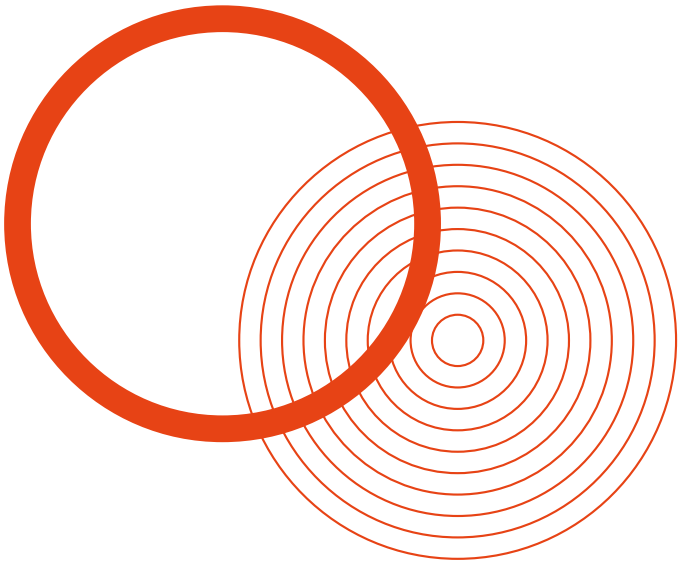

Summary

- For effective AI, develop a data strategy. A data strategy spans: data acquisition & processing; quality; context; storage; provisioning; and management & security. Define your data strategy at the outset of your AI initiative.

- Accelerate data acquisition by using multiple sources. Developers draw on several sources including: free resources (such as dataset aggregators); partnerships with third parties (companies, universities, data providers and government departments); and create new, proprietary data.

- A high-quality data set has appropriate characteristics to address your business challenge, minimises bias and offers training data labelled with a high degree of accuracy. Develop a balanced data set – if you possess significantly more samples of one type of output than another, your system will exhibit bias.

- Primary forms of bias are: unwarranted correlations (between inputs and output classifications); erroneous assumptions which cause relationships to be missed ('underfitting'); and modelling noise instead of valid outputs ('overfitting'). Adjust for overfitting and underfitting by using different data volumes and model structures. Remove unwarranted correlations through testing.

- Ensure that the results of your internal testing will be maintained when applied to real-world data. Test early, and frequently, on expected live data.

- Managing 'dirty data' is data scientists' most significant challenge (Kaggle). Smaller volumes of relevant, well-labelled data will typically enable better model accuracy than large volumes of poor-quality data. To label data effectively: consider developing a supporting system to accelerate data labelling and improve accuracy; draw on existing AI and data techniques; and seek data labelled by multiple individuals to mitigate mislabelling.

- Understand the data you use. Ensure you capture the human knowledge regarding how your data was gathered, so you can make downstream decisions regarding its use. Capture data provenance (where your data originated and how it was collected). Define your variables (differentiate between raw data, merged data, labels and inferences). Understand the systems and mappings through which your data pass to retain detail.

- Store and structure data optimally to support your objectives. Storage options include basic file-based, relational, NoSQL or a combination. When selecting storage plan for growth in data volume, updates, resilience and recoverability.

- One in three data scientists report that access to data is a primary inhibitor of productivity (Kaggle). Develop a provisioning strategy that: ensures data is accessible across your organisation when needed; contains safeguards to protect your company against accidents; optimises system input/output; and maintains data freshness.

- Implement robust data management and security procedures consistent with local and global regulations. Personal data is protected by UK and EU law and you must store it securely. Draw on principles of appropriate storage, transmission and minimum required access.

Data: The Checklist

Formulate a data strategy

- ☐ Develop a data strategy
- ☐ Review your data strategy quarterly

Optimise acquisition & processing

- ☐ Ensure your data collection is legal
- ☐ Preserve detailed fields
- ☐ Check you have included real world data

Develop a high-quality data set

- ☐ Confirm you have enough examples of data classes for fair predictions
- ☐ Understand variance in data required to solve your business challenge
- ☐ Identify sources of bias in your data
- ☐ Follow best practices for labelling data

Understand data context

- ☐ Document the sources of your data
- ☐ Add metadata to capture data collection methods

Store data optimally

- ☐ Forecast expected growth in data
- ☐ Evaluate methods of storage and access
- ☐ Develop and test a resilience plan

Provision data appropriately

- ☐ Ensure data requests do not block the addition of new data
- ☐ Develop a plan to archive stale data so access remains fast

Optimise management and security

- ☐ Ensure staff have the minimum access they require to perform their role
- ☐ Use multi-factor authentication
- ☐ Undertake regular penetration tests to validate your security
- ☐ Appoint an individual responsible for compliance with legislation

A data strategy will enable your company to acquire, process, govern and gain value from data effectively. Without a data strategy, your team's efforts will be greater than necessary, risks will be magnified and chances of success will be reduced.

Develop a data strategy for effective AI

"Data is the lifeblood of any AI system. Without it, nothing happens" (David Benigson, Signal). There are six components of an effective data strategy (Fig. 9.):

1. **Acquisition & Processing:** Obtain and process the data you need to develop effective prototypes and algorithms.
2. **Quality:** Develop a data set that has the appropriate characteristics to address your business challenge, minimises bias and offers training data labelled with a high degree of accuracy.
3. **Context:** Understand the provenance of your data and the mappings through which it passes so you use and share it effectively within your company.
4. **Storage:** Store and structure your data appropriately to support your objectives regarding access, speed, resilience and compliance.
5. **Provisioning:** Optimise the accessibility of data to your team and the implementation of safeguards.
6. **Management & Security:** Manage data security, access and permissioning to ensure appropriate use of your data stores.

Define your data strategy at the outset of your AI initiative. Review it quarterly and update it as product requirements change, your company grows or you are impacted by new legislation.

Fig. 9. The six components of an effective data strategy

Source: MMC Ventures

"Data is the lifeblood of any AI system. Without it, nothing happens."

David Benigson, Signal

"Build access to data at scale from day one."
David Benigson, Signal

Accelerate data acquisition by using multiple sources

Obtaining data to develop a prototype or train your models can be a lengthy process. Ideally, you will possess all the data you need at the outset and have a data strategy to govern its access and management. In the real world, neither is likely. Working on the project may highlight missing data.

"Build access to data at scale from day one" (David Benigson, Signal). Filling the gaps from your own initiatives can take months, so use multiple approaches to accelerate progress. Developers typically draw on several approaches to source data (Fig. 10) including free resources (such as dataset aggregators), partnerships with third parties and the creation of new, proprietary data.

- Use free resources: Evaluate data sources that already exist and are free to use. Kaggle (www.kaggle.com), a large community of data scientists and machine learning engineers, regularly posts data sources for competition experiments. These can be useful for prototyping and initial training of machine learning algorithms. Google Dataset Search (https://toolbox.google.com/datasetsearch) can help you find specific data sets – be they weather in London or public transport statistics for Manchester. Further, many authors of academic papers are now uploading sample code and data sets (either raw data or locations to acquire it) to platforms such as GitHub. These data sets are frequently used for benchmarking. Not all of the datasets from the above sources are free for business use, so check that your use of them is appropriate.
- Develop partnerships: Develop partnerships with other organisations – other companies, universities, data providers or government departments. Establishing a mutually beneficial relationship can offer your company exclusive data and associated benefits.
- Create data: The data you seek may be unavailable or prohibitively costly. You may need to invest time and resource to create the data you need – and a quarter of data scientists do so. The approach – embedding sensors, taking photos or videos, undertaking surveys or labelling existing datasets – will vary according to your industry and use case. Proprietary data is valuable – which is why so little is free. Developing your repository of proprietary data will yield value and defensibility over time.

You will need to de-duplicate and merge your data from multiple sources into a single, consistent store. New data must follow a comparable process so your data remains clean. If you merge fields, or decrease the precision of your data, retain the original data. Being able to analyse gaps in your data will enable you to plan future data acquisition and prioritise addressable business use cases.

Fig. 10. AI developers use multiple approaches to source training data

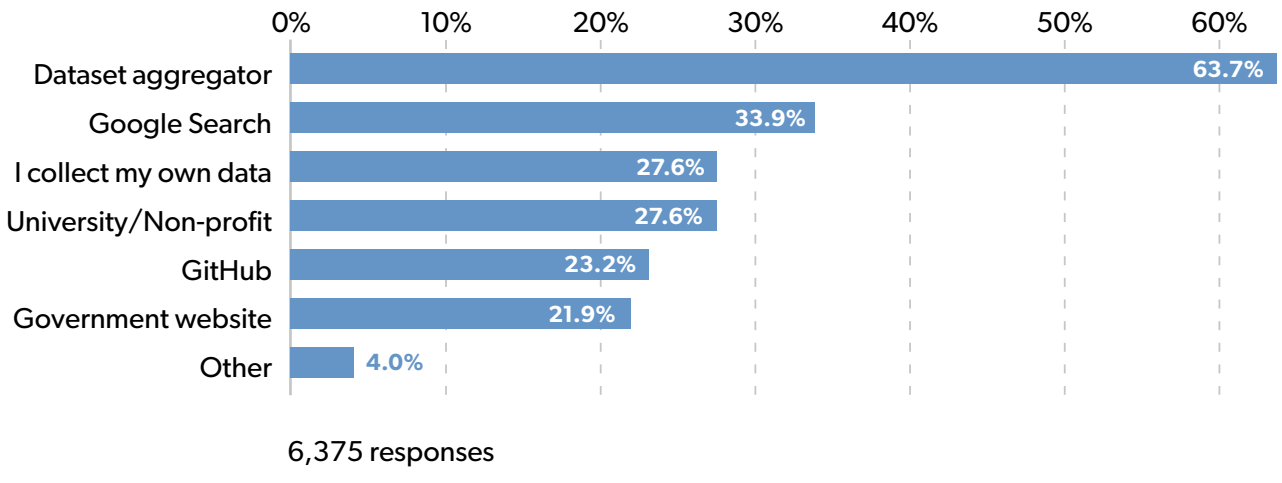

Source: Kaggle

Develop a balanced, well-labelled data set

A high quality data set has appropriate characteristics to address your business challenge, minimises bias and offers training data labelled with a high degree of accuracy.

It is important to develop a balanced data set. If you possess significantly more samples of one type of output than another, your AI is likely to exhibit bias. You can decide whether your system's bias will tend towards false positives or false negatives, but bias will be inevitable. There are three primary forms of bias in AI (Fig. 11):

Fig. 11. Three types of bias in AI

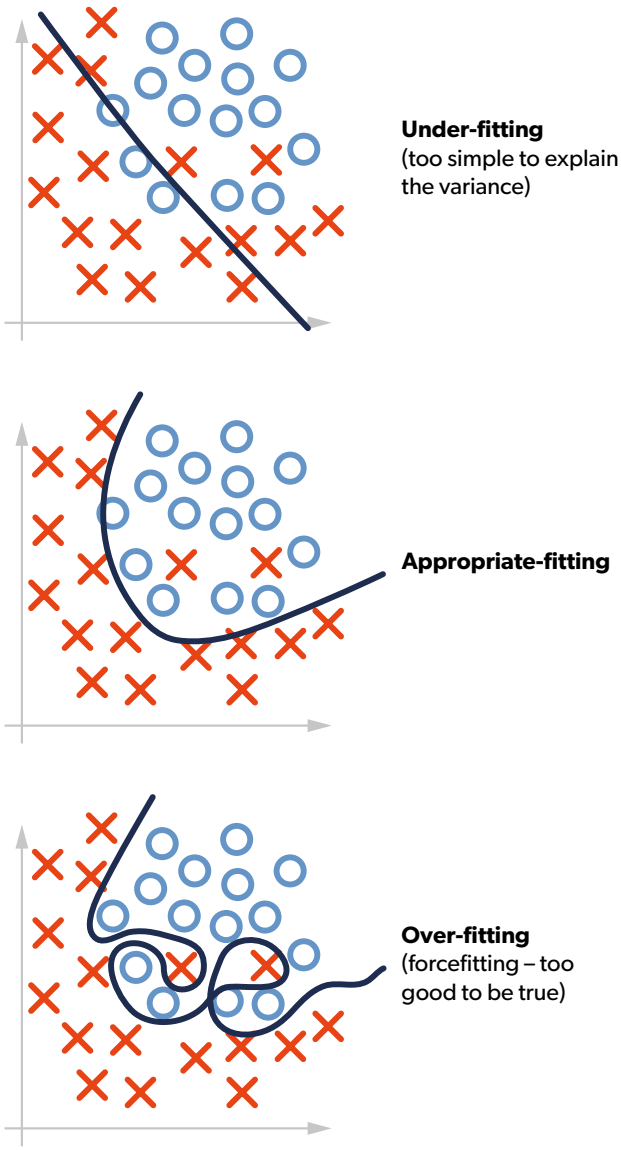

Source: Victor Lavrenko

1. Unwarranted correlations between inputs and output classification. Systems that offer jobs based on gender rather than skills, or provide or decline financial products based on ethnicity, are examples of unwarranted correlations resulting from unrepresentative input data.
2. Erroneous assumptions in learning algorithms, which result in relevant relationships being missed – so-called 'underfitting' (Fig. 12, overleaf). If underfitting, you have not sufficiently used the power of your data. If you seek to predict rental prices for properties, and base your model only on the number of bedrooms a property has, your predictions will perform poorly; your model will ignore important characteristics such as location, whether a property is furnished, and whether a property offers parking or a garden.
3. Modelling noise instead of valid outputs – "overfitting". An overfitted model takes account of so many details in the data that it cannot make accurate predictions. Considering all the health related data of a group of people, for example, will include so much natural variation in weights, blood pressures and general levels of fitness that predicting any characteristics or a new member of the group would be inaccurate.

Be aware of bias in your data and models to take appropriate action and minimise its impact. Overfitting and underfitting can be adjusted with different data volumes and model structures. Unwanted correlations are frequently more critical to the business; in addition to erroneous results they can lead to negative publicity. Test models thoroughly to ensure that variables that should not affect predictions do not do so. If possible, exclude these 'protected variables' from the models completely.

If you possess significantly more samples of one type of output than another, your AI system is likely to exhibit bias.

If the features you seek are rare, it can be challenging to achieve a balanced data set. You wish to develop a model that can deal with rare occurrences effectively, but not be overfit. You may be able to use artificial data, but not when the artefacts in the artificial data themselves impact the model. You may also choose to retain some overfit or underfit bias – and opt for a greater proportion of false positives or false negatives. If you err on the side of false positives, one solution is to let a human check the result. The bias you prefer – false positives or false negatives – is likely to depend on your domain. If your system is designed to recognise company logos, missing some classifications may be less problematic than incorrectly identifying others. If identifying cancerous cells in a scan, missing some classifications may be much more problematic than erroneously highlighting areas of concern.

