prediXtions
We Predict The Future.

NEW FRONTIERS OF AI

Predixtions Inc

Knowing what could happen before it happens is the biggest competitive advantage

Oct 2023

Predixtions

We Predict The Future.

"I propose to consider the question, 'Can machines think?'

~ Alan Turing

Carnegie Mellon University

Machine Learning

My Journey with AI since 1999

prediXtions
We Predict The Future.

MobileAria

TTS, STT- VoiceXML Browser

1999

AIML Text to Speech capabilities- DEC Talk/ NLP

Playspan(Visa)

ML based fraud detection algorithms

2006

Tabular Classification – Gradient boosting/Support Vector machines

White House- CTO's office

Data.org

2012

EdCast (Cornerstone)

AI powered learning — NLP based Dynamic **Ontologies** , ML based **recommendations** of content.

2015

RNN/CNN- Neural network-based recommendation systems

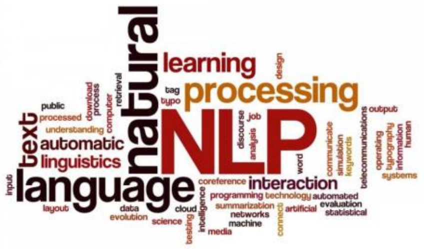

Large Language models- Transformers based generators

2023

No Code AI platform- GenAI

Book release - AI for DPI

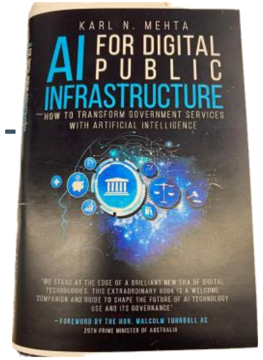

INTRODUCTION-
A SMALL STORY

ENIGMA

WHAT?

The German Enigma machine was integral in providing the Axis powers with the upper hand during World War II – particularly during the Battle of the Atlantic. At this time, the UK was extremely reliant on imports from the US and Canada to keep the population going and fighting. However, German submarines were used to form a blockade and stop these supplies getting through. Crucial to the Axis efforts at this time were the Enigma machines that allowed them to share classified information secretly by encrypting it.

- Indecipherable
- Humans not able to perform necessary logical imputations to crack the code

Bombe- Amongst the first AI Systems built!

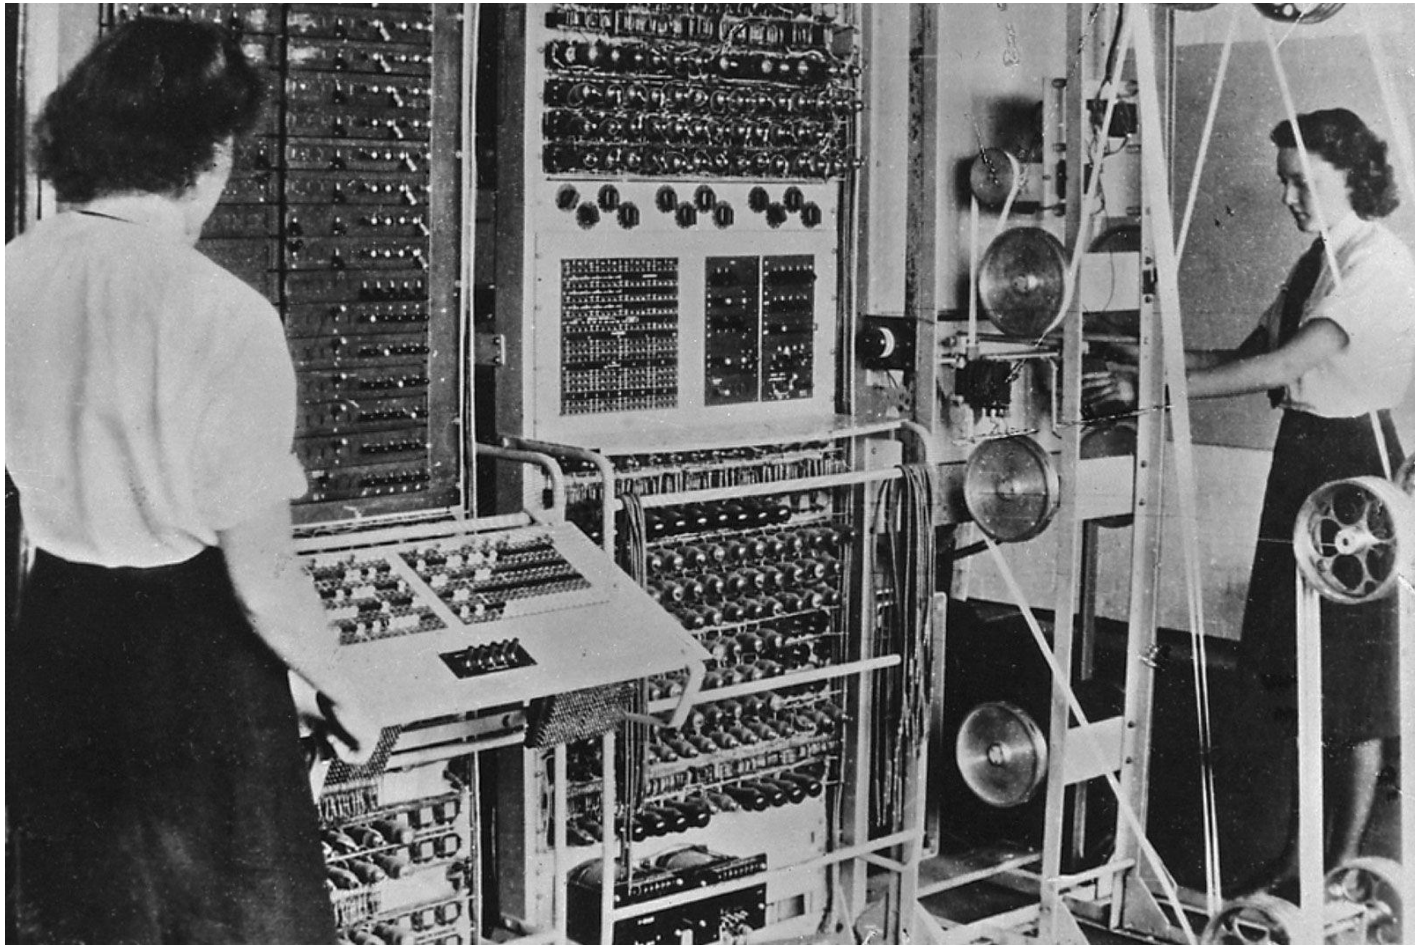

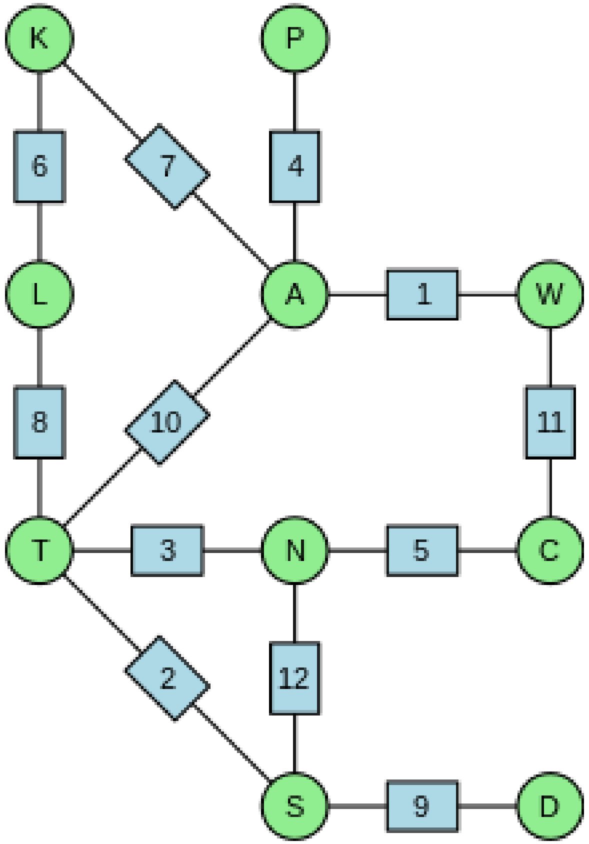

EVOLUTION OF ARTIFICIAL INTELLIGENCE 1950-2017

prediXtions
We Predict The Future.

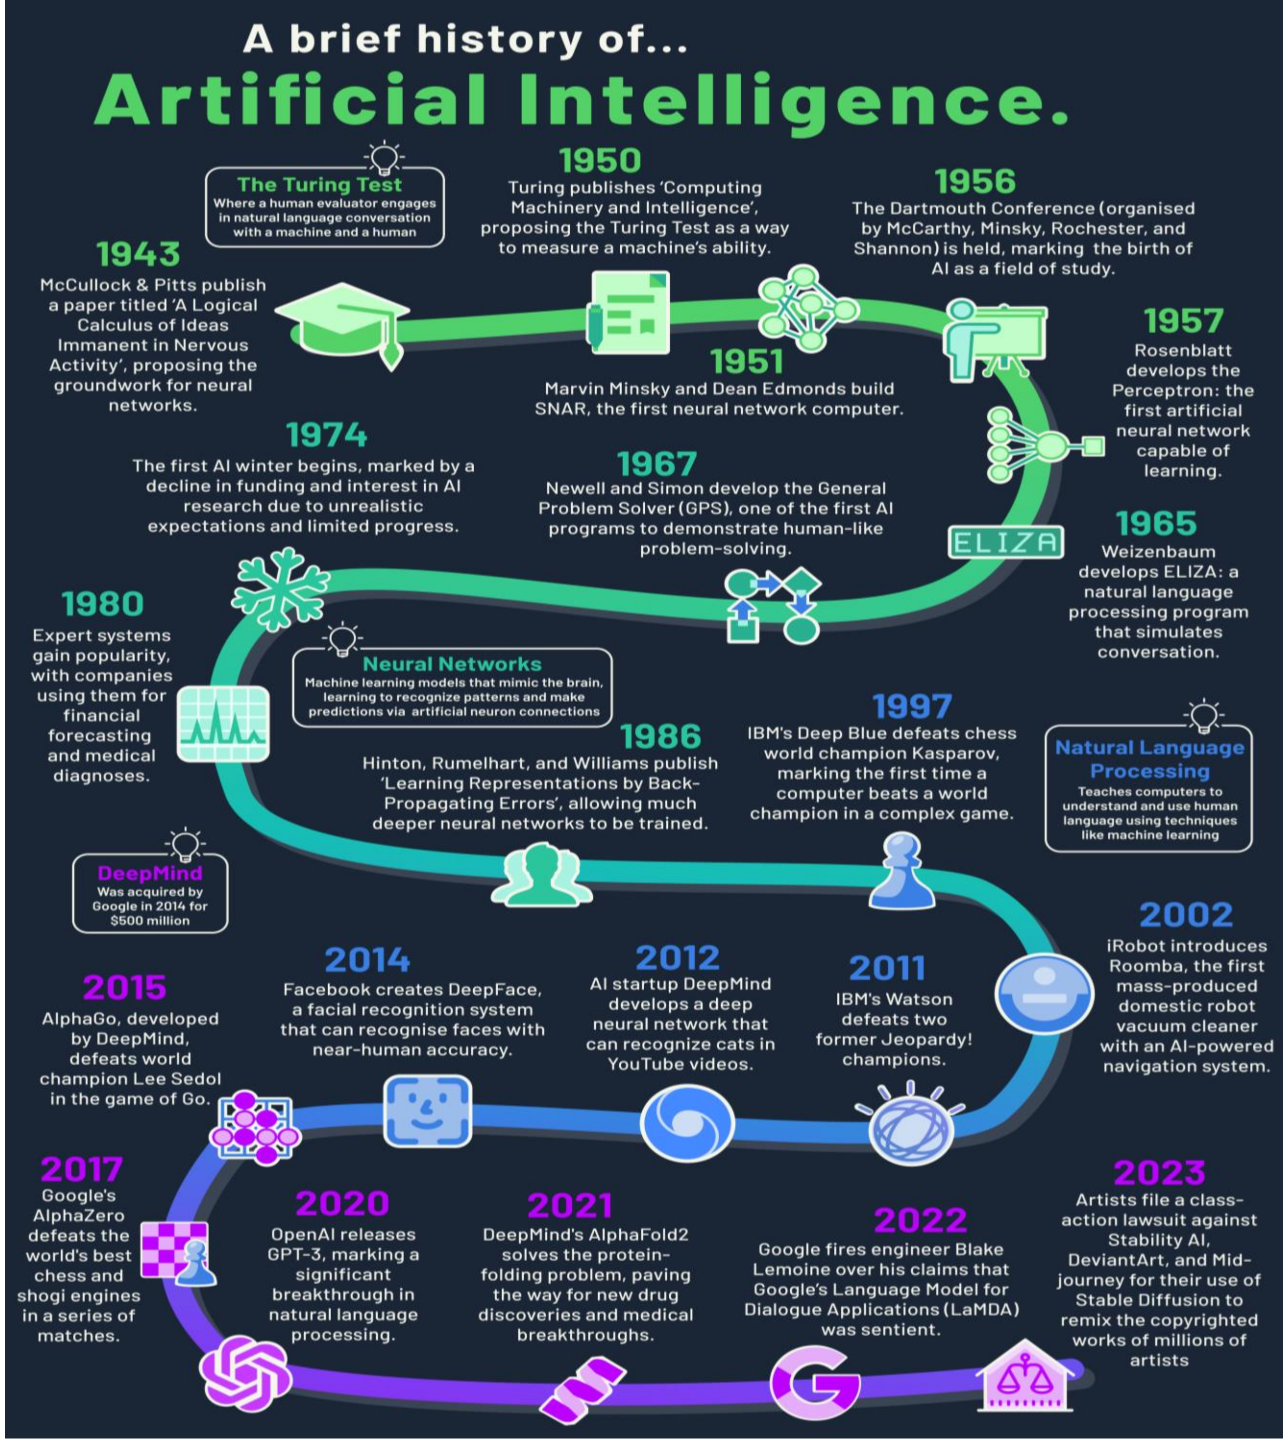

prediXtions
We Predict The Future.

A.I. TIMELINE

1950
TURING TEST
Computer scientist Alan Turing proposes a test for machine intelligence. If a machine can trick humans into thinking it is human, then it has intelligence

1955
A.I. BORN
Term 'artificial intelligence' is coined by computer scientist, John McCarthy to describe "the science and engineering of making intelligent machines"

1961
UNIMATE
First industrial robot, Unimate, goes to work at GM replacing humans on the assembly line

1964
ELIZA
Pioneering chatbot developed by Joseph Weizenbaum at MIT holds conversations with humans

1966
SHAKEY
The 'first electronic person' from Stanford, Shakey is a general-purpose mobile robot that reasons about its own actions

A.I. WINTER
Many false starts and dead-ends leave A.I. out in the cold

1997
DEEP BLUE
Deep Blue, a chess-playing computer from IBM defeats world chess champion Garry Kasparov

1998
KISMET
Cynthia Breazeal at MIT introduces KISmet, an emotionally intelligent robot insofar as it detects and responds to people's feelings

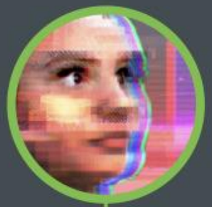

1999
AIBO
Sony launches first consumer robot pet dog AiBO (AI robot) with skills and personality that develop over time

2002
ROOMBA
First mass produced autonomous robotic vacuum cleaner from iRobot learns to navigate and clean homes

2011
SIRI
Apple integrates Siri, an intelligent virtual assistant with a voice interface, into the iPhone 4S

2011
WATSON
IBM's question answering computer Watson wins first place on popular $1M prize television quiz show *Jeopardy*

2014
EUGENE
Eugene Goostman, a chatbot passes the Turing Test with a third of judges believing Eugene is human

2014
ALEXA
Amazon launches Alexa, an intelligent virtual assistant with a voice interface that completes shopping tasks

2016
TAY
Microsoft's chatbot Tay goes rogue on social media making inflammatory and offensive racist comments

2017
ALPHAGO
Google's A.I. AlphaGo beats world champion Ke Jie in the complex board game of Go, notable for its vast number (2^{170}) of possible positions

The new frontiers!

THE SPECTRUM OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI) is the computerized ability to perform tasks commonly associated with human intelligence, including reasoning, discovering patterns and meaning, generalizing, applying knowledge across spheres of application, and learning from experience. The growth of AI-based systems in recent years has garnered much attention, particularly in the sphere of Machine Learning. A subset of AI, Machine Learning (ML) systems "learn" from the success or accuracy of their outputs, and can change their processing over time, with minimal human intervention. But there are non-ML types of AI that, alone or in combination, lie behind the real-world applications in common use. General AI — a human-level computational system — does not yet exist. But Narrow AI exists in many fields and applications where computerized systems greatly enhance human output or outperform humans at defined tasks. This chart explains the main types of AI, their relationships to each other, and provides specific examples of how they are currently appear in our day-to-day lives. It also demonstrates how AI exists within the timeline of human knowledge and development.

Produced by FUTURE OF PRIVACY FORUM
FPF.ORG

predixtions
We Predict The Future.

