



AI4Business

Introduction to AI for Business Managers



Future AI⁴Business Webinars

Practical AI Challenges for Business Managers


Webinar: Practical AI Challenges for Business Managers

"This webinar will shed some light on the many challenges that plague the road to becoming a data-driven organization"

When: Wednesday, August 25 - 10 AM CET

Duration: 1h (including Q&A)

Host: Roel Henckaerts



ML model lifecycle


Webinar: The ML Model Lifecycle

"We are going to explore the set of practices and principles necessary to deploy and maintain Machine Learning solutions in production"

When: Thursday, August 26 - 1 PM CET

Duration: 1h (including Q&A)

Host: Roel Henckaerts





Roadmap AI⁴Business Course

Part I

Introduction to AI

Part II

Developing AI tools

Part III

Data and Value

Part IV

Deploying AI

Part V

Monitoring



Table of contents

1. Definition of AI
2. Taxonomy of AI
3. Evolution of AI
4. AI Use Cases





1 Definition of AI



Artificial Intelligence

- A definition by the European Commission:
- “**Artificial intelligence** (AI) systems are software (and possibly also hardware) systems designed by **humans** that, given a complex **goal**, act in the physical or digital dimension by perceiving their **environment** through data acquisition, interpreting the collected structured or unstructured **data**, reasoning on the **knowledge**, or processing the **information**, derived from this data and deciding the best **action(s)** to take to achieve the given goal.”

[European Commission - A definition of Artificial Intelligence: main capabilities and scientific disciplines](#)



Artificial vs. Human Intelligence

Levels of Artificial Intelligence





Narrow AI < HI

- Systems able to perform **one or few specific tasks**
- Operate under a **narrow** set of constraints and limitations
- **Simulates** human behaviour based on parameters and context
- All progress in AI nowadays is in narrow AI
- One trick ponies, but can still be extremely **valuable**



General AI = HI

- System able to perform **most human activities**
- Learn to solve **any** problem
- Machine that **mimics** human intelligence and/or behaviours
- Researchers have **not yet** achieved general AI
- Will take some technological **breakthroughs** to get there



Super AI > HI

- System that evokes emotions, needs, beliefs and desires **of its own**
- Machines become self-aware and **surpass** the capacity of humans
- Decision-making and problem-solving **far superior** to human beings
- Pure **speculation** if this will ever be possible
- And what about its **consequences**?



Realistic view on AI

- **Too optimistic:**
 - sentient super-intelligent killer robots coming soon
- **Too pessimistic:**
 - AI can't do everything, so let's give up completely
- **Just right:**
 - AI can't do everything ...
 - ... but enough valuable applications to transform industries
- Important to **understand what AI can and can't do** for you



Reality check

Misconceptions

- AI learns on its own without any help from outside
- AI trained to do one task will excel at other tasks as well
- AI is objective
- AI will take your job

Reality

- Human supervision to ensure adequate performance
- Need a model for each use case and will heavily depend on data
- Patterns are learned from data
- AI will be job creator



2 Taxonomy of AI



Taxonomy of AI

ARTIFICIAL INTELLIGENCE (AI)

