Microsoft AI

Polytechnic Hack-Fest

Firas Bou Fakhereddine
Data and AI CSA

John Thabet
Data and AI SP

Microsoft AI

Microsoft AI

Building blocks for the future

Mixed reality

Quantum computing

Artificial intelligence

Blockchain

Digital feedback loop

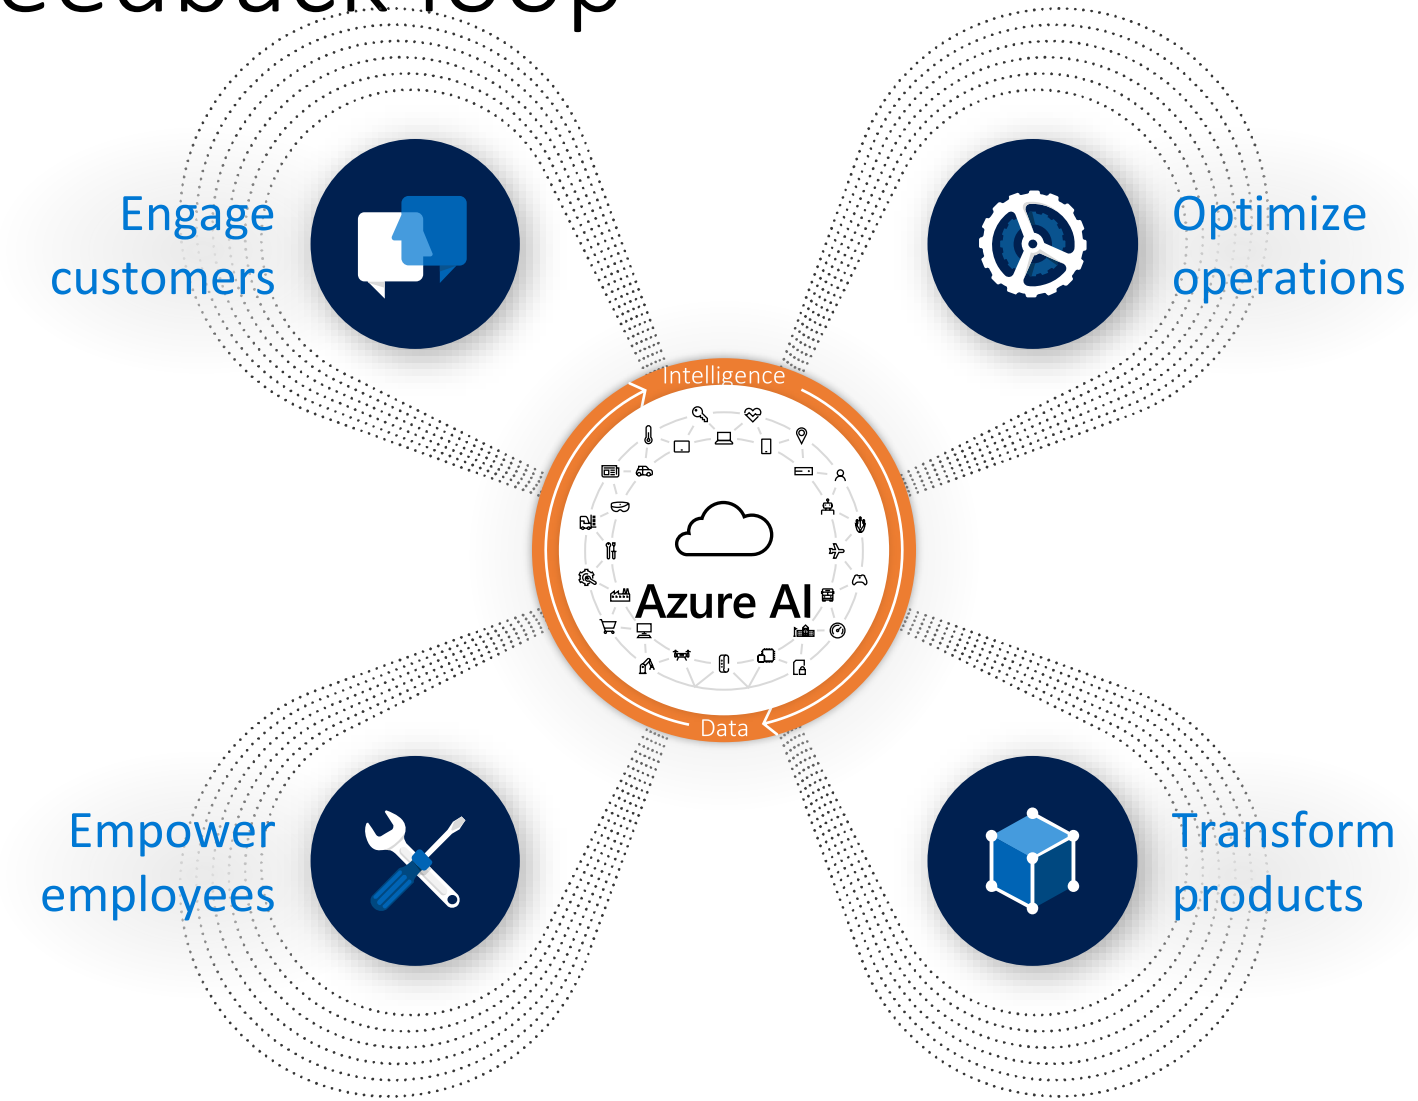

Microsoft AI investment areas

AI platform

Infusing AI

Business solutions

Azure AI

AI apps & agents

Azure Bot Service
Azure Cognitive Services

Machine learning

Azure Machine Learning

Knowledge mining

Azure Cognitive Search

Azure Data + AI Solution Areas

Data

Data Modernization on-premises

Data modernization to Azure

Globally distributed data

Cloud Scale Analytics

+

AI

AI apps & agents

Knowledge mining

Machine learning

Azure Cognitive Services

Perception

Vision

Speech

Comprehension

Language

Knowledge

Machine Learning on Azure

Sophisticated pretrained models
To simplify solution development

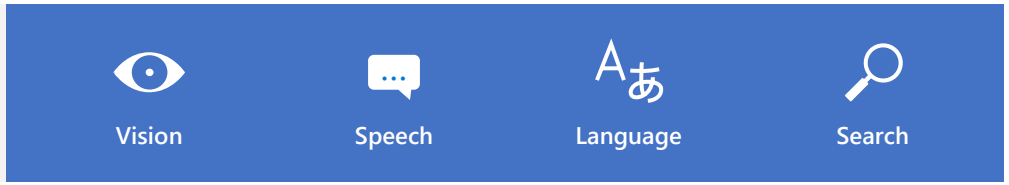

Popular frameworks
To build advanced deep learning solutions

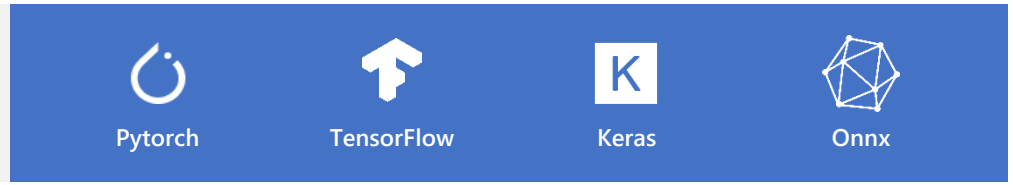

Productive services
To empower data science and development teams

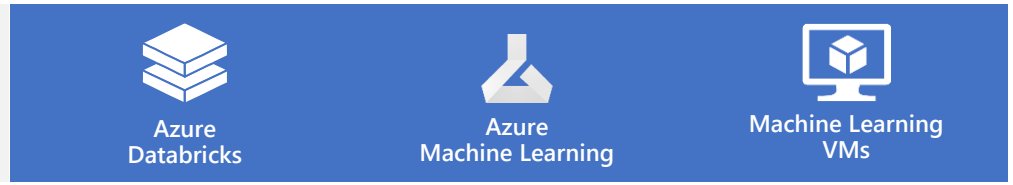

Powerful infrastructure
To accelerate deep learning

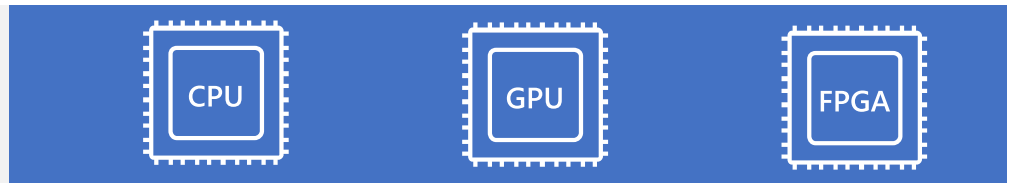

Flexible deployment
To deploy and manage models on intelligent cloud and edge

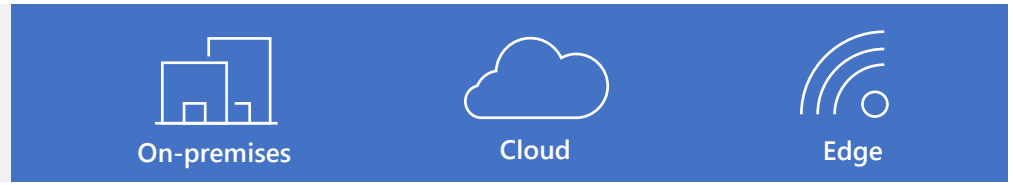

Microsoft AI investment areas

AI platform

Infusing AI

Business solutions

Infusing AI

Facial recognition | Cortana | Inking

Windows

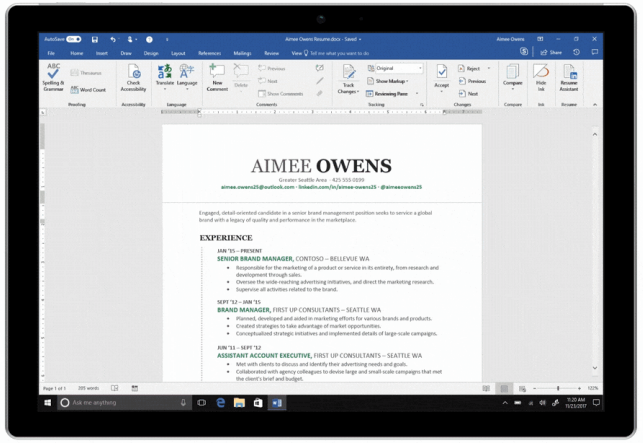

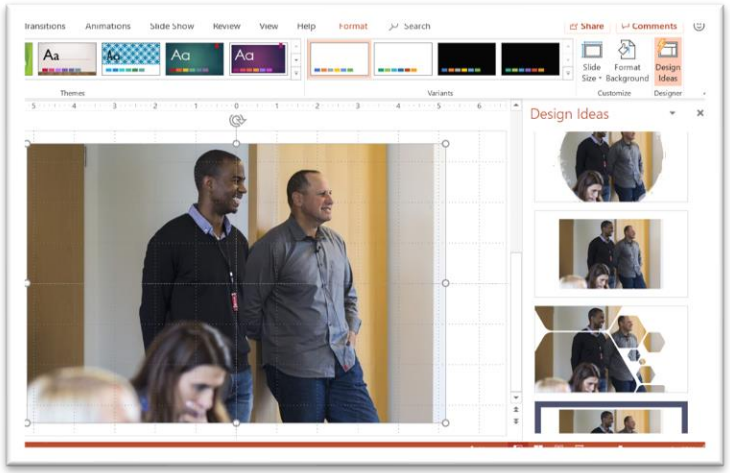

AI services

Threat detection

Autonomous services

Knowledge graph

Visual Studio