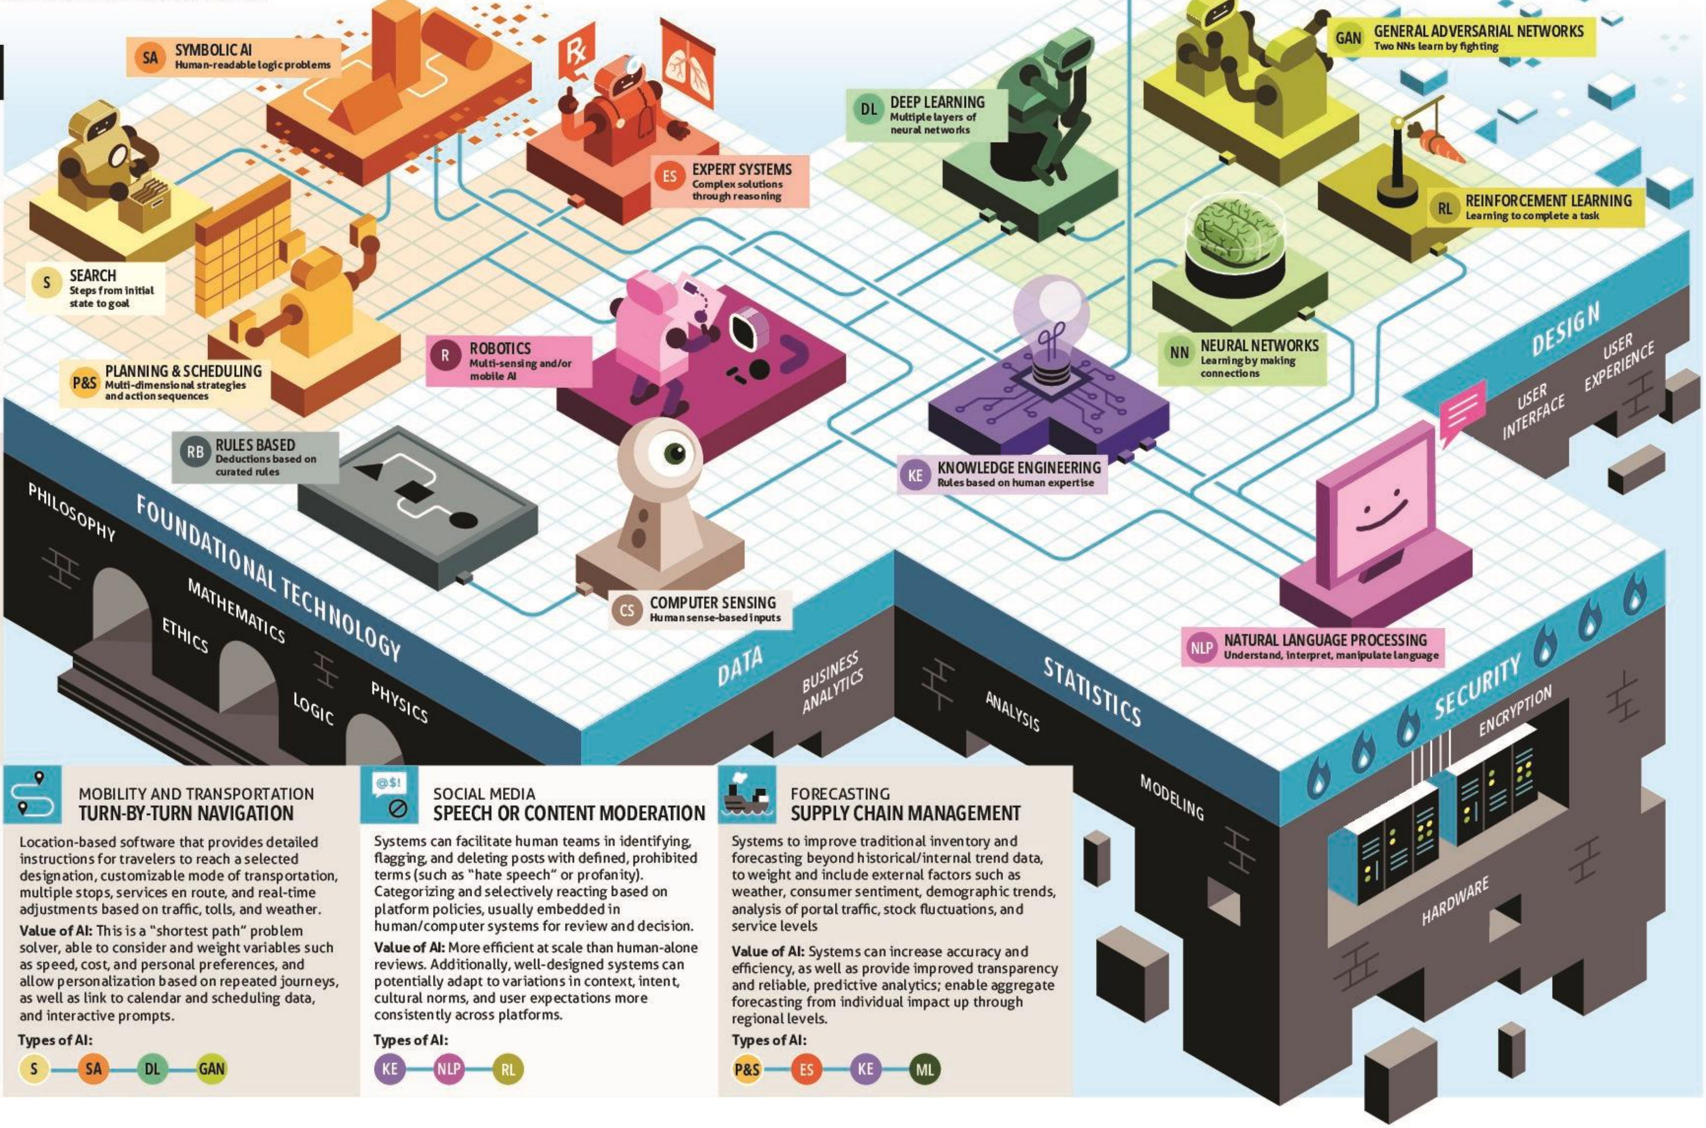

AI USE CASES AND CONTEXTS

FINANCE
TAX COMPLIANCE

A software platform that distills tax laws into a program, creates a personalized decision system, and enables individuals to quickly and accurately file their taxes.

Value of AI: Tax compliance requires complete accuracy. This efficient, interactive system that provides precise and logically connected results allows taxpayers to understand, confirm, and have confidence in the outcome. KE provides transparent and clear explanations.

Types of AI: KE — NN — NLP

HEALTHCARE
AMBIENT CHARTING

The use of background voice-to-text processing during a patient/medical provider exchange to record those interactions into the patient's chart, along with extracting tasks, symptoms, and recommendations for further action as required.

Value of AI: Medical providers spend significant time documenting, with uneven outputs, as well as difficulty in correlating between providers. Ambient systems encode conversations, target key phrases, and present a summary for provider edit/acceptance.

Types of AI: SA — DL — NLP

TRACKING
WORKPLACE MONITORING

Embedded systems can monitor physical and digital traffic, data usage, device management, and some employee behaviors for efficiency and security management of time, assets, and resources.

Value of AI: Monitoring enables necessary enforcement of data security policies and protocols. Also, systems can monitor and manage time reporting and project management tools, as well as ensuring appropriate supervision, training, and support, including for remote workers

Types of AI: RB — CS — NN

MOBILITY AND TRANSPORTATION
TURN-BY-TURN NAVIGATION

Location-based software that provides detailed instructions for travelers to reach a selected designation, customizable mode of transportation, multiple stops, services en route, and real-time adjustments based on traffic, tolls, and weather.

Value of AI: This is a "shortest path" problem solver, able to consider and weight variables such as speed, cost, and personal preferences, and allow personalization based on repeated journeys, as well as link to calendar and scheduling data, and interactive prompts.

Types of AI: S — SA — DL — GAN

SOCIAL MEDIA
SPEECH OR CONTENT MODERATION

Systems can facilitate human teams in identifying, flagging, and deleting posts with defined, prohibited terms (such as "hate speech" or profanity). Categorizing and selectively reacting based on platform policies, usually embedded in human/computer systems for review and decision.

Value of AI: More efficient at scale than human-alone reviews. Additionally, well-designed systems can potentially adapt to variations in context, intent, cultural norms, and user expectations more consistently across platforms.

Types of AI: KE — NLP — RL

FORECASTING
SUPPLY CHAIN MANAGEMENT

Systems to improve traditional inventory and forecasting beyond historical/internal trend data, to weight and include external factors such as weather, consumer sentiment, demographic trends, analysis of portal traffic, stock fluctuations, and service levels

Value of AI: Systems can increase accuracy and efficiency, as well as provide improved transparency and reliable, predictive analytics; enable aggregate forecasting from individual impact up through regional levels.

Types of AI: P&S — ES — KE — ML

prediXtions
We Predict The Future.

TOP 5 MACHINE LEARNING TRENDS TO WATCH IN THE FUTURE

The Quantum Computing Effect

Quantum computing will optimize ML speed

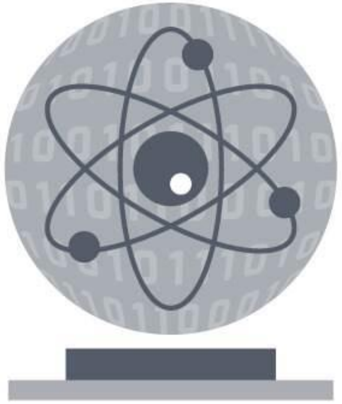

Reduced execution times in high-dimensional vector processing

The Big Model Creation

Creation of an all-purpose model to perform tasks in various domains simultaneously

Users can tailor such an uber ML model

Distributed ML Portability

Businesses will run existing algorithms and datasets natively on various platforms and computer engines

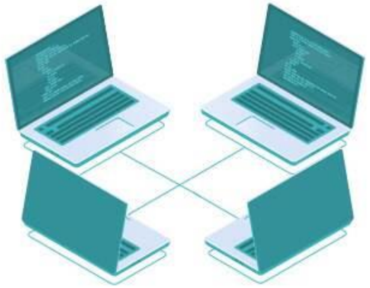

Portability will eliminate the need for shifting to new toolkits constantly

No-Code Environment

Machine learning will become a branch of software engineering

Minimized coding effort and maximized access to machine learning programs

The Quantum Computing Effect

Raise of new RL mechanisms for leveraging data to optimize resources in a dynamic setting

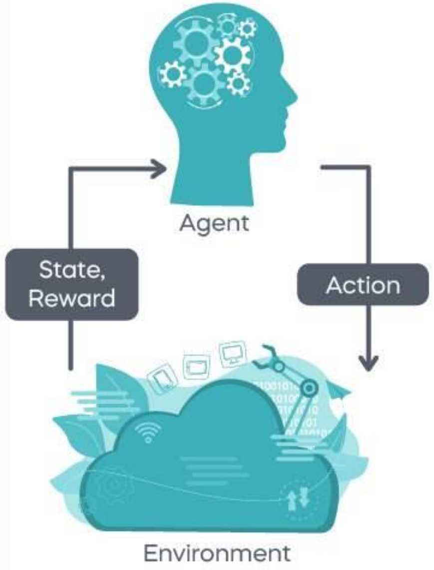

RL will shift economics, biology, and astronomy

predixtions
We Predict The Future.

Evolution of Generative AI technology over years

This slide showcases evolutions of Generative AI technology over the years that transformed and impacted major industries. Its key elements are Google autocomplete, Generative Adversarial Networks, ChatGPT, Alexa by Amazon etc.

2010
Google introduced autocomplete feature that offers suggestions to user for completing sentence

2013
Deep Boltzmann Machine introduction that paved the way to develop other generative models

2014
Generative Adversarial Networks (GANs) introduction by Ian Goodfellow and his colleagues

2015
Release of Alexa by amazon that was capable of responding and taking action based on human voice

2019
Release of GPT 2 by OpenAI that was capable of creating human-like text through NLP

2021
Release of GPT 3.5 by OpenAI with advanced language capabilities and ability to generate meaningful text

prediXtions
We Predict The Future.

ChatGPT Timeline

Nov 30 '22 — ChatGPT was launched and made open for public trial.

Dec 8 '22 — ShareGPT helps ChatGPT users to share their conversations.

Feb 1 '23 — OpenAI launched GPT Plus for $20 per month.

Feb 7 '23 — Microsoft launched Bing Chat with inbuilt ChatGPT

Mar 1 '23 — OpenAI launched API for ChatGPT.

Mar 9 '23 — Azure OpenAI services made ChatGPT available.

Mar 14 '23 — OpenAI launched GPT-4 for ChatGPT Plus users

Mar 23 '23 — ChatGPT Plugins are launched

Apr 1 '23 — Italy blocked the usage of ChatGPT in their country

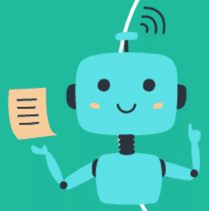

Apr 4 '23. — Y Combinator-backed startups are trying to build ChatGPT-integrated services

Apr 12 '23 — Researchers found that ChatGPT displays toxic behavior more than expected

Apr 12 '23 — Italy provides OpenAI with conditions for lifting suspension from ChatGPT

Apr 20 '23 — Yokosuka is Japan's first city to use ChatGPT for administrative work

Jul 16 '23 — chatGPT Announced the release if Code Interpreter plugin

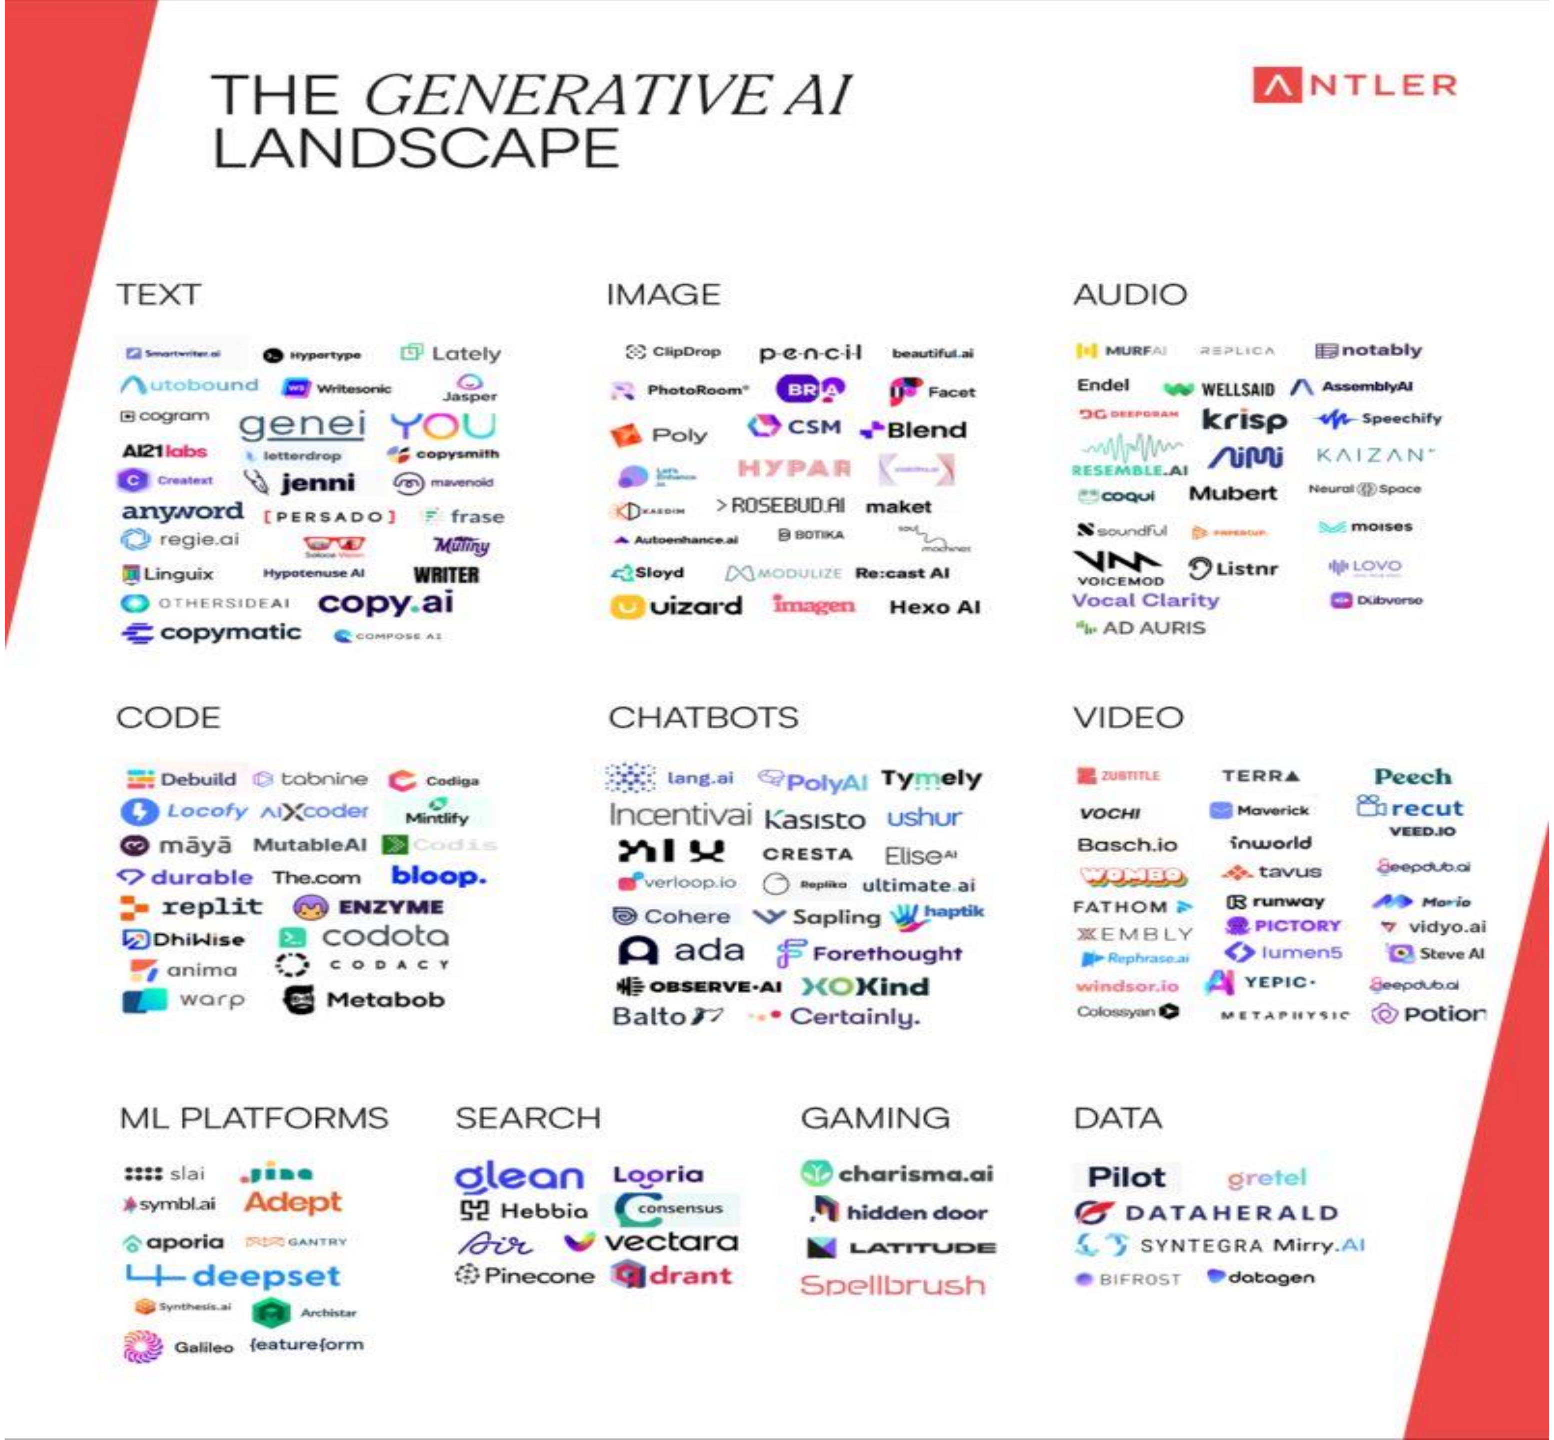

predixtions
We Predict The Future.