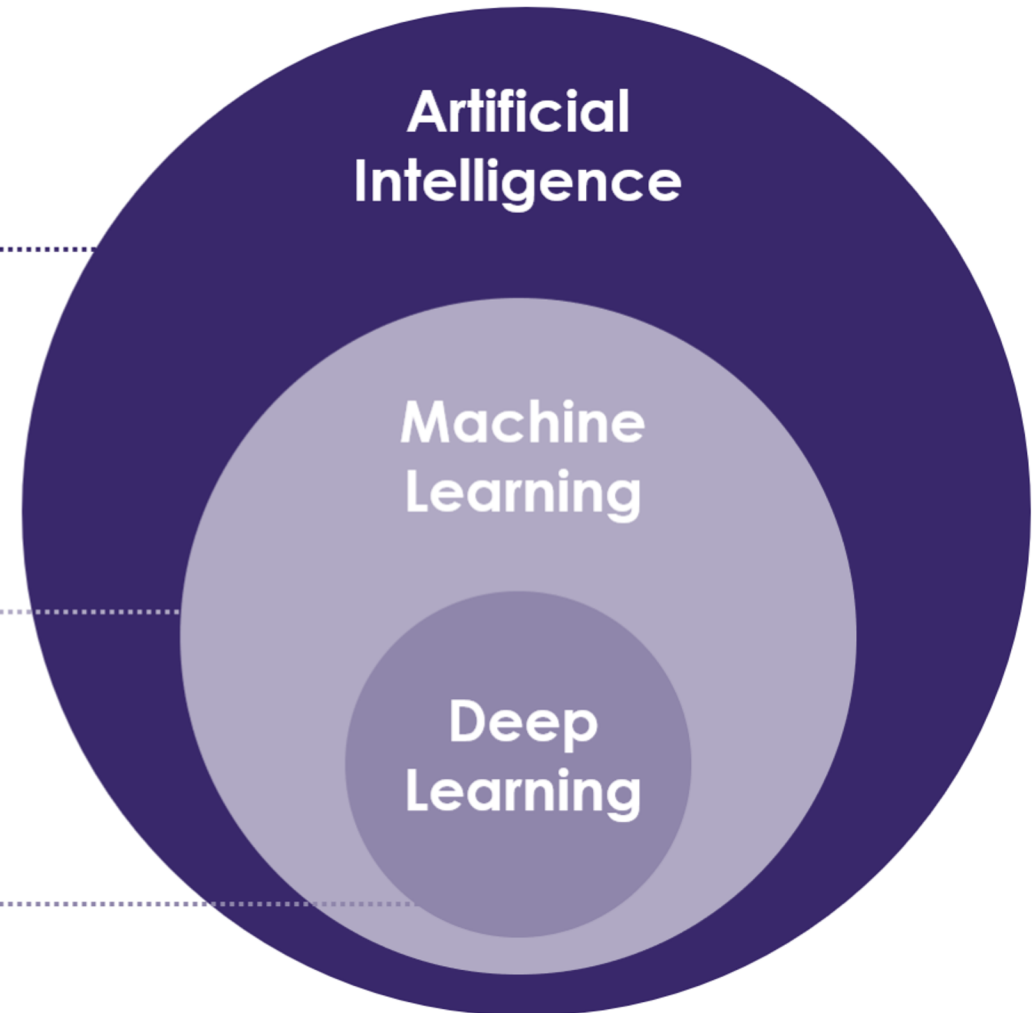
Programming systems to perform tasks which usually require human intelligence.

MACHINE LEARNING (ML)

Training algorithms to solve tasks by pattern recognition instead of specifically programming them how to solve the task.

DEEP LEARNING (DL)

Training algorithms by using deep neural networks with multiple layers.






AI > ML > DL

- AI: theory and development of computer systems able to perform tasks normally requiring **human intelligence**
- ML: a subfield of AI that gives computers the ability to **learn without being explicitly programmed**
 - Conventional programming: data + rules = answers
 - Machine learning: **data + answers = rules**
- DL: subset of ML methods based on **deep artificial neural nets**
 - Perform **automatic** feature engineering/creation