Fig. 12. The problem of overfitting

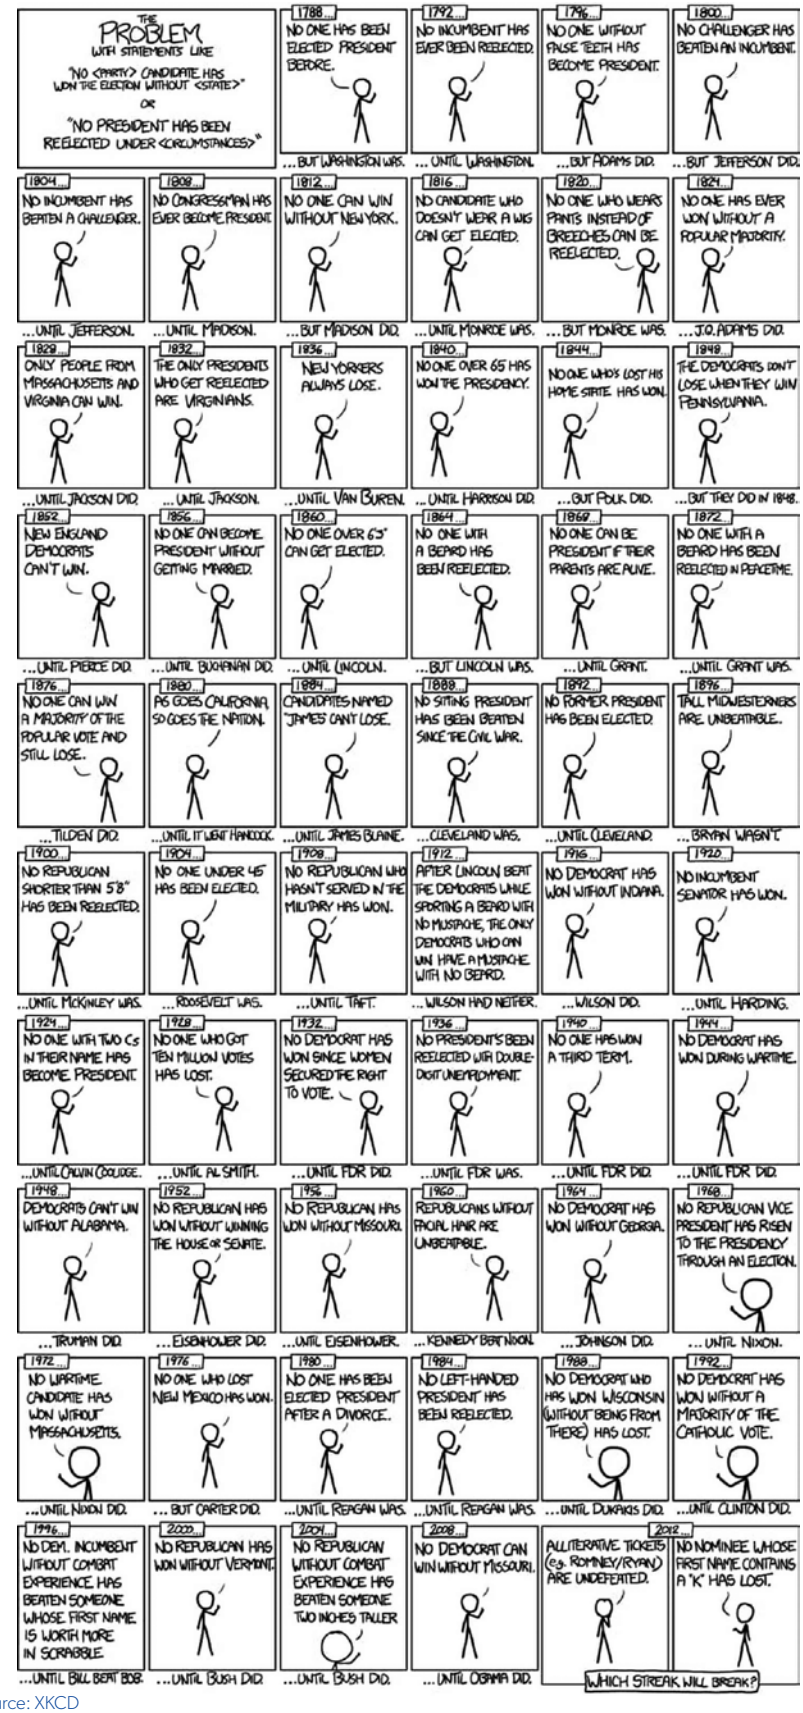

Source: XKCD

It is critical to ensure that the results of your internal testing are maintained when applied to real-world data. 99% accuracy on an internal test is of little value if accuracy falls to 20% when your model is in production. Test early, and frequently, on real-world data. "If you don't look at real-world data early then you'll never get something that works in production" (Dr. Janet Bastiman, Chief Science Officer, Storystream). Before you build your model, put aside a 'test set' of data that you can guarantee has never been included in the training of your AI system. Most training routines randomly select a percentage of your data to set aside for testing, but over multiple iterations, remaining data can become incorporated in your training set. A test set, that you are sure has never been used, can be reused for every new candidate release. "When we're looking at images of vehicles, I get the whole company involved. We all go out and take pictures on our phones and save these as our internal test set – so we can be sure they've never been in any of the sources we've used for training" (Dr. Janet Bastiman, Chief Science Officer, Storystream). Ensure, further, that your 'test set' data does not become stale. It should always be representative of the real-world data you are analysing. Update it regularly, and every time you see 'edge cases' or examples that your system misclassifies, add them to the test set to enable improvement.

Data scientists report that managing 'dirty data' is the most significant challenge they face (Kaggle). Smaller volumes of relevant, well-labelled data will typically enable better model accuracy than large volumes of poor quality data. Ideally, your AI team would be gifted data that is exhaustively labelled with 100% accuracy. In reality, data is typically unlabelled, sparsely labelled or labelled incorrectly. Human-labelled data can still be poorly labelled. Data labelling is frequently crowdsourced and undertaken by non-experts. In some contexts, labelling may also be intrinsically subjective. Further, individuals looking at large volumes of data may experience the phenomenon of visual saturation, missing elements that are present or seeing artefacts that are not. To mitigate these challenges, companies frequently seek data labelled by multiple individuals where a consensus or average has been taken.

To label data effectively, consider the problem you are solving. 'Identify the item of clothing in this image', 'identify the item of clothing in this image and locate its position' and 'extract the item of clothing described in this text' each require different labelling tools. Depending upon the expertise of your data labelling team, you may need a supporting system to accelerate data labelling and maximise its accuracy. Do you wish to limit the team's labelling options or provide a free choice? Will they locate words, numbers or objects and should they have a highlighter tool to do so?

Embrace existing AI and data techniques to ease the data labelling process:

- For visual classification use a generic object recognition tool, such as ImageNet, to identify relevant categories of images (such as cars) and the location of the object in an image. You can then show your labellers the image with a highlighted area and ask about the highlighted object to make a deeper classification (such as model).
- For natural language processing you may be able to draw on existing textual content and classifiers, such as sentiment analysers, to sort data into broad categories that a person can verify and build upon for downstream applications.
- Use clustering techniques to group large volumes of similar data that can be labelled together.

"If you don't look at real-world data early then you'll never get something that works in production."

Dr Janet Bastiman, StoryStream

Understand data context by capturing human knowledge

It is critical to understand the data you use. Using a number labelled "score" in your database is impractical – and may be impossible if you do not know how it was derived. Ensure you capture the human knowledge of how data was gathered, so you can make sound downstream decisions regarding data use.

Your data strategy should ensure you:

- **Understand data provenance:** It is imperative to understand where your data originated, how it was collected and the limitations of the collection process. Does data relate to current customers only, or a spread of the population? Are the images or audio you use raw or have they already been digitally edited?
- **Define your variables:** Defined variables should enable you to differentiate between raw data, merged data, labels and inferences (such as assuming an individual's gender from their title).
- **Understand systems and mappings** through which data have passed. As you process data through multiple systems and mappings, problems can arise – much as photocopies of a photocopy begin to degrade. For example, if a system has a date of birth field which is imported into a system that requires age instead, the mapping will be accurate at the time of processing but information has been lost and the quality of the data will degrade over time. If this is then mapped to a system that uses an age range, accuracy will be regained but at the expense of precision. Ensure your mappings retain detail.

Ensure you capture the human knowledge of how data was gathered, so you can make sound downstream decisions regarding data use.

Understanding the context of your data will depend upon process and documentation more than tooling. Without an understanding of the context in which data was collected, you may be missing nuances and introducing unintended bias. If you are predicting sales of a new soft drink, for example, and combine existing customer feedback with data from a survey you commission, you must ensure you understand how the survey was conducted. Does it reflect the views of a random sample, people in the soft drinks aisle, or people selecting similar drinks? It is important to understand the information not explicitly expressed in the data you use. Documenting this information will improve your understanding of results when you test your models. Investigating data context should prompt your employees to ask questions – and benefit from their differing perspectives. If you lack diversity in your team, you may lack perspectives you need to identify shortcomings in your data collection methodology. Ensure team members deeply understand your company's domain as well as its data. Without deeper knowledge of your domain, it can be challenging to know what variables to input to your system and results may be impaired. If predicting sales of computer games, for example, it may be important to consider controversy, uniqueness and strength of fan base in addition to conventional variables.

Store and structure data optimally to support your objectives

Your data storage strategy will impact the usability and performance of your data. The nature of your data, its rate of growth and accessibility requirements should inform your approach.

Types of storage include basic file-based, relational and No Structured Query Language (NoSQL):

- **Basic file-based:** Whether a cloud-based solution – such as Amazon Web Services (AWS) or HotBlob – or in-house, basic file-based storage has no limitations on file size – but is slow to search and search requests are typically based simply on file name, size or creation date.
- **Relational:** Relational databases (including MySQL or Oracle) can store extensive information in separate tables related to one another. Relational databases are well suited to defined information, with strict definitions, that can be grouped into tables. While powerful in their ability to enable complex queries, and offering security down to

Your data storage strategy will impact the usability and performance of your data.

the field level, relational databases can struggle with large data items (including images and documents) and prove challenging to scale.

- **NoSQL:** Recently, NoSQL databases (such as Mongo or Redis) have become popular because they do not demand the field restrictions associated with relational databases. NoSQL databases are effective for storing large volumes of hierarchical data. Accordingly, they are commonly associated with 'big data' initiatives. NoSQL databases can easily be scaled by adding extra machines to your system ('horizontal scaling'), but struggle to enable complex queries due to the way in which they store data.

The store you select will influence the performance and scalability of your system. Consider mixing and matching to meet your needs – for example, a relational database of individuals with sensitive information linking to data stored in a more accessible NoSQL database. The specific configuration you choose should depend upon the data types you will store and how you intend to interrogate your data.

To plan for growth and updates:

- Forecast increases in data volume. If starting with existing data, you will understand current data volumes and how much new data you are adding each day. If starting from scratch, you will need to estimate data growth based on a forecast of incoming data. Armed with an estimate of data growth, you can determine the volume of storage you will require for data for the first year of your project.
- Cloud solutions will enable you to store as much data as you wish – but balance the cost of immediate– and long-term storage (on AWS, the difference between S3 and Glacier). If operating your own hardware, you will also need to decide whether to archive data away from your primary store. You may need to maintain physically separate data stores for select personal data, to ensure its isolation.
- Monitor costs, remaining storage, and system performance so you can act before costs become prohibitive or you run out of storage space. For relational databases this is critical, because scaling is likely to require you to upgrade the hardware on which your database is operating. For NoSQL systems, it will be easier to scale horizontally.

For resilience and recoverability:

- Treat resilience as mission-critical. Data is the most valuable component of your AI strategy; if your data were lost, you could not rebuild your models and would lose a significant proportion of your company's uniqueness and value.
- While large companies will have dedicated resources and specialist skills, startups and scale-ups must also plan for resilience and recoverability.
- Ensure regular backups. Storage is inexpensive and accessible to every company.
- The degree of resilience you require will depend upon whether it is critical for your data store to be permanently available for read and write. Resilient systems will duplicate your data, so a replica can take over seamlessly if part of your system fails. Further, resilient systems typically load balance to ensure multiple requests do not cause delays.
- Many cloud providers offer resilient systems as part of their service. While most data centres have their own generators and redundant internet connectivity, significant events such as hurricanes and earthquakes can cause hours, or even days, of disruption. Other risks, including cascading software failures, can also crystallise. Depending upon the criticality of your data access you may also seek a separate provider, with a backup, that you can invoke in the event of a major disaster. If you manage your own data storage, you must manage recoverability as a minimum. Store backups in a separate geographic location and regularly test that you can restore them successfully. Your first disaster is not the time to learn that your backups have been failing silently.

When provisioning data consider access, safeguards and data freshness

One in three data scientists report that access to data is a primary inhibitor of productivity (Kaggle). Data provisioning – making data accessible to employees who need it in an orderly and secure fashion – should be a key component of your data strategy. While best practices vary according to circumstance, consider:

- Access: Your data science team will become frustrated if they are waiting for another team to provide them with data. Providing them with tools for direct access may be valuable. Most data stores offer only full administrative access or expert level tooling. You may need to allow time and resource to implement a specific solution for your team.
- Safeguards: Protect your company against accidents. Ensure data access is read-only. Except for an administrator, no-one should be able to delete or change data.
- Input/output: Reading data from your systems must not block the addition of new data. Similarly, if your data store is being continually updated, your team should not have to wait for a significant period before they can extract the data they require.

Stale data can be a significant challenge and is a key consideration when planning your provisioning strategy. If you are analysing rapidly-changing information, decide how much historical data is relevant. You might include all data, a specific volume of data points, or data from a moving window of time. Select an approach appropriate for the problem you are solving. Your strategy may evolve as your solution matures.

If you are correlating actions to time, consider carefully the window for your time series. If you are predicting stock levels, a few months of data will fail to capture seasonal variation. Conversely, if attempting to predict whether an individual's vital signs are deteriorating, to enable rapid intervention, an individual's blood pressure last month is likely to be less relevant. Understand whether periodic effects can impact your system and ensure that your models and predictions are based on several cycles of the typical period you are modelling. Pragmatically, ensure your access scripts consider the recency of data you require to minimise ongoing effort.

Implement robust data management and security procedures

Data management and security are critical components of a data strategy. Personal data is protected by UK and EU law and you must store it securely.

You may need to encrypt data at rest, as well as when transmitting data between systems. It may be beneficial to separate personal data from your primary data store, so you can apply a higher level of security to it without impacting your team's access to other data. Note, however, that personal data included in your models, or the inference of protected data through your systems, will fall under data protection legislation.

Establish effective data management by building upon the principles of appropriate storage and minimum required access.

- Physical Access: Direct access to your data store should be tightly limited to key, trusted individuals. Individuals with the highest level of access to your systems will frequently be targets for malicious third parties.
- Users: Employees' needs regarding data access will vary. If individuals do not need to view sensitive data, they should not have the ability to view or extract it.
- Applications: Other systems that connect to your data store should also be treated as virtual users and restricted. Many companies fail to restrict application access and suffer adverse consequences when there is an error in a connected application or the application's access credentials are compromised.

Additionally:

- Use multi-factor authentication as broadly as possible.
- Log every access request with the identity of the requester and the details of the data extracted.
- Hire a third party to undertake penetration testing to validate the security of your systems.

If an individual resigns, or has their employment terminated, immediately revoke access to all sensitive systems including your data. Ensure that employees who leave cannot retain a copy of your data. Data scientists are more likely to try to retain data to finish a problem on which they have been working, or because of their affinity for the data, than for industrial espionage. Neither is an appropriate reason, however, and both contravene data protection law. Ensure your team is aware of the law and that you have appropriate policies in place.

Chapter 3
Data

Chapter 4

Development

Summary

- There are many ways your company can engage with AI. Use third party AI APIs; outsource; use a managed service; build an in-house team; or adopt a 'hybrid' approach combining an in-house team with third party resources.