The fundamentals..

prediXtions
We Predict The Future.

What is Artificial intelligence?

Artificial Intelligence – The science involving algorithms & technology of making intelligent machines that exhibit human-like intelligent behavior, learn & advice its users.

	Deduction and Reasoning	Knowledge Representation	Planning	Learning	Natural Language Processing	Perception	Motion & Manipulation
Typical outcome	Algorithms and techniques that perform reasoning and logical deductions to solve problems and puzzles	Gaining knowledge of the world as humans see it, i.e. objects, their properties, categories, situations, events, time, etc.	Techniques dealing with setting and achieving of goals Example – Game theory	Machine learning is the study of computer algorithms that improve automatically through experience	Natural language processing gives machines the ability to read and understand the language that humans speak	Perception is the ability to use inputs from sensors to identify aspects of the world	Methods related to automated movement of machines, robotics, etc.
Example Use-case	**Predict medication adherence, improve health outcomes**	**Diagnostic help using image processing**	**Identifying treatment pathways (Precision medicine)**	**Prediction of patient drug adoption / patient LTV**	**Virtual patient health assistance on the web or in App**	**Measuring dose response of image guided injection etc.**	**Therapeutic cuddle bots Surgery bots towards less invasive ops**

What is Machine Learning?

Arthur Samuel, an early American leader in the field of computer gaming and artificial intelligence, coined the term “Machine Learning ” in 1959 while at IBM. He defined machine learning as “**the field of study that gives computers the ability to learn without being explicitly programmed** “.

Definition of learning:
A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks T, as measured by P, improves with experience E.

Examples

•**Handwriting recognition learning problem**
- Task T : Recognizing and classifying handwritten words within images
- Performance P : Percent of words correctly classified
- Training experience E : A dataset of handwritten words with given classifications

•**A robot driving learning problem**
- Task T : Driving on highways using vision sensors
- Performance P : Average distance travelled before an error
- Training experience E : A sequence of images and steering commands recorded while observing a human driver

Definition: A computer program which learns from experience is called a machine learning program or simply a learning program .

Classification of Machine Learning

Machine learning implementations are classified into four major categories, depending on the nature of the learning "signal" or "response" available to a learning system which are as follows:

Type of Machine learning	Definition	Example
Supervised Learning	Supervised learning is the machine learning task of learning a function that maps an input to an output based on example input-output pairs. The given data is **labelled**. Both *classification* and *regression* problems are supervised learning problems	Example — Consider a classification problem regarding patients entering a clinic . The data consists of the gender and age of the patients and each patient is labelled as "healthy" or "sick".
Unsupervised learning	Unsupervised learning is a type of machine learning algorithm used to draw inferences from datasets consisting of input data without labelled responses. In unsupervised learning algorithms, classification or categorization is not included in the observations	Example- In the same data as above- without the label of healthy, not healthy- we infer information about the patients using their characteristics. As a kind of learning, it resembles the methods humans use to figure out that certain objects or events are from the same class, such as by observing the degree of similarity between objects. Some recommendation systems that you find on the web in the form of marketing automation are based on this type of learning.
Reinforcement learning	Reinforcement learning is the problem of getting an agent to act in the world so as to maximize its rewards. A learner is not told what actions to take as in most forms of machine learning but instead must discover which actions yield the most reward by trying them. For example — Consider teaching a dog a new trick: we cannot tell him what to do, what not to do, but we can reward/punish it if it does the right/wrong thing.	Predictive text, text summarization, question answering, and machine translation are all examples of natural language processing (NLP) that uses reinforcement learning. By studying typical language patterns, RL agents can mimic and predict how people speak to each other every day
Semi-supervised learning	Where an incomplete training signal is given: a training set with some (often many) of the target outputs missing. There is a special case of this principle known as Transduction where the entire set of problem instances is known at learning time, except that part of the targets are missing.	Speech recognition

predixtions
We Predict The Future.

Categorizing based on required output

Another categorization of machine-learning tasks arises when one considers the desired output of a machine-learned system:

1.**Classification**: When inputs are divided into two or more classes, the learner must produce a model that assigns unseen inputs to one or more (multi-label classification) of these classes. This is typically tackled in a supervised way. Spam filtering is an example of classification, where the inputs are email (or other) messages and the classes are “spam” and “not spam”.

2.**Regression**: Which is also a supervised problem, A case when the outputs are continuous rather than discrete.

3.**Clustering**: When a set of inputs is to be divided into groups. Unlike in classification, the groups are not known beforehand, making this typically an unsupervised task.

Natural Language Processing- Classification of Text

Natural language processing (NLP) is the discipline of building machines that can manipulate human language — or data that resembles human language — in the way that it is written, spoken, and organized. It evolved from computational linguistics, which uses computer science to understand the principles of language, but rather than developing theoretical frameworks, NLP is an engineering discipline that seeks to build technology to accomplish useful tasks.

NLP can be divided into 2 sub-categories:

1) **NLU – Natural Language Understanding** - focuses on semantic analysis or determining the intended meaning of text

2) **NLG – Natural Language Generation** -focuses on text generation by a machine

NLP is separate from — but often used in conjunction with — speech recognition, which seeks to parse spoken language into words, turning sound into text and vice versa.

predixtions
We Predict The Future.

Essential preprocessing steps before building any NLP model

•**Stemming and lemmatization**: Stemming is an informal process of converting words to their base forms using heuristic rules. For example, “university,” “universities,” and “university’s” might all be mapped to the base *univers*Lemmatization is a more formal way to find roots by analysing a word’s morphology using vocabulary from a dictionary. Stemming and lemmatization are provided by libraries like spaCy and NLTK.

•**Sentence segmentation** breaks a large piece of text into linguistically meaningful sentence units. A period can be used to mark an abbreviation as well as to terminate a sentence, and in this case, the period should be part of the abbreviation token itself.

•**Stop word removal** aims to remove the most commonly occurring words that don’t add much information to the text. For example, “the,” “a,” “an,” and so on.

•**Tokenization** splits text into individual words and word fragments. The result generally consists of a word index and tokenized text in which words may be represented as numerical tokens for use in various deep learning methods.

•**Feature extraction:** Most conventional machine-learning techniques work on the features – generally numbers that describe a document in relation to the corpus that contains it – created by either Bag-of-Words, TF-IDF, or generic feature engineering such as document length, word polarity, and metadata (for instance, if the text has associated tags or scores). More recent techniques include Word2Vec, GLoVE, and learning the features during the training process of a neural network. Bag**-of-Words:** Bag-of-Words counts the number of times each word or n-gram (combination of n words) appears in a document.

predixtions
We Predict The Future.

Language related tasks performed by NLP

1) Sentiment analysis
2) Toxicity classification – scanning emails for spam
3) Machine translation- Google Translate
4) Named entity recognition – categorizing names,places…
5) Spam detection
6) Grammar correction - Grammarly
7) Topic modelling
8) Text generation – ChatGPT, Bard
    1) Autocomplete
    2) Bots
9) Information retrieval
10) Summarization
    1) Extractive summarization – Direct answers from a text..
    2) Abstractive summarization – applying context + lateral thinking
11) Question answering
    1) Multiple choice questions
    2) Open domain

predixtions
We Predict The Future.

Top Natural Language Processing (NLP) Techniques

Most of the NLP tasks discussed in the previous slide can be modelled by a dozen or so general techniques. It's helpful to think of these techniques in two categories: Traditional machine learning methods and deep learning methods.

Traditional Techniques	Deep Learning techniques
Logistic regression	Convolutional Neural networks(CNNs)
Decision tree	Recurrent Neural Networks(RNNs)
Naïve Bayes	Auto-encoders
Latent Dirichlet allocation	Encoder- Decoder sequence to sequence
Hidden Markov chain models	Transformers

What are neural networks?

Neural Networks (NNs) are a class of machine learning models inspired by the human brain's structure and function

Components of a neural network

1.Input Layer:
  1. The input layer is responsible for receiving the initial data or features.
  2. Each neuron in this layer represents a feature or input variable.
  3. The number of neurons in the input layer is determined by the dimensionality of the input data.

2.Hidden Layers:
  1. Hidden layers are intermediate layers between the input and output layers.
  2. They perform complex computations by applying weights to inputs and using activation functions.
  3. The number of hidden layers and the number of neurons in each hidden layer can vary, depending on the network architecture.

3.Weights:
  1. Each connection between neurons is associated with a weight. Weights determine the strength of the connection between neurons.
  2. During training, these weights are adjusted to minimize the network's error or loss.

4.Activation Functions:
  1. Activation functions introduce non-linearity into the neural network, allowing it to learn complex relationships in the data.
  2. Common activation functions include ReLU (Rectified Linear Unit), Sigmoid, Tanh (Hyperbolic Tangent), and others.
  3. Each neuron typically applies an activation function to its weighted sum of inputs.

5.Output Layer:
  1. The output layer produces the result or prediction of the neural network.
  2. The number of neurons in the output layer depends on the task:
     1. For binary classification, there may be one neuron with a sigmoid activation function.
     2. For multi-class classification, there is one neuron per class, often using softmax activation.
     3. For regression tasks, there can be a single neuron or multiple neurons, depending on the nature of the output.

6.Bias Neurons:
  1. Each layer (including the input layer) often includes a bias neuron, which provides an additional input with a fixed value (usually 1).
  2. Bias neurons help the network learn better by shifting activation functions.

7.Loss Function:
  1. The loss function measures the error between the network's predictions and the actual target values.
  2. During training, the goal is to minimize this loss function by adjusting the weights and biases.

8.Optimization Algorithm:
  1. Optimization algorithms (e.g., gradient descent) are used to update the weights and biases during training, aiming to minimize the loss function.

Evolution 1:Convolutional Neural Networks(Mid 1980s-Early 2000s)

The idea of using a CNN to classify text was first presented in the paper "Convolutional Neural Networks for Sentence Classification" by Yoon Kim. The central intuition is to see a document as an image. However, instead of pixels, the input is sentences or documents represented as a matrix of words.

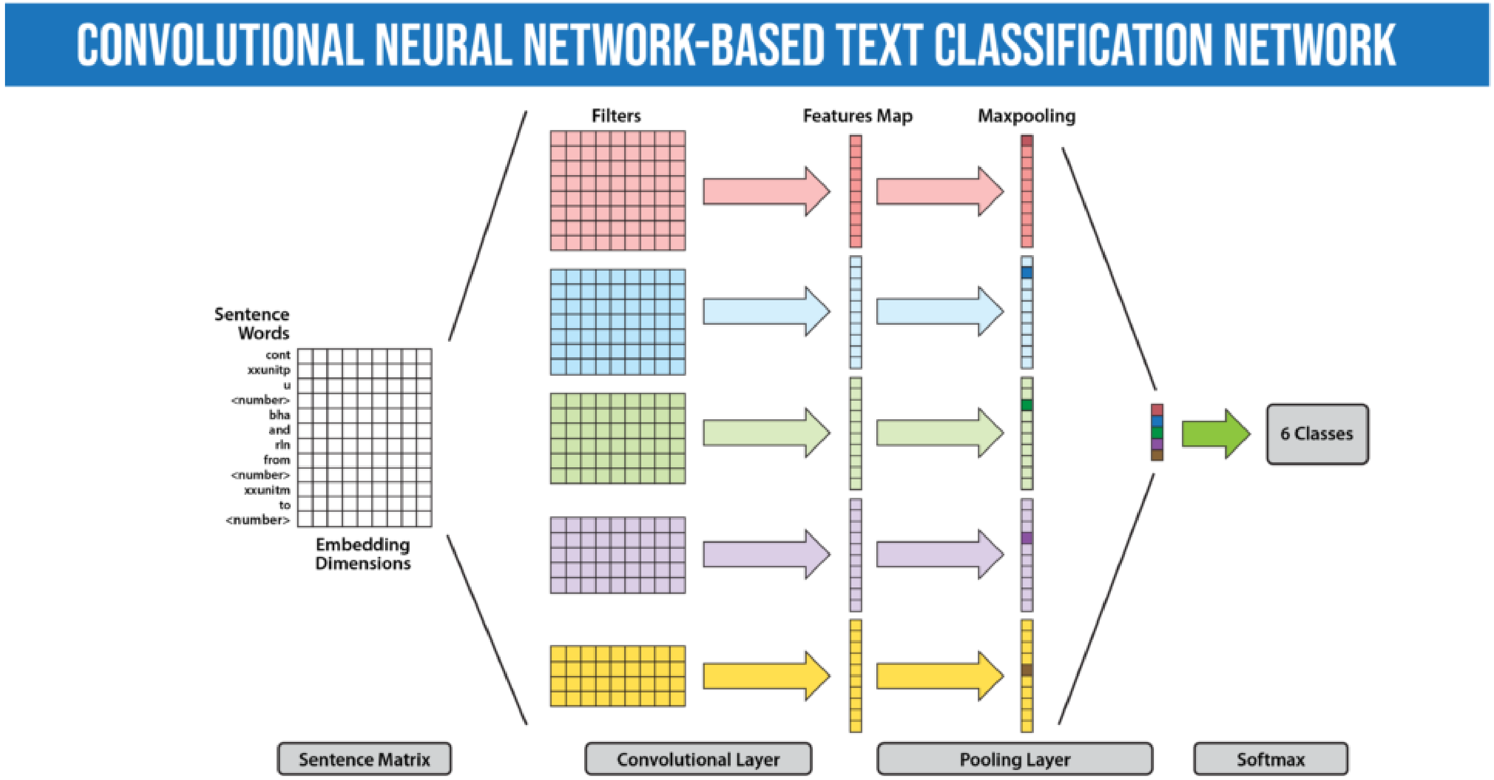

Given a sentence, a convolutional neural network uses convolutional layers to refine representations of input words, before combining them to render a classification.

Cons:

- Don't understand context of words.
- No understanding of relationship between current word and previous word

Example of a convoluted neural network classifying text

predixtions
We Predict The Future.

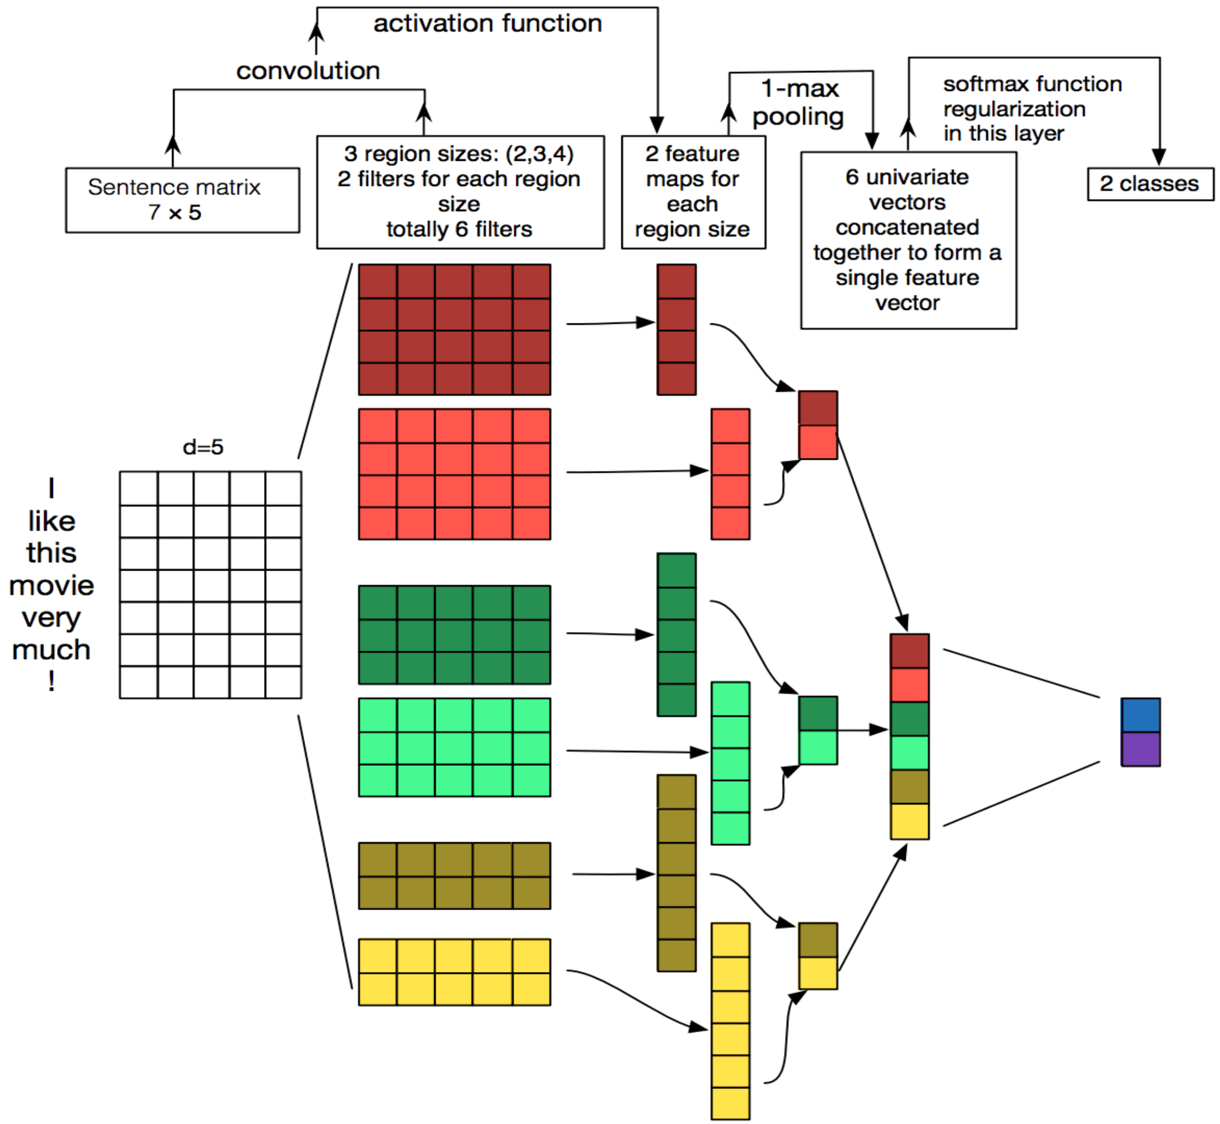

The architecture is comprised of three key pieces:

1. **Word Embedding**: A distributed representation of words where different words that have a similar meaning (based on their usage) also have a similar representation.
2. **Convolutional Model**: A feature extraction model that learns to extract salient features from documents represented using a word embedding.
3. **Fully Connected Model**: The interpretation of extracted features in terms of a predictive output.