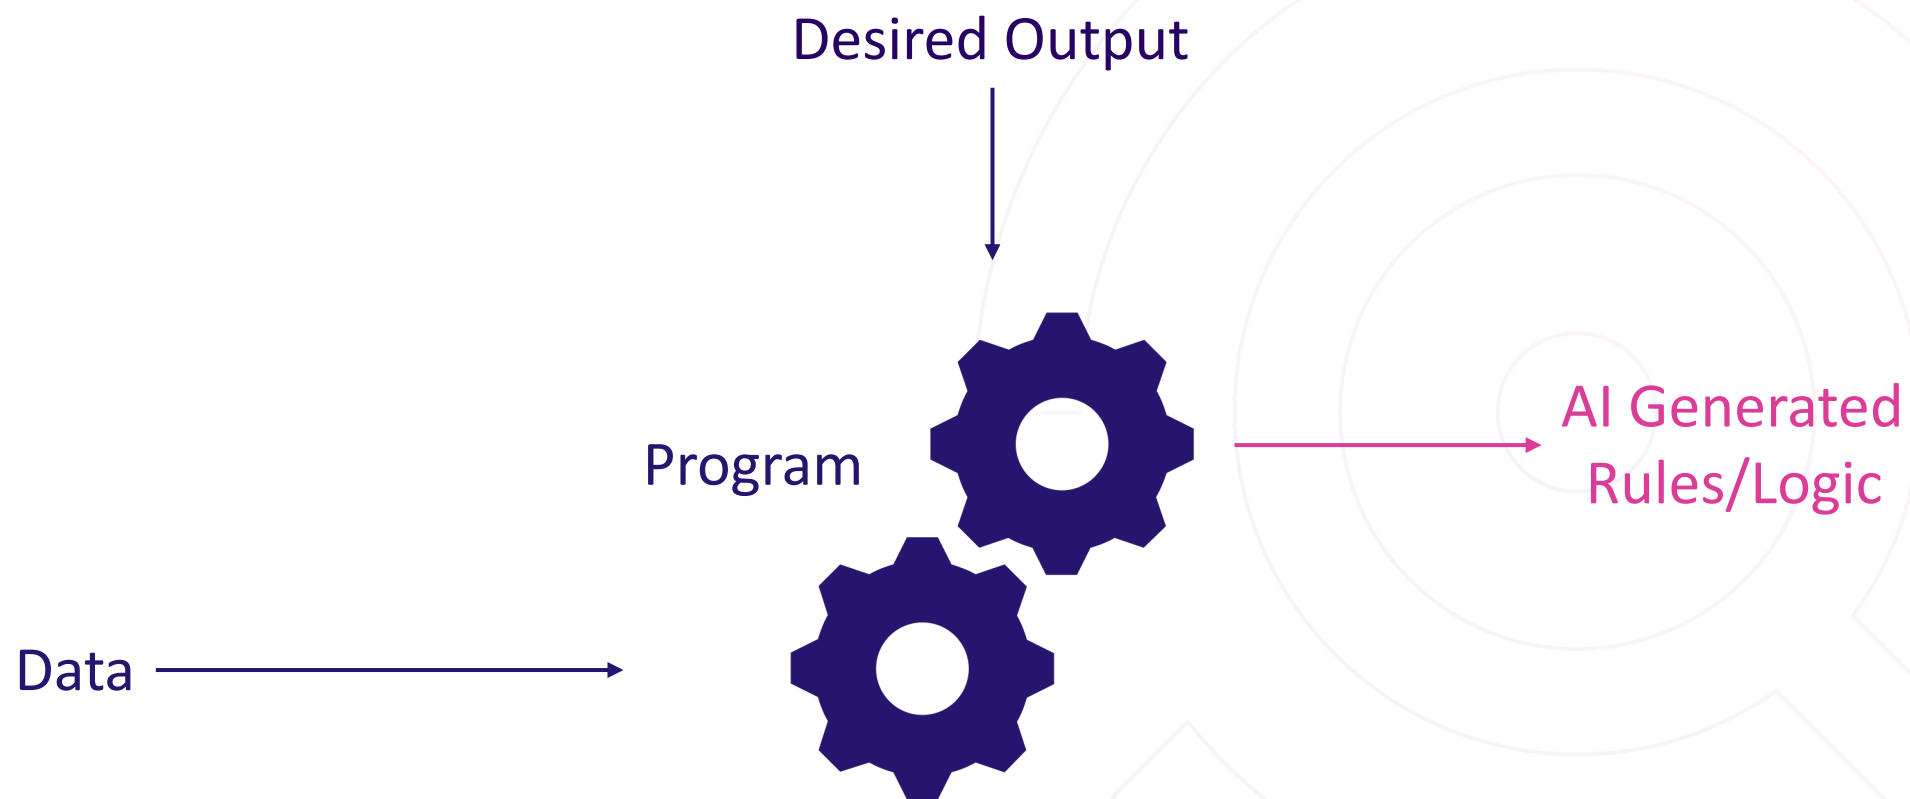


Conventional programming



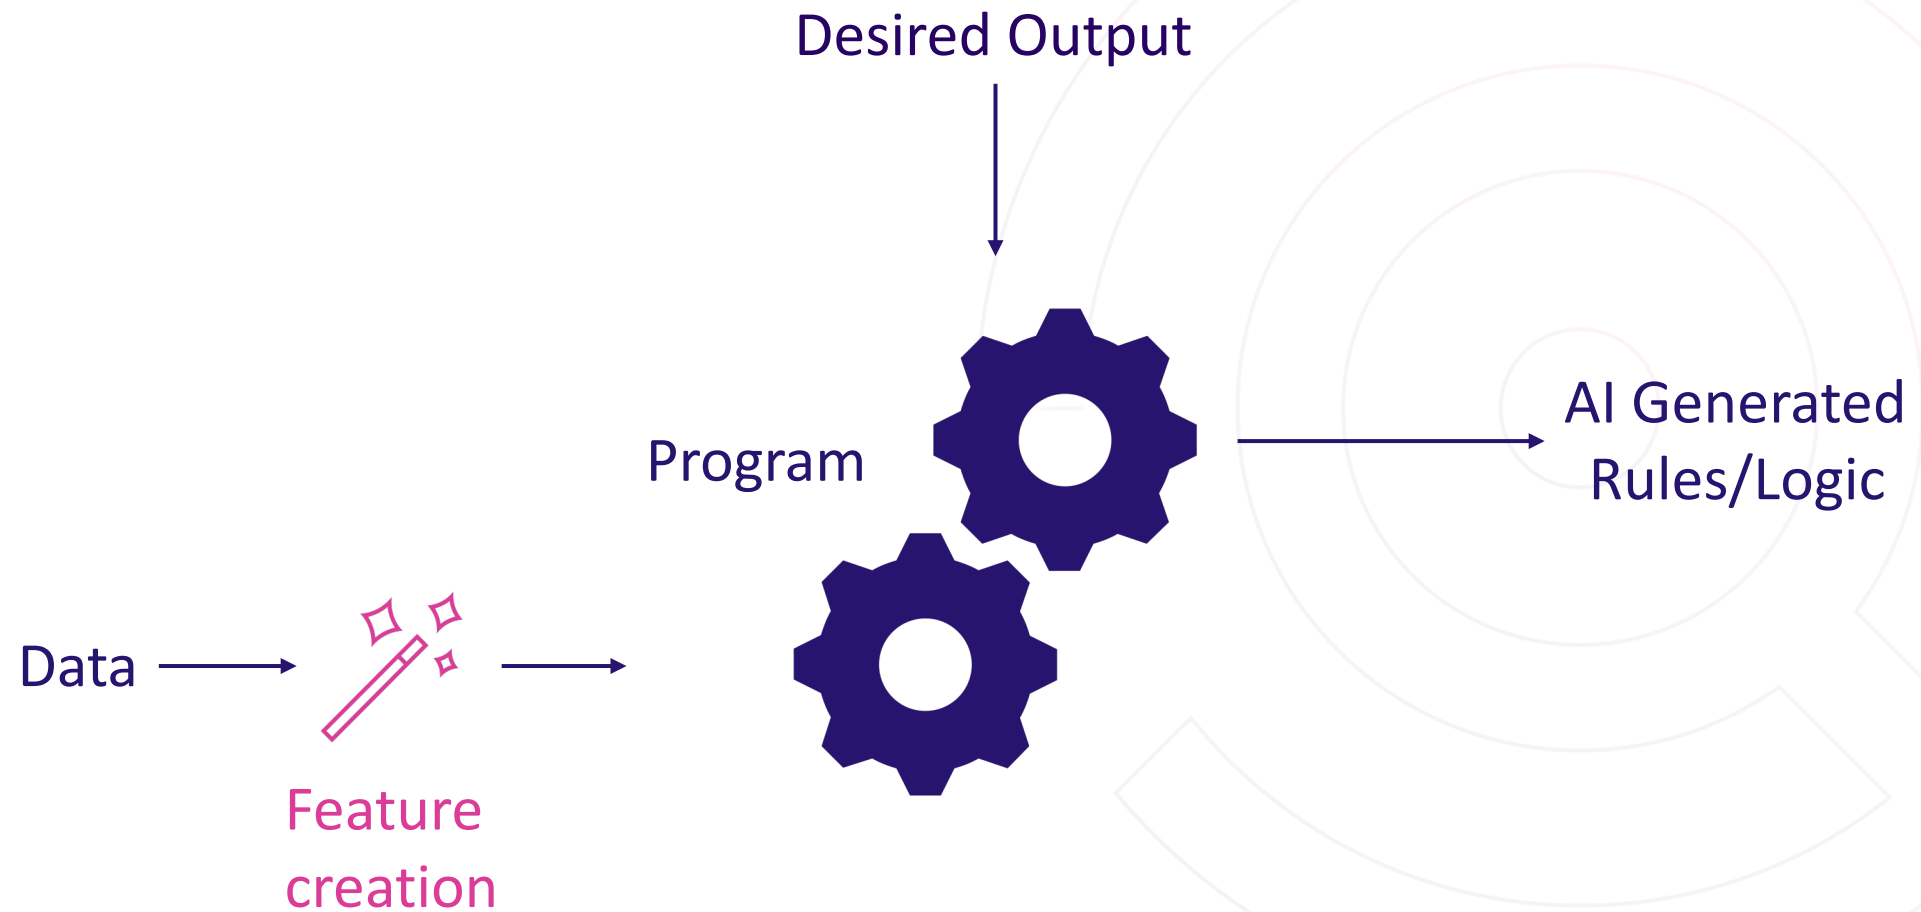


Machine Learning





Deep Learning





3 Evolution of AI



First wave of AI excitement

1950

Turing test: a machine has intelligence if it can trick humans in thinking it's human