- Third party AI APIs fulfil specific functions to a moderate or high standard at low cost. Most solve problems in the domains of vision and language. Numerous APIs are available from Amazon, Google, IBM, Microsoft and also other smaller companies. Features vary; we provide a summary. APIs offer immediate results without upfront investment, at the expense of configurability and differentiation. Use an API if you seek a solution to a generic problem for which an API is available. APIs are unsuitable if you seek solutions to narrow, domain-specific problems, wish to configure your AI, or seek long-term differentiation through AI.

- Managed services enable you to upload your data, configure and train models using a simple interface, and refine the results. Managed services abstract away much of the difficulty of developing AI and enable you to develop a custom solution rapidly. Managed services offer greater flexibility and control than APIs, but less flexibility than an in-house team, and also require you to transfer data to a third party and may create dependencies.

- If a third-party solution is unavailable and an in-house team is too expensive, you can outsource your AI development. Whether outsourcing is appropriate will depend upon your domain, expertise, required time to value and data sensitivity. If outsourcing, specify desired frameworks and standards, who will provide training data, costs, timescales and deployment considerations. Outsource if you require trusted expertise quickly and a cheaper option than permanent employees. Avoid outsourcing if your data permissions prohibit it, you require domain or sector knowledge that an outsourcer lacks, or you wish to build knowledge within your own company.

- An in-house AI team offers maximum control, capability and competitive differentiation – at a price. A small in-house team will cost at least £250,000 to £500,000 per year. A large team requires a multi-million-pound annual investment. To develop an in-house team your company must also: attract, manage and retain AI talent; select development frameworks and techniques; gather and cleanse data; learn how to productise AI into real-world systems; and comply with regulatory and ethical standards. Build an in-house team if you have a problem that cannot be solved with existing solutions, seek differentiation in the market, or seek to maintain control over your data.

- A 'hybrid' approach is ideal for many companies. Plan for an in-house team that will address your requirements to a high standard over time, but use third party APIs to solve an initial, simpler version of your challenge. A hybrid approach can be attractive if you seek rapid initial results, wish to limit spend until a business case is proven and want greater differentiation and resilience over time.

- To develop AI yourself you have choices to make regarding your AI 'technology stack'. The stack comprises six layers: hardware; operating systems; programming languages; libraries; frameworks; and abstractions. Not all problems require the full stack.

- Ensure your team has hardware with graphical processing units (GPUs) that support NVIDIA's CUDA libraries. Laptops with high performance graphics cards offer flexibility. For greater power, desktop machines with powerful GPUs are preferable. To train large models, use dedicated servers. Cloud-based servers offered by Amazon, Google or Microsoft are suitable for most early stage companies.

- Apply AI techniques suited to your problem domain. For assignment problems consider: Support Vector Classification; Naïve Bayes; K-Nearest Neighbour Classification; Convolutional Neural Networks; Support Vector Regression; or 'Lasso' techniques. We describe each and explain their advantages and limitations. For grouping problems, explore: Meanshift Clustering; K-Means; and Gaussian Mixture Models. For generation, consider: Probabilistic Prediction; Variational Auto-Encoders; and Generative Adversarial Networks.

Development: The Checklist

Create a development strategy

- ☐ Review the advantages and limitations of different development strategies
- ☐ For your AI initiatives, assess the relative importance to your company of time to value, capability, cost, differentiation, resilience and the development of in-house expertise
- ☐ Determine the availability of APIs that address your requirements
- ☐ Assess whether your data permissioning allows use of third party services
- ☐ Validate that your chosen development strategy offers the trade-offs, integrations with existing systems and resilience your organisation requires
- ☐ If developing an in-house team, review best practices with regard to strategy, people, data, development, production and regulation (Chapters 1 to 6)

Optimise system development

- ☐ Ensure your team has appropriate hardware for rapid iteration and review ongoing hardware requirements
- ☐ Match the language you use with the rest of your production activity for simplicity and speed
- ☐ Understand techniques appropriate for your problem domain (generation, assignment, grouping or forecasting)
- ☐ Experiment with multiple techniques to validate your challenge and highlight characteristics and limitations of your data
- ☐ Select a technique that offers the combination of accuracy, development speed and runtime efficiency you require
- ☐ Maintain awareness of alternative techniques and the pace of their improvement
- ☐ Select frameworks and libraries to accelerate development based upon your requirements for ease of use, development speed, size and speed of solution and level of abstraction and control

You may not require a large, in-house team to develop AI. There are many ways to engage with AI including third party AI APIs, outsourcing, managed services, creating an in-house AI team, or a 'hybrid' approach that combines an in-house team with third party resources. Extensive AI development, however, requires knowledge of AI hardware, development frameworks and techniques. Below, we provide a blueprint for AI development.

We begin by describing the advantages and disadvantages of different development strategies, so you can identify the ideal approach for your company.

The purpose and characteristics of AI frameworks (such as TensorFlow and PyTorch) and popular AI techniques (such as Support Vector Machines and Naïve Bayes) can be confusing. To catalyse your experimentation with AI, we then highlight and explain the AI frameworks and techniques best suited to solve a range of problems.

APIs offer specific functionality fast

You may be able to solve the problem you have identified by using an AI application programming interface (API) from a third party. These services fulfil specific, limited functions to a moderate or high standard at low cost. API calls can process your data and provide an immediate result.

Most AI APIs solve problems in the domains of vision and language. Language APIs include transcription, translation and topic extraction. Vision APIs include object recognition, scene detection and logo identification. Numerous AI APIs are available from Amazon, Google, IBM and Microsoft. Features vary (Fig. 13-15) and are advancing rapidly.

Fig. 13. Image Analysis APIs offer varying features

	Amazon	Microsoft	Google	IBM
Object detection	✔	✔	✔	✔
Scene detection	✔	✔	✔	✘
Face detection	✔	✔	✔	✔
Face recognition (human face identification)	✔	✔	✘	✘
Facial analysis	✔	✔	✔	✔
Inappropriate content detection	✔	✔	✔	✔
Celebrity recognition	✔	✔	✔	✘
Text recognition	✔	✔	✔	✔
Written text recognition	✘	✔	✔	✘
Search for similar images on web	✘	✘	✔	✘
Logo detection	✘	✘	✔	✘
Landmark detection	✘	✔	✔	✘
Food recognition	✘	✘	✘	✔
Dominant colours detection	✘	✔	✔	✘

Source: Altexsoft. Check with Amazon, Microsoft, Google and IBM to see their latest features beyond those shown above.

Fig. 14. Video APIs offer varying features

	Amazon	Microsoft	Google
Object detection	✔	✔	✔
Scene detection	✔	✔	✔
Activity detection	✔	✘	✘
Facial recognition	✔	✔	✘
Facial and sentiment analysis	✔	✔	✘
Inappropriate content detection	✔	✔	✔
Celebrity recognition	✔	✔	✘
Text recognition	✔	✔	✘
Person tracking on videos	✔	✔	✘
Audio transcription	✘	✔	✔
Speaker indexing	✘	✔	✘
Keyframe extraction	✘	✔	✘
Video translation	✘	9 languages	✘
Keywords extraction	✘	✔	✘
Brand recognition	✘	✔	✘
Annotation	✘	✔	✘
Dominant colour detection	✘	✘	✘
Real-time analysis	✔	✘	✘

Source: Altexsoft. Check with Amazon, Microsoft and Google to see their latest features beyond those shown above.

Fig. 15. Speech and text APIs offer varying features

	Amazon	Microsoft	Google	IBM
Speech recognition (speech into text)	✔	✔	✔	✔
Text into speech conversion	✔	✔	✔	✔
Entities extraction	✔	✔	✔	✔
Key phrase extraction	✔	✔	✔	✔
Language recognition	100+ languages	120 languages	120+ languages	60+ languages
Topics extraction	✔	✔	✔	✔
Spell check	✘	✔	✘	✘
Autocompletion	✘	✔	✘	✘
Voice verification	✔	✔	✘	✘
Intention analysis	✔	✔	✔	✔
Metadata extraction	✘	✘	✘	✔
Relations analysis	✘	✔	✘	✔
Sentiment analysis	✔	✔	✔	✔
Personality analysis	✘	✘	✘	✔
Syntax analysis	✘	✔	✔	✔
Tagging parts of speech	✘	✔	✔	✘
Filtering inappropriate content	✘	✔	✔	✘
Low-quality audio handling	✔	✔	✔	✔
Translation	6 languages	60+ languages	100+ languages	21 languages
Chatbot toolset	✔	✔	✔	✔

Source: Altexsoft. Check with Amazon, Microsoft, Google and IBM to see their latest features beyond those shown above.

Transferring large volumes of data can become expensive. If you are using Amazon, Google, IBM or Microsoft for other aspects of your platform, and your platform provider's APIs fulfil your requirements, your existing vendor may be an attractive option.

Many other companies, however, offer high-quality APIs in the fields of vision, language and forecasting (Fig. 16). Access a fuller list of nearly 200 APIs at www.programmableweb.com/category/artificial-intelligence/api (source: Programmable Web).

Fig. 16. Many additional companies provide AI APIs

Category	Company	Website
VISION	Clarifai	https://clarifai.com/developer/guide/
	EveryPixel	https://api.everypixel.com/
	Infinite Loop	http://imagerecognition.apixml.net/
	Prisma Labs	https://prismalabs.ai/api-sdk.html
	Reconess	https://reconess.com/
LANGUAGE	Aylien	https://aylien.com/
	Indata Labs	https://indatalabs.com/
	Meaning Cloud	https://www.meaningcloud.com/
	Spot Intelligence	https://www.spotintelligence.com/
	Tisane	https://tisane.ai/
	Automated Insights	https://automatedinsights.com/
FORECASTING	Ayasdi	https://www.ayasdi.com/platform/
	Infosys Nia	https://www.edgeverve.com/artificial-intelligence/nia/
	Nexosis	https://docs.nexosis.com/
	Unplugg	https://unplu.gg/test_api.html

Source: MMC Ventures

APIs offer immediate, useful results at the expense of niche functionality and differentiation. APIs deliver:

- **Time-to-value:** APIs provide immediate capability. By calling an API, your company can make immediate use of functions ranging from language translation to object recognition.
- **Low initial cost:** While extensive use can become expensive, APIs can cost as little as tens or hundreds of pounds to use – making AI accessible to companies of all sizes and organisations that seek proof of value before committing to greater budgets.
- **Quality:** Large companies, including Google and Microsoft, have invested billions of pounds in their AI services. Many are highly capable.
- **Ease of use:** AI APIs are accessible to developers without expertise in AI. Companies without knowledge of AI can immediately take advantages of AI via AI APIs.

Limitations of APIs include:

- **Functionality:** APIs offer specific functionality, often in the fields of vision and language. If your requirements fall outside of what is available, an alternative approach will be required.
- **Configurability**: APIs do not allow you to adjust the training data or models on which the services are based. If you wish to develop services based on unique training data you have, or tune underlying algorithms for improved results, APIs will be unsuitable.
- **Genericness:** APIs are designed for mass adoption; they tend to be generic and lack depth and domain specificity. Object recognition APIs can actually tell the difference between BMWs and Skodas but are unlikely to be able to tell the difference between a BMW 6 Series and 7 Series.
- **Commoditisation:** The APIs you use are available to your competitors. It will be challenging to create lasting competitive advantage, and associated market value, through use of third party APIs.
- **Lifetime cost:** Extensive use of APIs can attract a high cost relative to an in-house solution you own.
- **Dependence:** Large vendors have, on occasion, discontinued APIs. Smaller vendors can be acquired or cease to operate. Using third party APIs creates a dependency over which you have no control.
- **Privacy:** Using APIs involves passing your data to third parties. Does this comply with your data permissions? Does the third party retain a copy of your data or use it for any other purpose?

Overall, APIs are ideal if you seek an immediate, low cost solution to a generic problem. APIs will be insufficient, however, if you have a niche challenge, seek greater control and configurability, or seek long-term differentiation through AI (Fig. 17).

Fig. 17. APIs offer immediate results at the expense of differentiation

Use APIs if you:

- Seek a solution to a generic problem for which a relevant API is available
- Have limited budget
- Require immediate initial results
- Have limited in-house AI knowledge and resources.

Avoid APIs if you:

- Seek a solution to a domain-specific or niche problem for which an API is unavailable
- Have unique training data, or wish to control and configure your AI, for improved results
- Seek long-term differentiation through AI
- Do not wish to rely on third parties
- Have data permissions that prevent you passing data to third parties.

Source: MMC Ventures

Many companies adopt a 'hybrid' approach (page 54), using APIs for rapid proofs-of-concept while transitioning to an in-house team that can deliver improved, differentiated, domain-specific capabilities over time.

APIs can cost as little as tens or hundreds of pounds to use – making AI accessible to companies of all sizes and organisations that seek proof of value before committing to greater budgets.

Managed services offer increased capability at low cost

Several vendors offer managed AI services. A step beyond pre-tuned, function-specific APIs, managed services enable you to upload your data, configure and train your own AI models using a simple interface, and refine the results. These services abstract away much of the difficult of developing AI and enable you to develop a custom solution rapidly, via a simplified interface and limited coding.

Peak, a leading managed AI service company in which we have invested, offers an effective solution. Solutions are also available from Amazon (SageMaker), Google (AutoML), IBM (Watson), Microsoft (Azure) and Salesforce.

Managed services have several advantages:

- Capability: greater flexibility and control than simple APIs; managed services enable you to develop custom models and, potentially, bespoke IP.
- Cost: cheaper than building an in-house AI team or outsourcing development.
- Speed: faster time-to-value than building an in-house AI team.

Limitations include:

- Control: less control than in-house development; access to underlying models will be limited, reducing customisation and tuning.
- Permissioning: you must be comfortable transferring your data to a third party.
- Reliance: it may be expensive or unappealing to migrate away from a managed service provider, given dependencies and data transfer costs.
- Intellectual Property: Some vendors retain your data to improve algorithms for all; in other cases, pragmatically or contractually ownership of the model you develop may be limited.

If basic APIs will not satisfy your requirements, managed AI services offer a fast, flexible, way to develop bespoke solutions at a lower cost than building an in-house team. Managed services are also ideal for prototyping. If you require more control, flexibility, autonomy and ownership in time, however, significant re-development may be required.

Fig. 18. Managed services offer speed at the expense of control

Use managed services if:

- Your challenge is a solved problem but your data is key
- You wish to begin quickly
- Cost is a challenge.

Avoid managed services if:

- Your data permissions prohibit this approach
- You require extensive control and flexibility
- Speed of response is critical
- Your problem has unique demands.

Source: MMC Ventures

Outsourcing offers expertise for moderate initial investment

If a suitable API or third party product is unavailable, you will need to build an AI solution. However, investing in an in-house team is expensive – at least £500,000 per year, typically, even for a small team. There are cost-effective alternatives. Several companies provide outsourced AI capabilities, ranging from contractor developers, who work with your own engineers, to complete, outsourced AI development.

The nature of AI development enables reseachers to work on multiple workstreams simultaneously, so outsourcing can be cheaper than maintaining a permanent team. Conversely, transferring large volumes of data securely and frequently, and retraining models on an ongoing basis, can become expensive. Whether outsourcing is appropriate will depend upon a range of considerations including:

- Domain: will a third party offer the expertise you require in your problem domain and sector?
- Expertise: to what extent do you wish to build expertise in-house?
- Speed: do you require trusted expertise more rapidly than you can develop it in-house? Do you require a solution more quickly than you could build in-house?
- Data sensitivity: do you have permission to pass data to third parties?

- **Operation:** if an outsourcer builds your models, are you entitled to deploy them on your own infrastructure – or are you tied to your outsourcer on an ongoing basis?