Deep learning frameworks | Manage AI models | Sample gallery

Office

Designer | Morph transition | Resume assistant |Data types

Microsoft AI investment areas

AI platform

Infusing AI

Business solutions

Patterns for AI solutions

New generation of business agents

B2B, B2C, B2E

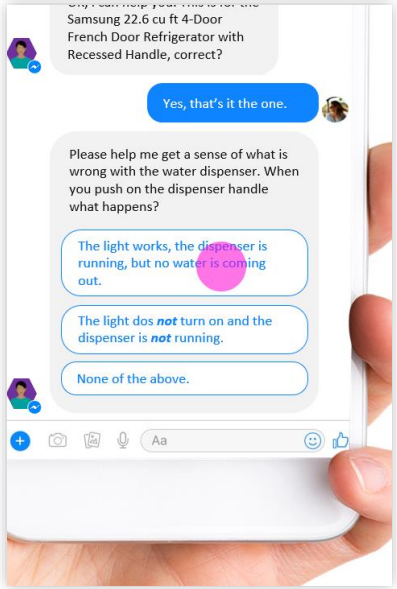

Person, object, and activity detection

Retail, manufacturing, security

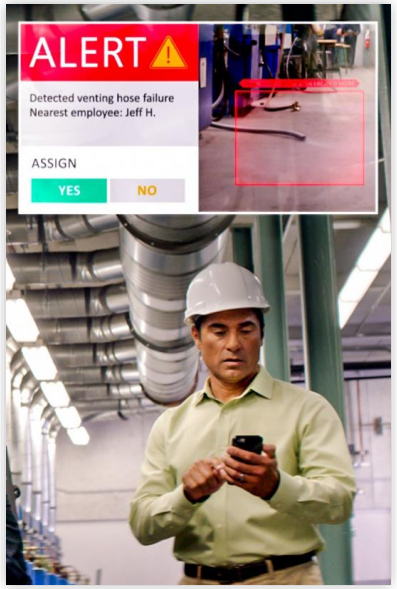

AI assisted professionals

Marketing, legal, financial

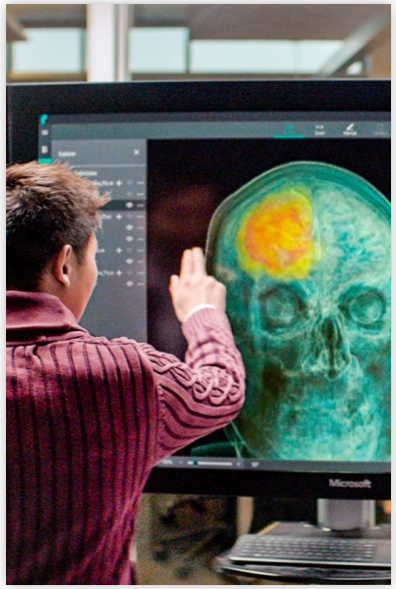

Knowledge mining

Documents, video

Autonomous systems

Vehicles, networks, RPA

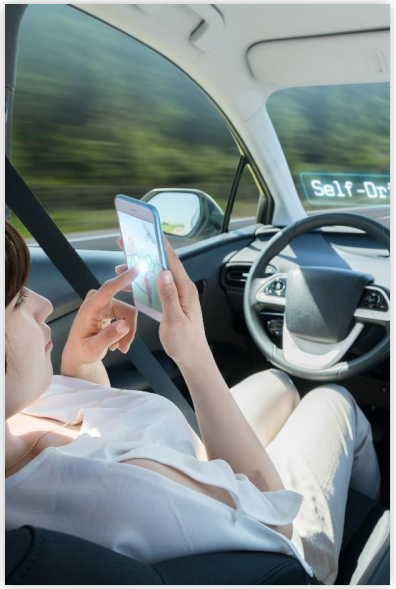


Develop applications that understand us

Cognitive Services

Microsoft Cognitive Services

Give your apps a human side

Microsoft Cognitive services are Artificial intelligence pre trained rest APIs. You don’t need to be a data scientist to use Cognitive Services. It allows developers to easily add cognitive features into their applications.

Returns Json as output.

More than three dozen languages are supported.

The catalog of services within Azure Cognitive Services can be categorized into five main pillars –

Vision, Speech, Language, Web Search, and Decision.

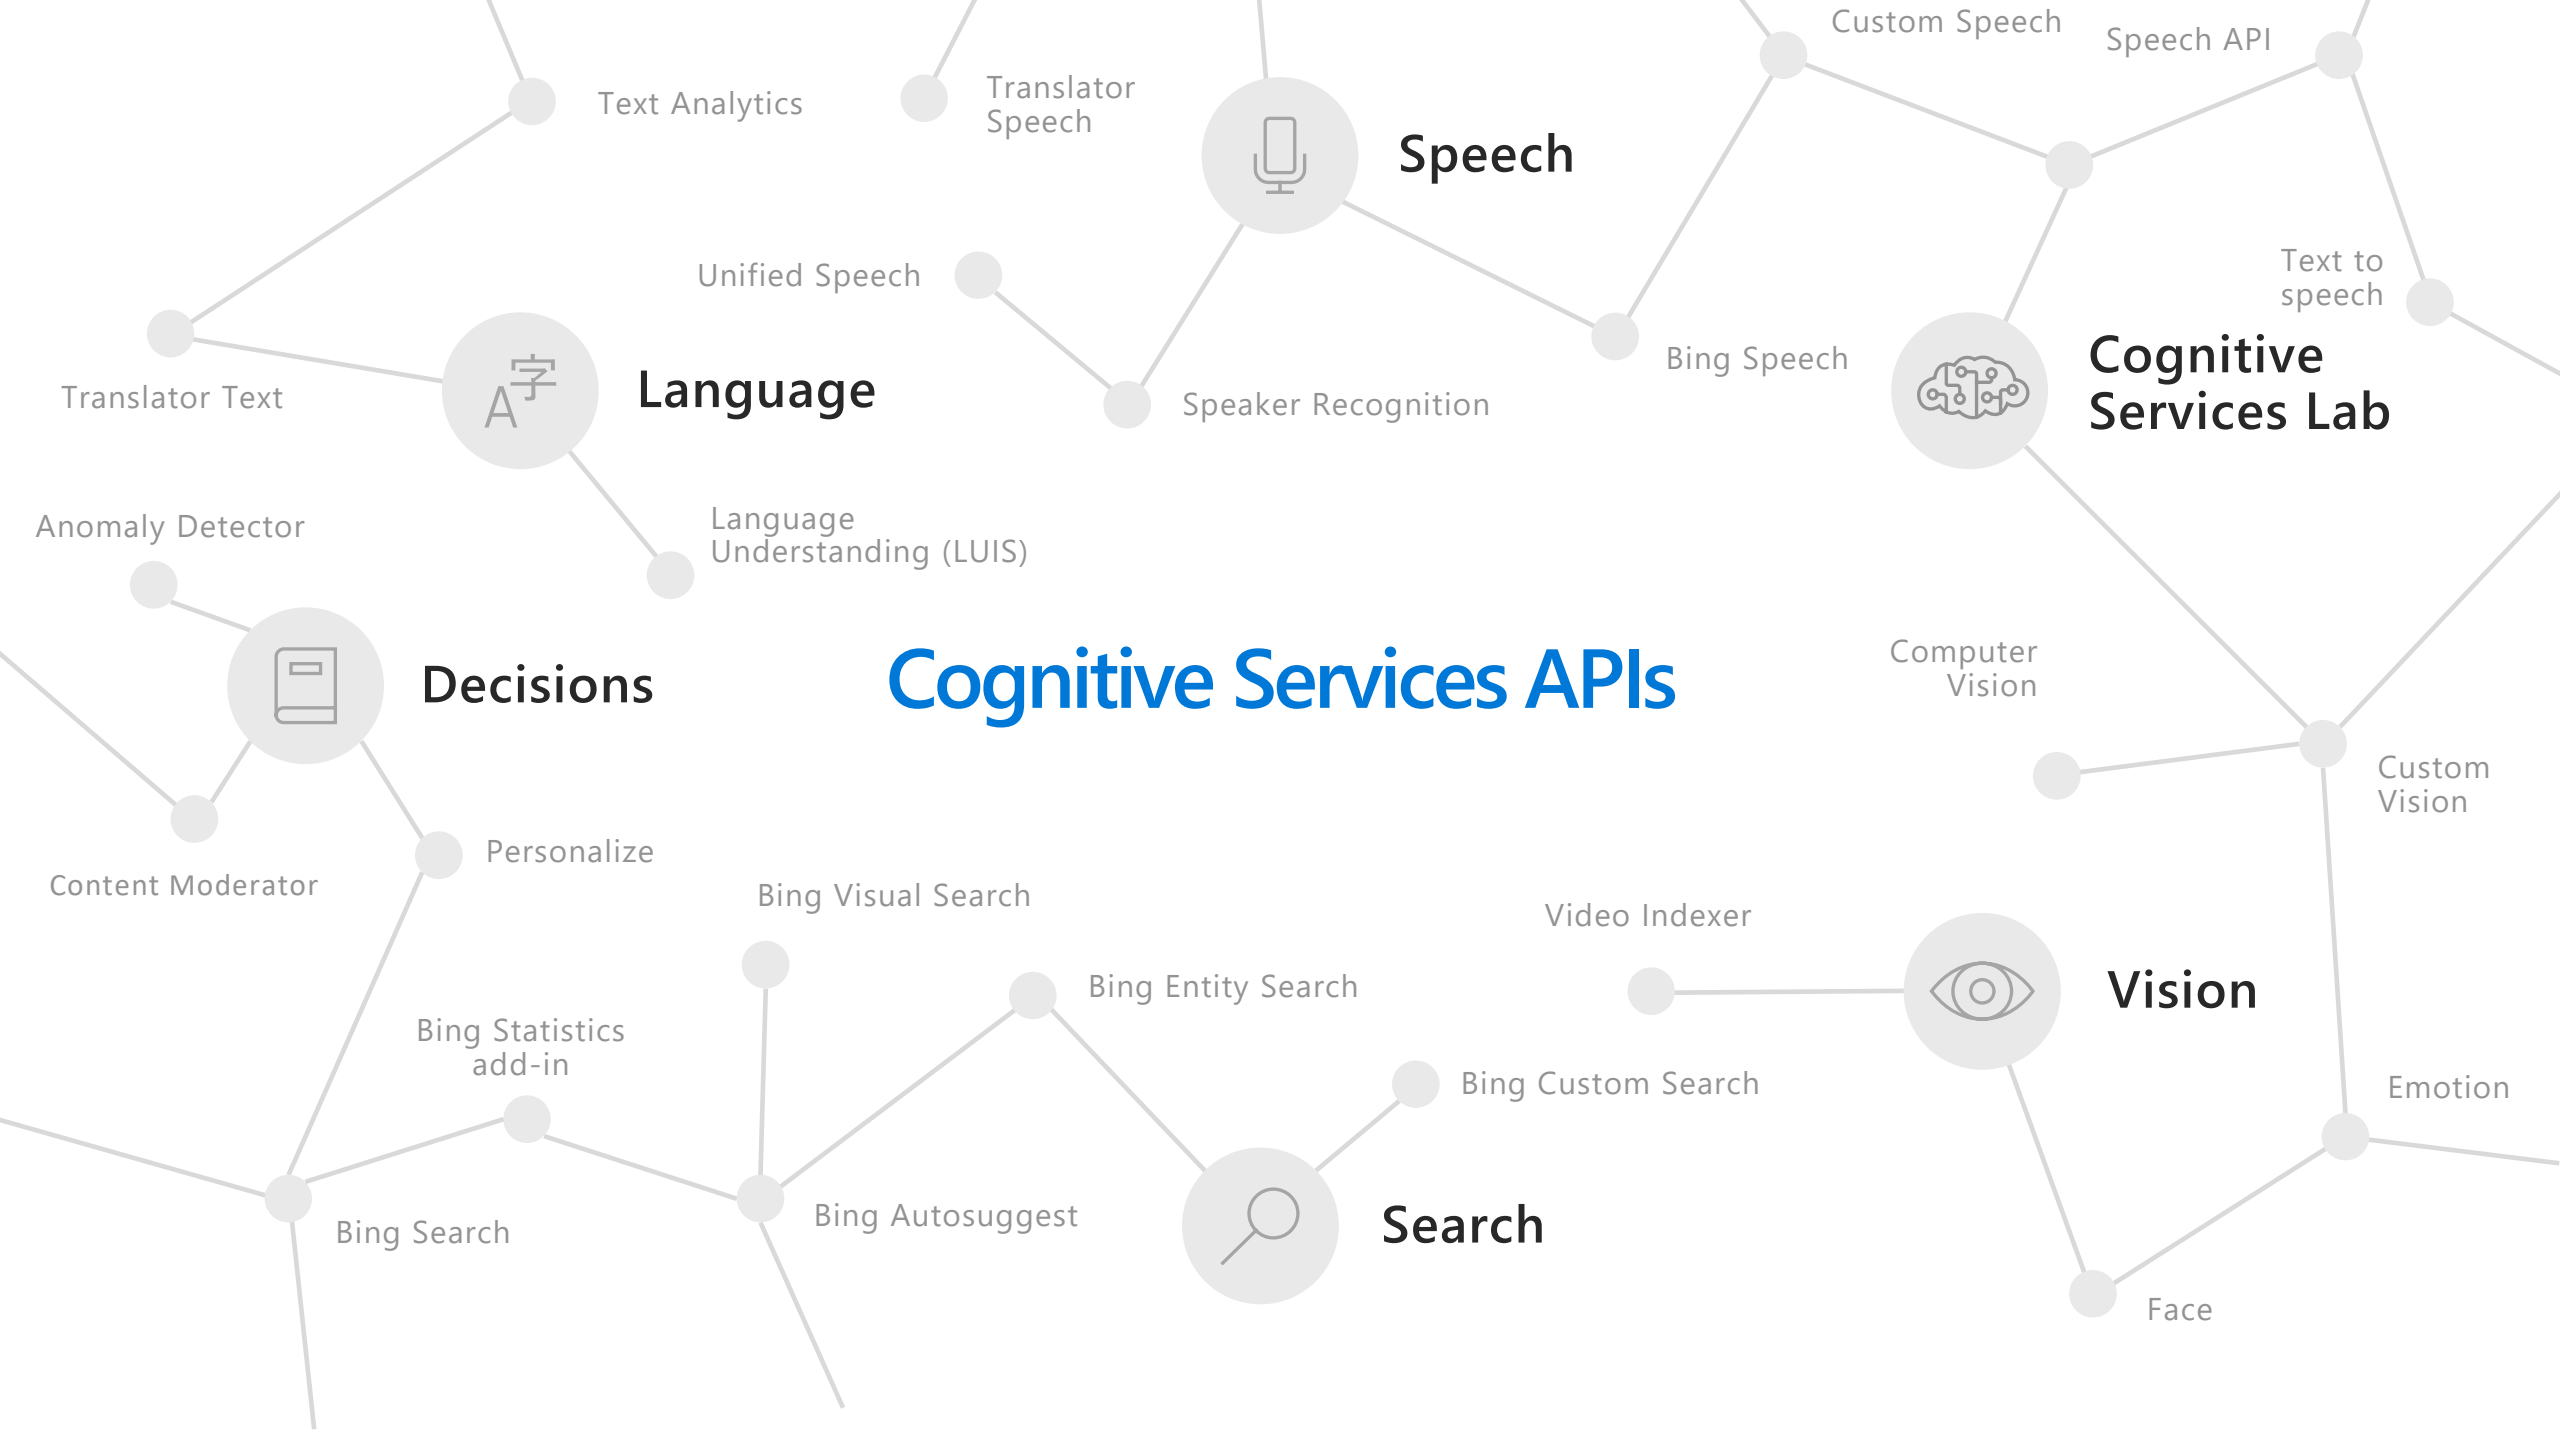
Text Analytics
Translator Speech
Custom Speech
Speech API
Speech
Unified Speech
Text to speech
Translator Text
Language
Bing Speech
Cognitive Services Lab
Speaker Recognition
Anomaly Detector
Language Understanding (LUIS)
Decisions
Cognitive Services APIs
Computer Vision
Custom Vision
Personalize
Content Moderator
Bing Visual Search
Video Indexer
Bing Entity Search
Bing Statistics add-in
Vision
Bing Custom Search
Emotion
Bing Search
Bing Autosuggest
Search
Face