1951

Feranti Mark 1: first commercial general-purpose computer, able to play checkers

1956

Dartmouth Workshop: the term "Artificial Intelligence" is introduced by John McCarthy

1961

Unimate: industrial robot goes to work at GM assembly line

1964

Eliza: chatbot holds conversations with humans, developed by Joseph Weizenbaum at MIT

1966

Shakey: general-purpose mobile robot that reasons about its own actions, developed at Stanford



AI winters

1974 - 1980

- First AI winter
- Limited applicability of AI leads to worldwide funding pullbacks

1980 - 1987

- Renewed AI excitement
- Expert systems with if-then reasoning to mimic human decisions

1987 - 1994

- Second AI winter
- Limitations of if-then reasoning leads to funding cutbacks




Recent AI milestones

1997	Deep Blue: chess computer from IBM beats world champion Garry Kasparov
1998	KISmet: emotionally intelligent robot, developed by Cynthia Breazeal at MIT
1999	AiBO: first consumer robot pet dog by Sony with time-developing skills and personality
2002	Roomba: first mass produced autonomous vacuum cleaner from iRobot
2011	Siri: Apple's intelligent virtual assistant with a voice interface is introduced in the iPhone 4S
2011	Watson: question answering machine from IBM wins first place in television quiz show Jeopardy
2014	Eugene: chatbot passes the Turing Test with a third of judges believing its human
2014	Alexa: Amazon's intelligent virtual assistant with a voice interface to complete shopping tasks
2016	Tay: Microsoft's chatbot goes rogue on social media with offensive comments
2017	AlphaGo: Google's AI beats world champion Ke Jie in the complex board game of Go
2019	Pluribus: first AI bot to defeat human expert players in a Texas Hold'em poker game