Overall if maximising speed and minimising initial costs are your highest priorities, and APIs are unavailable, consider outsourcing (Fig. 19).

If outsourcing, specify:

- **Frameworks:** is there a specific AI framework you require the outsourcer to use?
- **Standards:** what accuracy (precision and recall – see Chapter 5) must models meet?
- **Data:** will you provide cleaned, labelled training data? Or is data to be created by the outsourcer?
- **Costs:** what costs have been agreed?
- **Timescales:** what timescales must be met? This can be more challenging than for traditional software development because improving a model may require experimentation.
- **Deployment:** how production-ready must the output be?

Fig. 19. Outsourcing offers speed at the expense of in-house knowledge

Use outsourcing if you:
• Require trusted expertise quickly • Have clarity regarding the solution you require • Require a cheaper alternative to permanent employees.
Avoid outsourcing if you:
• Have data permissions that prohibit outsourcing • Require knowledge regarding your problem domain or sector that an outsourcer cannot offer • Wish to build knowledge wihin your company.

Source: MMC Ventures

If maximising speed and minimising initial costs are your highest priorities, and APIs are unavailable, consider outsourcing.

An in-house team offers differentiation – at a price

Investing in an in-house AI team offers maximum control, capability and competitive differentiation – at a price.

An AI team of your own can deliver:

- **Flexibility:** Control over the hardware, programming languages, frameworks, techniques and data you employ offers the flexibility to iterate and expand your solutions as your needs evolve.
- **Capability:** APIs offer defined functionality. Managed service environments limit your ability to tune underlying algorithms. Outsourced talent will lack your team's domain expertise. With an in-house team you have the opportunity to create optimised solutions, potentially beyond the current state of the art.
- **Differentiation:** An in-house team can develop a unique AI offering that delivers competitive advantage, credibility in the market and associated value for your company.
- **Resilience:** Without reliance on third party APIs or outsourcers, your AI initiatives can enjoy greater resilience and longevity.
- **Security:** Retain control over your own data; none of it needs to be passed to third parties.

Drawbacks of an in-house team include:

- **Cost:** A small in-house team, comprising two to four people and the hardare they require, will cost at least £250,000 to £500,000 per year – potentially more to productise the resulting system. A large team, recruited to solve problems at the edge of research, will require a multi-million pound annual investment in personnel and hardware.
- **Complexity:** To develop an in-house AI team you must attract, structure, manage and retain AI talent; select the development languages, frameworks and techniques you will employ; undertake data gathering and cleansing; learn how to productise AI into real-world systems; and ensure compliance with regulatory and ethical standards.
- **Speed:** It will require months to build a productive in-house AI team, and potentially longer to collect the data you require and develop customised solutions that deliver results to the standard you require.

An in-house team may be necessary if your challenge cannot be solved with existing AI techniques and solutions, if you face significant restrictions on your ability to pass data to third parties, or if you seek competitive differentiation through AI. Before incurring the cost and complexity of an AI team, however, explore whether alternative methods can deliver your requirements faster and for a lower cost (Fig. 20). A hybrid strategy, described below, may be ideal.

To develop an in-house AI team, review all chapters of this Playbook for best practices regarding strategy, talent, data, development, production and regulation & ethics.

Fig. 20. An in-house team offers differentiation – at a price

Use an in-house team if you:
• Have a niche problem that cannot be solved with existing solutions or techniques • Seek differentiation in the market and associated value • Wish to retain control over your own data.
Avoid an in-house team if you:
• Have a simple problem for which solutions are available • Require an initial solution quickly • Have a modest budget.

Source: MMC Ventures

A hybrid approach can offer the 'best of both worlds'

For many companies, a 'hybrid' approach to AI is ideal. Plan for an in-house team that will address your requirements to a high standard over time, but use third party APIs (or even a non-AI solution) to solve an initial, simpler version of your challenge.

A hybrid approach may enable you to prove the viability or value of your idea and justify in-house spend. It can also serve as a cost-effective way to identify the aspects of your challenge can be readily addressed and those that will require bespoke work. "A hybrid approach gives me flexibility. I don't need to reinvent the wheel and can focus on doing very specific tasks better than anyone else in the world" (Dr Janet Bastiman, StoryStream).

A hybrid strategy offers a rapid, low cost start that suits many companies (Fig. 21). Initial investment in hardware, team and software can be minimal. Many APIs offer free trial periods in which you can assess scope for results. Even if your data restrictions prohibit use of third party APIs, you can adopt a hybrid approach with in-house developers using pre-trained AIs. Further, many academic papers and coding competition entries have code uploaded to GitHub and many have unrestricted licenses.

If you adopt a hybrid approach, develop a data strategy (Chapter 1) and pipeline of training data upfront. You can continue to use third-party APIs if they fulfil your needs unless costs become prohibitive, you wish to build resilience, or you seek improved results and differentiation with your own technology. As you gather additional data, you can create more accurate and complex models in-house, as needed and when the business case has been proven.

While the risk of interruption to services from Amazon, Google, IBM and Microsoft is low, vendors do occasionally remove APIs. Smaller vendors offering APIs may be acquired, or their services changed or discontinued. If you adopt a hybrid approach, develop a strategy for resilience. Once elements of your product are in place, examine the pre-trained models and consider moving these services in-house if you can achieve results comparable with the API. You may be able to use your chosen APIs in perpetuity and continually develop niche AI to complement these – a popular approach.

"A hybrid approach gives me flexibility. I don't need to reinvent the wheel and can focus on doing very specific tasks better than anyone else in the world".

Dr Janet Bastiman, StoryStream

Fig. 21. A hybrid approach can offer the 'best of both worlds'

Use a hybrid approach if you:

- Require rapid initial results
- Wish to limit spend until a business case is proven
- Have an evolving problem and desire for greater differentiation and resilience over time.

Avoid a hybrid approach if you:

- Have a generic problem solved with existing APIs
- Have a complex problem, to which a simple solution will cause more harm than no solution
- Have data permission challenges that prevent use of APIs.

Source: MMC Ventures

To develop AI, optimise your technology stack

To develop AI – via a managed service provider, outsourcer or in-house team – you have choices to make regarding your AI technology stack. The stack comprises six layers: hardware; operating systems; programming languages; libraries; frameworks and abstractions (Fig. 22).

We offer hardware recommendations overleaf. The problem domain you are tackling (assignment, grouping, generation or forecasting) will then favour particular machine learning techniques and associated libraries and frameworks. Select components for your development stack accordingly.

The degree of abstraction you select will depend upon the skill of your development team, the speed of development you require and the degree of control you seek over the models you develop. Greater abstraction offers faster development and requires less skill, but limits your ability to tune models to your requirements. The size and speed of your models may also be limited.

Not all problems require the full stack; some solutions can be achieved rapidly, without frameworks or abstractions.

Fig. 22. The six layers of the AI technology stack

Abstractions	*(e.g. Keras, Digits)*
Frameworks	*(e.g. TensorFlow, PyTorch)*
Libraries	*(e.g. NumPy, Pandas)*
Languages	*(e.g. Python, R)*
Operating System/CUDA	*(e.g. Linux, Windows)*
Hardware	*(e.g. GPUs, CPUs)*

Source: MMC Ventures

For effective R&D, use appropriate hardware

Research and development requires hardware. To process models quickly, ensure your team has hardware with graphical processing units (GPUs) that support NVIDIA's Compute Unified Device Architecture (CUDA) libraries. These allow your AI programmes to use the specialised mathematics of the GPUs and run at least ten times faster than on a CPU. For many, a laptop with a high performance graphics card is ideal. Current, potentially suitable cards include the NVIDIA GTX 1050 Ti, 1070, 1080 and the RTX 2080.

For greater power, desktop machines with more powerful GPUs are preferable – but at the expense of flexibility for your team. If you are a multi-premise team, or expect your personnel to travel, your team may expect a laptop in addition to a desktop machine you provide for research.

For large models, or rapid training, you will require dedicated servers. Buying and hosting servers yourself, either on-premise or in a data centre, is the most cost-effective over the long term but requires considerable upfront capital expenditure. The majority of early stage companies will find it more appropriate to use cloud-based servers offered by large providers including Google, Amazon and Microsoft. All offer GPU servers, costed according to usage time. Using the cloud providers, at least at the early stages of your AI initiatives, will enable you to control costs more effectively and push the decision regarding buying hardware to later in the process when you have a minimum viable product.

Apply AI techniques suited to the problem domain

For each problem domain (assignment, grouping, generation and forecasting – see Chapter 1, 'Strategy') – there are numerous established machine learning techniques.

Techniques vary in their data requirements, training dynamics, deployment characteristics, advantages and limitations. While deep learning methods are powerful, other techniques may be sufficient or better suited. Experiment with multiple techniques. Below, we highlight techniques popular for each domain.

For assignment problems consider SVCs, Bayes, KNNs and CNNs

Classification problems, which offer a defined, correct output to ease development, are frequently an attractive starting point for AI. While convolutional neural networks became popular, in part, due to their efficacy in solving classification problems there are many alternative techniques you can apply – many of which offer effective results and are quicker to implement.

Fig. 23. For assignment problems consider SVCs, Bayes, KNNs and CNNs

Technique	Approach	Advantages	Challenges
Support Vector Classification (SVC)	SVC is effective when classifying images or text and you have fewer than 10,000 examples. Plot data in multi-dimensional space, based upon the number of variables in each example, and the SVC algorithm will determine the boundaries of each class (Fig. 24). New examples are classified based upon their relationship to the calculated boundaries.	Effective when there are many variables. Memory-efficient.	Prone to overfitting. Cannot directly provide probability estimates to evaluate results.
Naïve Bayes	Naïve Bayes assumes that variables are independent and is particularly effective for text classification. Classifications are developed based upon the probability of each variable being contained within a specific class. Probabilities are then combined to provide an overall prediction.	Fast to train and run. Effective for text and variables.	Highly sensitive to training data. Probability for classifications is unreliable.
K–Nearest Neighbours Classification (KNN)	KNN is a productive statistical technique when you possess a complete data set. All training data is mapped into vectors, from an origin based on the variables in the data. Each point in space is assigned a label. New data is then classified by mapping it to the same space and returning the label of the closest existing datapoints (Fig. 26).	Effective when boundaries between classes are poorly defined.	All data must be stored in memory for classification; predictions require additional resources and time.

Continued on next page.

Technique	Approach	Advantages	Challenges
Convolutional Neural Networks (CNNs)	CNNs comprise multiple layers of neurons. Data passing through the network is transformed, by examining overlaps between neighbouring regions to create areas of interest. The final layer of the network is then mapped to target classes.	Excels with complex data and multiple output classes.	Compuationally expensive. Slow to train.

Source: MMC Ventures

Fig. 24. SVCs maximise the boundaries between classes

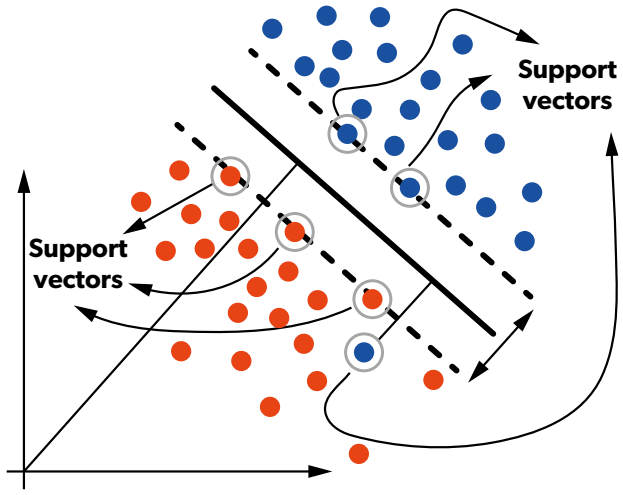

Source: Haydar Ali Ismail, (https://bit.ly/2vcSDLf)

Fig. 25. Naïve Bayes classifies based on the probability of a variable being contained in a class

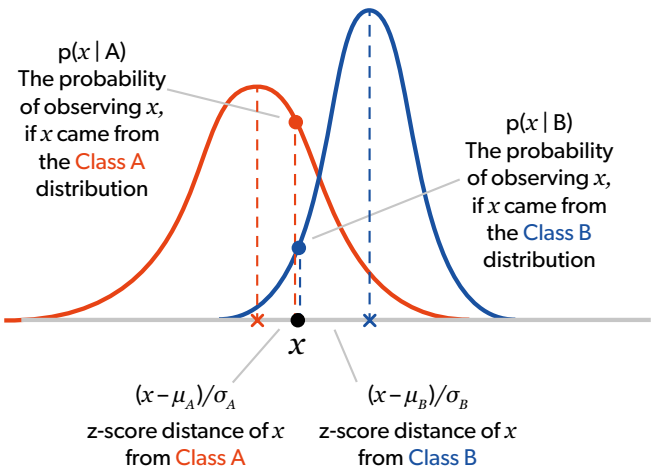

Source: Rajeev D. S. Raizada, Yune-Sang Lee (https://doi.org/10.1371/journal.pone.0069566)

Fig. 26. KNNs return the label of the closest datapoint

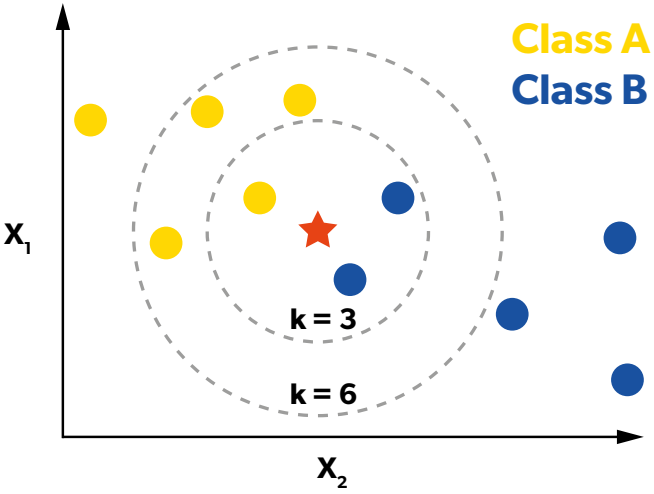

Source: Savan Patel (https://bit.ly/2GAYWR5)

While convolutional neural networks are popular, there are many alternative techniques you can apply – many of which are quicker to implement.

Regression problems quantify the extent to which a feature exists. Because they are also assignment problems, the techniques used for assignment frequently overlap with those used for regression.

Fig. 27. For regression problems, explore SVRs, Lasso and CNNs

Technique	Approach	Advantages	Challenges
Support Vector Regression (SVR)	SVR is similar to SVC; training data plotted in multi-dimensional space. However, unlike SVC (where hyperplanes are generated to maximise distance from the data), with SVR hyperplanes are matched as closely as possible to the data.	Effective with large numbers of variables. Versatile. Can extrapolate for new data.	Prone to overfitting. The prediction is provided without confidence in its correctness; confidence must be determined through indirect methods.
Least Absolute Shrinkage and Selection Operator (Lasso)	Lasso minimises the number of variables used to make a prediction. If there are multiple, correlated variables Lasso will select one at random.	Fast predictions. Well suited to situations in which few variables are important for a prediction.	Minimising input variables may cause overfitting to training data. Selected variables may oversimplify the problem.
Convolutional Neural Networks (CNNs)	CNNs can also be used for regression assignment tasks. Unlike when used for classification, the CNN provides a single neuron, with the prediction value as an output.	Effective for complex problems.	Difficult to determine which inputs contribute to a prediction. Difficult to determine confidence in the prediction.