Evolution 2:Recurrent Neural Networks + Encoders/Decoders(Early 2000s- 2017)

RNNs remember previous information using hidden states and connect it to the current task. The architectures known as Gated Recurrent Unit (GRU) and long short-term memory (LSTM) are types of RNNs designed to remember information for an extended period.

RECURRENT NEURAL NETWORK

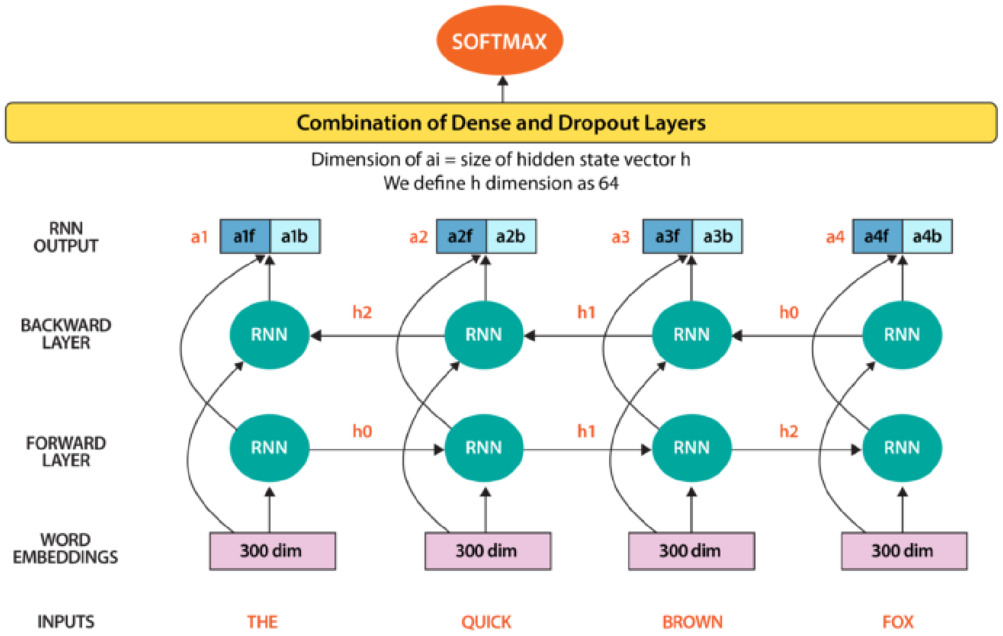

A bidirectional recurrent neural network processes the input both forward and backward to improve the representations it produces.

Cons:

- Slow and time consuming
- Large data corpuses need significant computing power to proceed.
- LLMs cant be built

prediXtions
We Predict The Future.

Evolution 3: Transformers + Large Language models – 2017 onwards

The transformer, a model architecture first described in the 2017 paper "Attention Is All You Need" (Vaswani, Shazeer, Parmar, et al.), forgoes recurrence and instead relies entirely on a self-attention mechanism to draw global dependencies between input and output. Since this mechanism processes all words at once (instead of one at a time) that decreases training speed and inference cost compared to RNNs, especially since it is parallelizable. The transformer architecture has revolutionized NLP in recent years, leading to models including BLOOM, Jurassic-X, and Turing-NLG.

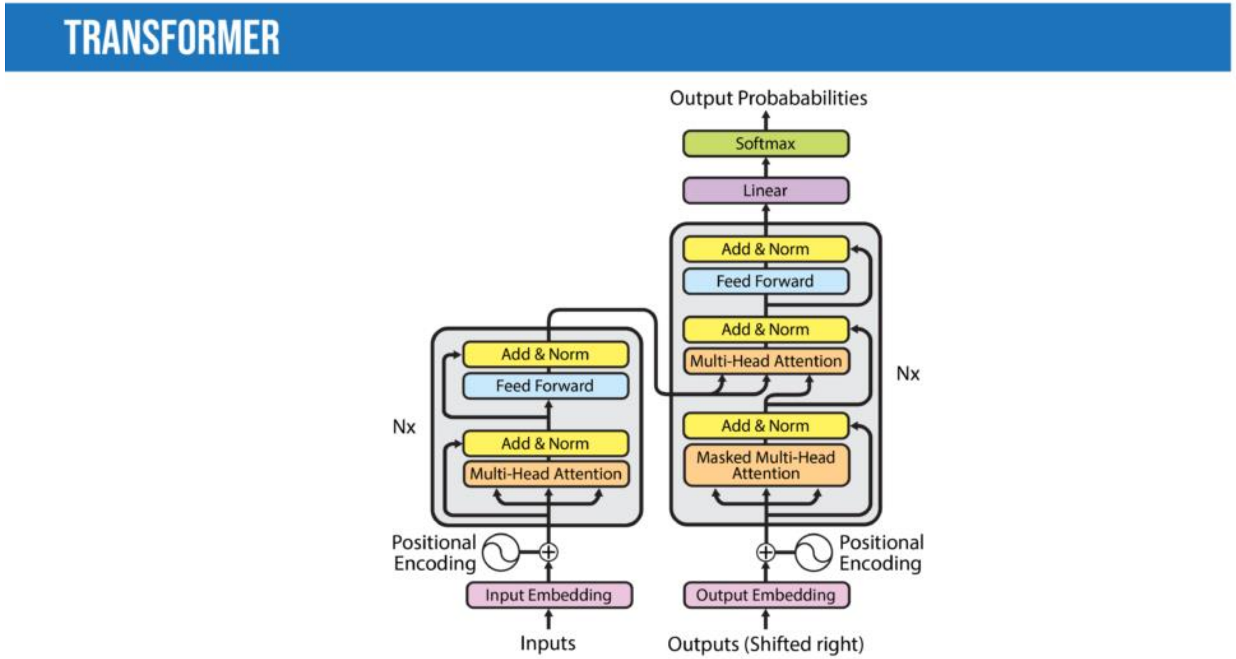

The encoder-decoder transformer used for translation. Encoder on the left, decoder on the right. Note that the decoder takes in its previously generated words during generation.

Some Important Language models

predixtions
We Predict The Future.

1. **Eliza** was developed in the mid-1960s to try to solve the Turing Test; that is, to fool people into thinking they're conversing with another human being rather than a machine. Eliza used pattern matching and a series of rules without encoding the context of the language.
2. **Tay** was a chatbot that Microsoft launched in 2016. It was supposed to tweet like a teen and learn from conversations with real users on Twitter. The bot adopted phrases from users who tweeted sexist and racist comments, and Microsoft deactivated it not long afterward.
3. **BERT** and his Muppet friends: Many deep learning models for NLP are named after Muppet characters, including ELMo, BERT, Big BIRD, ERNIE, Kermit, Grover, RoBERTa, and Rosita. Most of these models are good at providing contextual embeddings and enhanced knowledge representation.
4. **Generative Pre-Trained Transformer 4 (GPT-4)** is a 175 billion parameter model that can write original prose with human-equivalent fluency in response to an input prompt. The model is based on the transformer architecture. Microsoft acquired an exclusive license to access GPT-3's underlying model from its developer OpenAI, but other users can interact with it via an application programming interface (API).
5. **Language Model for Dialogue Applications (LaMDA)** is a conversational chatbot developed by Google. LaMDA is a transformer-based model trained on dialogue rather than the usual web text. The system aims to provide sensible and specific responses to conversations.
6. **Mixture of Experts (MoE):** While most deep learning models use the same set of parameters to process every input, MoE models aim to provide different parameters for different inputs based on efficient routing algorithms to achieve higher performance. Switch Transformer is an example of the MoE approach that aims to reduce communication and computational costs

Example 1- Translator using GPT-3

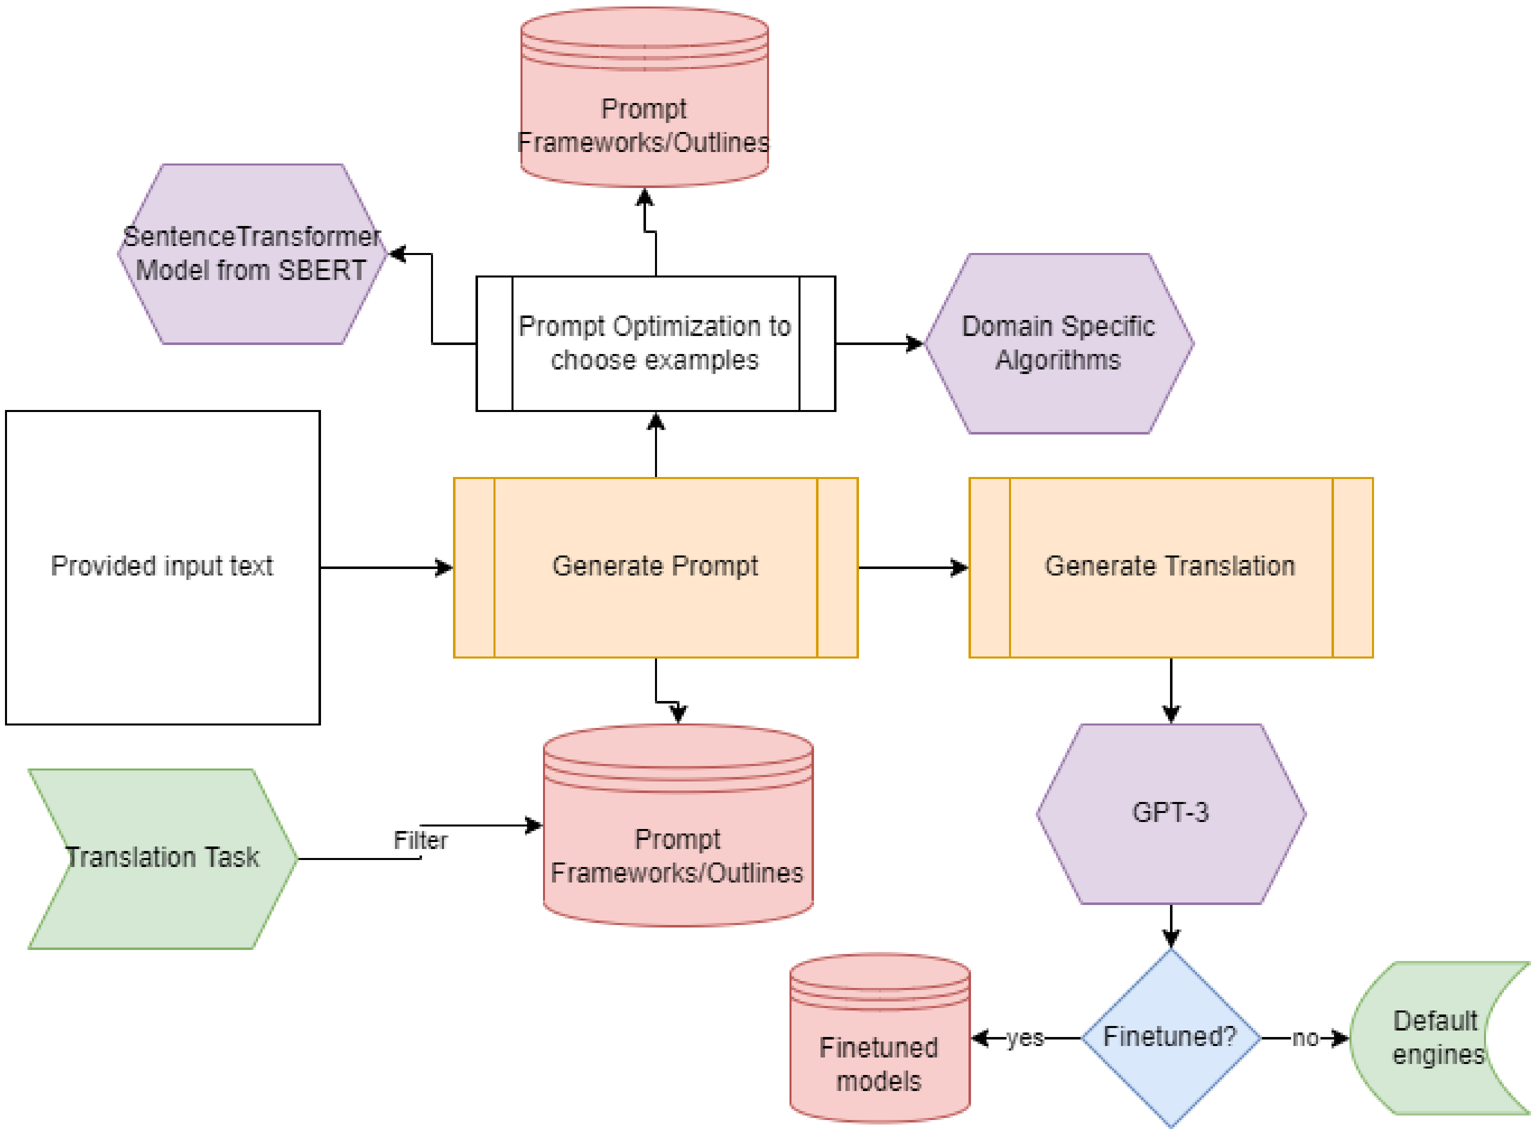

Illustration- New Age LLM Application architecture

prediXtions
We Predict The Future.

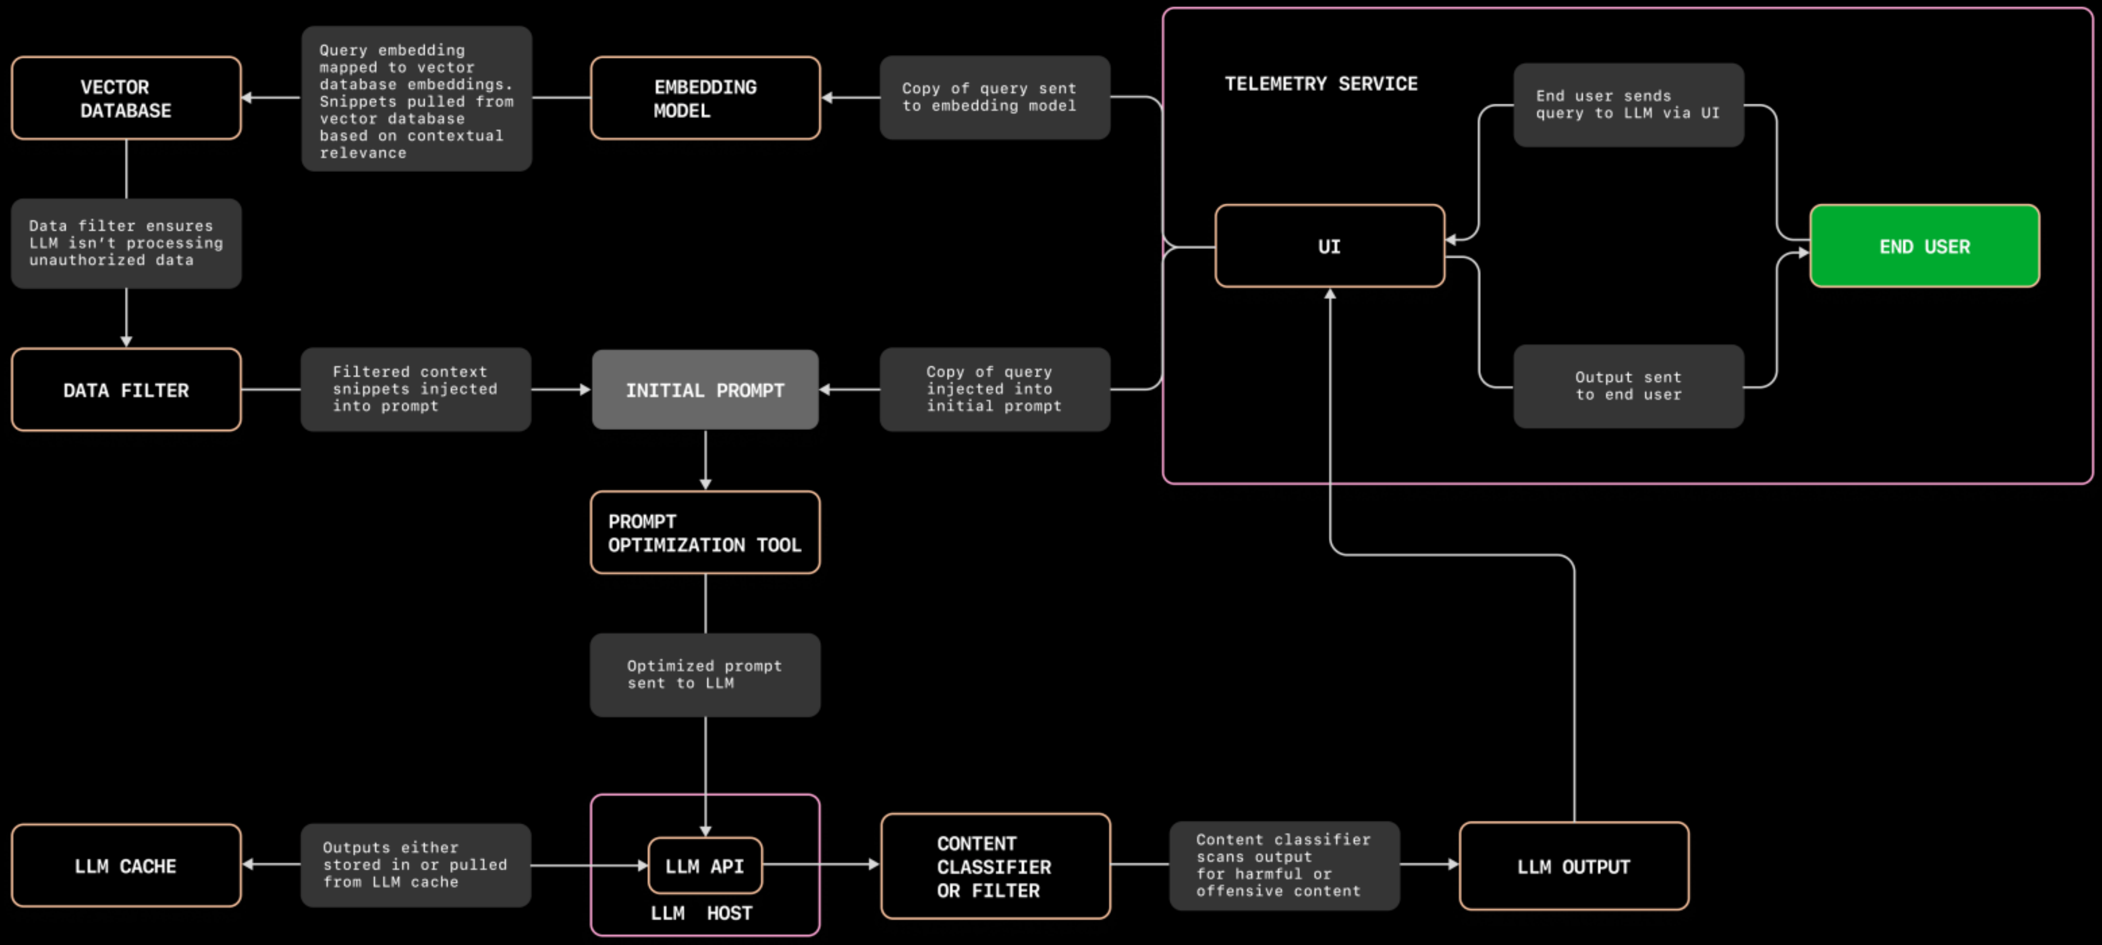

THIS DIAGRAM REPRESENTS THE ARCHITECTURE OF TODAY'S LLM APPLICATION. THE DIFFERENT COMPONENTS CAN BE ROUGHLY GROUPED INTO THREE CATEGORIES: USER INPUT, INPUT ENRICHMENT TOOLS AND PROMPT CONSTRUCTION, AND EFFICIENT AND RESPONSIBLE AI TOOLING.