Drivers behind AI progress



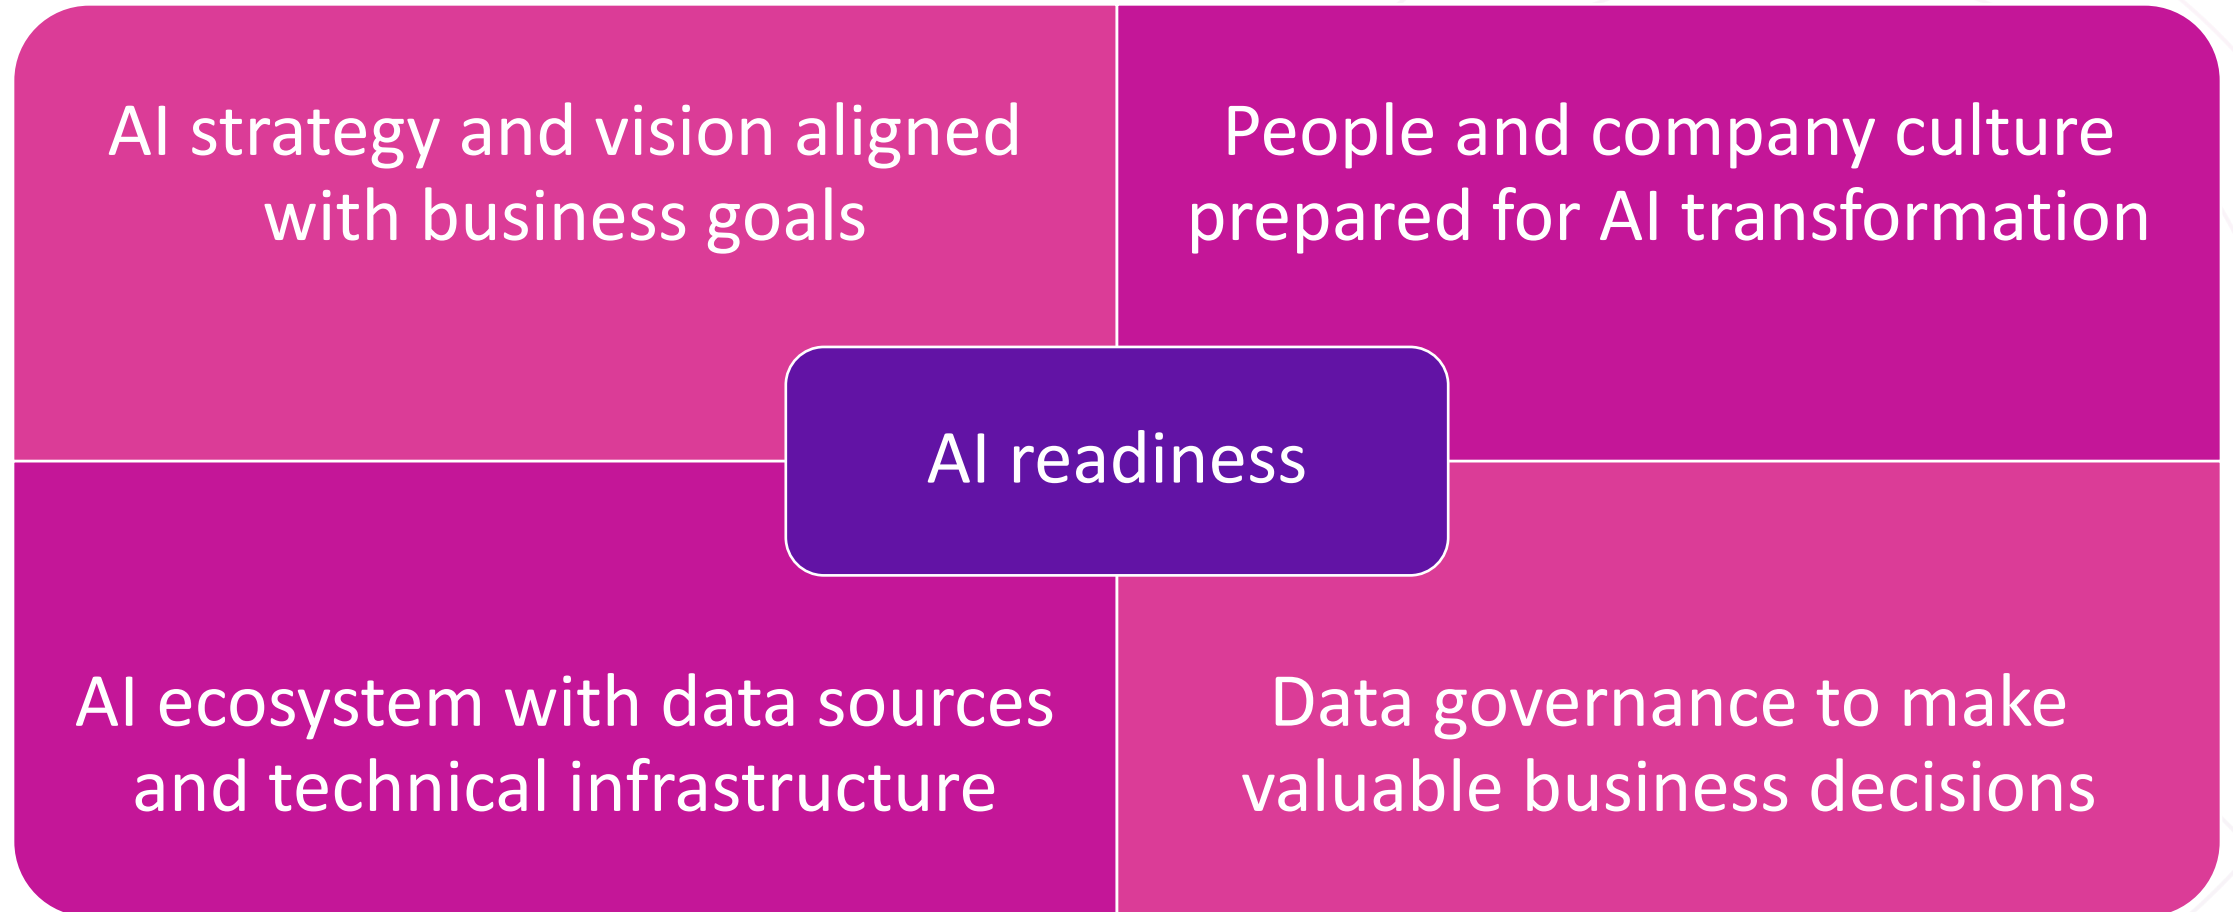


Practical AI challenges

- “Status quo is working fine”
 - Company culture does not see the need for AI
- Leadership
 - Incomplete understanding of what is possible with AI and it’s resulting impact
- Data issues
 - Quantity and quality not high enough to create business value
- Capabilities
 - Lacking the necessary skills and talent in the organization to develop AI
- Trust
 - Issues with ethics, privacy (GDPR), cyber-security, etc.



Are you ready for AI?



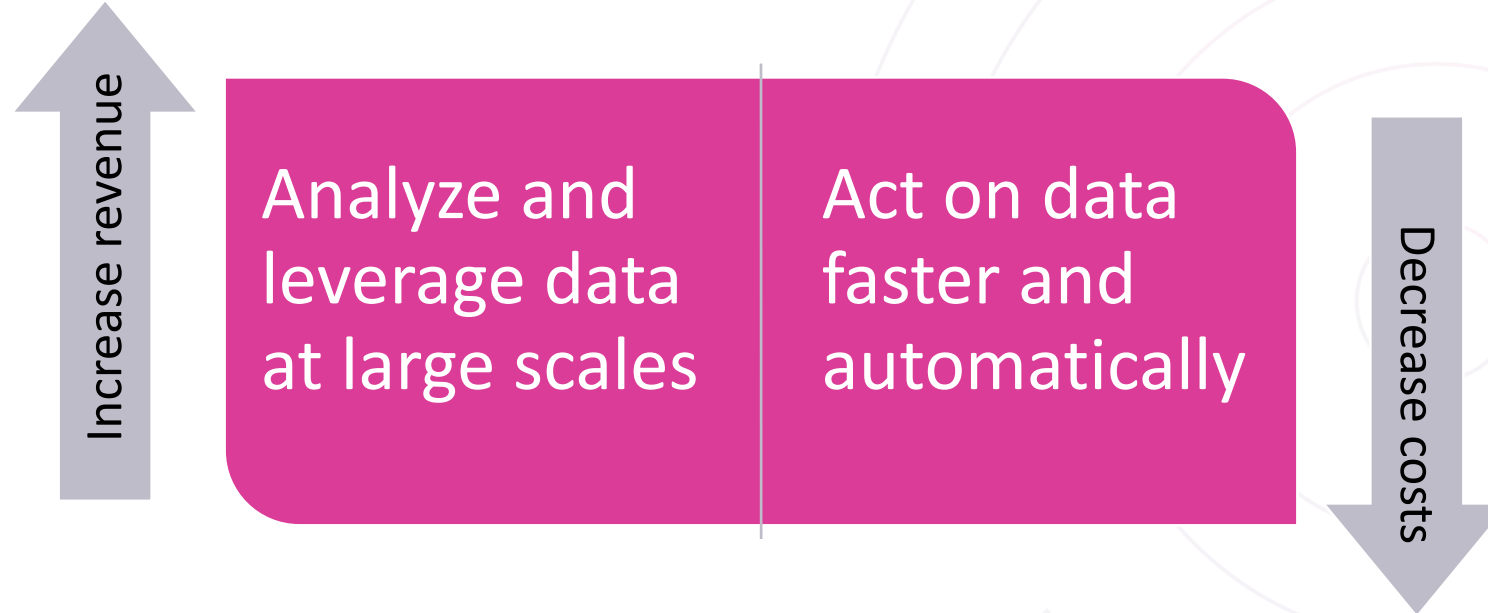


4 AI Use Cases



Why invest in AI?

$$\text{Profit} = \text{Revenue} - \text{Costs}$$





Use Case Life Cycle



- Find relevant use cases consistent with AI strategy



- Compare the expected value with implementation complexity



- Rank cases based on high value and low complexity



- Start with the most valuable cases first



Identify use cases

- Talk to the **right people**
 - Bring together domain experts, business stakeholders and AI experts
 - Ensure that initiatives address broad organizational priorities
 - Increase adoption chances by involving end users in the application design
- **Brainstorm** sessions to keep communication lines open
 - Defer judgement and encourage wild ideas
 - Build on ideas but stay on target
 - Go for quantity, more is better at this stage
- Not AI-ready?
 - Bring in external expertise



Questions to ask

- Strategy
 - What **goals** are driving the company right now?
 - Which **challenges** keep you up at night?
- Processes
 - What is driving current **bottlenecks** or preventing progress?
 - Which things are done **over and over** again?
 - Where would you benefit from knowing the **future**?
- Data
 - Where do you have a lot of **relevant** data?
 - Where do you already use some data to drive decision-making?