Source: MMC Ventures

For grouping explore Meanshift Clustering, K-Means and GMMs

If you have unlabelled data and seek to cluster it into similar groups, you will require techniques that expose similarity. Defininining similarity can be challenging when there are many dimensions to the data.

Fig. 28. For grouping explore Meanshift Clustering, K-Means and GMMs

Technique	Approach	Advantages	Challenges
Meanshift Clustering	Meanshift clustering discovers groups within a data set by selecting candidates for the centre of a group from the arithmetic mean of the datapoints in the region. The process continues until there are a distinct set of groups, each with a centre marker (Fig. 29).	You do not need to know in advance how many clusters you expect.	The algorithm's scalability is limited due to the number of calculations between neighbours in each iteration.
K-Means (Lloyd's algorithm)	K-Means groups data into a pre-defined number of clusters of equal variance (data spread within the group).	Scalable to large data sets.	Defining the number of clusters in advance can be difficult because it requires some knowledge of the probable answers. If data is irregularly shaped, when plotting in multi-dimensional space the algorithm can become confused and suggest peculiar distributions.
Gaussian Mixture Models (GMMs)	GMMs can offer more flexibility than K-Means. Instead of assuming that points are clustered around the mean of each group, GMMs assume a Gaussian distribution and can offer ellipse shapes (Fig. 30).	Because they draw upon probabilities, GMMs can label datapoints as belonging to multiple classes – which may be valuable for edge cases.	If the Gaussian distribution assumption is invalid, the clustering may perform poorly with real data.

Source: MMC Ventures

Fig. 29. Meanshift Clustering produces distinct groups with centre markers

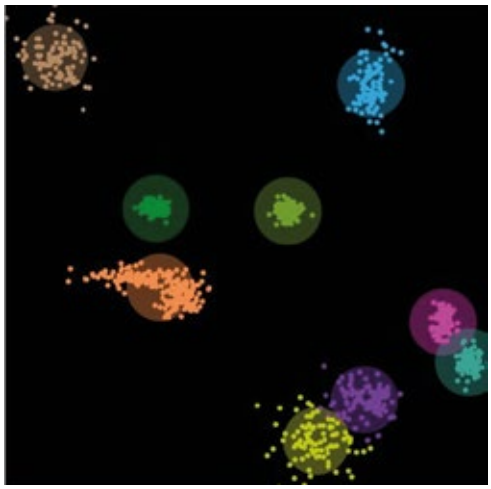

Source: Miroslav Radojević (https://bit.ly/2GWfm5W)

Fig. 30. GMMs offer elliptical groupings instead of assuming points are clustered round a mean

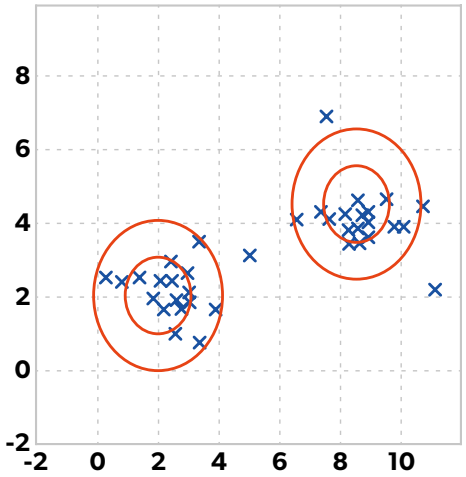

Source: John McGonagle, Geoff Pilling, Vincent Tembo (https://bit.ly/2tzlc5k)

For generation, VAEs and GANs can be effective

Since its inception, AI has been used to synthesise text; MIT's ELIZA natural language processing programme, created from 1964 to 1966, offered the illusion of understanding in psychology and other domains. In the decades since, the quality of generation techniques has been transformed – particularly following the introduction of Generative Adversarial Networks (GANs) – while domains of application have broadened to include visual imagery and sound.

Fig. 31. For generation, VAEs and GANs can be effective

Technique	Approach	Advantages	Challenges
Pattern matching	Pattern matching is among the most naïve of techniques but offers the illusion of intelligence in text generation. Using a dictionary of phrases and key words to recognise input statements, it is possible to create moderately effective responses with little effort.	Useful for repetitive situations that may be fully mapped – such as sports reporting or basic customer support.	Rapidly becomes nonsensical when inputs are outside a predefined area.

Continued on next page.

Technique	Approach	Advantages	Challenges
Probabilistic prediction	Probabilistic prediction can be effective for text generation. Given a word or words from a sentence, probabilistic models determine a word or phrase to follow and recommend the text with the highest probability.	Improve quickly with use.	Addresses a set of problems limited in scope.
Variational Auto-Encoders (VAEs)	VAEs train from real-world data. VAEs use a convolutional neural network to encode data into a vector and a second network to deconvolve the vector back to the original image (Fig. 32). After training the network, varying the input vector will provide realistic outputs.	Compare the output directly to the original.	The likelihood of a realistic output decreases if the difference between the original data vector and new input vector becomes too great. Image outputs can be blurry.
Generative Adversarial Networks (GANs)	Generative Adversarial Networks (GANs) comprise a generator network such as DCGAN (Deep Convolutional GAN) and a discriminator network (a standard classification CNN) (Fig. 33). The generator attempts to create an output that will fool the discriminator, while the discriminator becomes increasingly sophisticated at identifying outputs that are unreal. With sufficient training, the generator network learns to create images or text that are indistinguishable from real examples.	Create realistic outputs from random input noise.	Cannot generate outputs with specific features unless the GAN searches the entire input space. Random inputs give random (although realistic) outputs; you cannot force a specific output condition. The discriminator identifies only real images and fakes, not whether the output includes elements of interest. The more complex the image or text being created, the harder to create realistic output. Current research centres on splitting the challenge into multiple generative steps.

Source: MMC Ventures

Fig. 32. VAEs encode images into a vector and add noise before regenerating

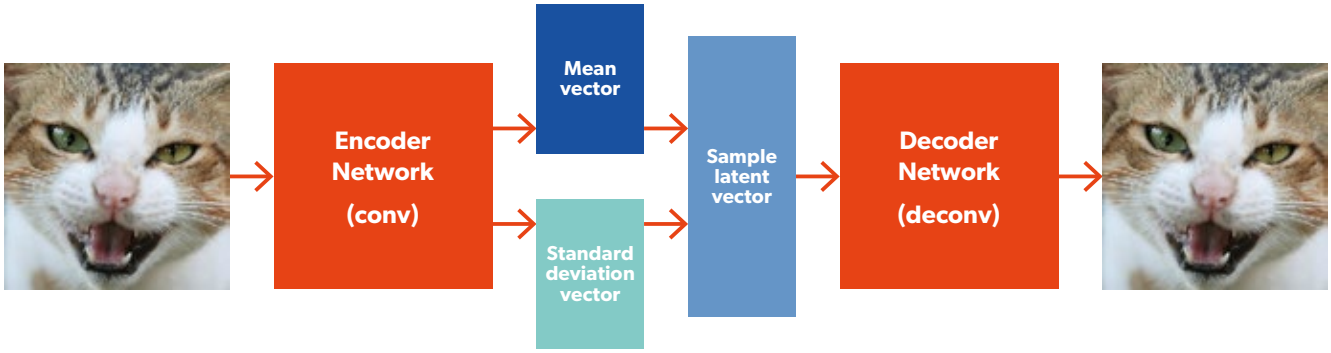

Source: Kevin Frans (https://bit.ly/2GDnUij)

Fig. 33. With one network, GANs generate output from random noise; a second network serves as a discriminator

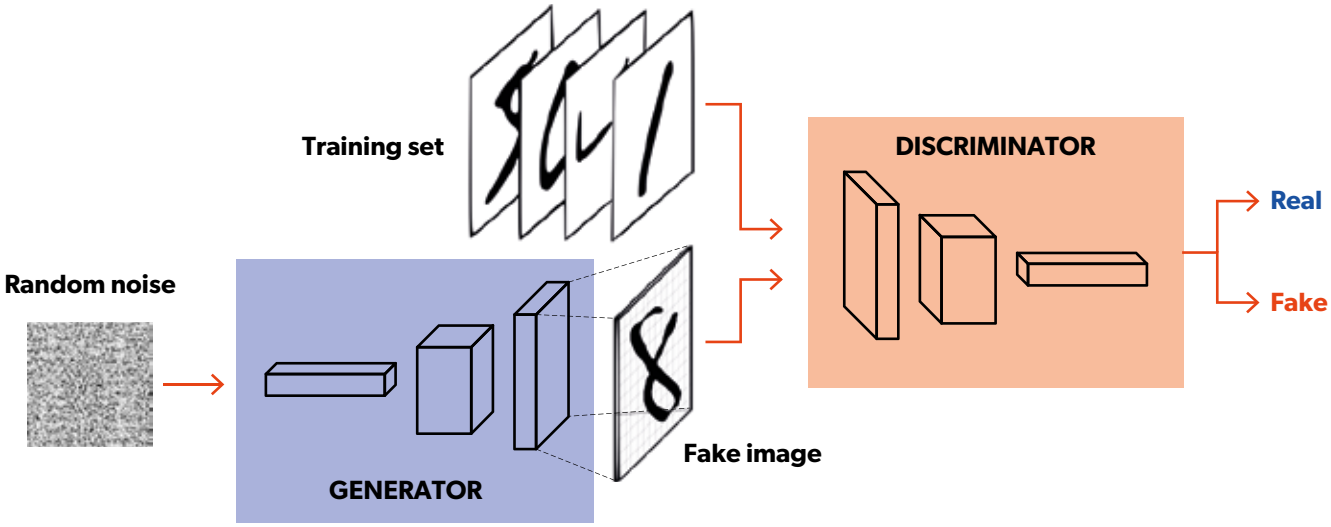

Source: Thalles Silva (https://bit.ly/2MZGKAs)

For forecasting, causal models and HMMs are popular

Applying techniques to predict the future is challenging; forecasts may be affected by variables outside the data available to you. While the past is not always indicative of the future, AI forecasting techniques are effective when there are causal or periodic effects. Understanding the volume of data you require may need initial knowledge of causal and periodic effects, in the absence of which your model may miss these relations.

While the past is not always indicative of the future, AI forecasting techniques are effective when there are causal or periodic effects.

Fig. 34. For forecasting problems, experiment with causal models, HMMs and ARMA

Technique	Approach	Advantages	Challenges
Causal models	A sub-class of assignment problem, causal models can use the same techniques – with the additional consideration of variables' rate of change – to predict new values.	Straightforward to implement.	Consider a point in time; may fail to take into account longer-term trends.
Hidden Markov Models (HMMs)	Markov models provide a sequence of events based upon the previous time step. HMMs assume that predictions of the future can be based solely upon the present state; further history is irrelevant.	Well suited to learning and predicting sequeneces within data based upon probability distributions.	Challenging to train. Rapidly become inaccurate if sequences change.
Auto-Regression Moving Average (ARMA)	Despite dating from the 1950s, ARMA remains useful. ARMA considers past values and uses regression to model and predict a new value, while a moving average calculates the error. A further algorithm determine the best fit for future predictions.	Considers past values and prediction error, offering greater adaption than HMMs.	Can oversimplify problems that have complex periodicity, or randomness, in the time series.

Source: MMC Ventures

Use frameworks to accelerate development

If your team is developing an AI solution, use libraries to accelerate development. The algorithms described above have been coded into efficient libraries for Python and R. Implementing an algorithm directly in Python will be slower – in some cases 60 times slower (Fig. 35).

Fig. 35. Libraries offer improved performance

Implementation	Run time
Pure Python (with list comprehensions)	18.65 seconds
TensorFlow on CPU	1.20 seconds
NumPy	0.32 seconds
Run time for a linear regression problem implemented in pure Python, using TensorFlow (on CPU for comparability) and using in-built functions in NumPy (a Python library for numerical analysis).	

Source: https://realpython.com/numpy-tensorflow-performance/ (Renato Candido)

For numerical analysis, NumPy is a library of choice

There are many libraries available to support numerical analysis. The most popular include:

- NumPy: A library of choice for numerical analysis in Python. Functions are optimised in C so run quickly, matrices and vectors are well handled, and there are many in-built statistical algorithms to support AI.
- Scikit-learn: Released in 2010 as a lightweight library for deep learning, Scikit-learn is built on NumPy and offers considerable overlap, although the two complement each other well.
- Matplotlib: Predominantly a plotting library, to support visual analysis of plots Matplotlib requires its own numerical calculations. These are limited and further libraries are required for broader analytical techniques.
- R packages are not as extensive as those for Python but there are many for numerical analysis on top of core R functions – including caret, gimnet, randomForest and nmle.

In addition to libraries there are specific applications, such as Matlab and Mathematica, which offer extensive functions. While popular in academic settings, they are rarely used in industry given the high cost of software licenses compared with the free libraries available.

For deep learning, TensorFlow and Caffe are popular frameworks

Deep learning frameworks are typically more extensive than numerical analysis libraries and serve as 'scaffolding' for your projects. Your choice of framework will impact speed of development as well as the features and scalability of your solution.

With numerous frameworks available, take time to evaluate your project priorities and the framework best suited to your goals. The most popular framework may not be optimal for your initiative. When selecting a framework consider its advantages and limitations, the skills it requires, availability of skills, and scaling and speed requirements (both for development and production).

Unless you are using a pre-trained network, if you have implemented models in a single framework then reimplementing them in another will involve retraining from scratch. You may elect to use multiple frameworks for different problems, particularly if doing so allows consistency with existing development languages.

Frameworks evolve at different speeds and, particularly when maintained by a single business or university, may be discontinued with limited notice. In a rapidly evolving field, frameworks with high levels of community support can be attractive.

Fig. 36. Different deep learning frameworks offer advantages and challenges

Framework	Features	Maintained by	Community Support	Availability of Talent	Advantages	Challenges
TensorFlow	One of the most widely used frameworks, TensorFlow is implemented as a Python library, enabling rapid development of a wide variety of projects. There are many example projects for TensorFlow, and numerous code samples (available with an open source license) for different classes of problem that can be adapted rapidly for your own tasks.	Google	High	High	Numerous example projects are available. Becoming a standard as many training courses use TensorFlow. Allows lower-level data manipulation for tuning.	Significant computational overhead. Less efficient than numerical libraries for certain calculations. Challenging to optimise.

Continued on next page.