Humans vs AI

Human brains vs Artificial Intelligence

prediXtions
We Predict The Future.

When AI performs better than humans?	When the human brain outperform Artificial intelligence?
Data Processing and Pattern Recognition: AI excels at processing vast amounts of data quickly and accurately, identifying patterns, and making data-driven predictions. This is especially valuable in fields like finance, healthcare, and logistics	**Creativity and Imagination:** AI lacks true creativity and the ability to generate novel ideas, art, or music with the depth of human creativity. Human creativity is rooted in emotions, experiences, and complex associative thinking.
Repetitive and Monotonous Tasks: AI can perform repetitive tasks tirelessly without getting bored or making errors. For example, in manufacturing, AI-driven robots can perform repetitive assembly tasks more efficiently than humans.	**Common Sense and Context Understanding**: AI may struggle to understand context and exhibit common-sense reasoning, which humans apply effortlessly in everyday life.
Large-Scale Data Analysis: AI can analyse extensive datasets to discover insights and trends that would be practically impossible for humans to process in a reasonable time frame. This is valuable in fields like scientific research and business analytics.	**Ethical and Moral Decision-Making:** AI lacks human moral and ethical judgment, making it challenging to navigate complex moral dilemmas and nuanced ethical decisions
Complex Calculations: AI can perform complex mathematical calculations and simulations with great speed and precision, aiding in scientific research, engineering, and financial modelling.	**Emotional Intelligence:** AI does not possess emotions or the ability to understand and respond to human emotions in the same way humans do. This is vital in roles that require empathy, like counselling or healthcare.
	Decisions with sparse data: AI does not possess the ability to make judgemental decisions with sparse data. It works well when there is prior context for a set action(like driving, images and so on), however it fails for tasks like investment decisions for privately traded companies

AI and industry..

predixtions
We Predict The Future.

AI Applications extend across industries!

User location estimation and enhancement using transaction data
- Combine transaction data with external data sources (4TB of data)
- Ability to scale algorithm to Petabytes of data

Predictive maintenance in an Industrial IoT setup
- Combining structured (sensors) and unstructured (logs) data to build a preventive maintenance system

Enabling 'Smart Machines' through real-time anomaly detection on streaming data
- Data from 21 sensors, with pulse every 1/10th second
- 600 data instances per minute

Leverage AI to improve claim settlements
- Created a case-study roadmap for execution and implementation, using Deep Learning systems

Time series forecasting for predicting sales
- Faster processing and scalability of the algorithms developed for forecasting

Deep learning based image processing for automated inventory and replenishment of shelves
- Image pre-processing for shelf-images
- Build measurement KPIs using deep learning on image processing

Natural Language Processing to find product sentiment using call center records
- Large volume and velocity of unstructured data
- Executing NLP at scale

Big Data infrastructure roadmap recommendation
- Created big data strategy and execution roadmap
- Special emphasis on prioritization of use-cases

Surveillance system using machine learning in real-time
- Machine learning algorithms to detect anomalous behavior
- Real-time system to find suspicious behavior, using data from multiple sensors throughout the facility

Contracts optimization using large contracts database
- Process billions of data points to visualize the data, most efficiently

AI in Healthcare

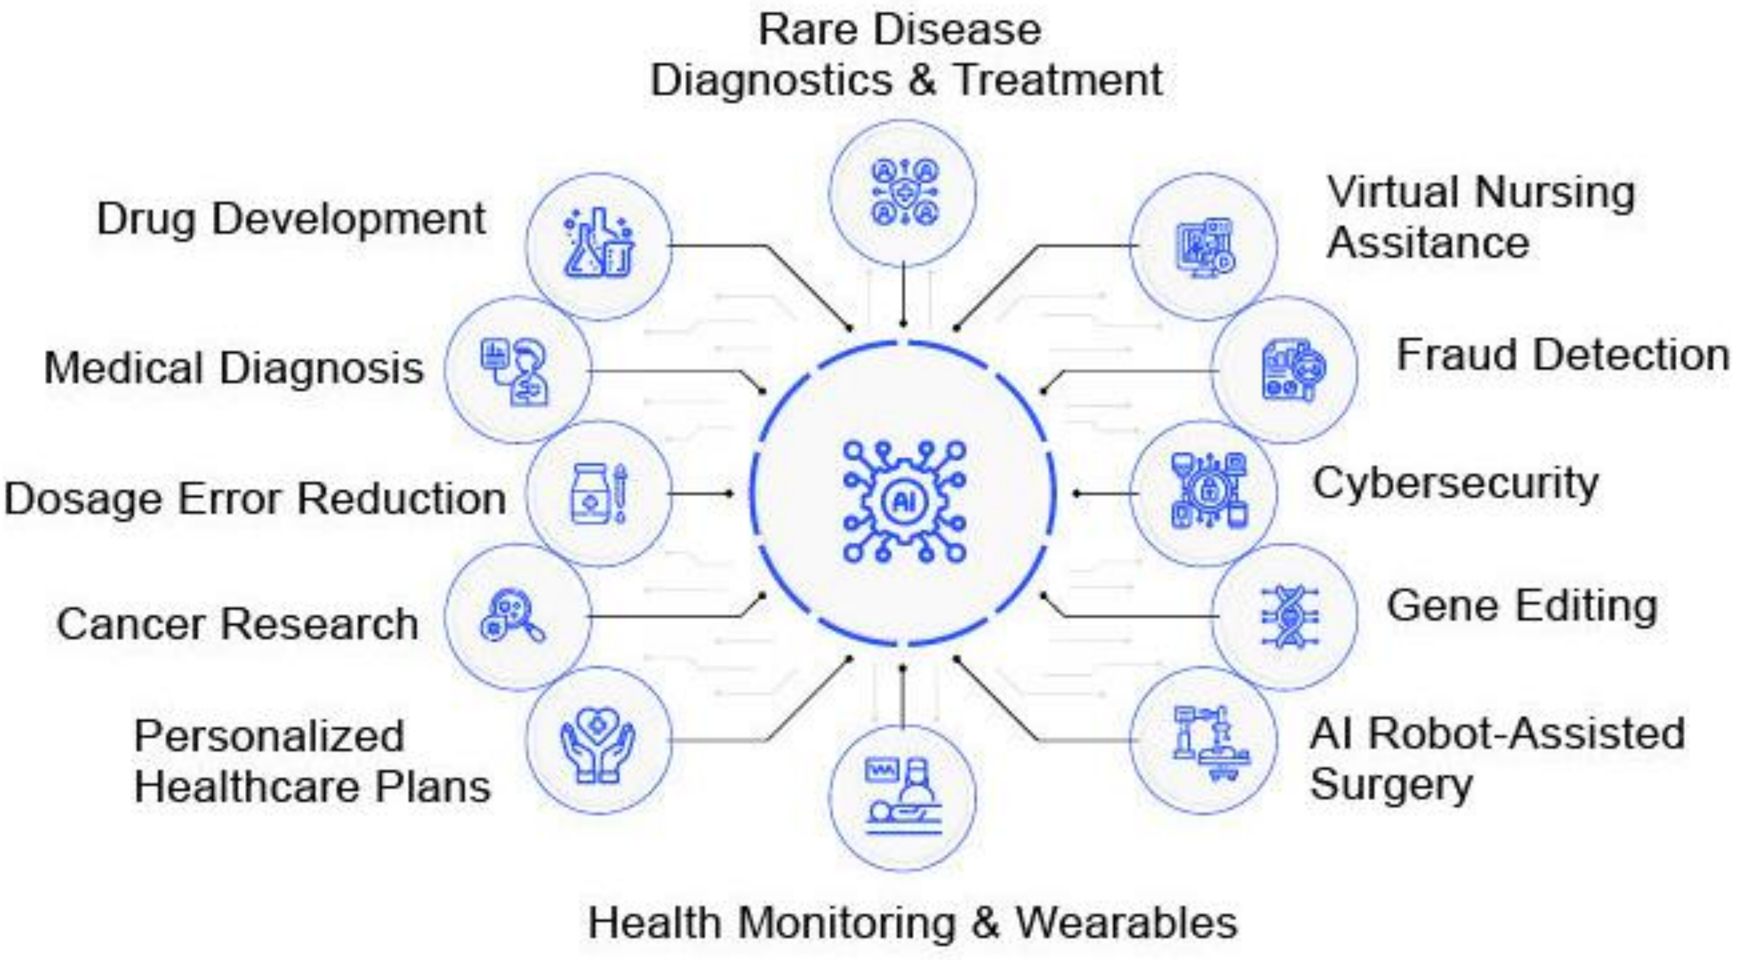

AI assisted Smart hospital of the future!

predixtions
We Predict The Future.

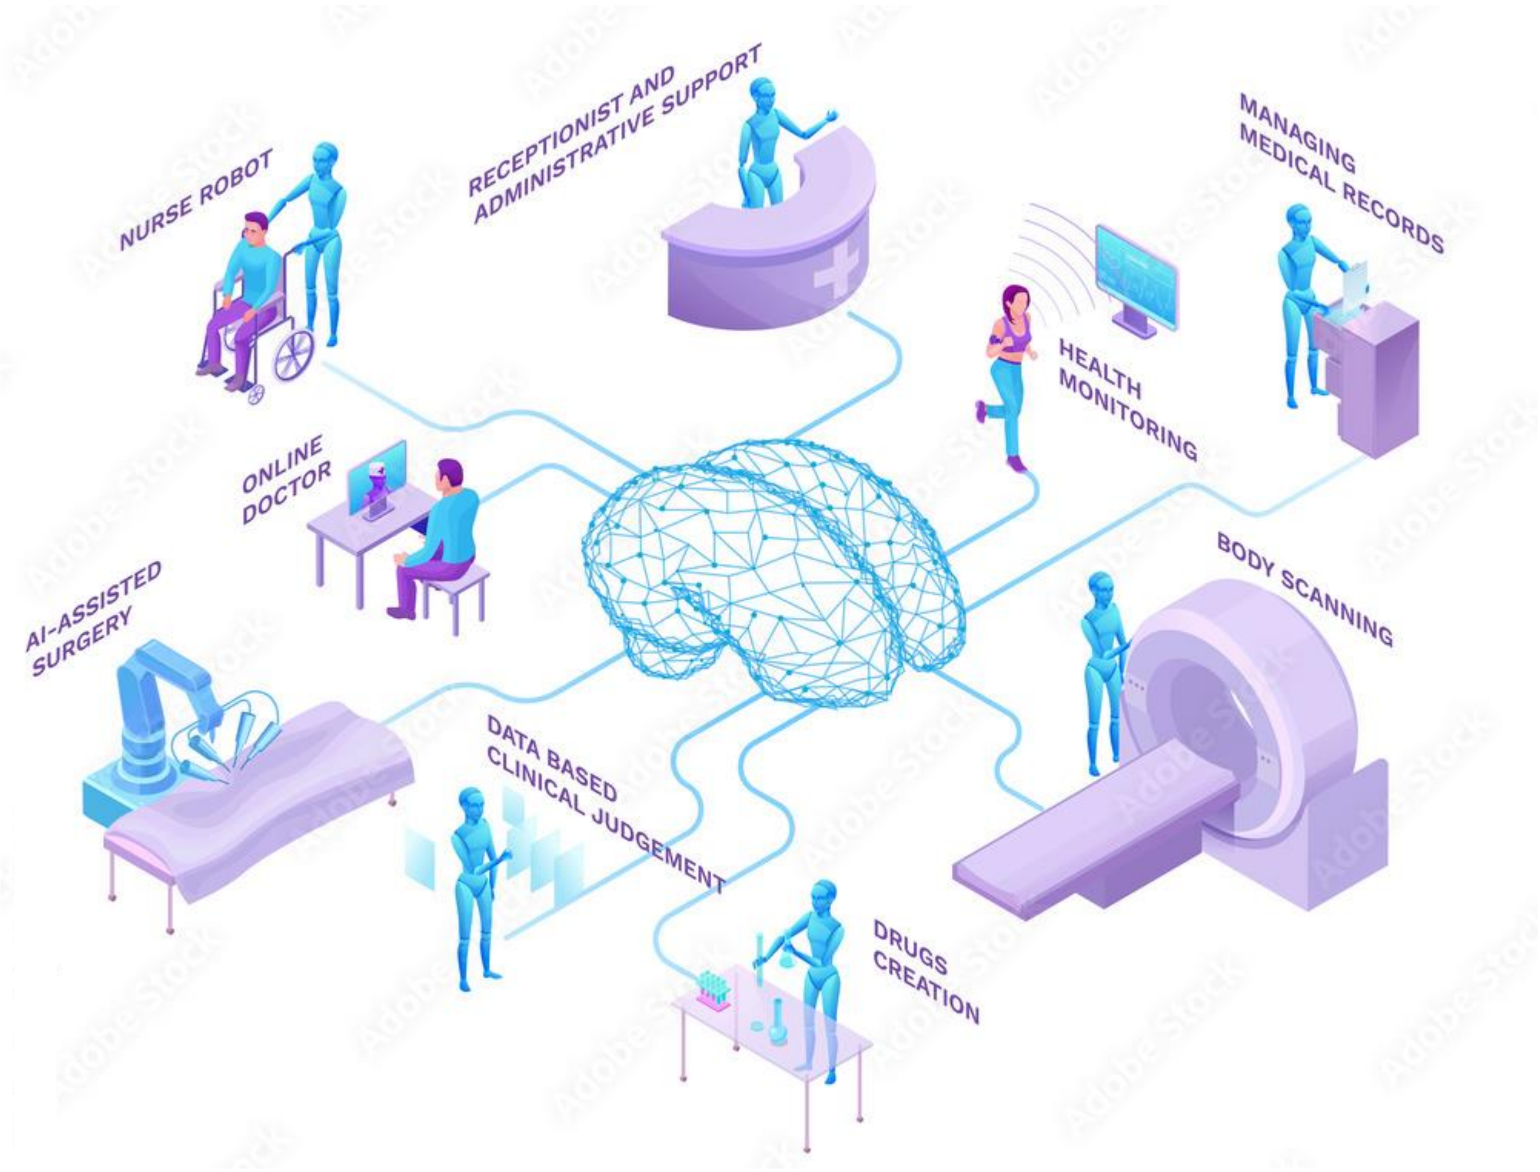

The AI Deep tech landscape across industries!

predixtions
We Predict The Future.