Cognitive Services capabilities

Infuse your apps, websites, and bots with human-like intelligence

Vision

Object, scene, and activity detection

Face recognition and identification

Celebrity and landmark recognition

Emotion recognition

Text and handwriting recognition (OCR)

Customizable image recognition

Video metadata, audio, and keyframe extraction and analysis

Explicit or offensive content moderation

Speech

Speech transcription (speech-to-text)

Custom speech models for unique vocabularies or complex environment

Text-to-speech

Custom Voice

Real-time speech translation

Customizable speech transcription and translation

Speaker identification and verification

Language

Language detection

Named entity recognition

Key phrase extraction

Text sentiment analysis

Multilingual and contextual spell checking

Explicit or offensive text content moderation

PII detection for text moderation

Text translation

Customizable text translation

Contextual language understanding

Knowledge

Q&A extraction from unstructured text

Knowledge base creation from collections of Q&As

Semantic matching for knowledge bases

Customizable content personalization learning

Search

Ad-free web, news, image, and video search results

Trends for video, news

Image identification, classification and knowledge extraction

Identification of similar images and products

Named entity recognition and classification

Knowledge acquisition for named entities

Search query autosuggest

Ad-free custom search engine creation

Text-to-Speech service overview

Standard voices	49 languages/locales, REST APIs, SDKs, containers
Neural voices	4 languages/locales GA with more in preview, REST APIs, SDKs
Customization	9 languages/locales, REST APIs, portal

Microsoft is leading the industry with its Neural Text-to-Speech technology

The latest breakthrough that produces near human-parity digital voices

Concatenative

Human

Neural TTS

Form Recognizer

Azure Form Recognizer is a cognitive service that uses machine learning technology to identify and extract key/value pairs and table data from form documents. It then outputs structured data that includes the relationships in the original file.

Form Recognizer uses unsupervised learning to create the model.

Form Recognizer supports printed and handwritten forms, PDFs, and images.

Form Recognizer is available in English, with additional language availability growing.

The AI-powered document extraction service that understands your forms

Form Recognizer

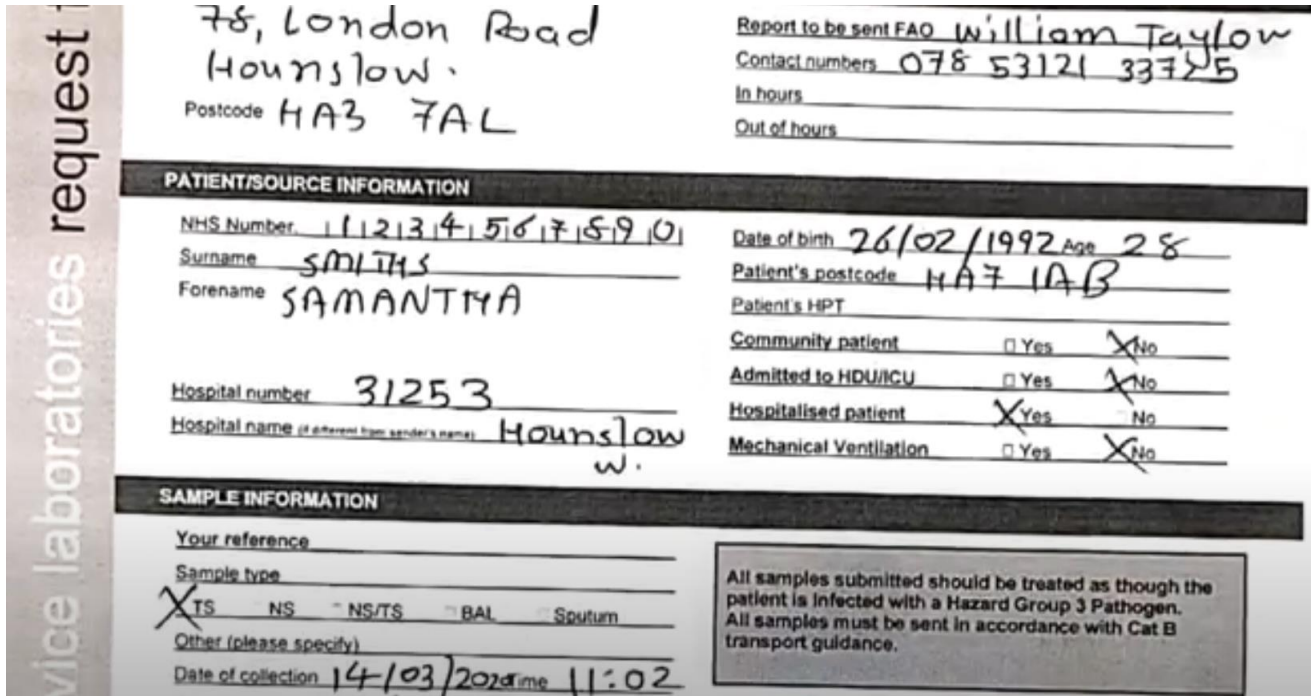

Customize extraction to your forms

Cloud to Edge

Privacy & Security

Cognitive Search

AI Builder

ZELROS

AUTOMATION ANYWHERE

UiPath

Chevron

blueprism

FINANCIAL FABRIC

Capabilities that extract printed or handwritten text from images

Optical Character Recognition (OCR)

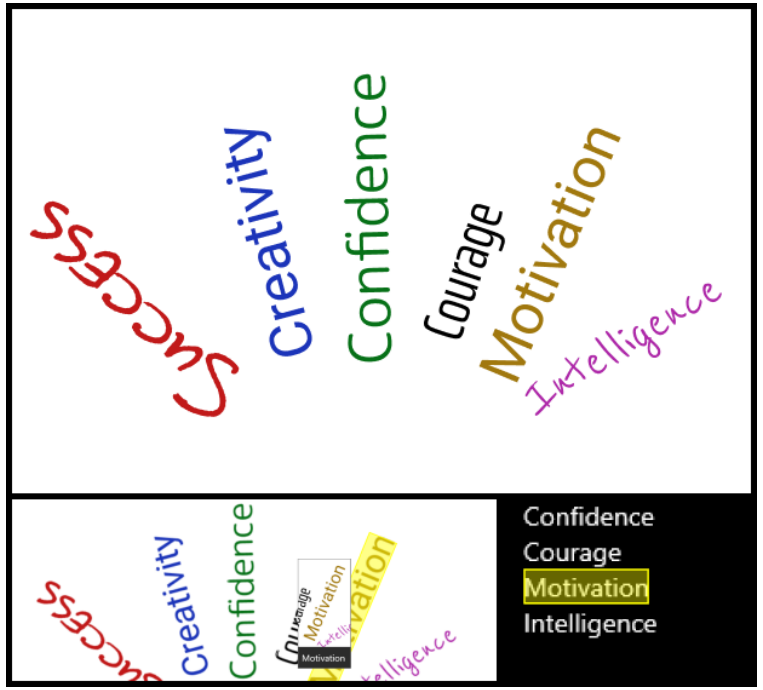

Text from images

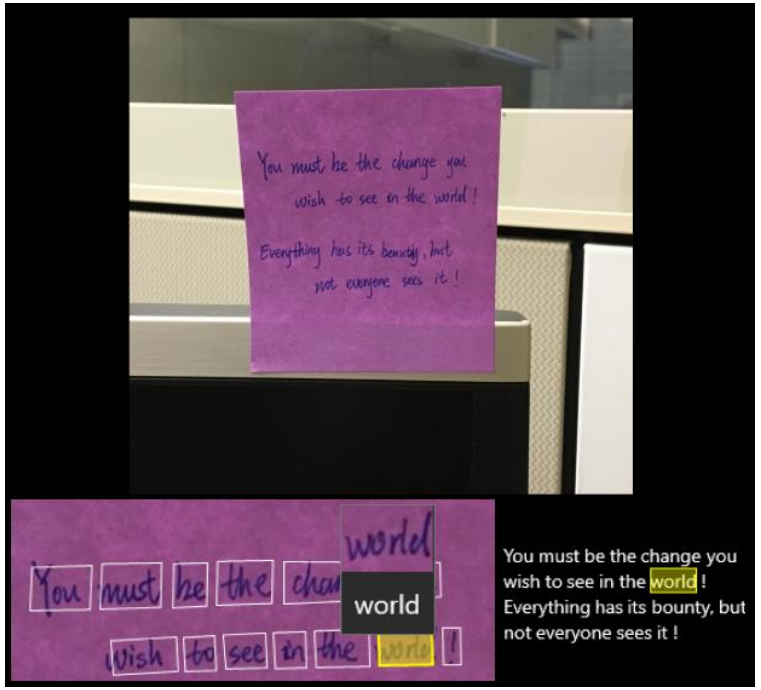

Handwritten text

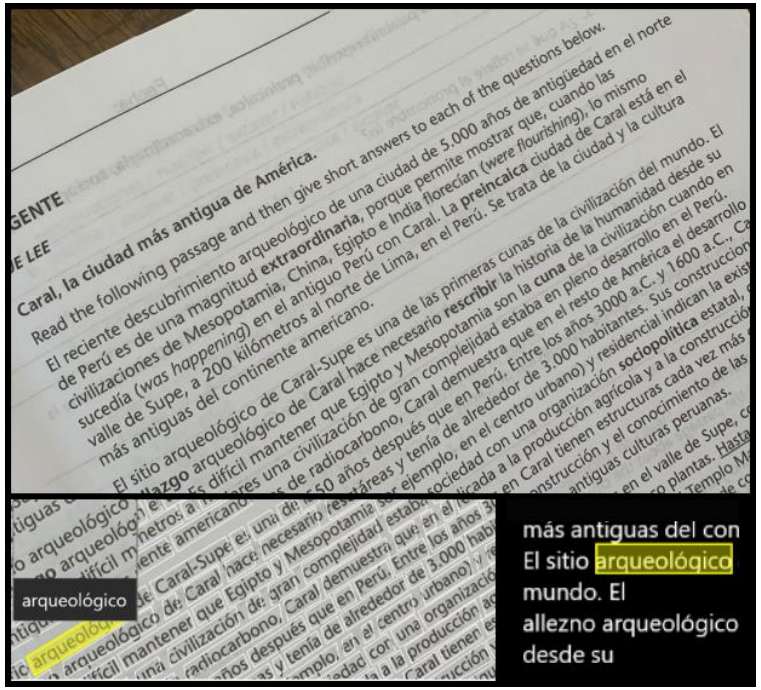

Printed text

READ API

Language Understanding (LUIS)

- Designed to identify valuable information in conversations, LUIS interprets user goals (intents) and distills valuable information from sentences (entities), for a high quality, nuanced language model. LUIS integrates seamlessly with the Azure Bot Service, making it easy to create a sophisticated bot
- Active learning is used to continuously improve the quality of the natural language models. Once the model starts processing input, LUIS begins active learning, allowing you to constantly update and improve the model.

Container support in Azure Cognitive Services

Container support in Azure Cognitive Services allows developers to use the same rich APIs that are available in Azure, and enables flexibility in where to deploy and host the services that come with Docker containers. Container support is currently available for a subset of Azure Cognitive Services, including parts of:

- Anomaly Detector
- Computer Vision
- Face
- Form Recognizer
- Language Understanding (LUIS)
- Speech Service API
- Text Analytics

Cognitive Services container configurations of resources are controlled by customers, so Microsoft will not offer an SLA for general availability (GA).

Speech Services

Human Parity (Speech & Translation)

45 Languages

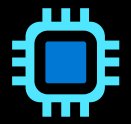

6 Indic Languages

English, Hindi/ hinglish, Tamil, Telugu, Gujrati, Marathi

Speech cognitive services

- Speech to Text
- Transcribe audible speech into readable, searchable text.