Framework	Features	Maintained by	Community Support	Availability of Talent	Advantages	Challenges
Caffe/Caffe2	Caffe is one of the earlier frameworks implemented in C++ with a Python interface. Originally designed for convolutional neural networks, Caffe grew to support feed-forward networks. Facebook recently introduced Caffe2 which is built for mobile, includes pre-trained models, and is likely to be merged with PyTorch (also from Facebook).	Berkeley Vision (Caffe) Facebook (Caffe2)	Medium	Medium	Widely used in the academic community.	Challenging to compile. Limited support.
Theano	Among the oldest deep learning libraries, Theano is more mathematical than many and positioned between analytical libraries such as NumPy and abstract frameworks such as TensorFlow. Much academic research was undertaken in Theano, with Python, and many early adopters of AI employed it.	University of Montreal	Low	Medium	Efficient and scalable. Straightforward to implement new algorithms. Used for many early machine learning courses.	No longer developed or supported. Developers with experience may advocate for its use.
MXNet	MXNet supports a wide range of programming languages including C++, R, Python and Javascript, and is maintained by the open source community.	Apache	Medium	Low	Fast and scalable; designed for high performance.	Little support in academia or industry except niche, high performance use cases.
Torch/PyTorch	Torch provides numerous pre-trained models and development that is similar to traditional programming. While Torch supported only the Lua language, PyTorch supports Python.	Facebook	Low	Low	Uses standard debugging techniques Supports distributed training.	PyTorch 1.0 was recently released (October 2018); change may be rapid. Limited available Lua talent for Torch.
DeepLearning4J	DeepLearning4J is written for Java and Scala, and supports a wide variety of networks.	Eclipse Foundation	Low	Low	Fast and scalable. Can operate with an existing Java stack.	Lacking support for Python, its use is uncommon. Few examples.
Chainer	Chainer is a smaller library for Python, used extensively for natural language tasks (speech recognition, sentiment analysis and, translation).	Preferred Networks	Low	Low	Networks can be modified while running.	Challenging to debug. Few examples.
Digits	NVIDIA's framework is freely available to participants in the Company's developer programme or users of the NVIDIA cloud. Digits abstracts much of the programming into a visual interface, which allows researchers to focus on network design instead of coding data import routines or network architecture components. Digits will operate with Tensorflow or on a standalone basis.	NVIDIA	Low	Low	Enables rapid prototyping. Highly optimised.	Low levels of support in academia and industry. Few available examples. Restrictive abstraction.
Keras	Keras is a Python library that allows rapid prototyping of neural networks. Not a framework in its own right, we consider it here as an extension to Theano and TensorFlow.	François Chollet	Medium	Medium	Enables rapid prototyping and experimentation. Accessible for beginners and useful for all levels of experience.	Requires additional frameworks. Challenging to customise networks beyond the abstraction layer; re-working may be required to utilise underlying frameworks.

Source: MMC Ventures

Chapter 5

Production

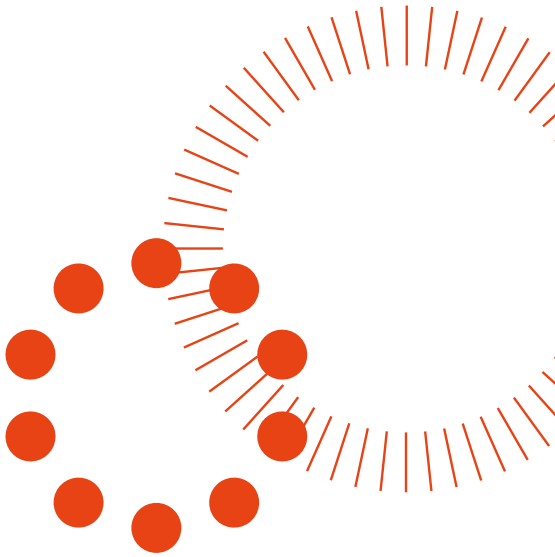

Summary

- An unused AI system delivers no value. Develop a production process that smoothly transitions AI systems you have in development to live use.
- AI production follows a conventional development process and requires you to undertake research, develop a prototype and create a minimum viable product (MVP). Once in production, undertake cycles of ideation, research, development and quality assurance.
- Effective R&D requires rapid iteration. Initially, optimise for speed over quality. Releasing an early model into production for feedback is preferable to waiting until a research model is perfect.
- During the R&D phase, solicit feedback about prototypes from beyond the AI and production teams to minimise expensive, later redevelopment.
- When moving from MVP to production, select an appropriate hosting environment. On-premise hosting is suitable for those with highly sensitive data and existing on-premise hardware, but is rarely preferred by early stage companies given high upfront costs, unpredictable activity levels and required security expertise. Hosting your own hardware in a data centre offers control and value over the long term. Upfront costs can be high, however, and managing a data centre can prove a distraction for young companies. Cloud hosting, which offers low upfront costs and high levels of flexibility, is well suited to many early stage companies – although annual costs can be double that of a self-managed data centre and cloud hosting may be unsuitable for highly sensitive data. Consider the physical location in which your cloud servers are hosted. Different countries have varying rules regarding data and you may be required to keep your data within its area of origin.
- Proving that AI systems are effective differs from the typical software quality assurance (QA) process. Test your AI system at multiple stages – during training, validation and continuously through its life. Efficiency is critical; automate testing to as great an extent as possible.
- Understand the three common measures of 'accuracy' in AI – recall, precision and accuracy – and monitor all three to capture performance. Balancing precision and recall is challenging. Whether you elect to minimise false positives or false negatives should depend upon the nature of your sector and the problem you are solving.
- An effective maintenance programme will sustain your AI's intelligence. Beyond the maintenance you would typically perform on a software system, you should verify and update your AI system on an ongoing basis. AI technology is developing at pace. Invest in continual improvement to ensure your system avoids obsolescence.

Production: The Checklist

Optimise Research & Development

- ☐ Clarify the required characteristics of your planned system
- ☐ Limit abstract research – focus resources on solutions to business problems
- ☐ Identify system demands and associated RAM/GPU requirements
- ☐ Leverage existing infrastructure, if present

Prototype effectively

- ☐ Identify relative system priorities (speed, precision, recall)
- ☐ Create production-ready code, even for prototypes
- ☐ Solicit feedback from people outside the AI team
- ☐ Understand the additional development required for an MVP

Develop efficient Production and Deployment workflows

- ☐ Establish a controlled release process
- ☐ Plan for rapid deployment of code and models
- ☐ Co-deploy models and related code
- ☐ Anticipate continual cycles of improvement
- ☐ Select an appropriate hosting environment
- ☐ Establish increasing automation over time
- ☐ Ensure problematic deployments can be reversed

Create a rigorous testing process

- ☐ Report against all measures of accuracy (precision, recall, accuracy)
- ☐ Establish definitions for 'better' models
- ☐ Automate testing to as great an extent possible

Establish an effective maintenance programme

- ☐ Validate live results frequently
- ☐ Test edge cases
- ☐ Establish defined downtime and update procedures

An unused AI system delivers no value. It's essential to develop a production process that smoothly transitions the AI systems you have in development to live use. Below, we describe an optimal production pipeline – in which rapid iteration, appropriate hardware, suitable hosting, rigorous testing and ongoing maintenance deliver high quality results.

AI production follows a conventional development process

By the time you are considering how to take your solution live you should:

- Understand how you are going to solve the business problem
- Possess the required data
- Know the development languages and frameworks you will use
- Appreciate whether hardware with Graphical Processing Units (GPUs) will be required
- Understand whether your live system will be powered by local hardware or the cloud.

The aforementioned will enable you to determine an optimal process to move from development to a live solution – the production environment.

Progressing an AI system from idea to reality should follow a broadly conventional development practise – although timescales may be less certain. After ideation, undertake research, prototype and then develop a minimum viable product (MVP). Once in production undertake cycles of ideation, research, development and quality assurance.

Fig. 37. The AI production pipeline is similar to a normal development practice

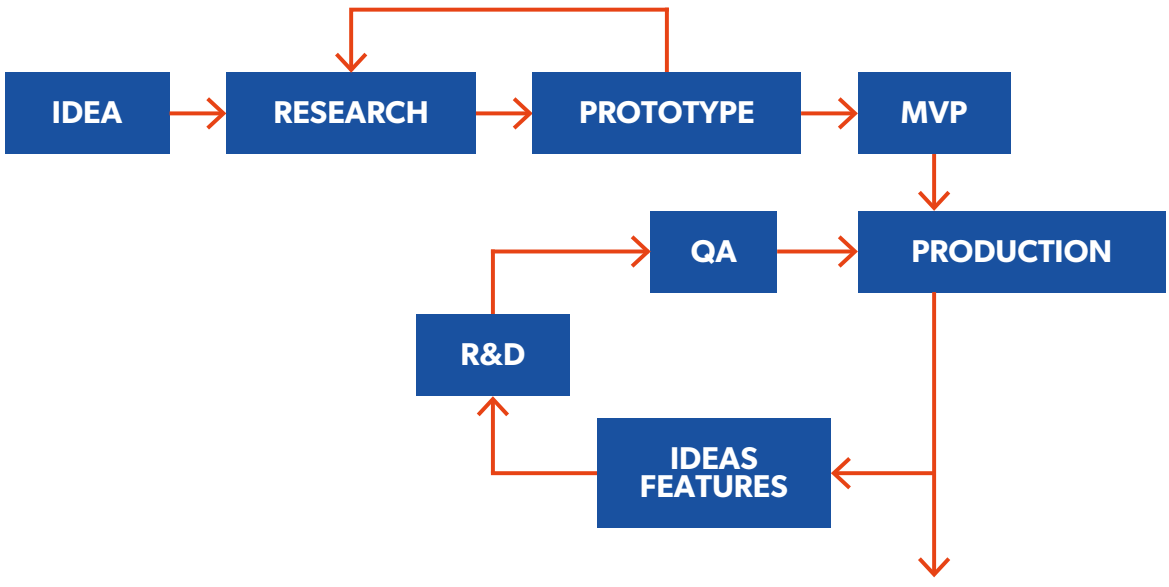

Source: MMC Ventures

For effective R&D, iterate rapidly and use appropriate hardware

Whether you have an in-house team or are outsourcing, ensure the team understands the characteristics required from the end system. Which considerations are flexible? Which are not? Different decisions will be made if speed is more important than accuracy – and vice versa.

Even in the research phase, ensure that development is undertaken in the language used for deployment. If you plan to deploy in Python, for example, avoid the overhead of rewriting models created in MatLab or R.

Initially, optimise for speed over quality. It's better to release an early version of a model from the research environment into production, and then to solicit feedback in the live environment, than to wait until the research model is perfect. "Spend a month to get a weak model and then iterate to make it great" (Eddie Bell, Director of Machine Learning, Ravelin). Isolating models within the research environment will push considerations of deployment, usability, performance and scalability to the end of the project instead of addressing them early. In addition, it increases the risk of a model performing poorly with unexpected real-world data. Many data scientists resist releasing models that are not 'good enough'. Overcome this hurdle by developing a culture in which the dynamics of AI development are understood and people are not blamed for early, poor quality results.

Effective research & development requires appropriate hardware – see page 55 for guidance.

Ensure your AI team has its code in a source control system – Git, Mercurial and Subversion are popular – and update it regularly. The size of trained models can exceed file size limits on these systems. If file size is a constraint, find an alternative way of versioning and storing your files. A simple solution (for example, creating zip files on a shared drive) can be effective but ensure these files are regularly backed up to prevent accidental deletion or changes breaking your AI models.

Your research team may find that it is creating many similar models – for comparable problems or multiple clients. Automate repetitive tasks to as great an extent as possible, with your Research team validating the results and using their specialised skills to adjust the network architectures.

Develop prototypes and solicit feedback beyond the AI team

During the research and development phase, your non-AI development and production teams should take the AI models you have in development and insert them into environments in which the models may be used.

These early prototypes will be incomplete and unfriendly for users, but will show the capacity for AI to solve the problem. Before your system can become a minimum viable product (MVP), prototypes will highlight related development work required – including website changes, the creation of database connections, mobile application modifications or development of application programming interfaces (APIs). Prototyping will engage stakeholders, allow other applications to call the model, enable initial scrutiny of results, and serve as a starting point for improvement.

During the prototype phase it is critical to solicit feedback from people outside the AI and production teams. Begin with internal stakeholders and, with each improvement, move closer to feedback from end users. Are your models:

- adequately performing their intended function?
- as fast as they need to be?
- scaling with usage as required?

Answering these questions early will avoid expensive redevelopment later. As with developing non-AI systems, frequent and iterative changes offer flexibility to address difficulties as they emerge.

Before your team completes the research and development iterations that feed your prototypes, finalise plans for a release process and for deploying code to its final environment. The number of stages in this process, and its complexity, will depend on factors including: the importance of controlling the code (processes for code review, testing, code merging, build, and versioning); the implications of system downtime; and the level of automation you require.

Considerations are company-specific – but evaluate:

- During development, will your code and models be tested with every update or only when there is a viable release candidate?
- Will testing be automated? How frequently will tests be updated?
- Does a successful test trigger a live deployment? What manual steps are required for a system to be made live?
- Will you deploy code directly or create containers? How will large AI model files be deployed? Will system downtime be required to make new versions live?

If you have existing development practises, adopt these to the extent possible to ensure that AI is not considered separately from the rest of your team's development efforts.

While automating release based on certain metrics may be straightforward, understanding whether a new AI system is an improvement overall may be difficult. A new version of your AI system may offer improved accuracy at the expense of speed, or vice versa. Whether you are automating deployment or verifying it manually, prioritise what is important to your use case.

"Spend a month to get a weak model and then iterate to make it great."

Eddie Bell, Ravelin

When moving from MVP to Production, select an appropriate hosting environment

With an initial model, supporting code and an established deployment process you should have a minimum viable product (MVP) ready for release to your production (live) environment. The MVP is distinct from your prototypes – while imperfect, it will contain all the elements required to solve the problem including peripheral components (web interfaces, APIs, storage and versioning control).

Having a model in production does not mean it needs to be publicly visible or impact live results. It should, however, be exposed to live data so your team can make refinements until it meets the requirements for a major release. With live data you can undertake longer-running tests and provide your data science team with feedback on what is working well and what is not. At this stage, prioritise establishing a controlled release process with thorough code testing, and the stability of your solution. You should also monitor the performance and scalability of your system.

Plan continual cycles of improvement – investigate and implement ideas for iterating the model, changing the interface and responding to feedback. New models must be demonstrably superior to old ones. Test all changes before updates are released to the production environment, allocating work between the AI team and the general development team. These cycles will continue for the life of the system.

If you've yet to decide where your system will run – on premise, in a data centre or in the cloud – at this point you will have the information you need to select an environment, and hardware, that are suitable for your needs.

On-premise: If your data is highly sensitive and cannot leave your network, or you wish to keep data and inferencing entirely under your control, you may wish to host your AI systems within your own premises. Usually, this is possible only for companies that have their own internal hardware infrastructure already. This can be a highly cost-effective option if the volume of requests to manage is known and relatively stable. However, all new hardware must be ordered and provisioned, which will limit scalability. Further, security will be entirely your responsibility. As such, on-premise deployment is a less preferred option for early stage companies that will lack these specialised skills.

Fig. 38. When to use on-premise deployment

Use on-premise if you:	Avoid on-premise if you:
Need to fix your costs	Do not have robust in-house security expertise
Have existing on-premise hardware	Cannot guarantee volumes of requests
Are working on highly sensitive data	Need your models to be accessed from outside your network

Source: MMC Ventures

Data centre: If you can afford the capital expense of buying servers, and have limited need to scale rapidly, hosting your own hardware in a data centre – either your own or a third party – can be an attractive option. The cost, even over a year, can be far lower than using a cloud service and you will maintain control over your system's performance. Using a data centre can also be sensible when you already have large volumes of data on your own servers and wish to avoid the cost of uploading the data to a cloud service.

The capital expense of establishing and managing a data centre can, however, be high – although for early stage companies there are programmes, such as NVIDIA Inception, which offer discounted hardware. As with the on-premise option, only consider a data centre approach if your company already has its own data centre for other servers, as well as staff with skills to install and configure your hardware. In addition to the upfront cost, the distraction of managing a data centre may prove an inappropriate and unwelcome distraction for your early stage company focused on its core initiatives.

Cloud: For good reason, many early stage companies choose cloud hosting from the outset. Amazon AWS, Google Cloud, Microsoft Azure and Rackspace are popular cloud providers. The majority of cloud providers offer specialised options so you can begin quickly and need set up little more than a security layer and links to other systems in their cloud.