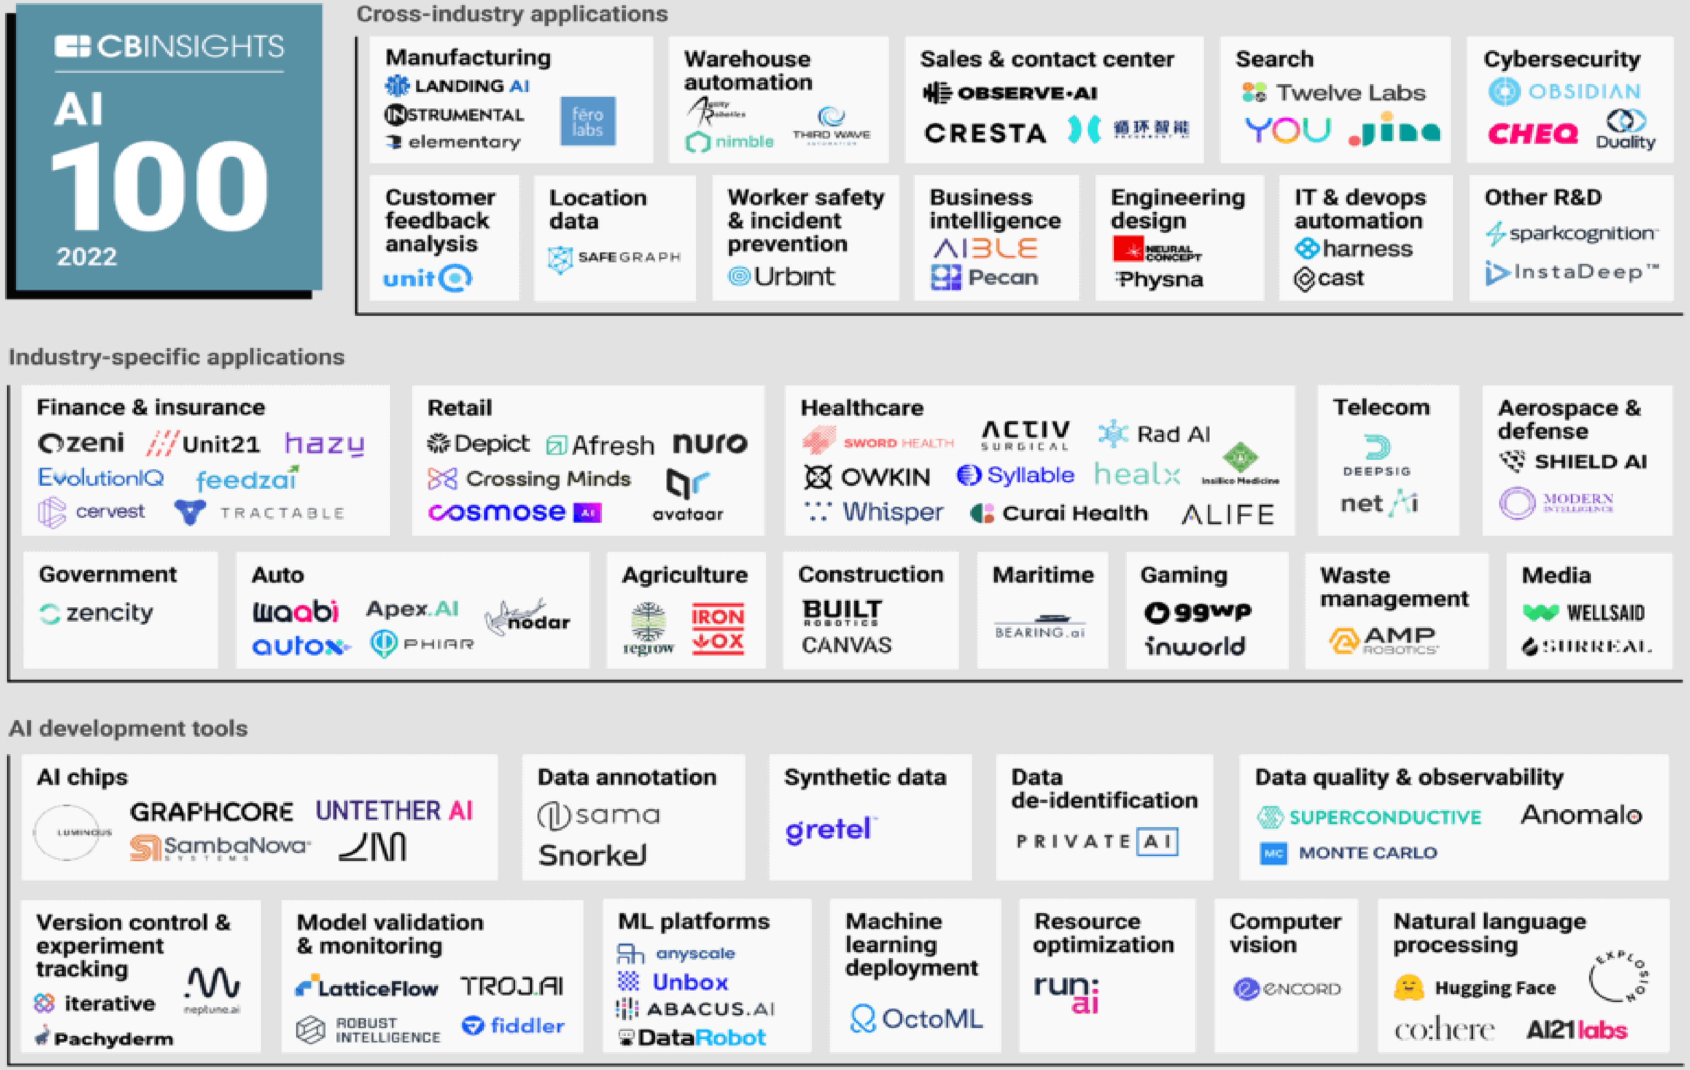

The deep learning tech landscape!

prediXtions
We Predict The Future.

AI and robotics..

prediXtions
We Predict The Future.

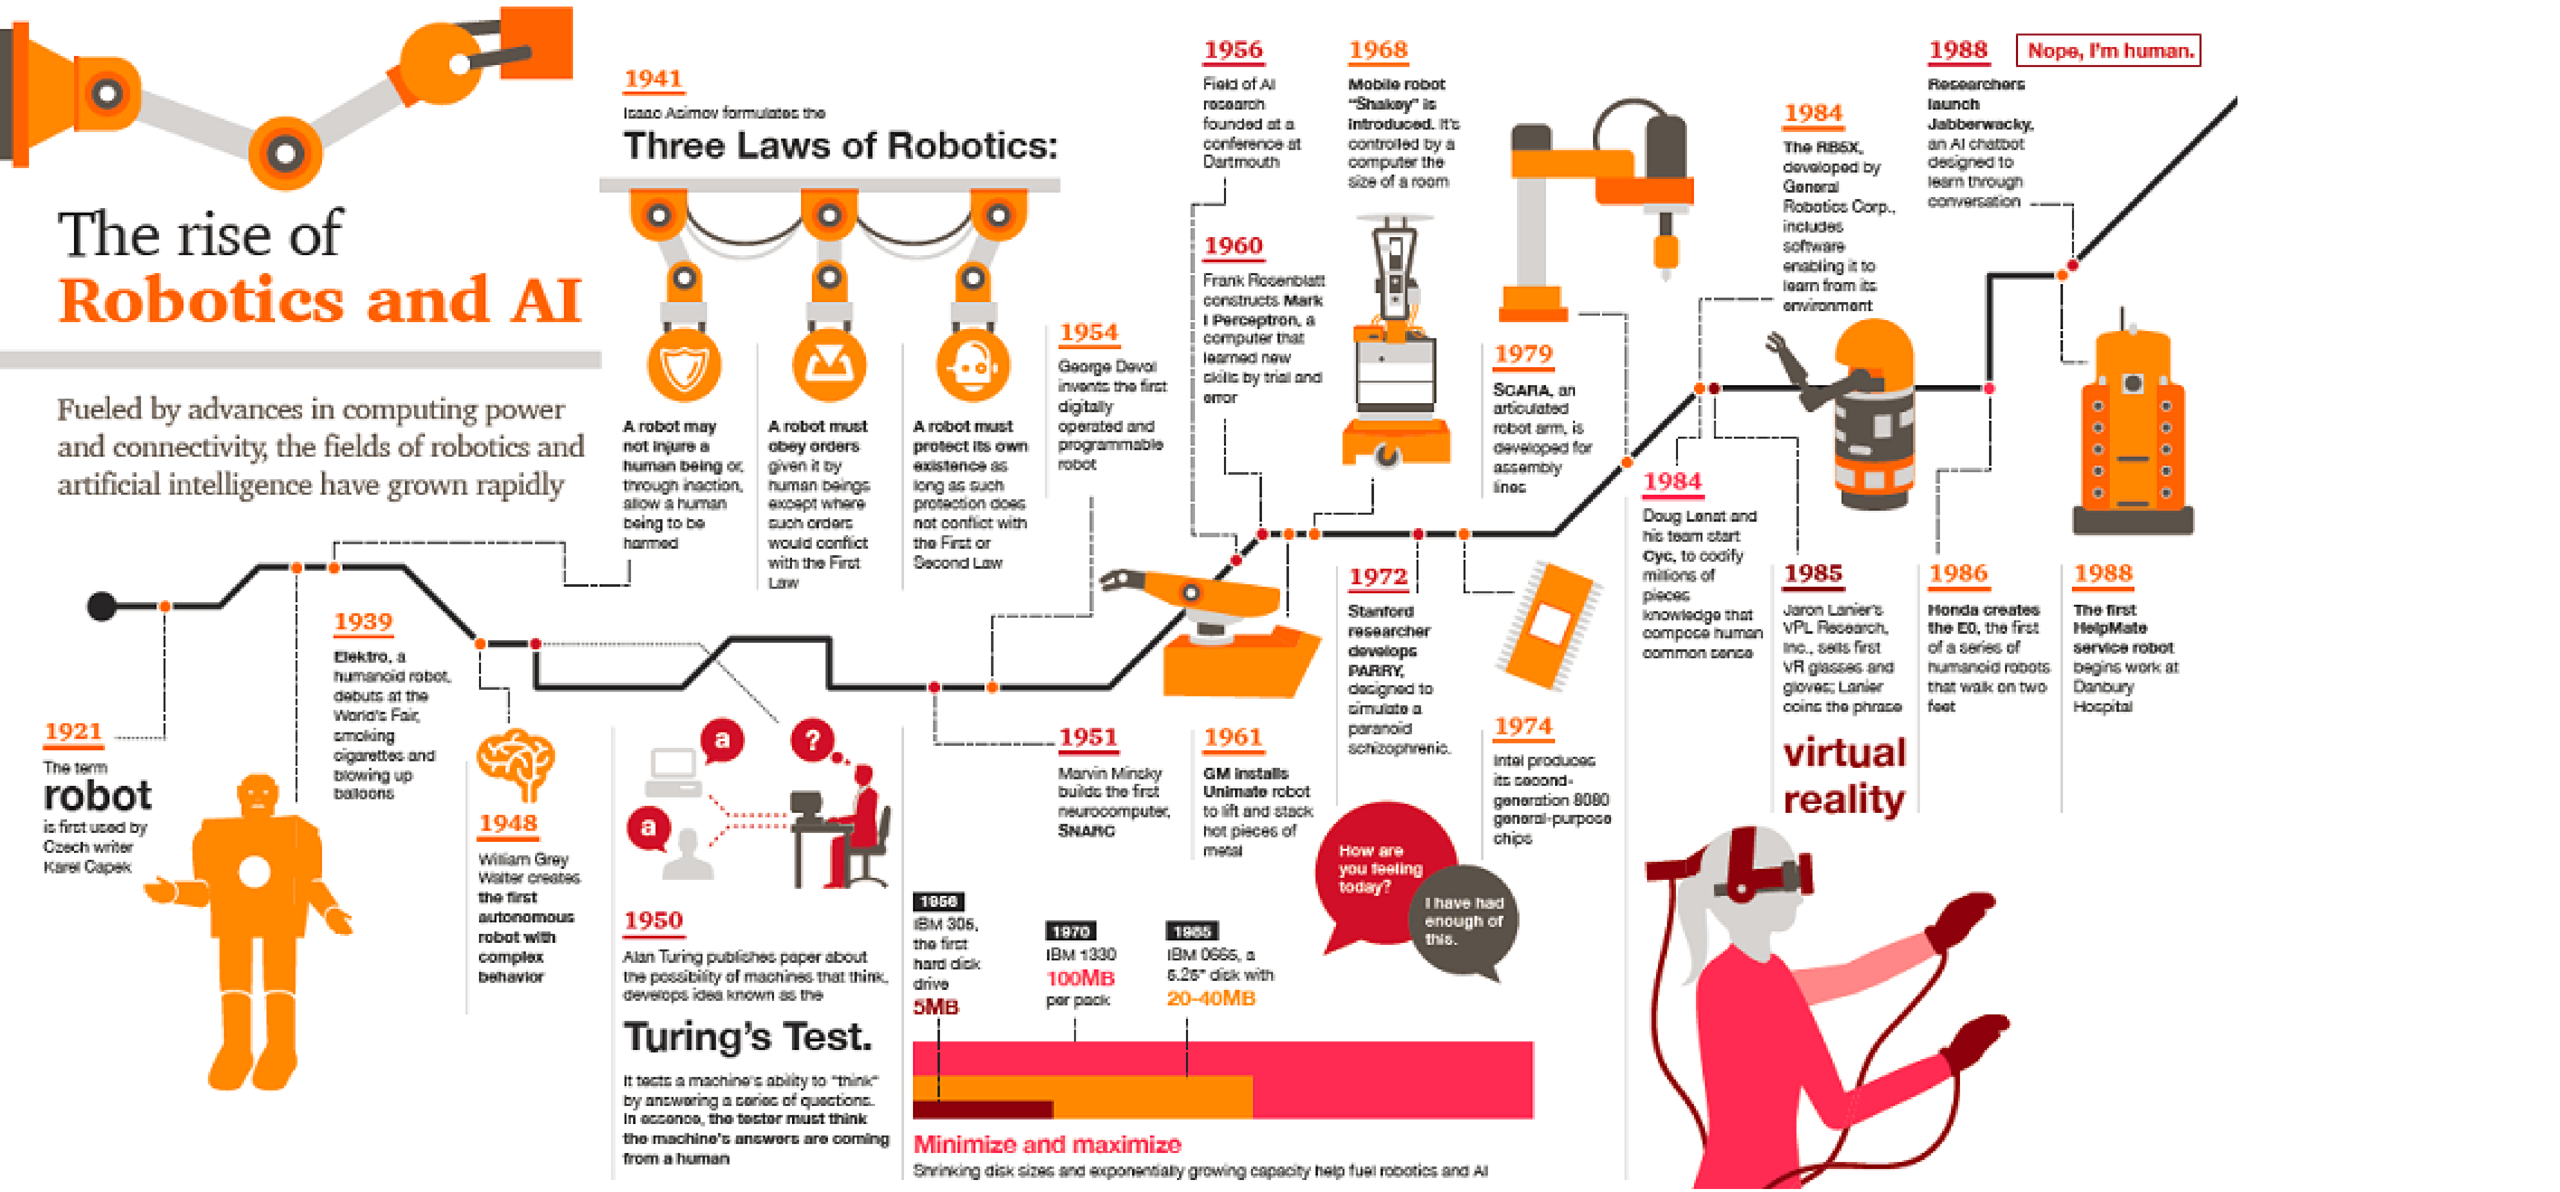

predixtions
We Predict The Future.