- Text to Speech
- Convert text to lifelike speech for more natural interfaces.

- Speech Translation
- Integrate real-time speech translation into your apps.

- Speaker Recognition (PREVIEW)
- Identify and verify the people speaking based on audio.

Text to Speech

- Create delightful voice experience for your customers with Microsoft Text-to-Speech Service.

Text-to-Speech service overview

Standard voices	49 languages/locales, REST APIs, SDKs, containers
Neural voices	4 languages/locales GA with more in preview, REST APIs, SDKs
Customization	9 languages/locales, REST APIs, portal

Speech Services (STT, TTS)

Call analytics (both post call and real-time)

- Unlock the information contained in every customer call by automatically transcribing and analyzing calls
- Extract: sentiment, entities (people, products, etc.), key phrases, call duration, trending topics, customer opinions, etc.
- Agent Assist (IVR Bots)

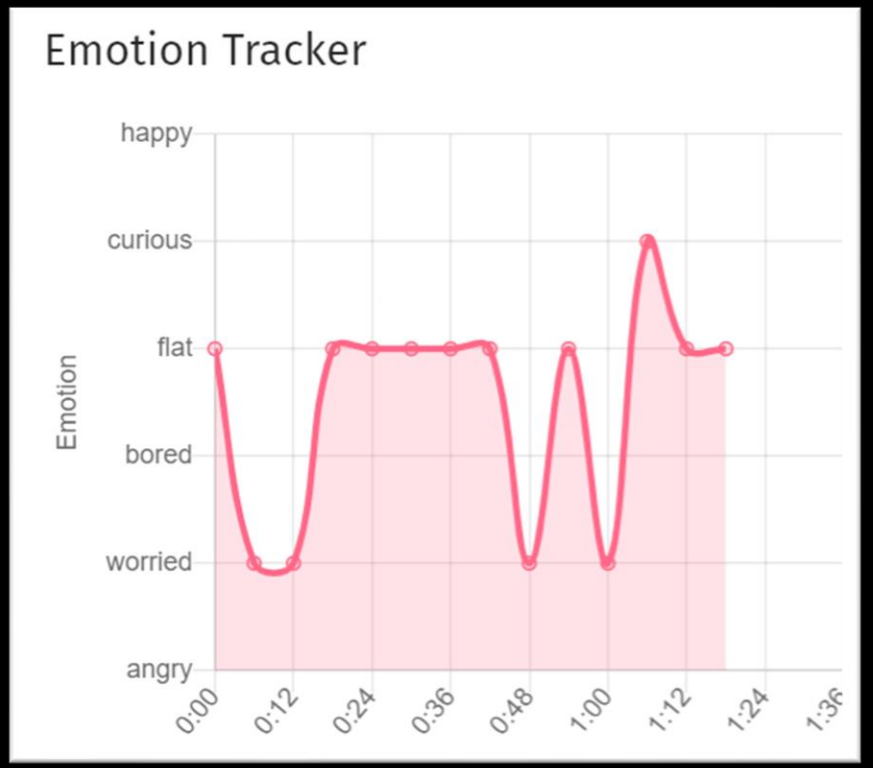

Conversational Voice Agents

Challenge

Interacting with a bot can be inaccurate, impersonal, and robotic

Solution

Speak to users naturally while improving accessibility to a global audience so customers can interact with a brand on their terms

Azure AI Services

Neural Text to Speech

Direct Line Speech

Keyword Spotter

LUIS, BOT Framework, ASR

Impact

Increased satisfaction by providing real-time customer service while freeing up employee time to focus on tasks requiring human intervention