Further, time-based costings allow rapid upscaling and downscaling of resources as required. For companies without dedicated system administrators, cloud may be an easy choice. You will, however, pay a premium for the service. A year of cloud hosting can cost twice as much as hosting in a data centre. You will also pay to transfer data in and out. Nonetheless, costs are payable monthly, rather than as a single large capital expenditure, and you will also avoid the cost of staff to manage the hardware.

Fig. 39. When to use a data centre for deployment

Use this approach if you:	Avoid this approach if you:
Wish to fix your costs	Require flexibility in your resourcing
Have existing data centre hardware	Wish to avoid high up-front capital costs
Seek control over your data	

Source: MMC Ventures

Unless there is a compelling reason to do so – cost, location or you are replacing a supplier – it is usually desirable to use the same cloud provider for your AI hosting that you use for your other infrastructure. This will limit data transfer costs and provide a single infrastructure to manage for security and resilience.

Although cloud systems offer extensive physical security, be aware that you are placing your software and data on the internet. If you do not secure the cloud servers that you establish, you will be at risk. Your cloud provider should ensure that: access to their data centre is secure; their data centre has multiple power sources and internet connections; and there is resilience in all supporting infrastructure such that the provider can resist any direct security challenge either in person or via attempted hacks into their network. They should also ensure that the data images they provide to you are secured from the rest of their infrastructure and other customers'. It is your responsibility, however, to ensure that your systems on their infrastructure are secure. Direct access to your account should only be via multi-factor authentication – not a simple username and password. Data stored should be private and any external data access or calls to your AI must be established using best practices for authentication. There are many malicious individuals who scan the IP addresses registered to cloud providers, looking for unsecured systems they can exploit. Finally, consider the physical location in which your cloud servers are hosted. Different countries have varying rules regarding data and hardware. You may need to keep your data within its area of origin. Be aware of local laws that could allow the cloud servers to be restricted. US law for example, allows hardware from a cloud provider to be seized if authorities suspect its use for criminal activity. If you are unlucky enough to have data on the same physical system, you could lose access to your systems without notice. This risk can readily be mitigated with appropriate monitoring of your remote systems and images of your servers that you can start in other zones if required. Finally, different regions may have varying performance at different times of day – a dynamic you can use to your advantage.

Fig. 40. When to use cloud deployment

Use this approach if you:	Avoid this approach if you:
Need flexibility in resource	Already have systems and personnel established in a data centre
Have existing systems and data in the cloud	Use highly sensitive data
Have limited capital to get started	

Source: MMC Ventures

For good reason, many early stage companies choose cloud hosting from the outset.

Test for precision, recall and accuracy at multiple stages

Proving that new AI releases are effective, and an improvement on prior versions, differs from the typical software quality assurance (QA) process. Test your AI system at multiple stages:

- During training: While your model is being trained, constantly test it against a subset of training data to validate its accuracy. The results will not represent the performance of the model fully, because the randomised test data will have influenced the model. As a result, this testing will overstate the model's accuracy.
- During validation: Set aside a part of your training data for validation. This test set – known as the validation set – is never used for training. Accordingly, the predictions your AI system makes from the validation set will better represent the predictions it makes in the real world. Validation accuracy is usually lower than training accuracy. If your data set does not represent real world data well, however, validation accuracy will still over-report the accuracy of your model.
- Continuously: Once your model has been created, test it against live data for a more appropriate measure of accuracy.

"Accuracy" has a specific meaning in AI – but, confusingly, is also used as a general term to cover several measures. There are three commonly-used measures of accuracy in AI: recall, precision and accuracy. Understand these measures to decide which are important for your systems so you can validate them appropriately.

Consider an AI that determines whether an apple is 'good' or 'bad' based on a picture of the apple. There are four possible outcomes:

1. True positive: The apple is good – and the AI predicts 'good'.
2. True negative: The apple is bad – and the AI predicts 'bad'.
3. False positive: The apple is bad – but the AI predicts 'good'.
4. False negative: The apple is good – but the AI predicts 'bad'.

Using the example above:

- Recall: ***What proportion of the good apples did I find correctly?***
  The number of correctly identified good apples divided by the total number of good apples (whether correctly identified or not).
- Precision: ***What proportion of the apples I said are good, did I get right?***
  The number of correctly identified good apples divided by the total number of apples labelled as good (whether correctly identified or not).
- Accuracy: ***What proportion of the apples did I label correctly?***
  The number of apples correctly identified as good or bad, divided by the total number of apples.

Avoid the temptation to use a single measure that flatters results. You will obtain a truer picture by using all three measures.

Balancing precision and recall can be difficult. As you tune your system for higher recall – fewer false negatives – you will increase false positives, and vice versa. Whether you elect to minimise false negatives or false positives will depend on the problem you are solving and your domain. If developing a marketing solution, you may wish to minimise false positives. To avoid the embarrassment of showing an incorrect logo, missing some marketing opportunities may be acceptable. If developing medical diagnostics, on the other hand, you may wish to minimise false negatives to avoid missing a diagnosis.

Automate testing to as great an extent as possible. Every new model should be tested automatically. "Efficiency is critical. If you have to do something more than once, automate it." (Dr. Janet Bastiman, Chief Science Officer, Storystream). If all measures of accuracy are higher, the decision to deploy the new model will be straightforward. If measures of accuracy decrease, you may need to verify the new model manually. A decrease in one measure of accuracy may not be problematic – you might have re-tuned your model for precision or recall, or decided to change the entire model to improve performance. If your models produce results that are concerning, speak to your AI team to discuss why. It may be that your training data set does not contain enough appropriate data. If you encounter problems, add examples of these types of data to your test set so you can monitor improvements.

An effective maintenance programme will sustain your AI's intelligence

A deployed AI solution reflects a point in time; available data, business requirements, market feedback and available techniques will change. Beyond the typical maintenance you would perform on any software system, you need to verify and update your AI system on an ongoing basis. Once your solution is live, ensure it continues to perform well by:

- Continuously sampling its result and verifying that the outcome from your model is as you expect from live data.
- Adding problematic data to your test set to ensure your team has addressed issues.
- Exploring whether new, third-party APIs are available which outperform your model, which will enable you to focus your AI team's efforts onto harder problems.
- Ensuring your system issues alerts if incoming data fails, so problems can be addressed.

AI technology is developing at pace. Further, the varieties and volume of available training data continue to evolve. Invest in continual improvement to ensure the system you develop today avoids obsolescence.

Available data, business requirements and techniques will change over time. Invest in continual improvement to avoid obsolescence.

Chapter 6

Regulation & Ethics

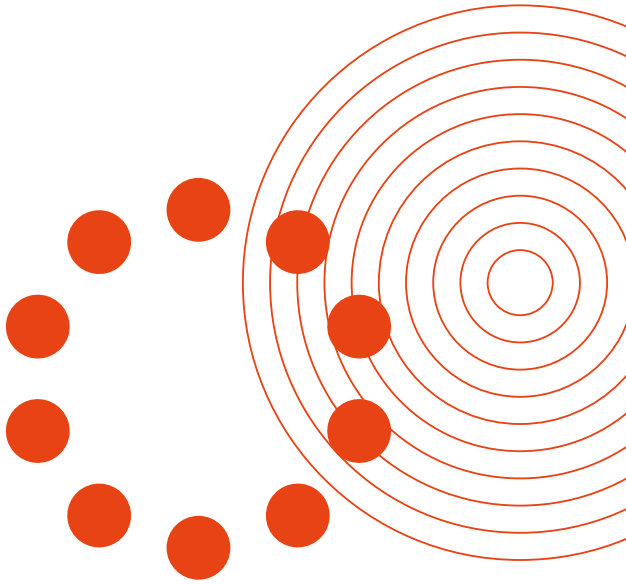

Summary

- As consideration of data privacy grows, and with the General Data Protection Regulation (GDPR) in force across the European Union (EU), it is vital to ensure you are using data appropriately. The GDPR applies to all companies processing the personal data of people in the EU, regardless of a company's location.

- Companies that are 'controllers' or 'processors' of personal information are accountable for their handling of individuals' personal information. Demonstrate compliance with GDPR data handling requirements and the principles of protection, fairness and transparency.

- Minimise the personal data you require, to reduce regulatory risk, and pseudonymise all personal data through anonymisation, encryption or tokenisation.

- In addition to standardising data handling requirements and penalties for misuse, the GDPR introduced considerations that can impact AI systems specifically. Verify that automated systems meet GDPR stipulations. Article 22 of the GDPR prohibits legal effects that result solely from automated processing being undertaken without an individual's explicit consent, when consent is required. Several legislative terms are subject to interpretation at this time. It may be prudent to make your system advisory only, and include a human check, if you are developing a system that could materially impact an individual's life.

- 'Explainability' – explaining how the outputs of your AI system are derived – is growing in importance. Convention 108 of the Council of Europe, adopted into UK and EU law in May 2018, provides individuals with the right to obtain knowledge of the reasoning underlying data processing systems applied to them. Explainability can be challenging in relation to deep learning systems. Explore varying approaches to explainability including Inferred Explanation, Feature Extrapolation and Key Variable Analysis. Each offers trade-offs regarding difficulty, speed and explanatory power.

- Develop a framework for ethical data use to avoid reputational and financial costs. The ALGOCARE framework, developed by the Durham Police Constabulary in partnership with academics, highlights issues you should consider when managing data. It incorporates: the nature of system output (Advisory); whether data is gathered lawfully (Lawful); whether you understand the meaning of the data you use (Granularity); who owns the IP associated with the data (Ownership); whether the outcomes of your system need to be available for individuals to challenge (Challenge); how your system is tested (Accuracy); whether ethical considerations are deliberated and stated (Responsible); and whether your model has been explained accessibly to as great an extent as possible (Explainable).

Regulation & Ethics: The Checklist

Comply with regulations and license requirements

- ☐ Review your compliance with current legislation including the UK Data Protection Act, the EU GDPR and EU Convention 108
- ☐ Monitor proposed legislation to anticipate implications
- ☐ Review permissions for the customer data you collect
- ☐ Check that the data sets you use are available for commercial use
- ☐ Validate licenses for open source models you use

Deliver explainable, ethical AI

- ☐ Understand industry-specific explainability requirements
- ☐ Define an explainability framework
- ☐ Select and apply an approach to explainability
- ☐ Update your framework documentation as your models and data change
- ☐ Validate that your use of data is ethical

As consideration of data privacy grows, and with the new General Data Protection Regulation (GDPR) in force across the European Union, it is important to ensure you are using data appropriately. Today, data permissioning and management are critical aspects of any AI-driven company. Below, we describe regulatory and ethical considerations to help you manage safely the data you use to build your models. Seek legal advice to ensure compliance with any applicable legislation; the information below is introductory in nature and will not reflect your company's individual circumstances.

Ensure compliance with GDPR data handling requirements

The GDPR came into force across the European Union on 25th May 2018. It applies to all companies processing the personal data of people in the EU, regardless of a company's location. Among other considerations, it standardises data handling requirements and penalties for data misuse. Article 83 of the GDPR specifies fines of up to 4% of a company's global revenue or €20m – whichever is greater – for non-compliance.

Individuals, organisations and companies which, according to the GDPR, are either "controllers" or "processors" of personal information are accountable for their handling of individuals' personal information. Companies must "implement measures which meet the principles of data protection by design and by default". Transparency and fairness are also key concepts within the GDPR. You must be clear and honest regarding how you will use individuals' personal data – and must not use personal data in a way that is unduly detrimental, unexpected or misleading for the individuals concerned.

Demonstrate compliance with the GDPR principles of protection, fairness and transparency in multiple ways, including by:

- Collecting only the data you require
- Being transparent regarding why data is collected, what it will be used for and who will have access to it
- Ensuring you have appropriate permissions to store and process your data
- Removing unnecessary personal data
- Deleting data when its agreed purpose has been fulfilled
- Anonymising data, where possible, to remove personal identifiers
- Encrypting personal data
- Securing physical access to your data storage
- Limiting access to your data
- Monitoring data access and saving an audit trail of individuals who have viewed or changed data
- Using data only for the purposes agreed
- Implementing a process to provide an individual with a copy of all the data you hold about him or her
- Implementing a process to remove all the data you hold about a specific individual.

The GDPR has expanded the definition of personal data, which broadly refers to information relating to an identified or identifiable person, to include information "specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that person". This includes job titles, employers and social media handles – even if the individual has made these public on selected websites. While some information may directly identify an individual, other information may do so indirectly if combined with other information. In both circumstances, the data is deemed personal information. Further, if you are inferring personal information – such as gender, age or salary – using your system, you must treat the information as if it were gathered from the user directly.

Demonstrate compliance with GDPR principles of protection, fairness and transparency in multiple ways.

Certain personal data – in categories including racial origin, religious beliefs, genetic data and data concerning a person's sexual orientation – may be considered "sensitive data" and require special care. Pseudonymise all personal data through anonymisation, encryption or tokenisation:

- **Anonymisation:** Remove or replace personal data with random information. Even if unauthorised individuals read the data, they will be unable to identify the data subject.
- **Encryption:** Encrypt personal data fields. The decryption key will be required to identify an individual from the encrypted data. The decryption key must be stored safely. AI techniques remain effective on encrypted data, enabling you to identify patterns even if the input data is not human-readable. This offers a way to incorporate personal attributes more safely.
- **Tokenisation:** Remove personal data from the main data set and replace it with numerical tokens that relate to each aspect of personal data. The process may be as simple as providing each individual with a unique identifier. The personal data and corresponding token are stored on a separate, more secure, system, allowing the data to be reconnected at a later date. Tokenisation is effective when one party has permission to view personal data and needs to interpret the results of the AI system, but the company providing the AI system does not need to view the personal data – for example, a medical analysis system.

Consider the security of data not just when it is stored but when it enters, moves within and leaves your environment.

Even with security best practices in place, holding personal data remains a risk. Minimise the personal data you require. If you can fully anonymise your data and avoid the need to store any personal information, do so. If you must store personal data, consider the security of data not just when it is stored but when it enters, moves within and leaves your environment. Examine every point where personal data could be read by an employee or malicious third party and ensure you have pursued every measure within your control to protect it. Delete data when it has been processed according to its agreed purpose.

Verify that automated systems meet GDPR stipulations

In addition to standardising data handling requirements and penalties for misuse, the GDPR introduced considerations that can impact AI solutions:

- **Article 22 (Paragraph 1):** *"The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her."*
- **Article 22 (Paragraph 2):** *"Paragraph 1 shall not apply if the decision:*
  - *is necessary for entering into, or performance of, a contract…[or]*
  - *is based on the data subject's explicit consent."*
- **Article 22 (Paragraph 3):** *"In the cases referred to in [Paragraph 2], the data controller shall implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision … and, in Recital 71 only, the right to an explanation of the decision."*

These articles are yet to be comprehensively tested in court. However, they explicitly prohibit legal effects – such as sentencing and parole decisions – that result solely from automated processing undertaken without the individual's explicit consent, when consent is required.

What constitutes "similarly significant" effects, "explicit consent" (beyond acceptance of an extensive set of online conditions containing a relevant paragraph) and whether something is "necessary" to perform a contract are subject to interpretation at this time.

If you are developing an AI system that could materially impact an individual's life, therefore, it is prudent to consider making your system advisory only and including a human check. Once case law has better established the meanings of the terms above, there will be greater clarity regarding the implications of the legislation.