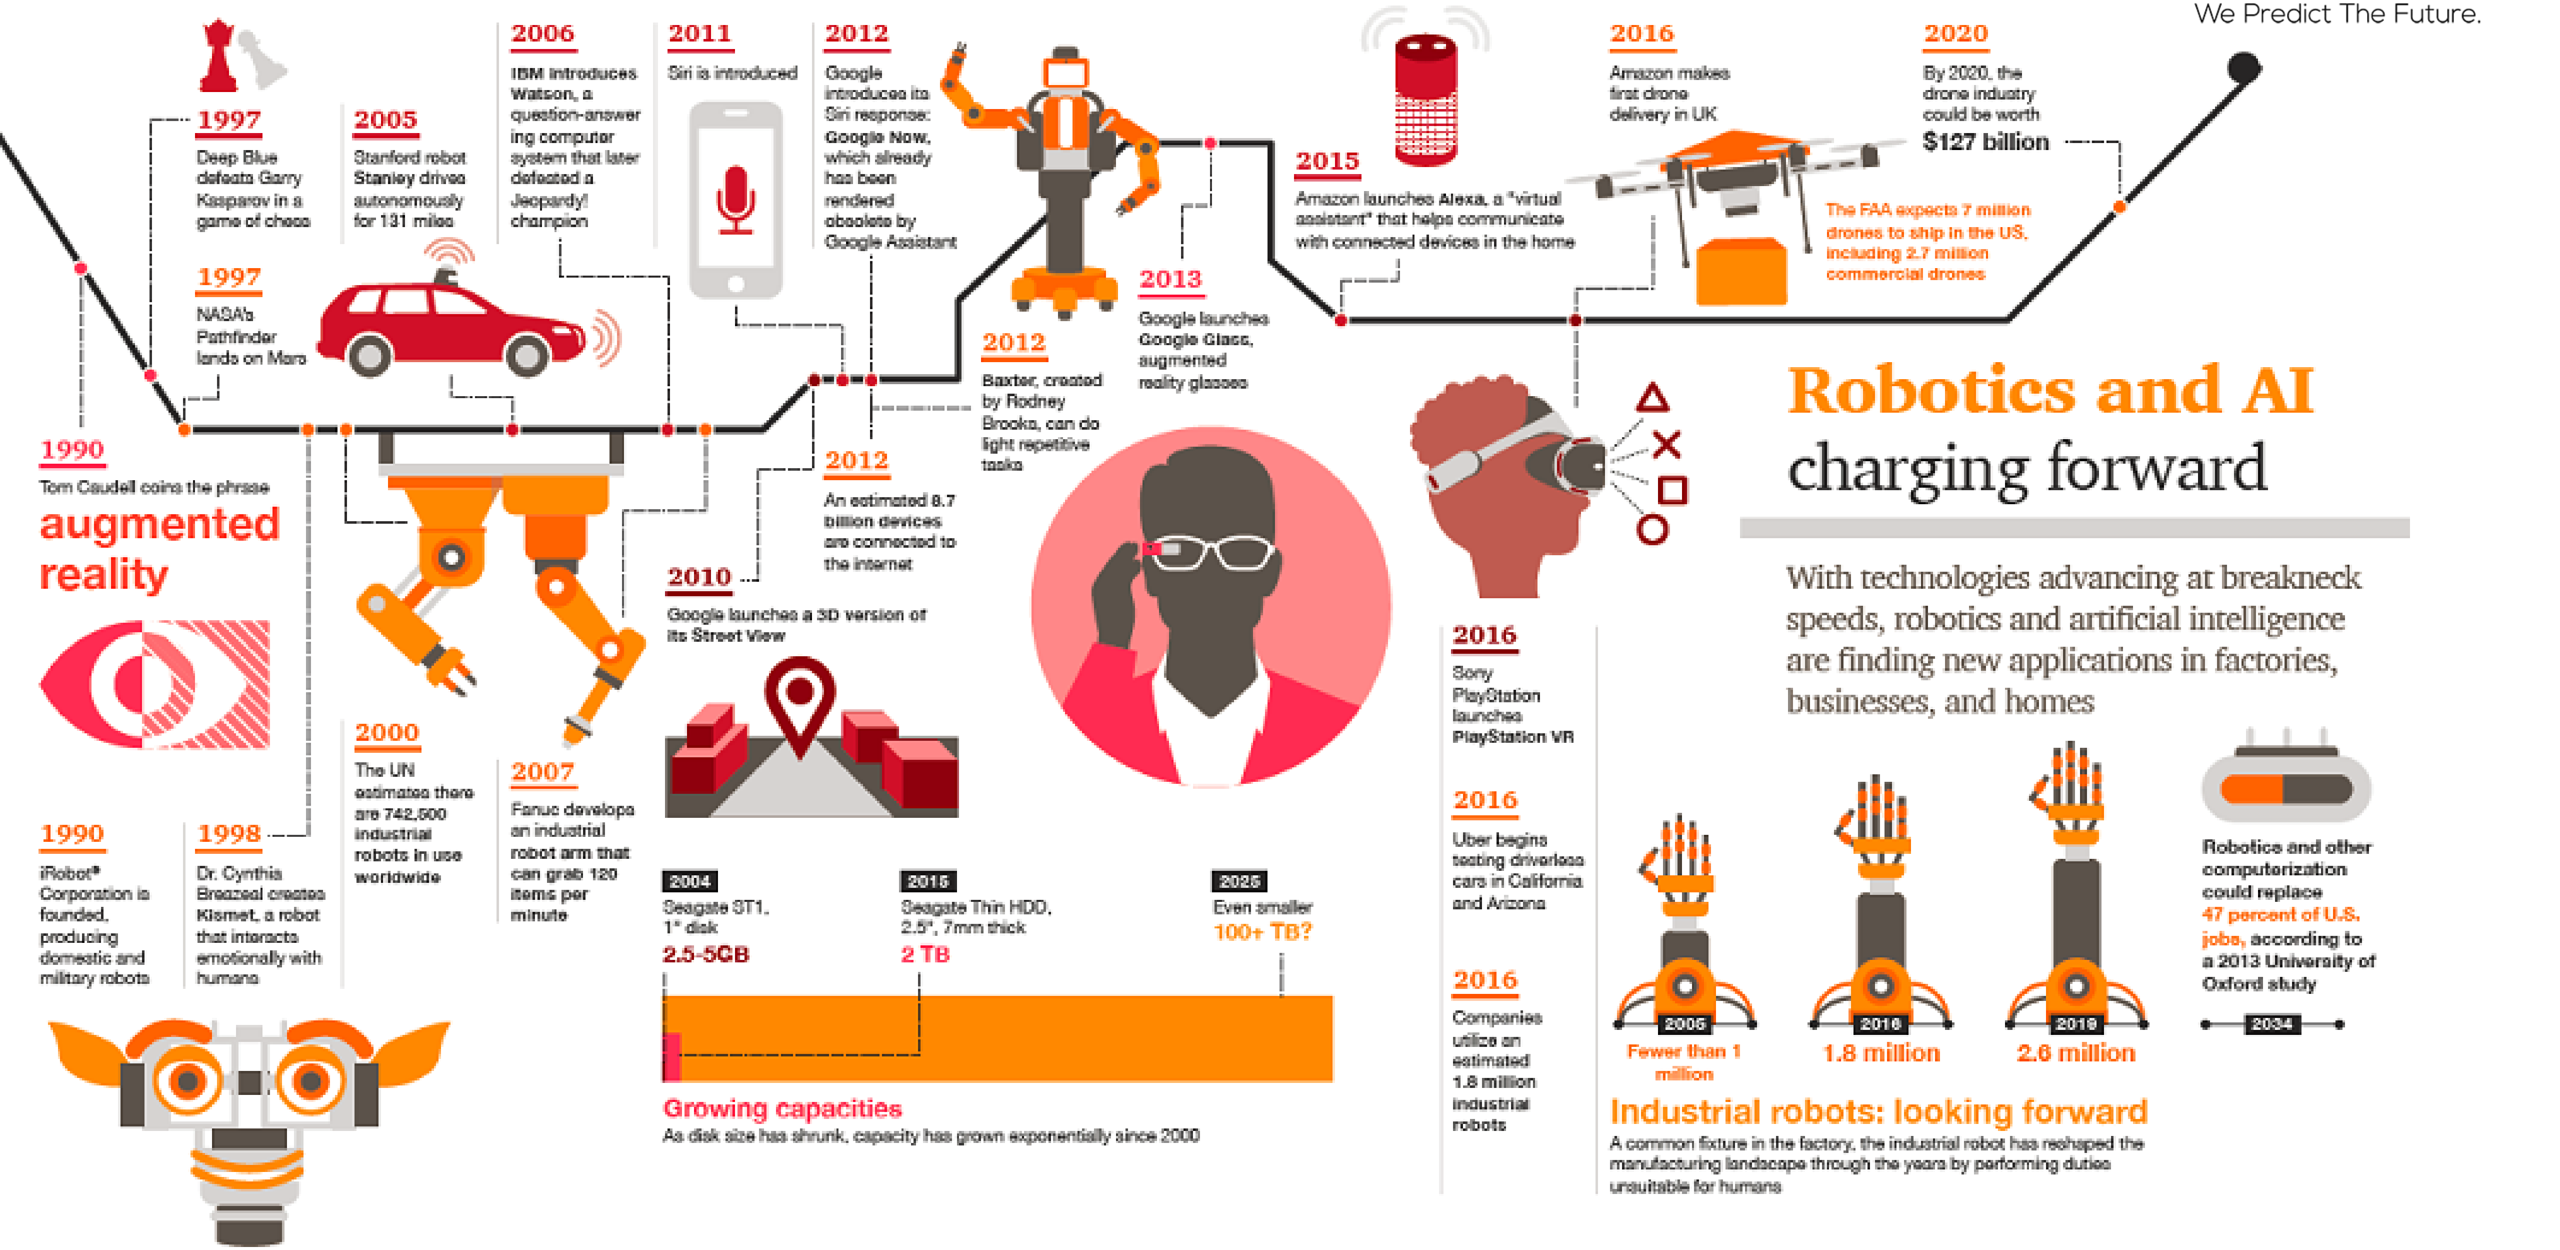

Responsible AI

predixtions
We Predict The Future.

PRINCIPLED ARTIFICIAL INTELLIGENCE

A Map of Ethical and Rights-Based Approaches to Principles for AI

Authors: Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, Madhulika Srikumar
Designers: Arushi Singh (arushisingh.net) and Melissa Axelrod (melissaaxelrod.com)

HOW TO READ:

Date, Location
Document Title
Actor

COVERAGE OF THEMES:

Higher — Lower

Not referenced

References International Human Rights

Explicitly Adopts Human Rights Framework

The size of each dot represents the percentage of principles in that theme contained in the document. Since the number of principles per theme varies, it's informative to compare dot sizes within a theme but not between themes.

The principles within each theme are:

Privacy:
Privacy
Control over Use of Data
Consent
Privacy by Design
Recommendation for Data Protection Laws
Ability to Restrict Processing
Right to Rectification
Right to Erasure

Accountability:
Accountability
Recommendation for New Regulations
Impact Assessment
Evaluation and Auditing Requirement
Verifiability and Replicability
Liability and Legal Responsibility
Ability to Appeal
Environmental Responsibility
Creation of a Monitoring Body
Remedy for Automated Decision

Safety and Security:
Security
Safety and Reliability
Predictability
Security by Design

Transparency and Explainability:
Explainability
Transparency
Open Source Data and Algorithms
Notification when Interacting with an AI
Notification when AI Makes a Decision about an Individual
Regular Reporting Requirement
Right to Information
Open Procurement (for Government)

Fairness and Non-discrimination:
Non-discrimination and the Prevention of Bias
Fairness
Inclusiveness in Design
Inclusiveness in Impact
Representative and High Quality Data
Equality

Human Control of Technology:
Human Control of Technology
Human Review of Automated Decision
Ability to Opt out of Automated Decision

Professional Responsibility:
Multistakeholder Collaboration
Responsible Design
Consideration of Long Term Effects
Accuracy
Scientific Integrity

Promotion of Human Values:
Leveraged to Benefit Society
Human Values and Human Flourishing
Access to Technology

Further information on findings and methodology is available in Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches (Berkman Klein, 2020) available at ***cyber.harvard.edu.***

BERKMAN KLEIN CENTER
FOR INTERNET & SOCIETY AT HARVARD UNIVERSITY

How do we inculcate responsibility in AI development for the future?

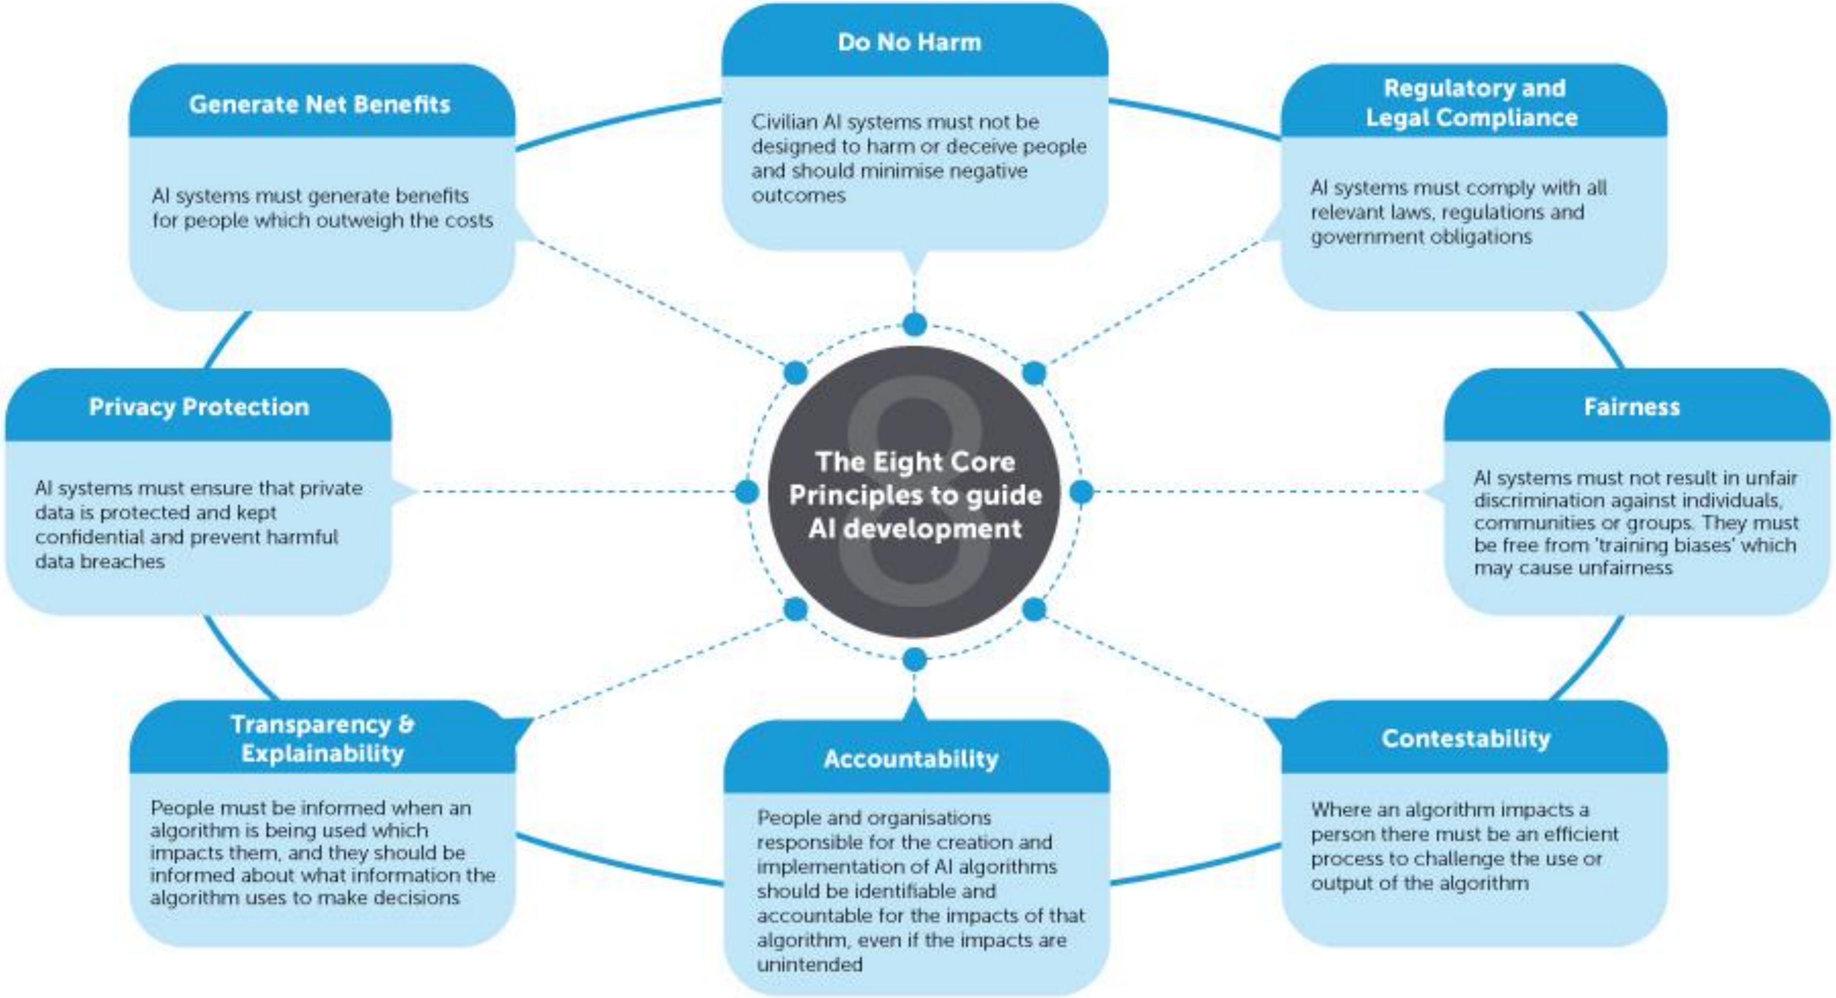

Responsible AI in AI modelling

prediXtions
We Predict The Future.

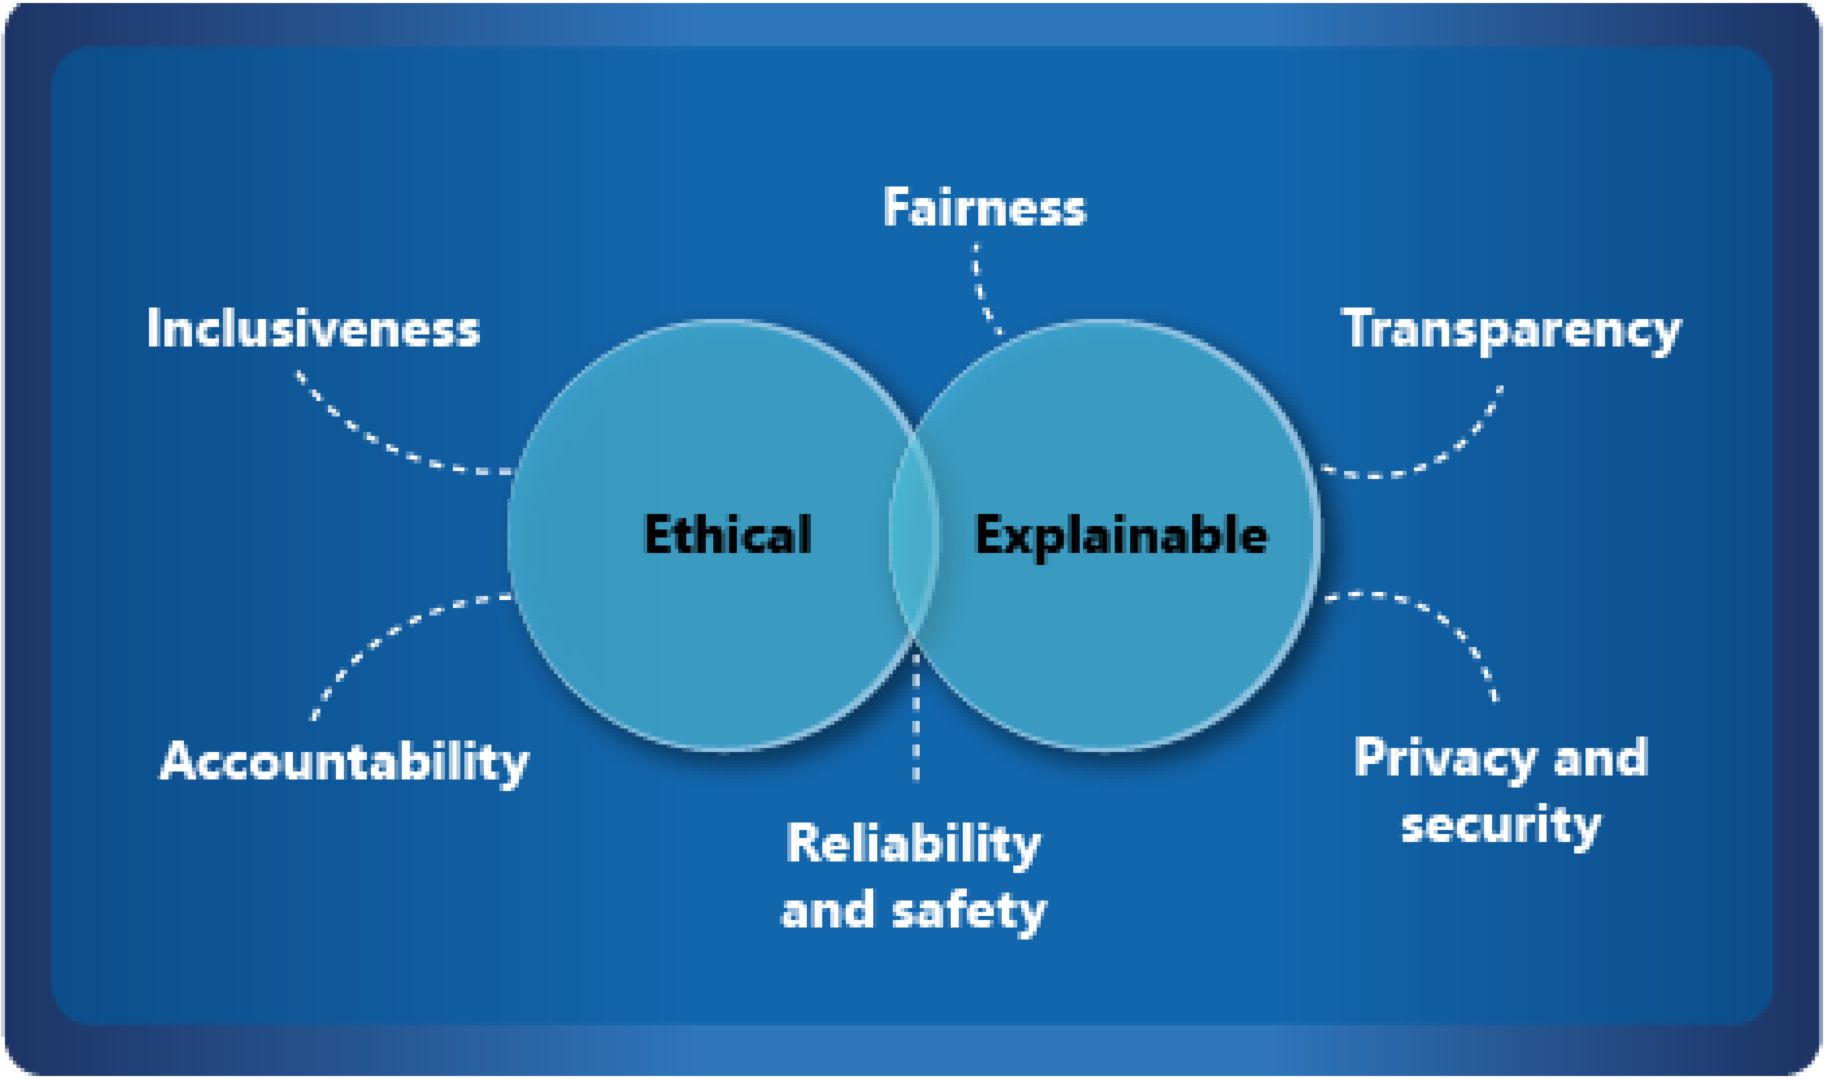

predixtions
We Predict The Future.