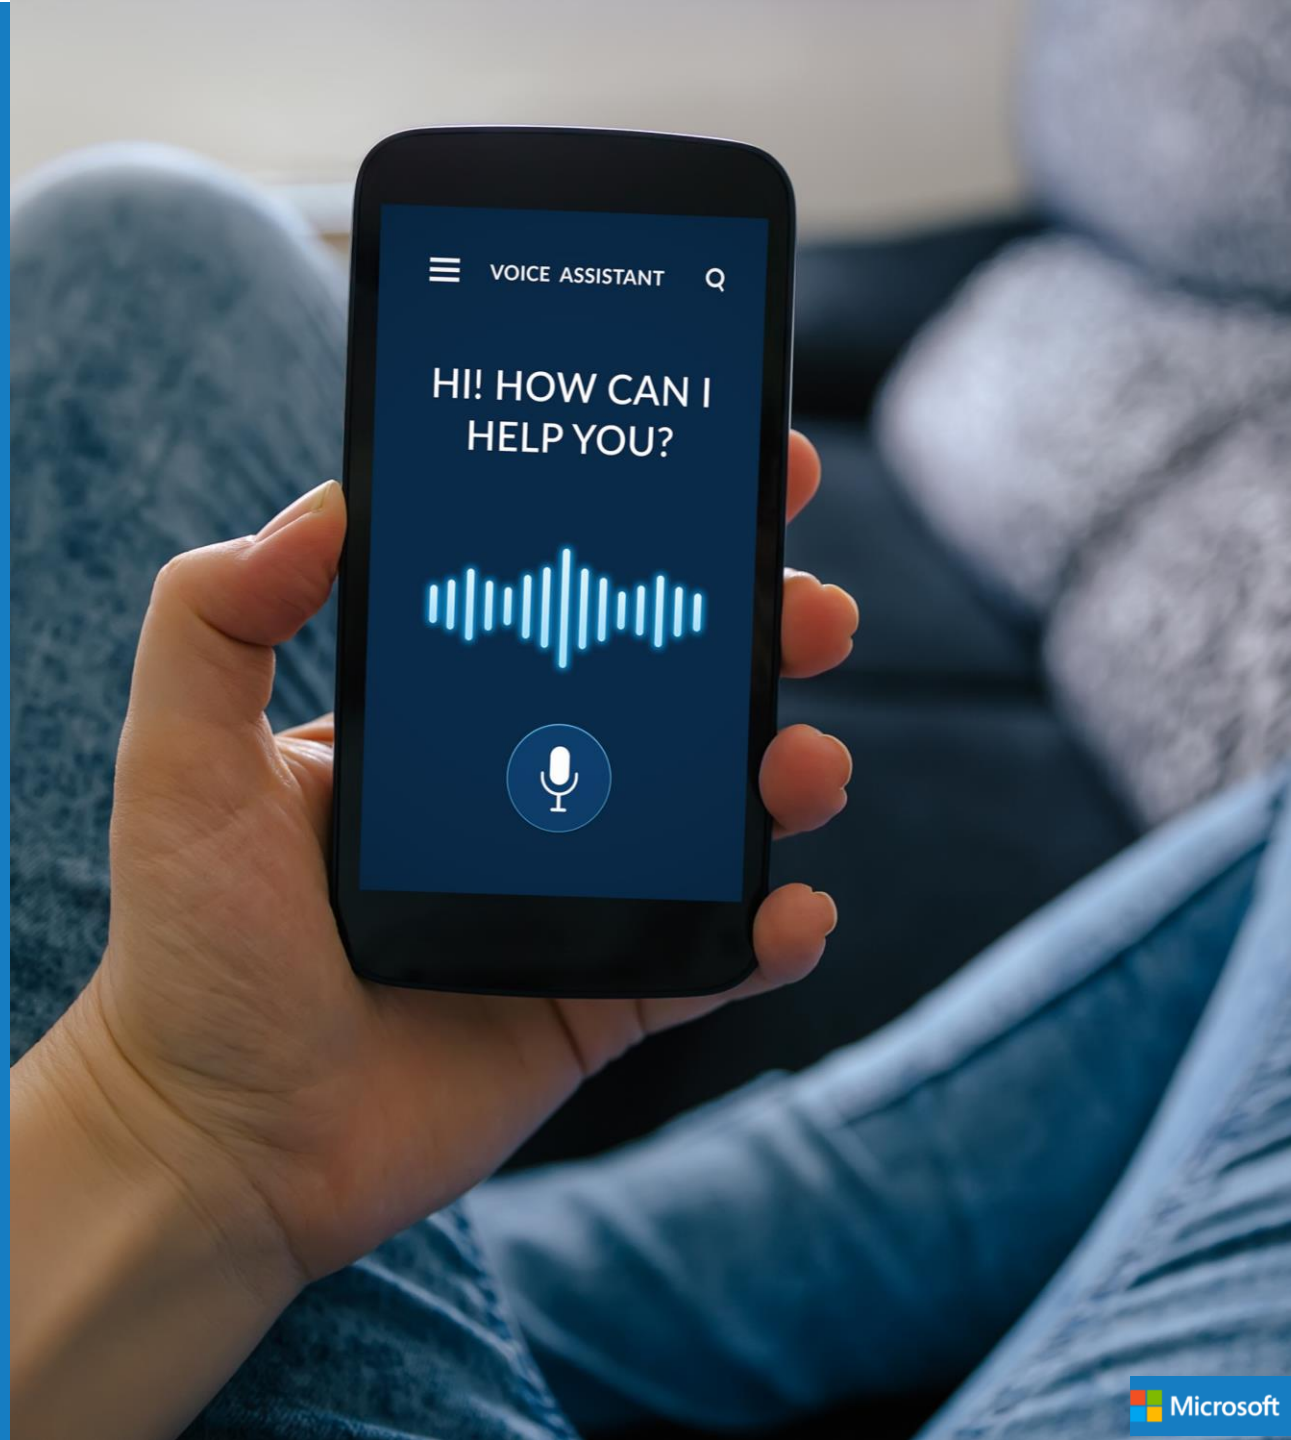

Bot Frame-work

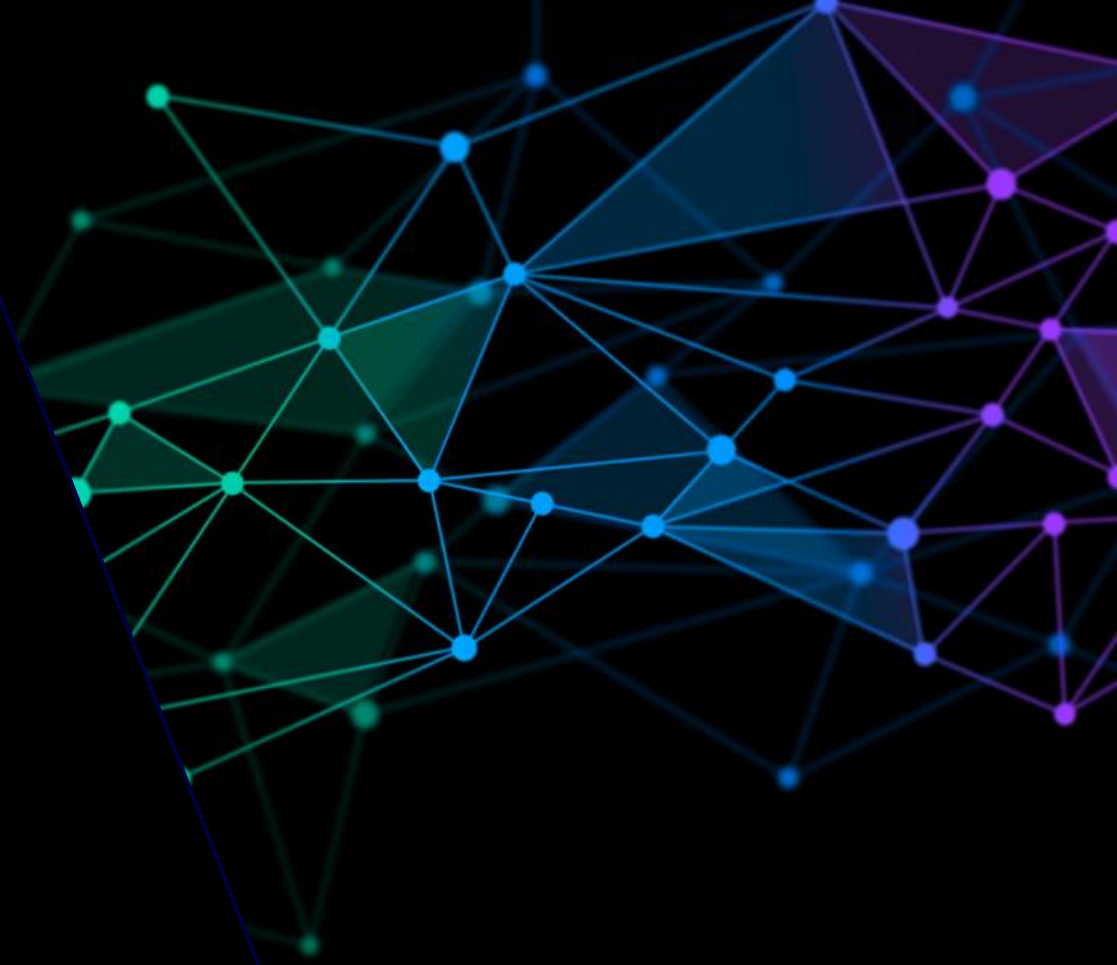

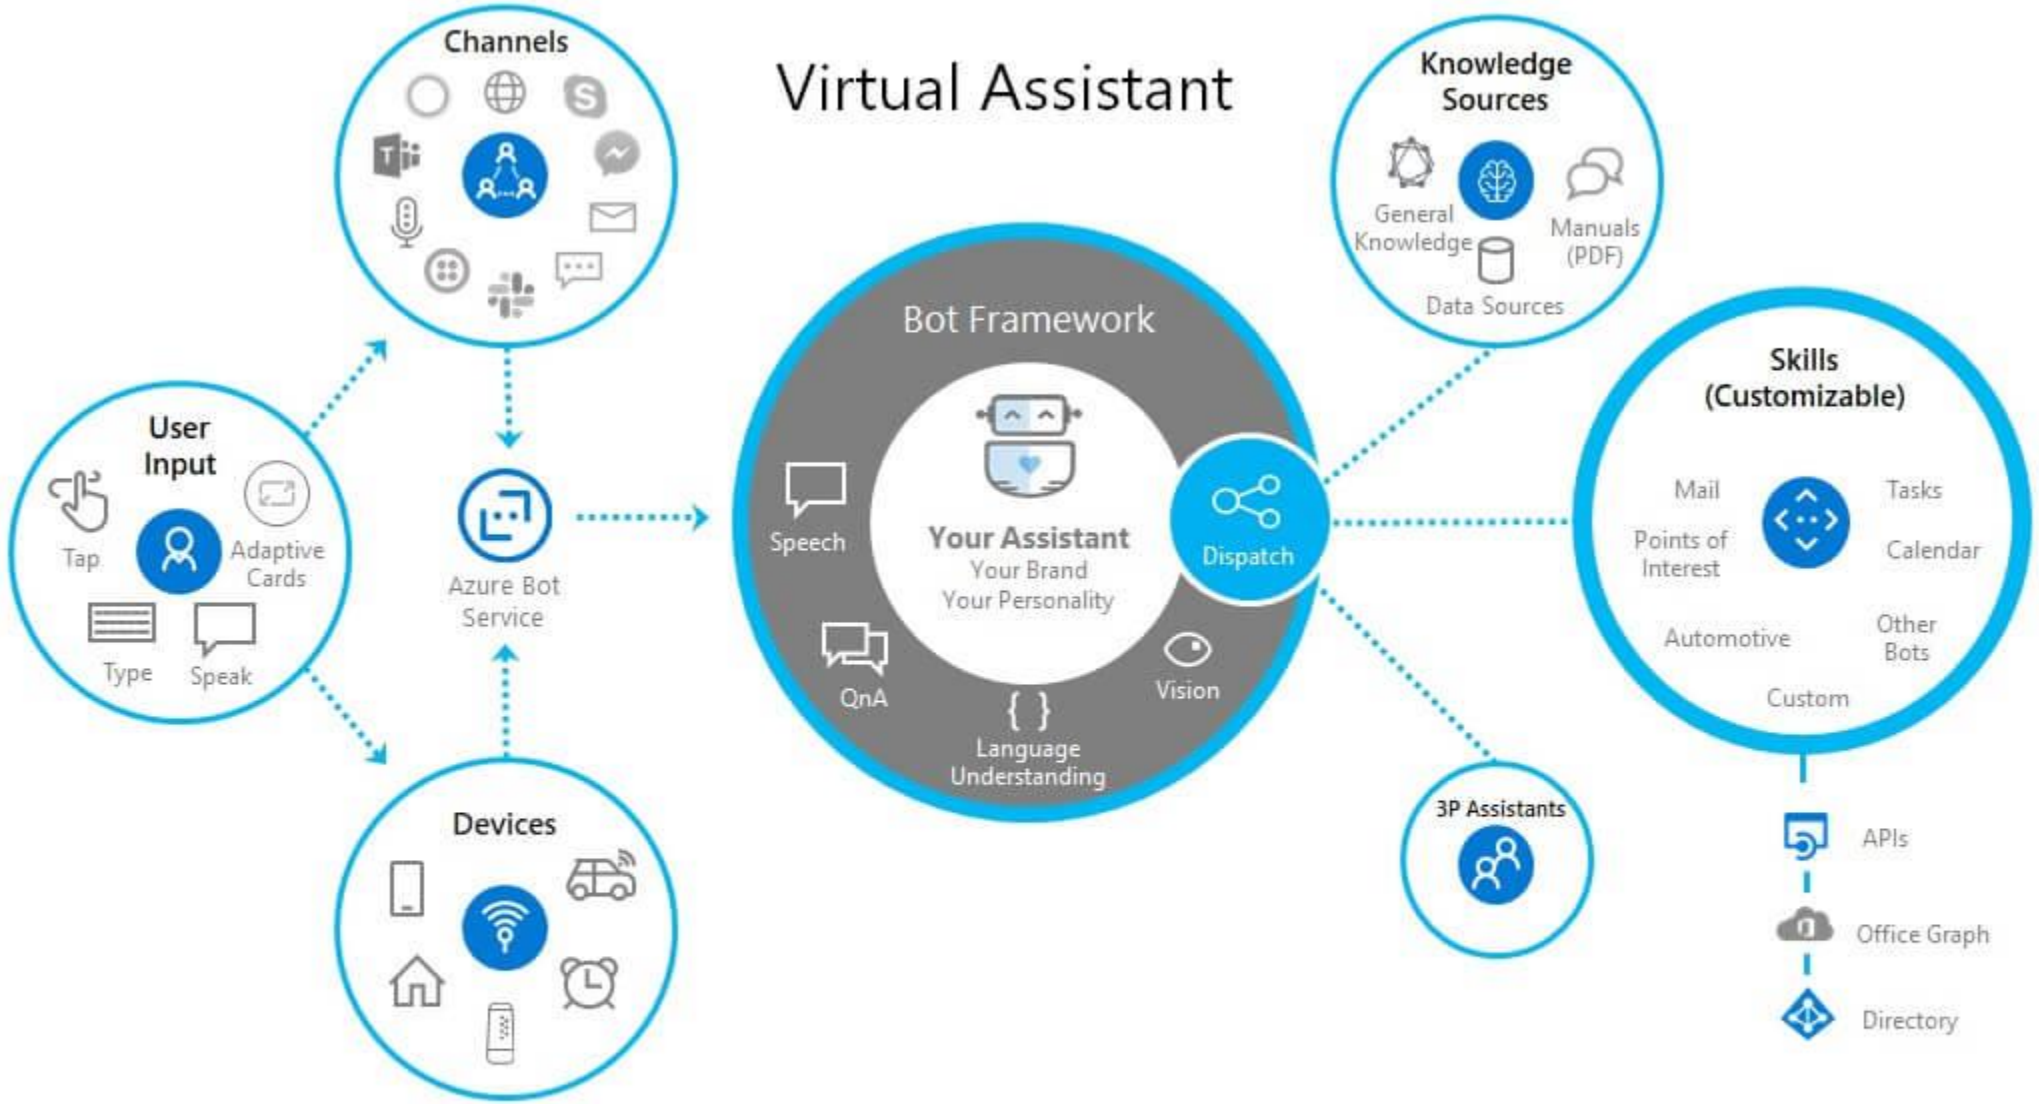
Virtual Assistant
Channels
User Input
Tap
Adaptive Cards
Type
Speak
Azure Bot Service
Devices
Bot Framework
Speech
Your Assistant
Your Brand
Your Personality
QnA
Language Understanding
Vision
Dispatch
Knowledge Sources
General Knowledge
Manuals (PDF)
Data Sources
Skills (Customizable)
Mail
Tasks
Points of Interest
Calendar
Automotive
Other Bots
Custom
3P Assistants
APIs
Office Graph
Directory

Kinds of bots

1,000+ companies engaging us

Scenario	Retail	Finance	Insurance	Telecoms	Government	Automotive	Manufacturing	Healthcare	Media	Events
Customer service	✓	✓	✓	✓	✓	✓	✓	✓		✓
Customer retail	✓	✓	✓	✓				✓		
Audio/speech analysis	✓	✓	✓	✓	✓				✓	
Translation	✓	✓	✓		✓					
Surveillance		✓			✓					
Knowledge extraction	✓	✓	✓	✓			✓			
Video/photo analysis		✓			✓				✓	
Product identification	✓						✓	✓		
Digital assistant			✓			✓		✓		
Footfall analysis	✓									✓
HD maps and object detection	✓					✓				