The GDPR and Convention 108 impose obligations of explainability

Article 22 (Paragraph 3) of the GDPR, which requires companies to protect the data they control and allows individuals to challenge an automated system they believe is treating them unfairly, demands a robust explanatory framework for the outputs of your systems. Convention 108 of the Council of Europe (https://bit.ly/2n6POrT), adopted into UK and EU law in May 2018, imposes related requirements:

- **Convention 108 (Article 8):** *"Every individual shall have a right… to obtain, on request, knowledge of the reasoning underlying data processing where the results of such processing are applied to him or her."*

Convention 108 affords individuals the right to understand how decisions made about them, using data processing, have been developed. Because every individual possesses this right, you must be able to explain, in lay terms, how decisions that affect individuals are made.

Explore varying approaches to explainability

Beyond the stipulations of Convention 108 of the Council of Europe, there is growing demand more broadly for greater explainability of AI systems. Improved explainability was a recommendation, for example, from the UK Parliamentary Science and Technology Select Committee on AI. Regulatory or pragmatic demands may force you to consider the explainability of your systems.

For a system that uses a decision tree, it will be straightforward to explain how data maps to the system's decision. For machine learning-based systems, and particularly deep learning-based systems, this will not be possible. There may be thousands of abstract numbers corresponding to the connections in the network that contribute to its output. These numbers will be meaningless to individuals who seek to understand the system, which is why many AI systems are considered to be 'black box' and inexplicable.

There are, however, means of explanation that do not involve a system's mathematics. These approaches consider the impact of variables inputted to a system and their influence on the output. There are several techniques you can apply including Inferred Explanation, Feature Extrapolation and Key Variable Analysis (Fig. 41). Which you favour will depend on the level of explainability you require and the challenge of providing it.

Fig. 41. Three approaches to explainability

Approach	Difficulty	Speed	Advantages	Disadvantages
Inferred Explanation	Low	Fast	Easy to understand	Limited explanatory power
Feature Extrapolation	Moderate	Slow	Easy to understand	Limited applicability
Key Variable Analysis	Very high	Very slow	Thorough	Challenging to understand

Source: MMC Ventures

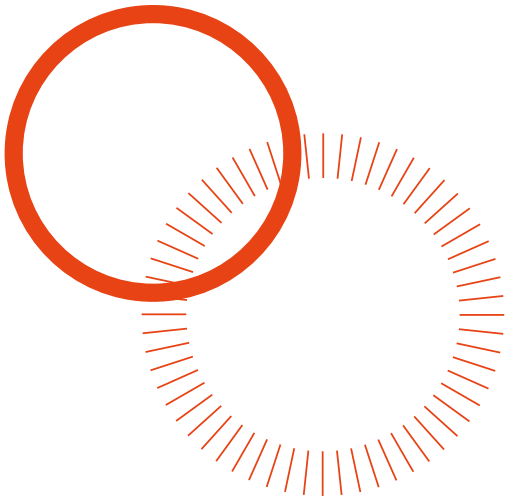

1. **Inferred Explanation:** Inferred Explanation is the easiest way to explain AI. The algorithm is not described and a 'black box' is retained around it. Correlations are considered between system inputs and outputs, without explaining the steps between.

By demonstrating examples of decisions, individuals can see the correlations between input data and output decisions (Fig. 42), without detail regarding how the inputs and outputs are connected. Inferred explanation does not provide complete clarity regarding a model, but will demonstrate how decisions relate to inputs in a manner that will be satisfactory in many situations.

Fig. 42. Inferred Explanation

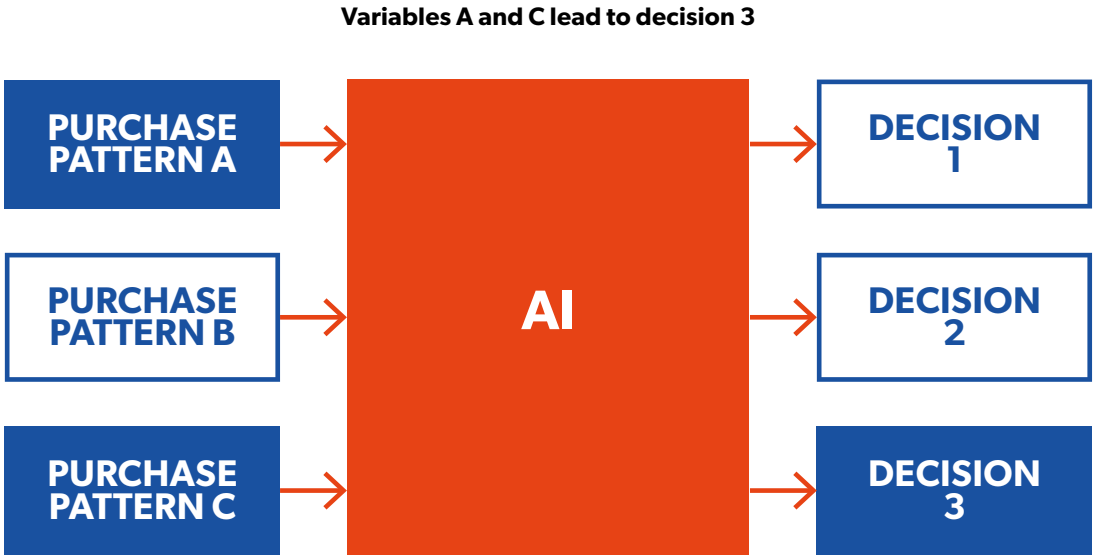

Source: MMC Ventures

There are means of explanation that do not involve a system's mathematics.

2. **Feature Extrapolation:** Some systems, including financial models and systems that materially impact individuals, require an explanation – beyond correlation – of how models reach their conclusions. While more effort, it is possible to evaluate features in data that are activating parts of a network. This is particularly fruitful for image classification systems. Using test data, and reversing the flow of data in a network, you can create images that demonstrate the features that activate a particular layer in the network (Fig. 43). Further, Google recently released a library for TensorFlow to undertake this visualisation automatically, within a browser, during training (bitly.com/2R6XeZu). While not suitable for all AI systems, feature extrapolation provides a degree of explainability in a manner that non-technical individuals can appreciate.

Some systems require an explanation – beyond correlation – of how models reach their conclusions.

It is possible to evaluate features in data that are activating parts of a network.

Fig. 43. Feature Extrapolation

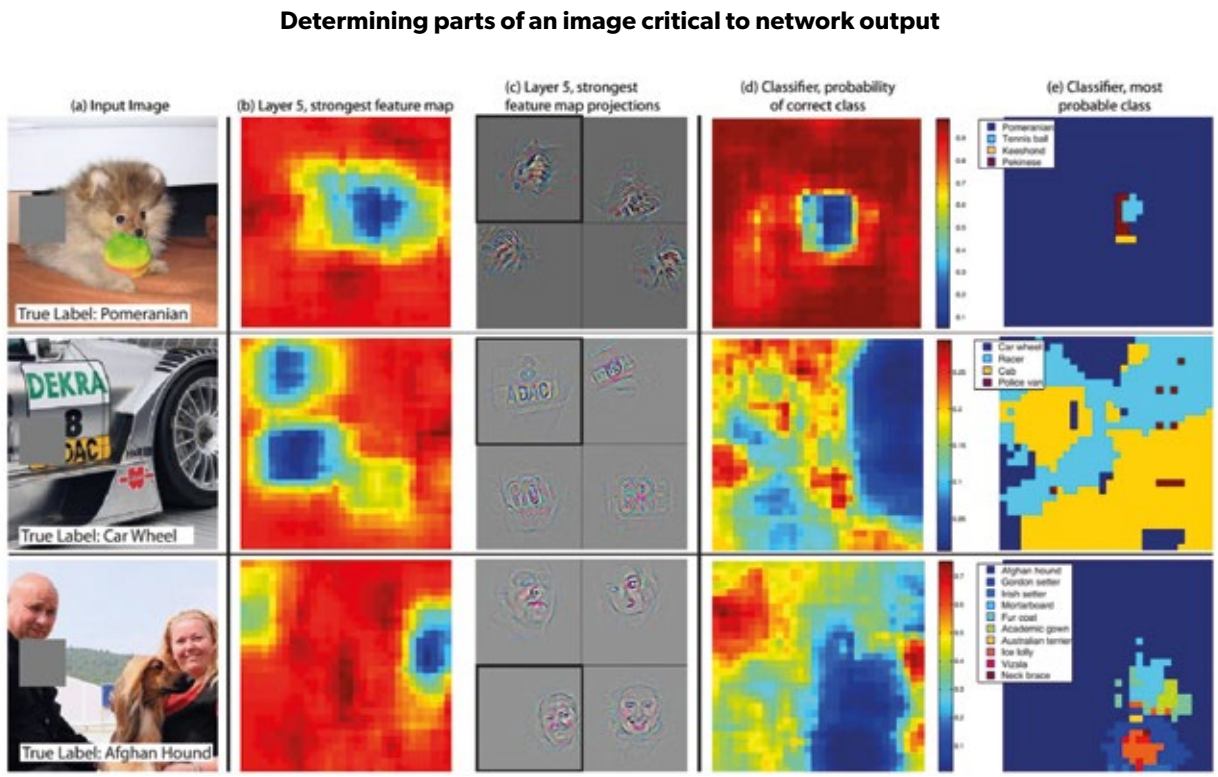

Source: Zeiler and Fergus, https://bit.ly/2JjF4R0

3. **Key Variable Analysis:** If you wish to provide the most precise explanation of your system, you must analyse the impact of each input on the system's decision-making process and overall decision. This will require a full statistical analysis and is a significant undertaking. For each output decision, you will require an example of input data that strongly results in that decision. Change each variable, in turn, from the value that leads to Decision 1 to the value that leads to Decision 2. Then, change the variables in combination. The effort required will increase exponentially according to the number of variables you have, and the process will be time-consuming. However, you will be able to determine whether any system inputs, singularly or in combination, have a disproportionate effect on your system's output decision ("variable bias"). You may find, for example, that your model places a high importance on gender when providing an output decision. This is possible, even if gender is not explicit in your data, if you have other closely associated variables (such as vocation in a gender biased industry).

Key variable analysis has drawbacks as well as advantages. In addition to being complex and resource-intensive, it can be difficult to explain results accessibly. Further, if you explain your model in a high degree of detail, malicious third parties can use this information to force results from your model that they wish to see.

If you wish to provide the most precise explanation of your system, you must analyse the impact of each input on the system's decision-making process and overall.

Fig. 44. How to select an approach to explainability

Use this approach if you:	Avoid this approach if you:
Inferred Explanation	
– Seek a high-level overview of your AI system – Believe correlation offers sufficient explainability	– Require detail regarding how variables lead to decisions
Feature Extraction	
– Require detail from within the network – Have a network type (e.g. images) where abstractions can be mapped onto input data	– Have limited time – Require precise impact of input variables, not general features – Are not using an assignment–based or generative AI network
Key Variable Analysis	
– Require detail about the importance of variables – Seek to prevent unwanted bias in your variables	– Have limited time – Seek to publish your results – Wish to offer a layperson's guide to your model

Source: MMC Ventures

Develop a framework for ethical data use

When developing and deploying AI systems, as well as providing sufficient explainability it is important to use data ethically. "Plan for ethical AI from the outset and underpinning all initiatives. It's got to be foundational, not an afterthought" (Steven Roberts, Barclays). In addition to the intrinsic importance of doing so, a growing number of companies are incurring reputational and financial costs from failing to do so.

The Durham Police Constabulary, in conjunction with computer science academics, is trialling a framework – ALGOCARE – to ensure its AI system uses data transparently within an explainable, ethical process. Many companies with AI systems also have frameworks in place, albeit privately and often loosely defined. While every company's framework differs, ALGOCARE highlights issues you should consider when managing data.

- **Advisory:** Is the output of the algorithm used as a suggestion or a fact? How to interpret the output of a system is a key consideration. Does a human, or automated system, act on the output without further thought or investigation? Do you wish your car, for example, to brake automatically if danger is perceived (even when the system is incorrect) or to warn you so you can make the decision? To an extent, what is optimal may be determined by the domain of the problem.

  Too often, the numbers returned alongside a label are interpreted and used as a confidence score. Usually, however, they will be a probability distribution that has been tuned to give an output above a specific level for the predicted result. Even incorrect results can have high probabilities. Decisions based on the "confidence" of the network decision, therefore, can be disastrous. While there are tuning techniques that better align a network's prediction probabilities with confidence levels (Guo et al, https://bit.ly/2JiRNTS) they are rarely used given the time required and teams' focus on measures of accuracy.

- **Lawful:** Is your data gathered lawfully and do you have the right to use it? Under the GDPR, individuals may not have consented to their data being processed in the way you intend. Data gathered without the informed consent of the individual should not be used.

- **Granularity:** Do you understand the meaning of the data you feed into your model? To avoid biased results and models that fail when exposed to real-world data, it is important to do so. What variables are missing, combined or inferred? How varied is your data? Do you have sufficient time series data? Data scientists can excel in this regard, questioning data and anticipating problems before building models.

- **Ownership:** Who owns the intellectual property associated with the data and algorithms you use? Many companies use academic data sets and models to validate concepts in the early stages of development. Confirm that the licenses, for both the data and models you use, are suitable for commercial use. The availability of something on the internet does not confer on your company the right to use it.

- **Challengeable:** Do the outcomes of your system need to be available to individuals? For example, under GDPR may the system be challenged? In some sectors, there will be a greater need than others to be open about the basis of your results. If you undertake all projects assuming that your results will be challenged, you will be prepared.

"Plan for ethical AI from the outset and underpinning all initiatives. It's got to be foundational, not an afterthought."

Steven Roberts, Barclays

- **Accuracy:** How is your system tested? How is changing data or inaccuracies fed back into the system? Is dated data removed? Many companies report 'accuracy' by cherry-picking precision or recall (Chapter 5), which will overstate model performance. Continuous testing with real world data is the only way to verify your model's performance.

- **Responsible:** Are ethical considerations deliberated and stated? The impact of your system will depend on its domain and the consequences of false positives and false negatives. If your system could adversely impact an individual's life, you must understand and consider the implications.

- **Explainable:** Has your model been explained, in jargon-free language, to as great an extent as possible without exposing your intellectual property? Explanations that avoid technical terminology will enable everyone in your business to understand what you are creating and to challenge it. "Having a framework to explain your models is valuable for all stakeholders" (Dr Janet Bastiman, Chief Science Officer, StoryStream). What does your model do? How do inputs correlate with outputs?

EU and UK Parliamentary committees, including the Science and Technology Select Committee and the House of Lords Artificial Intelligence Select Committee on AI, are engaged on issues of AI, explainability and data privacy. The UK Science and Technology Select Committee, for example, launched an inquiry into the use of algorithms in public and business decision-making. Further, more specific, legislation is probable. Ensure that a senior member of your team (Chief Science Officer, Head of AI or Head of Data) is responsible for staying up-to-date regarding proposed legislation and the impact it could have on your business.

EU and UK Parliamentary committees are engaged on the issues of AI and explainability.

Ensure that a senior member of your team is responsible for staying up-to-date regarding proposed legislation.

MMC Ventures
24 High Holborn
London WC1V 6AZ
www.mmc.vc

Barclays UK (BUK) Ventures
Registered office:
1 Churchill Place
London
E14 5HP
www.home.barclays

Report design by
Aviary Creative
www.aviarycreative.co.uk

The AI Playbook

The step-by-step guide to taking advantage of AI in your business