Churn modeling

- Why **important**?
 - Losing clients affects company revenue numbers and profits
- How does it **benefit business**?
 - Customer retention increases revenue and decreases costs
 - Understanding churn behavior leads to more effective retention strategies
- What **data** is needed?
 - Customer behavior, transactions, demographics, product usage/patterns, etc.



Demand forecasting

- Why **important**?
 - Used for strategic business plans (e.g., budgeting, financial planning, sales and marketing plans, capacity planning, risk assessment and mitigation plans)
- How does it **benefit business**?
 - Improved inventory availability can increase revenue
 - Reducing storage waste can decrease costs
- What **data** is needed?
 - Sales data, product demand, market conditions, ecommerce, etc.

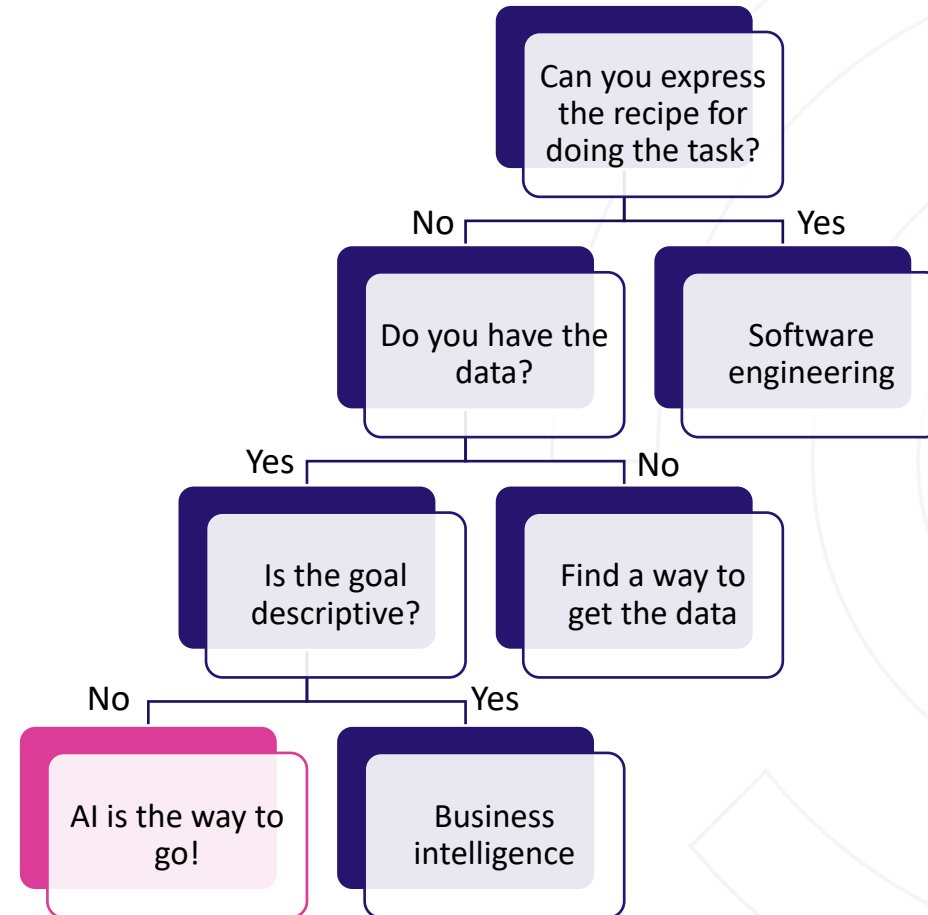


Targeted advertising

- Why **important**?
 - Cost-effective approach as it minimizes wasted advertising
- How does it **benefit business**?
 - Sales go up and customer satisfaction increases (less annoyed by random ads)
 - Targeting successfully leads to increased revenue and decreased costs
- What **data** is needed?
 - Purchase history & client personality, attitude, opinions, lifestyle and interests



Is AI the answer to your problem?





AI4Business