LINE

Information Chatbot

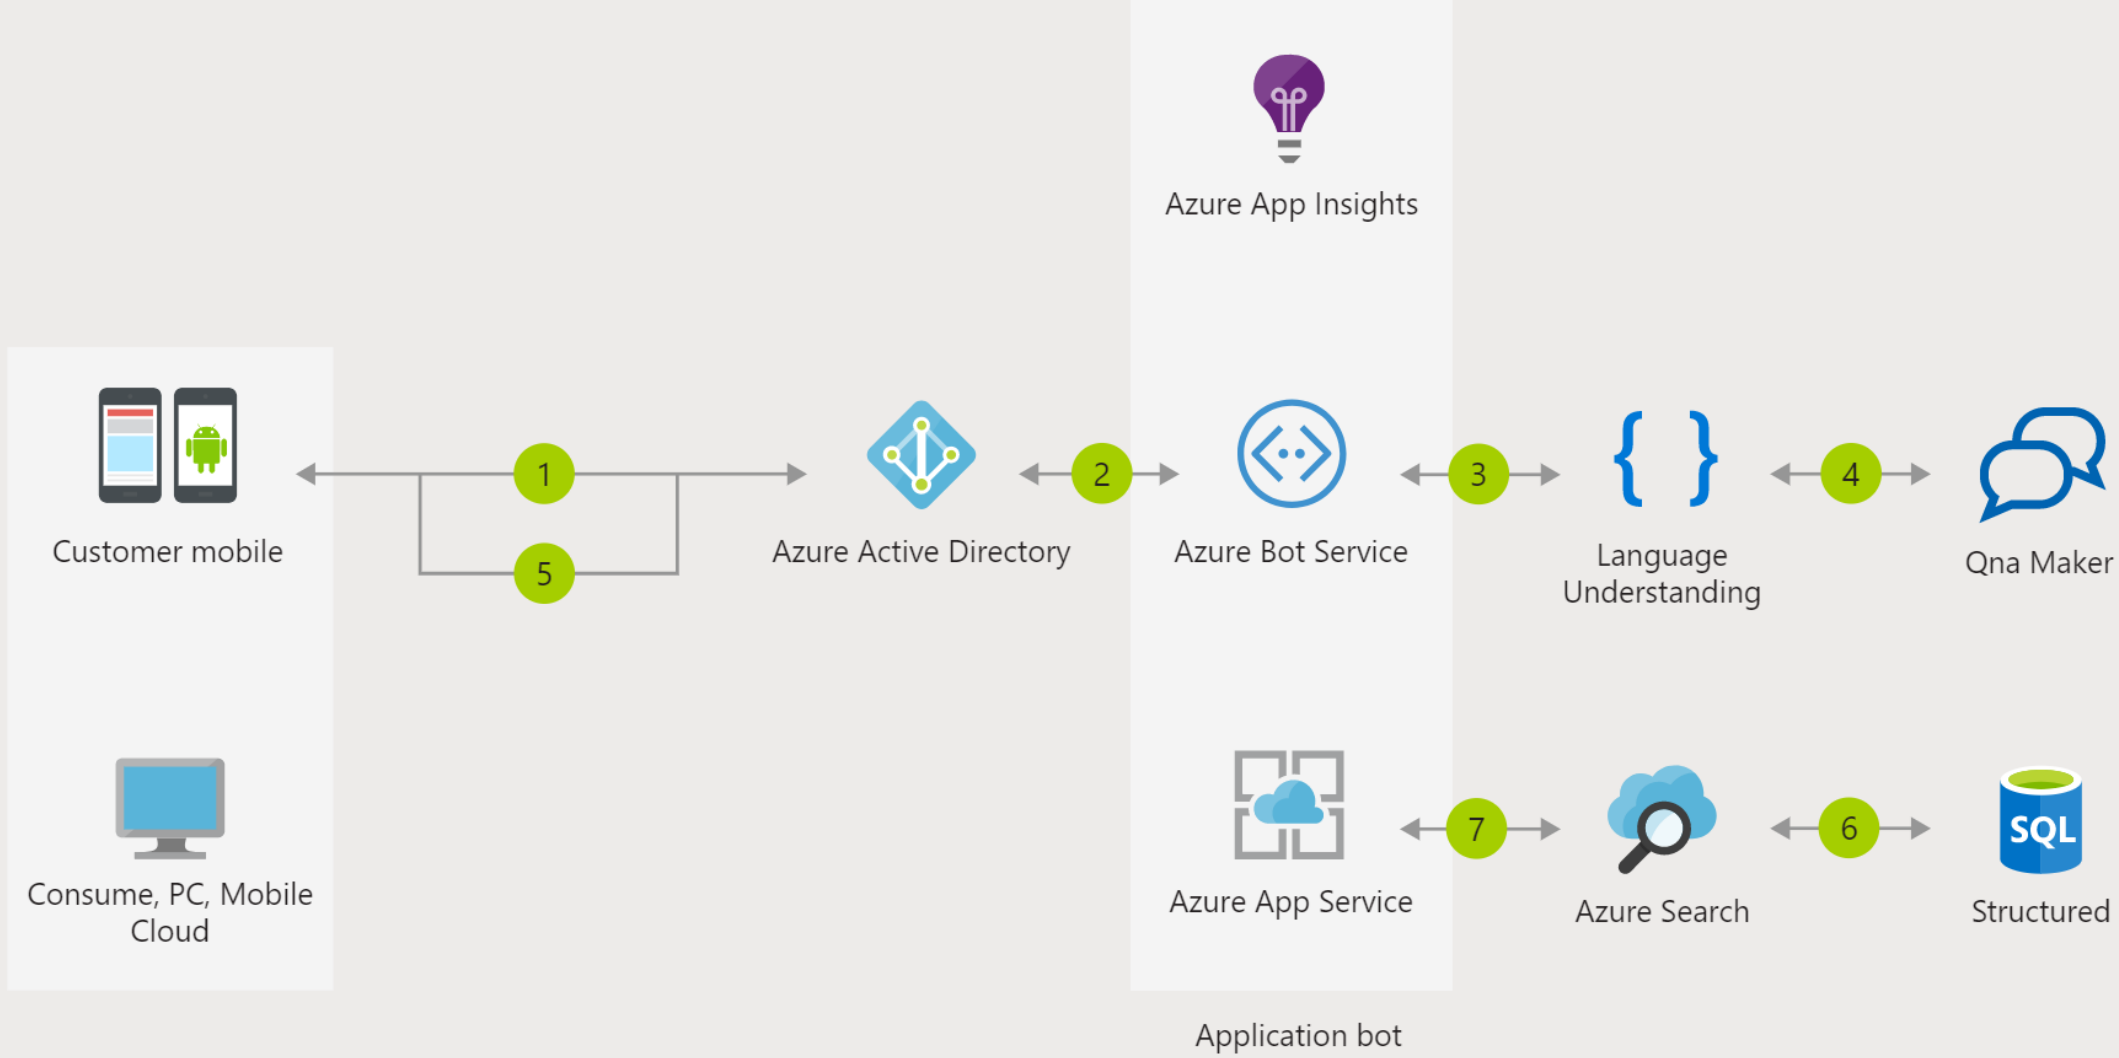

1. Employee starts the Application Bot
2. Azure Active Directory validates the employee's identity
3. The employee can ask the bot what type of queries are supported
4. Cognitive Services returns a FAQ built with the QnA Maker
5. The employee defines a valid query
6. The Bot submits the query to Azure Search which returns information about the application data
7. Application insights gathers runtime telemetry to help development with Bot performance and usage