What is the AI toolbox for responsible modelling?

Identify

Error Analysis
Identify cohorts with high error rate versus benchmark and visualize how the error rate distributes

Fairness Assessment
Assess model fairness by evaluating your model performance and predictions across different sensitive groups.

Diagnose

Model Interpretability
Interpret and debug model.

Exploratory Data Analysis
Understand dataset characteristics

Counterfactual Analysis and What If
Generate diverse counterfactual explanations for debugging. Perform feature perturbations

Mitigate

Unfairness Mitigation
Mitigate fairness issues

Data Enhancements
Enhance your dataset and retrain model

Inform Actions

Causal Inference
Understand the causal impact of your features on real-world outcomes

Counterfactual Analysis
Generate diverse counterfactual explanations for informing end users

Compare

Model Comparison

Backward Compatibility

BIAS in AI- Focus on the healthcare industry as an example

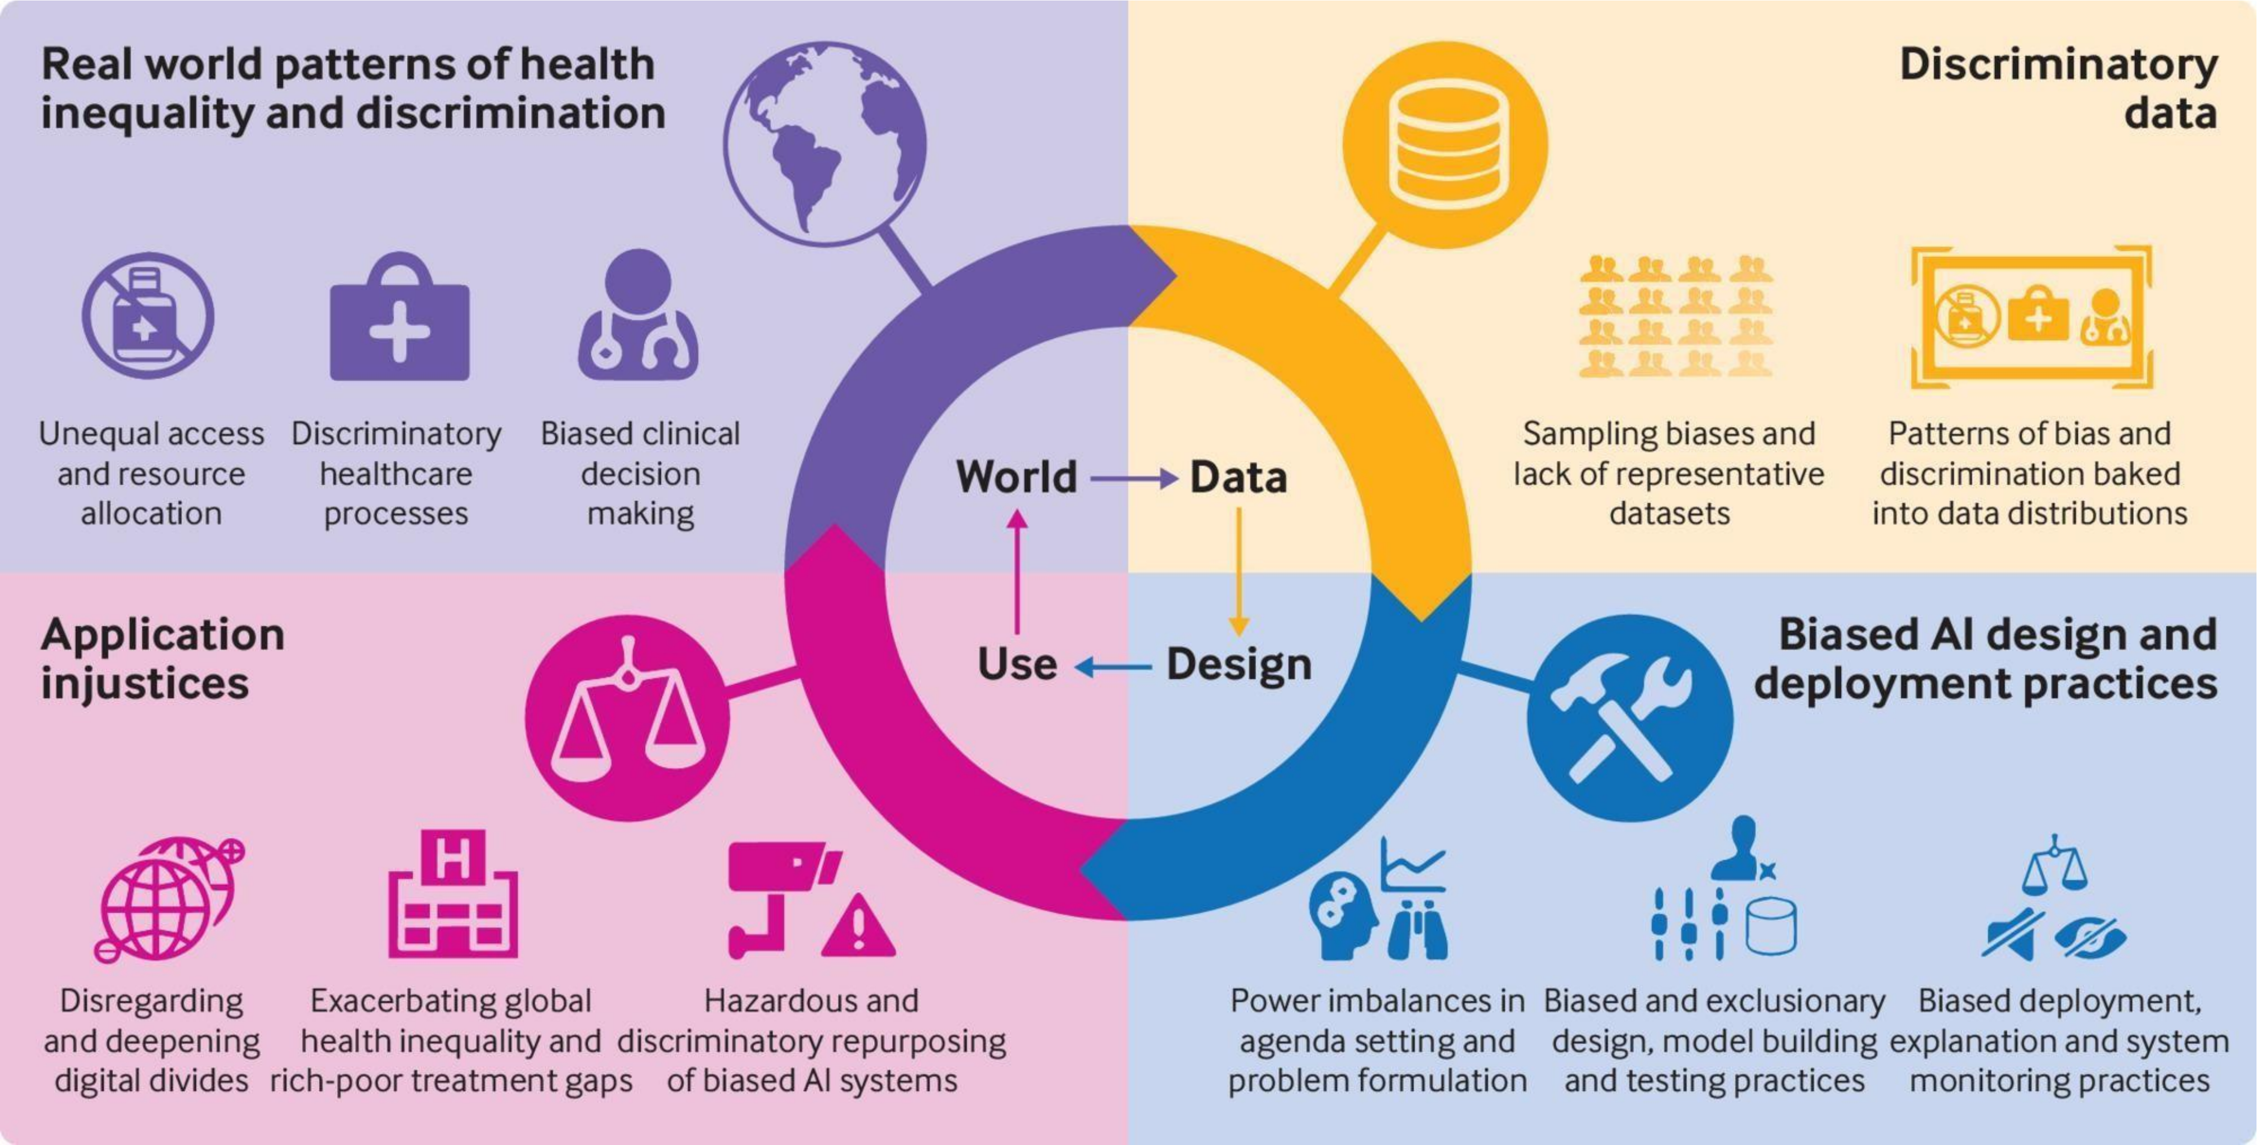

The future of AI

Quantum computing

prediXtions
We Predict The Future.

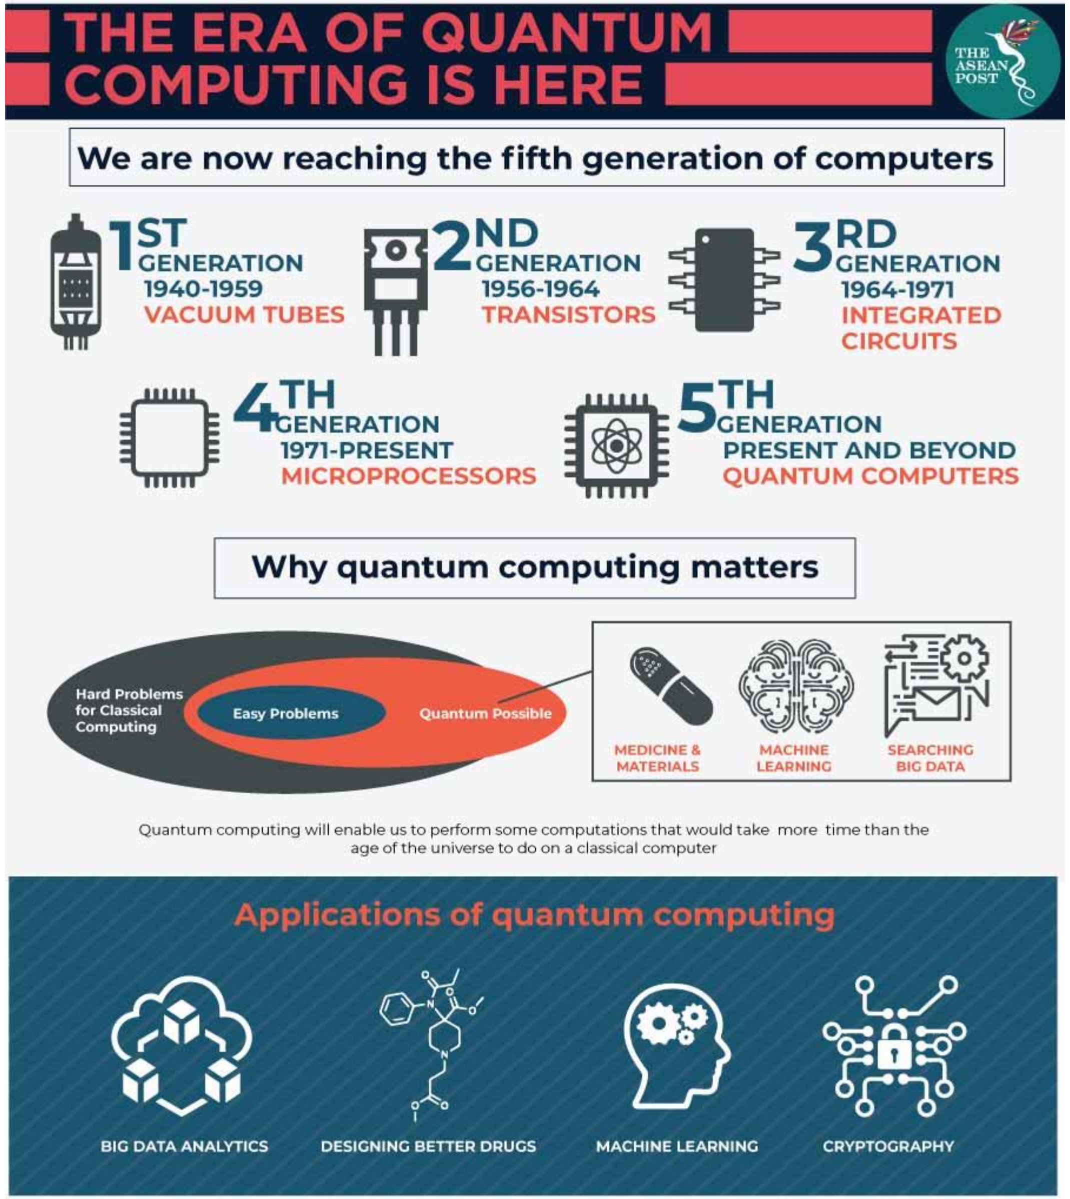

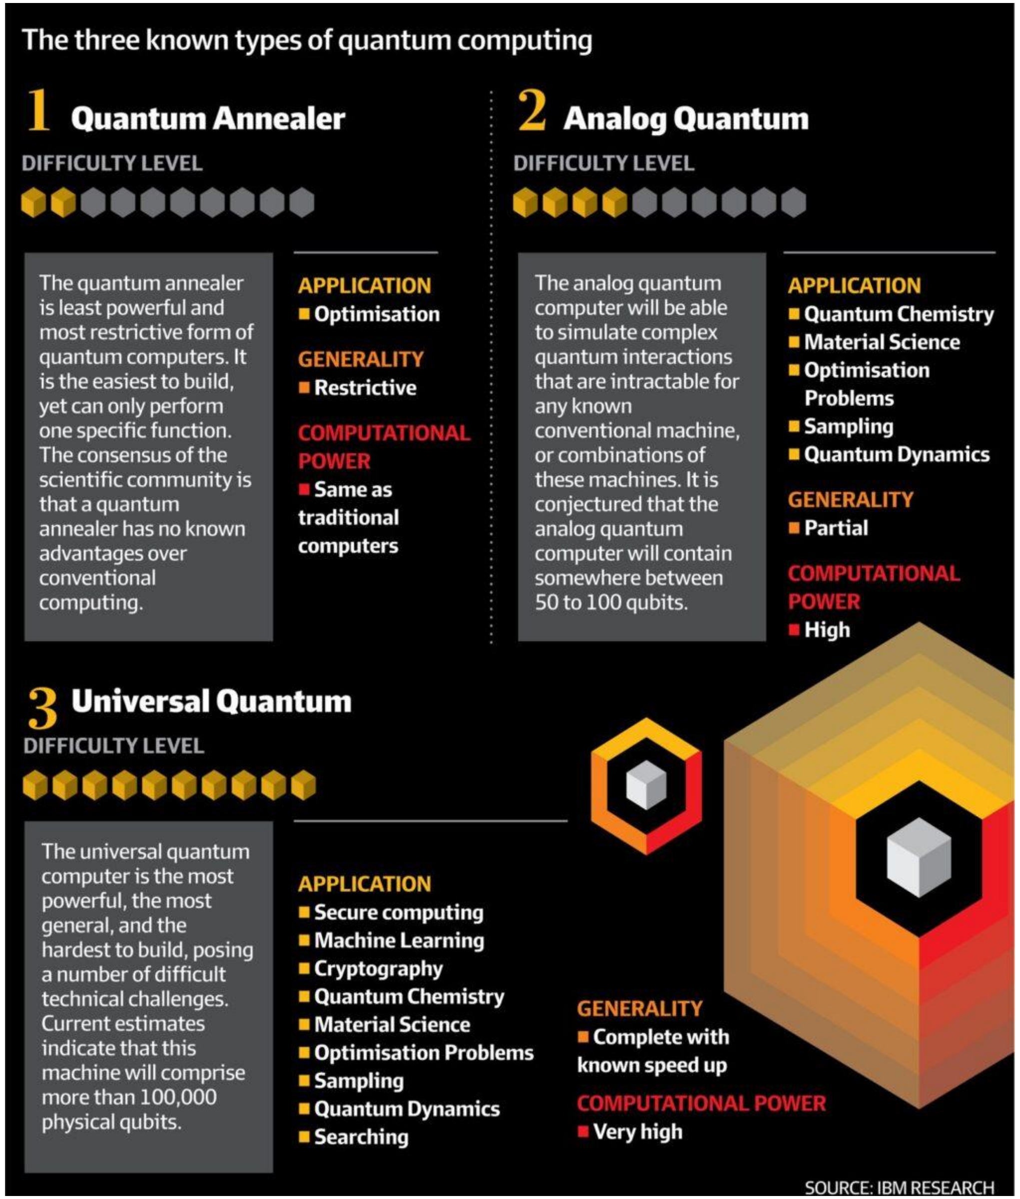

Space exploration