AI Driven Knowledge Mining

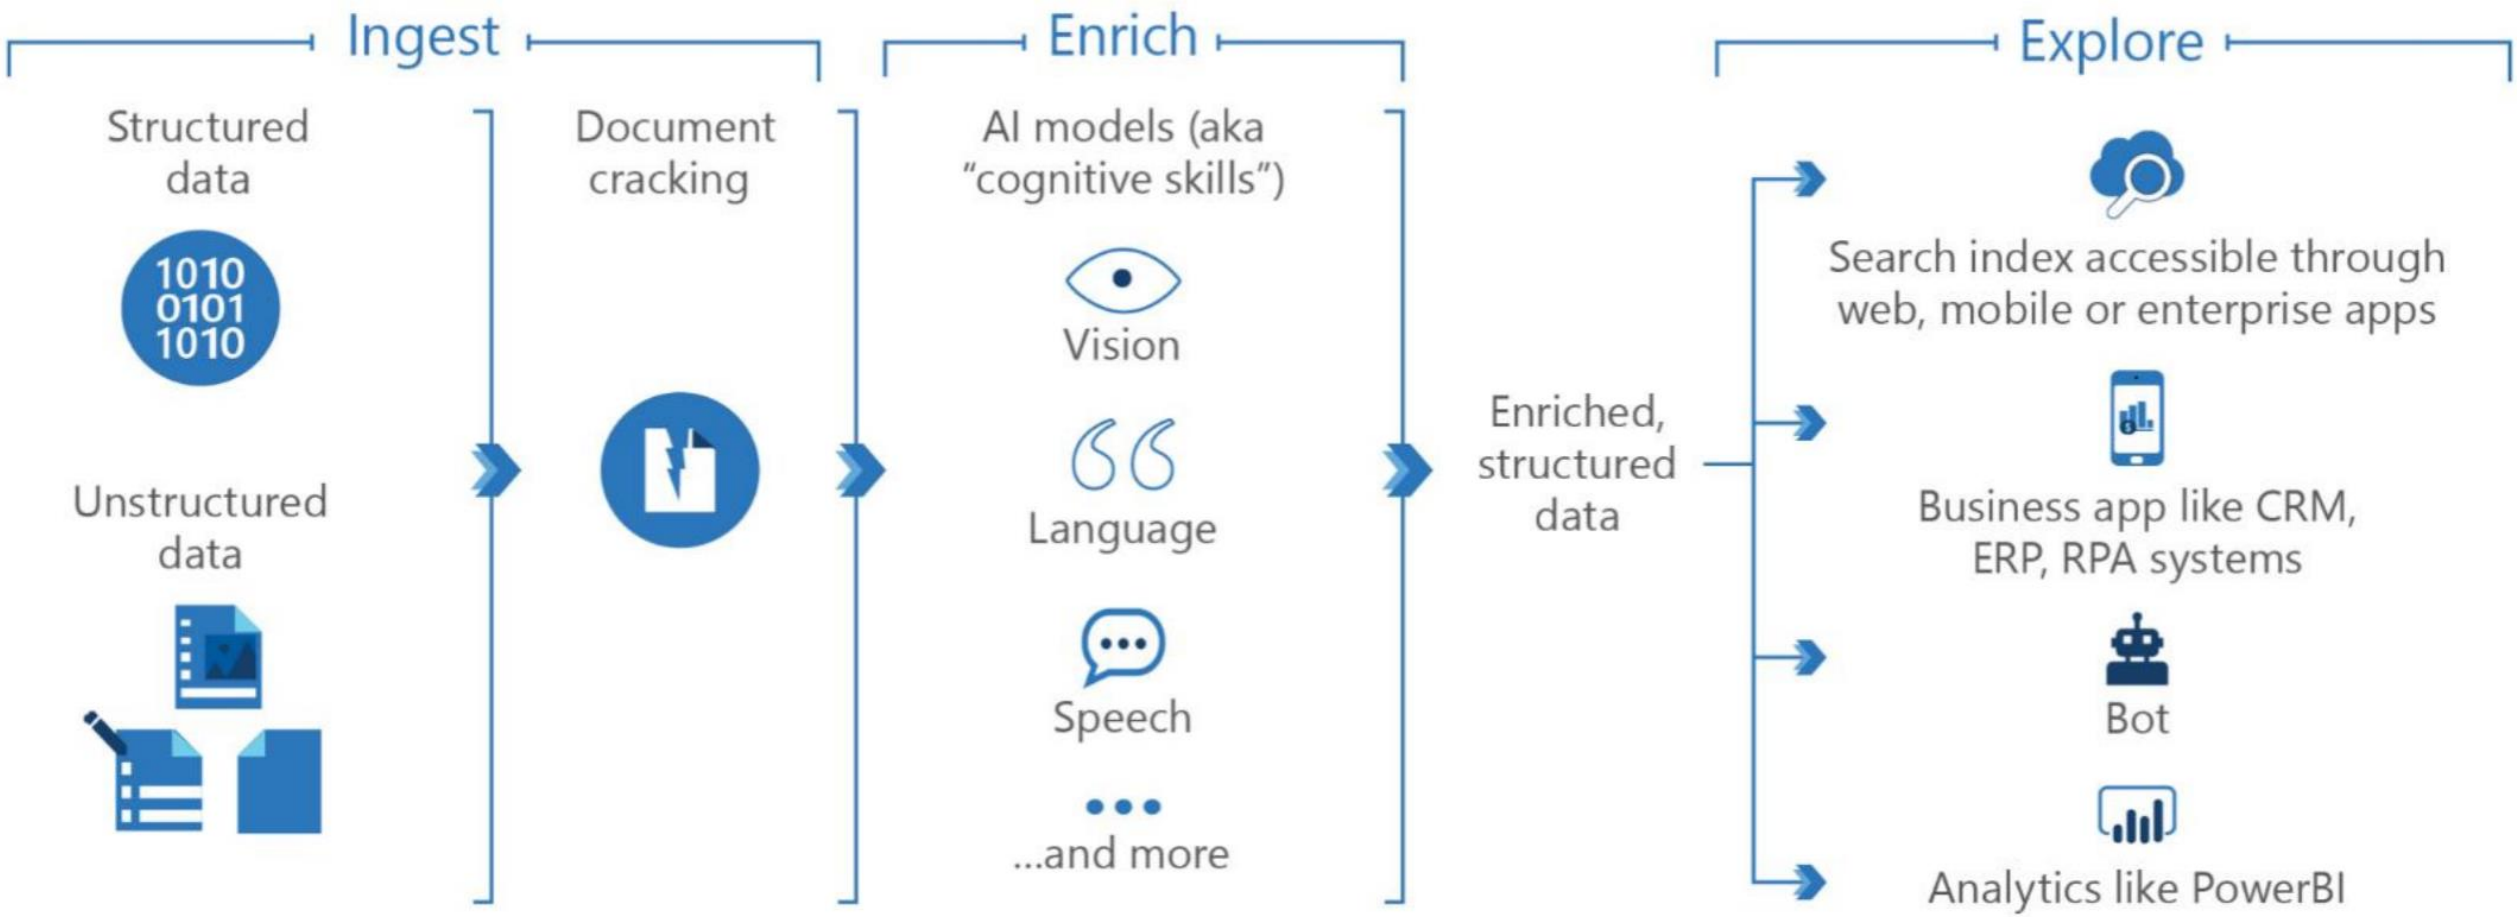

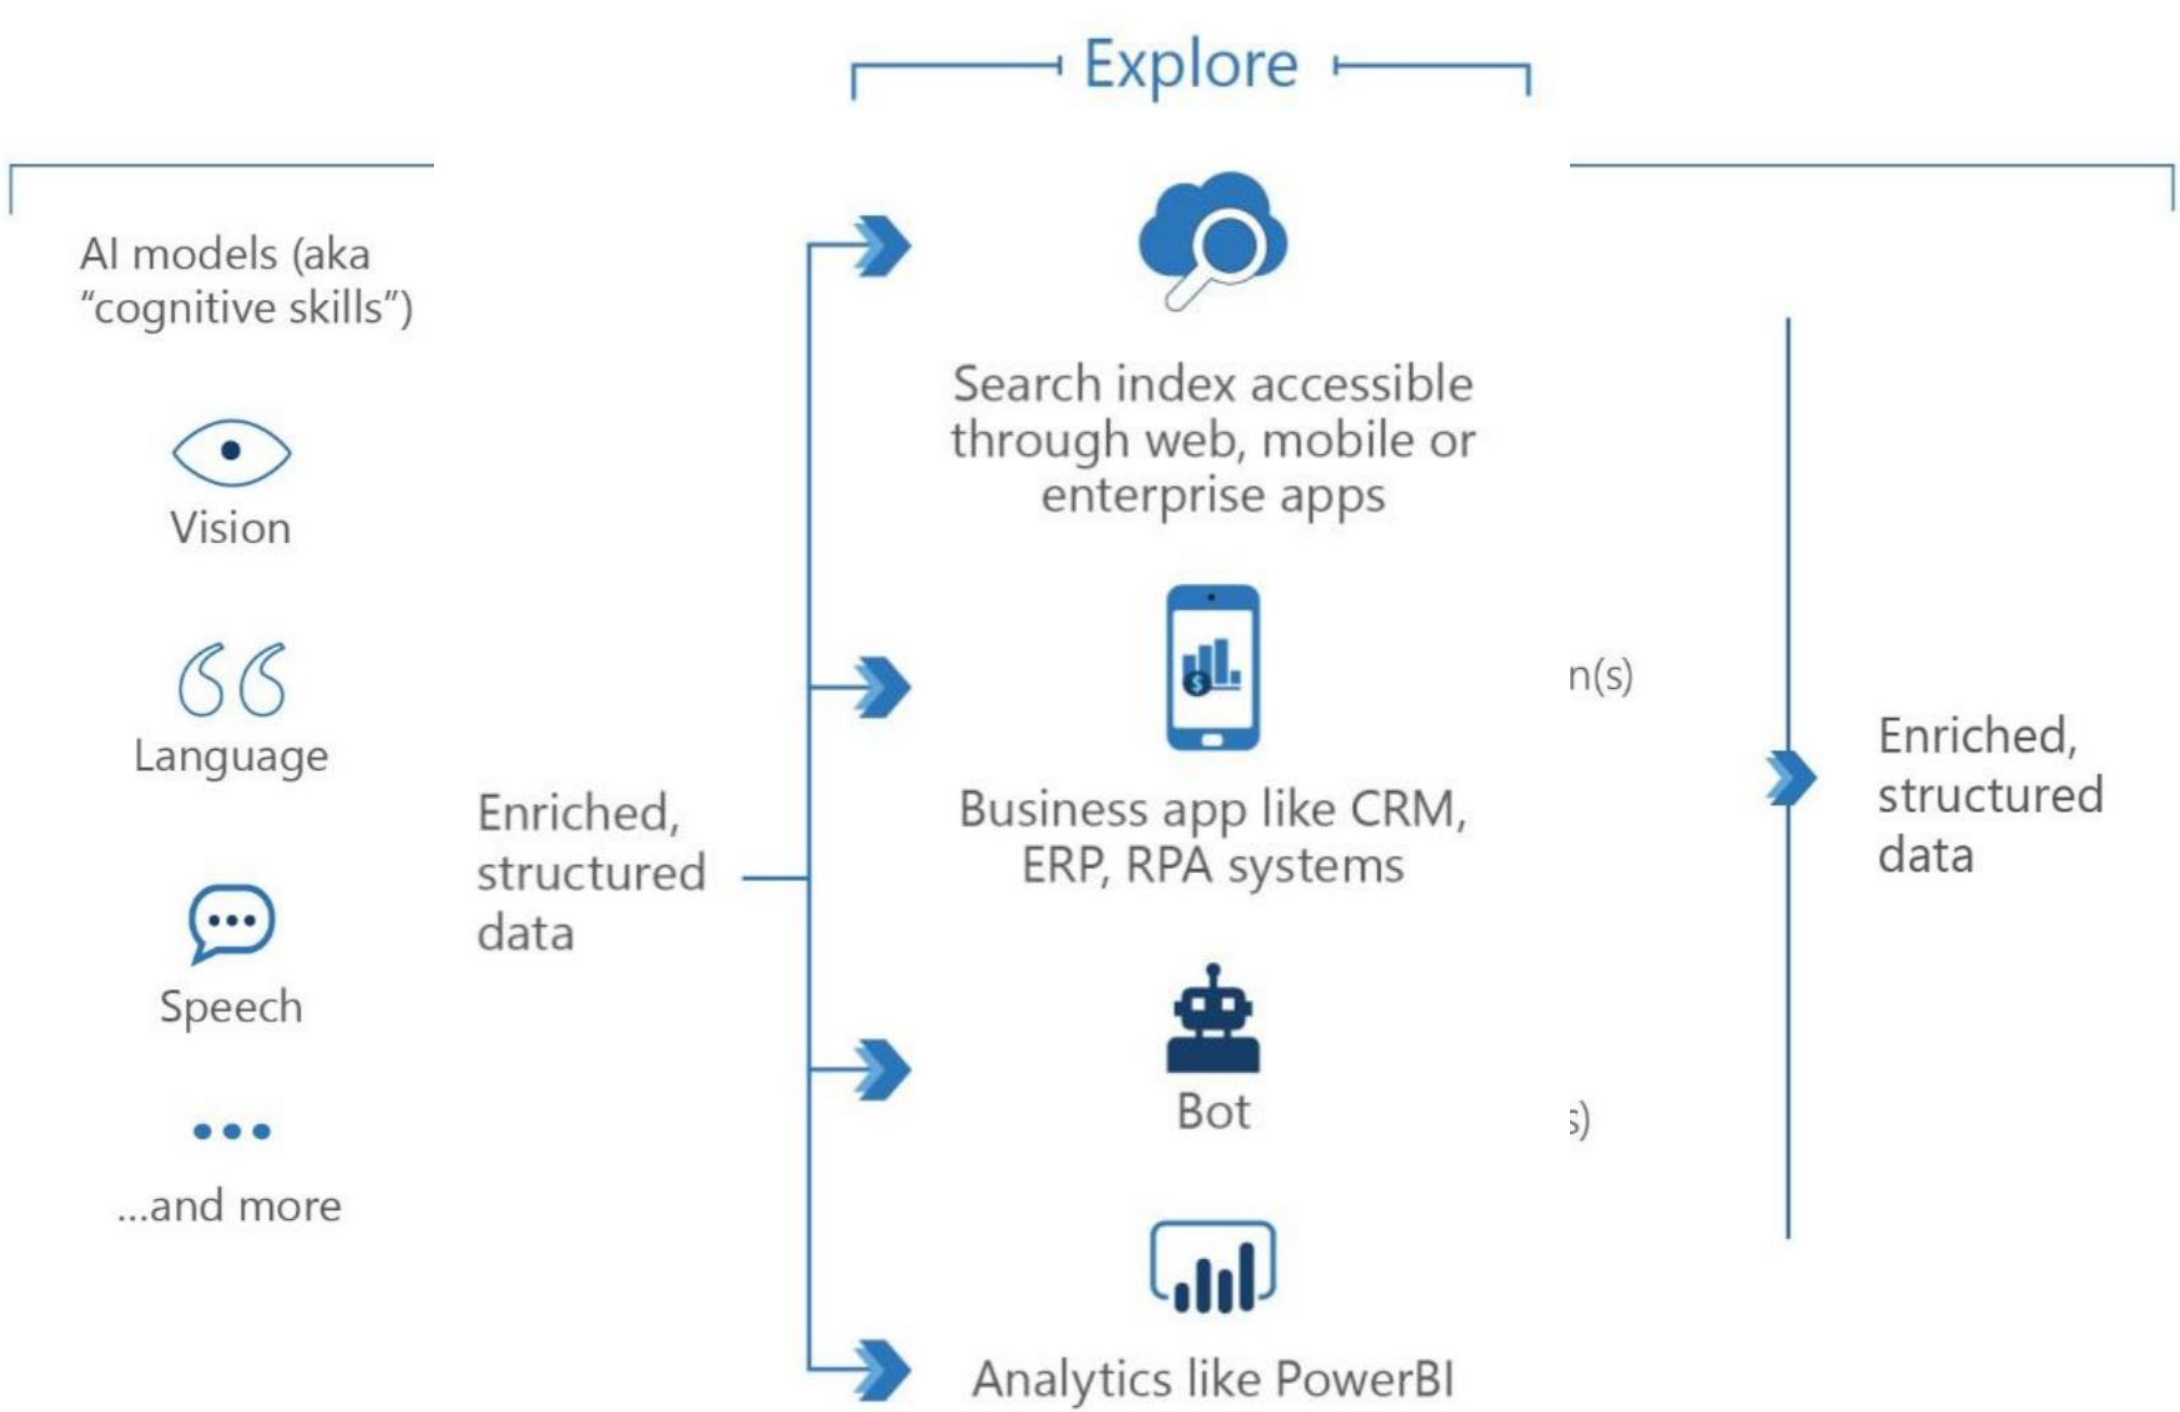
Explore
AI models (aka "cognitive skills")
Vision
Language
Speech
...and more
Enriched, structured data
Search index accessible through web, mobile or enterprise apps
Business app like CRM, ERP, RPA systems
Bot
Analytics like PowerBI
n(s)
s)
Enriched, structured data

Poll Two

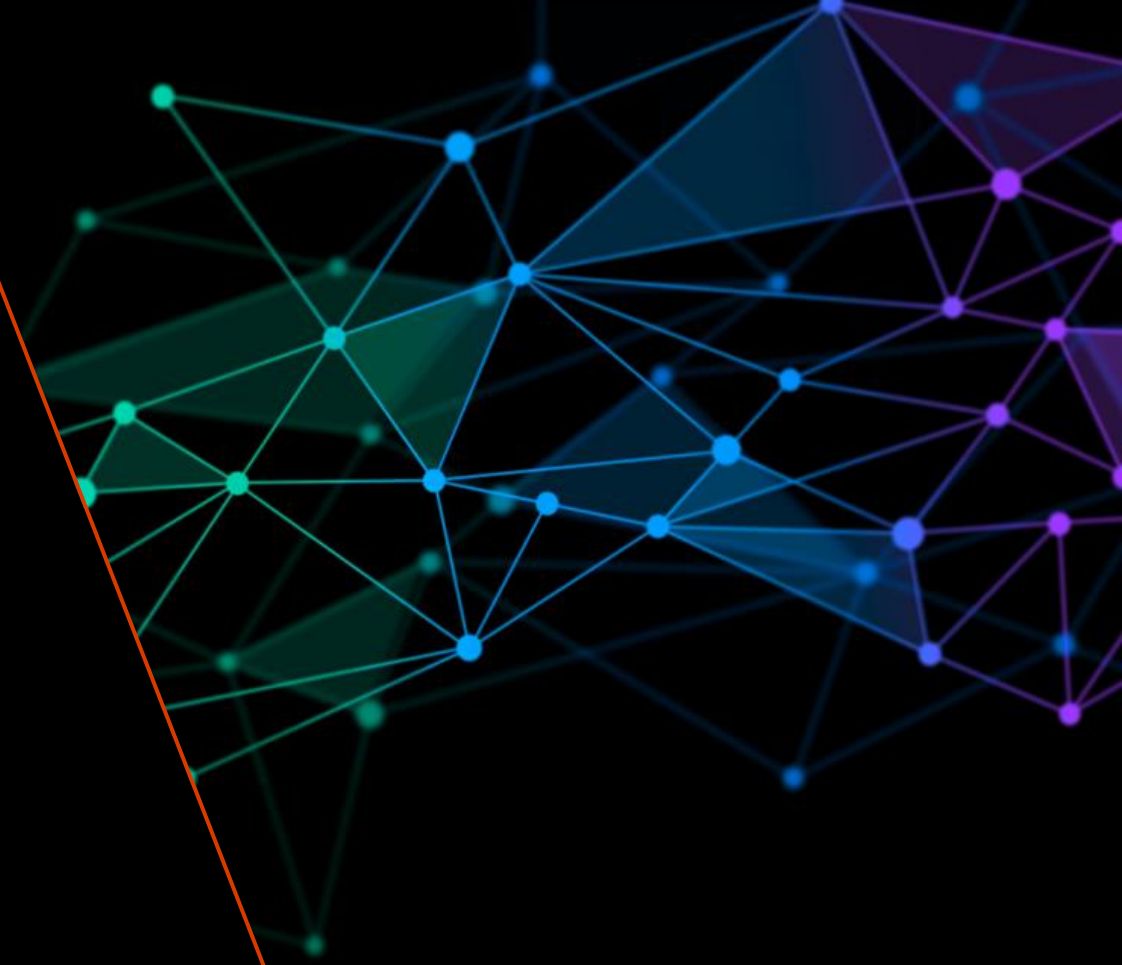

Machine Learning

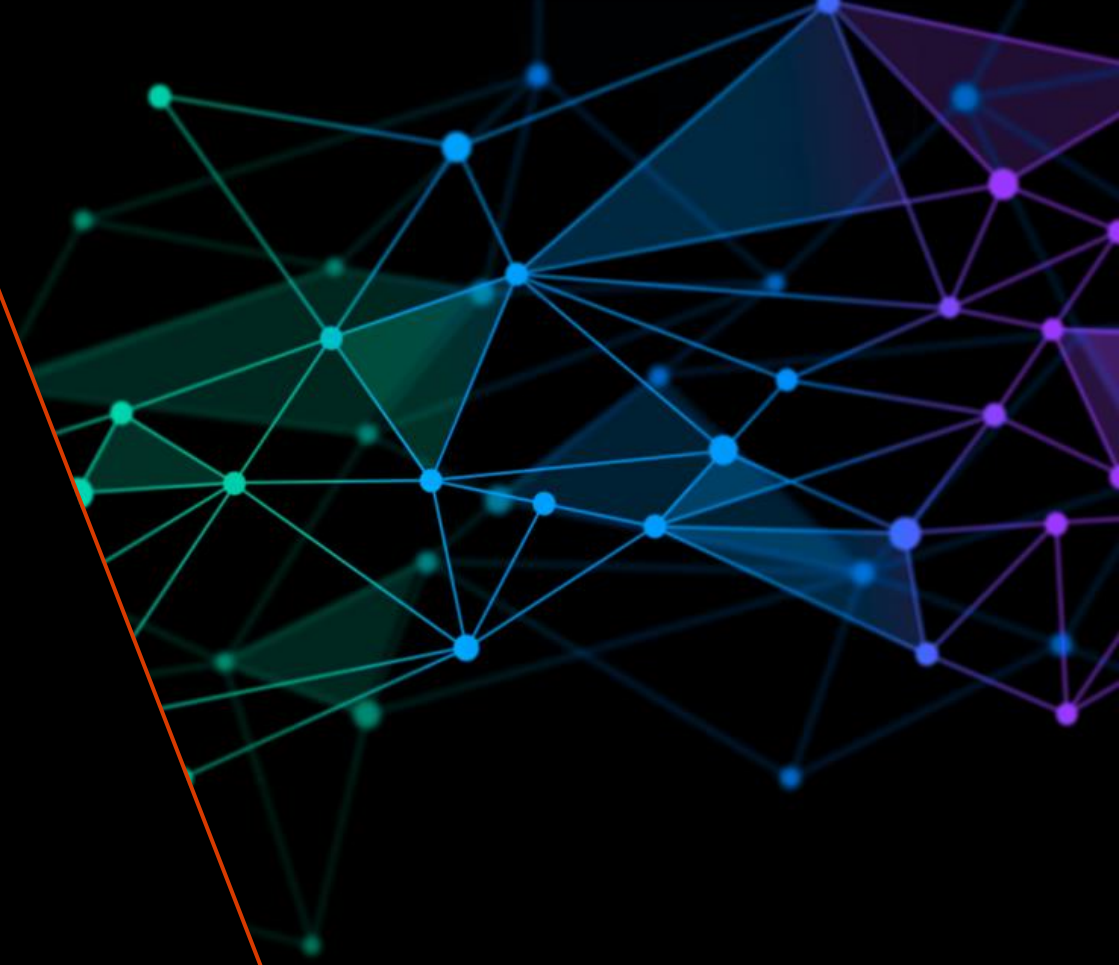

Azure AI

AI apps & agents

Azure Bot Service
Azure Cognitive Services

Machine learning

Azure Machine Learning

Knowledge mining

Azure Cognitive Search

Requirements of an advanced ML Platform

DevOps loop for data science

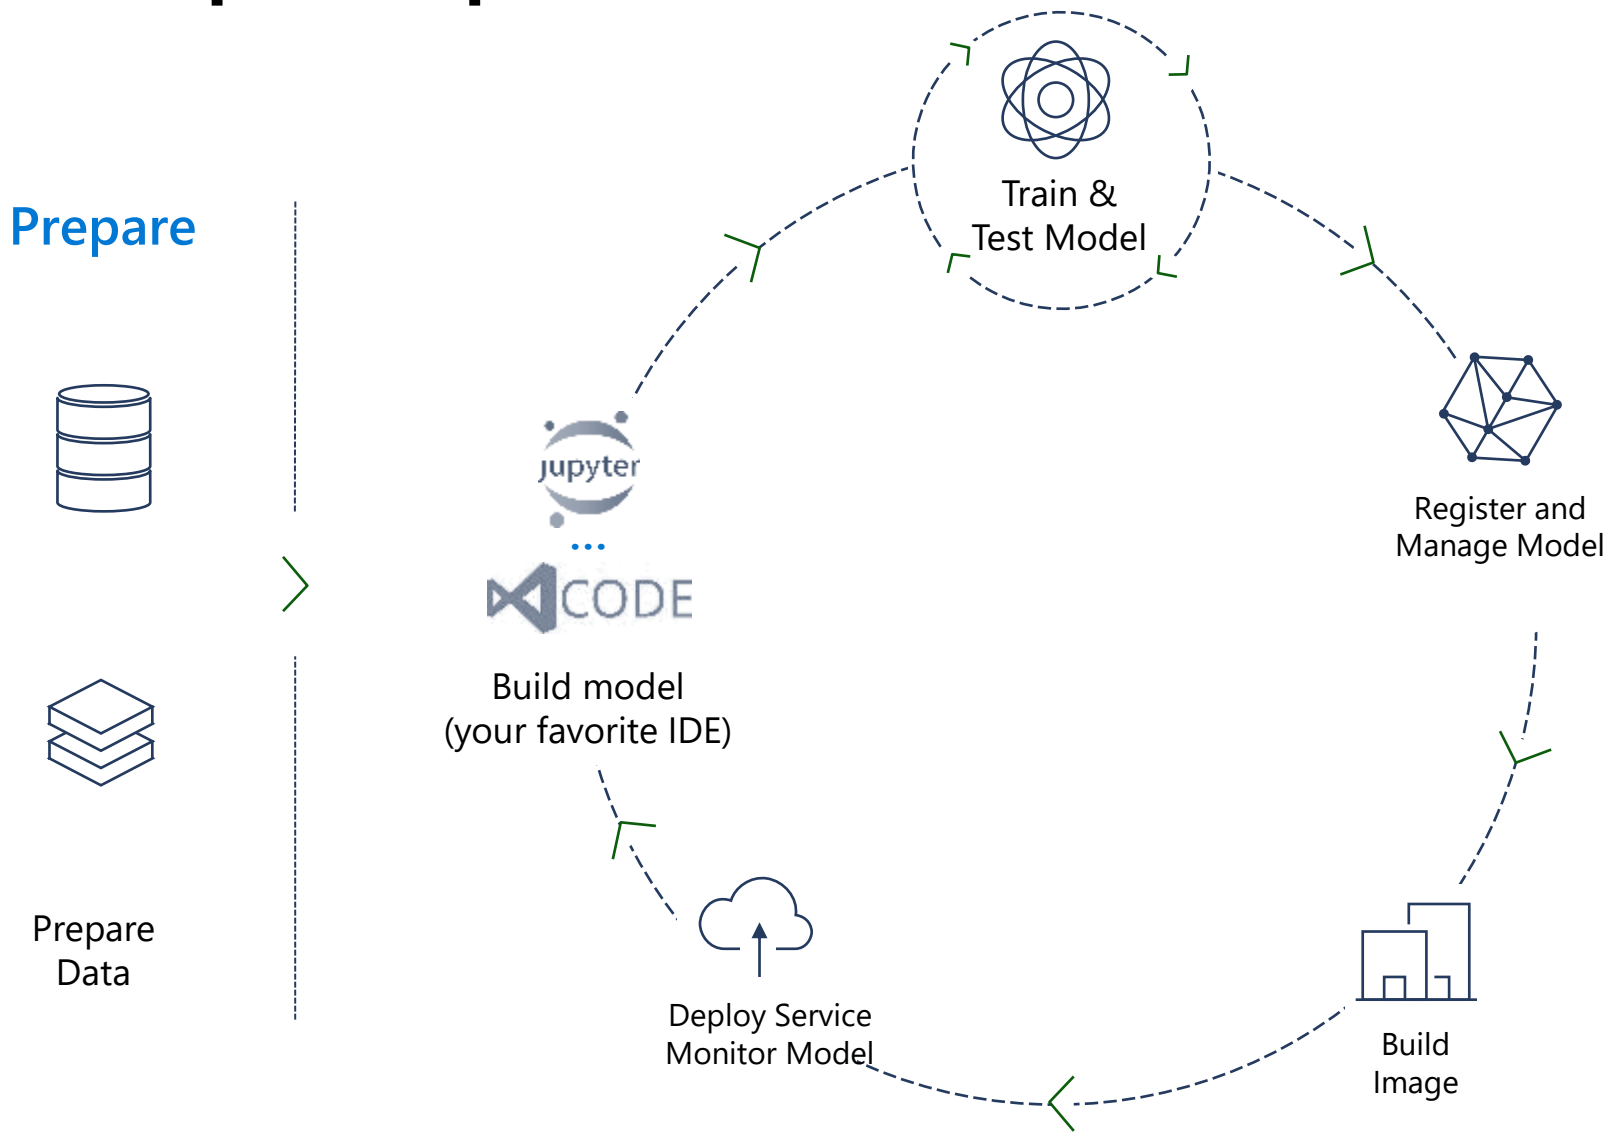

Data Preparation

Requirements

Multiple Data Sources

SQL and NoSQL databases, file systems, network attached storage and cloud stores (such as Azure Blob Storage) and HDFS.

Multiple Formats

Binary, text, CSV, TS, ARFF, etc.

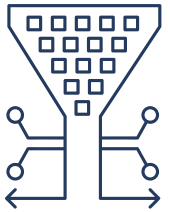

Cleansing

Detect and fix NULL values, outliers, out-of-range values, duplicate rows.

Transformation

General data transformation (transforming types) and ML-specific transformations (indexing, encoding, assembling into vectors, normalizing the vectors, binning, normalization and categorization).

Model Building

Requirements

Choice of algorithms

Choice of language

Python, R

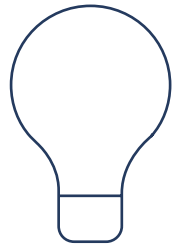

Choice of development tools

Browser-based, REPL-oriented, notebooks such as Jupyter, PyCharm and Spark Notebooks.

Desktop IDEs such as Visual Studio and R-Studio for R development.

Local Testing

To verify correctness before submitting to a more powerful (and expensive) training infrastructure.

Model Training

Requirements

Powerful Compute Environment

Choice should include scale-up VMs, auto-scaling scale-out clusters

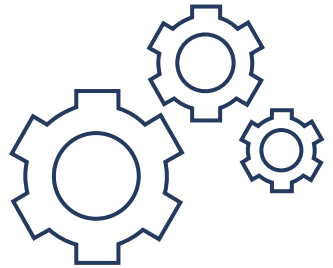

Preconfigured

The compute environment should be pre-setup with all the correct versions ML frameworks, libraries, executables and container images.

Job Management

Data scientists should be able to easily start, stop, monitor and manage Jobs.

Automated Model and Parameter Selection

Solution should automatically select the best algorithms, and the corresponding best hyperparameters, for the desired outcome.

Model Registration and Management

Requirements

Containerization

Automatically convert models to Docker containers so that they can be deployed into an execution environment.

Versioning

Assign versions numbers to models, to track changes over time, to identify and retrieve a specific version for deployment, for A/B testing, rolling back changes etc.

Model Repository

For storing and sharing models, to enable integration into CI/CD pipelines.

Track Experiments

For auditing, see changes over time and enable collaboration between team members.

Model Deployment

Requirements

Choice of Deployment Environments

Single VM, Cluster of VMs, Spark Clusters, Hadoop Clusters, In the cloud, On-premises

Edge Deployment

To enable predictions close to the event source-for quicker response and avoid unnecessary data transfer.

Security

Even when deployed at the edge, the e2e security must be maintained. Models should be deployed and data transmitted only to secure, authenticated devices.

Monitoring

Monitor the status, performance and security.

Machine Learning on Azure

Domain specific pretrained models
To reduce time to market


Vision


Speech


Language


Search

Familiar Data Science tools
To simplify model development


PyCharm


Jupyter


Visual Studio Code


Command line

Popular frameworks
To build advanced deep learning solutions


Pytorch


TensorFlow


Scikit-Learn


Onnx

Productive services
To empower data science and development teams


Azure Databricks


Azure Machine Learning


Machine Learning VMs

Powerful infrastructure
To accelerate deep learning


CPU


GPU


FPGA


From the Intelligent Cloud to the Intelligent Edge


Simplify machine learning for any skill level

Azure Machine Learning service

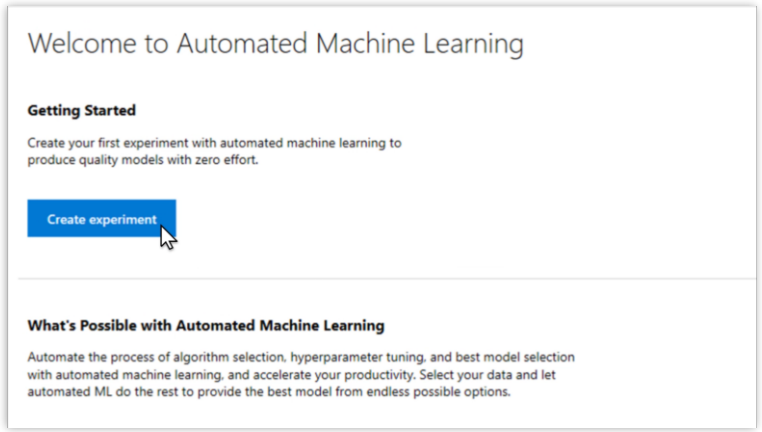

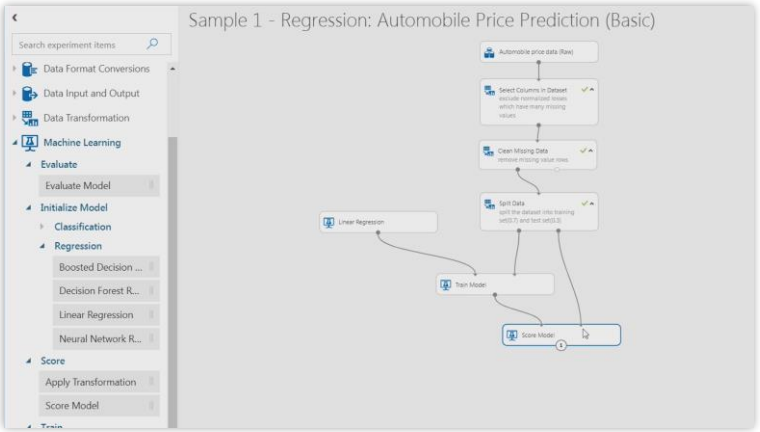

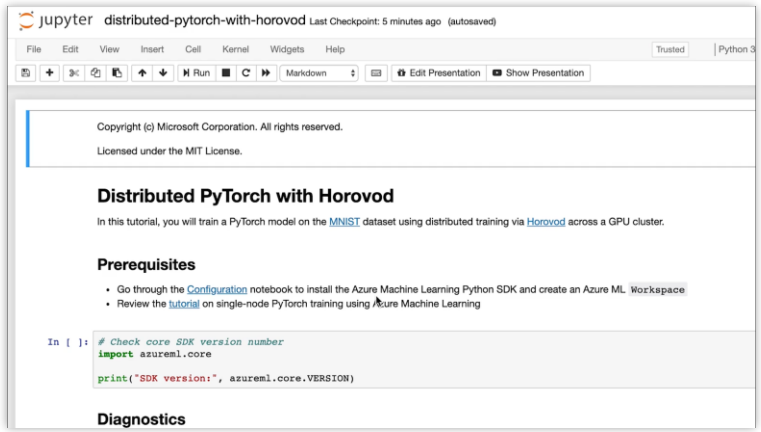

Knowledge Mining

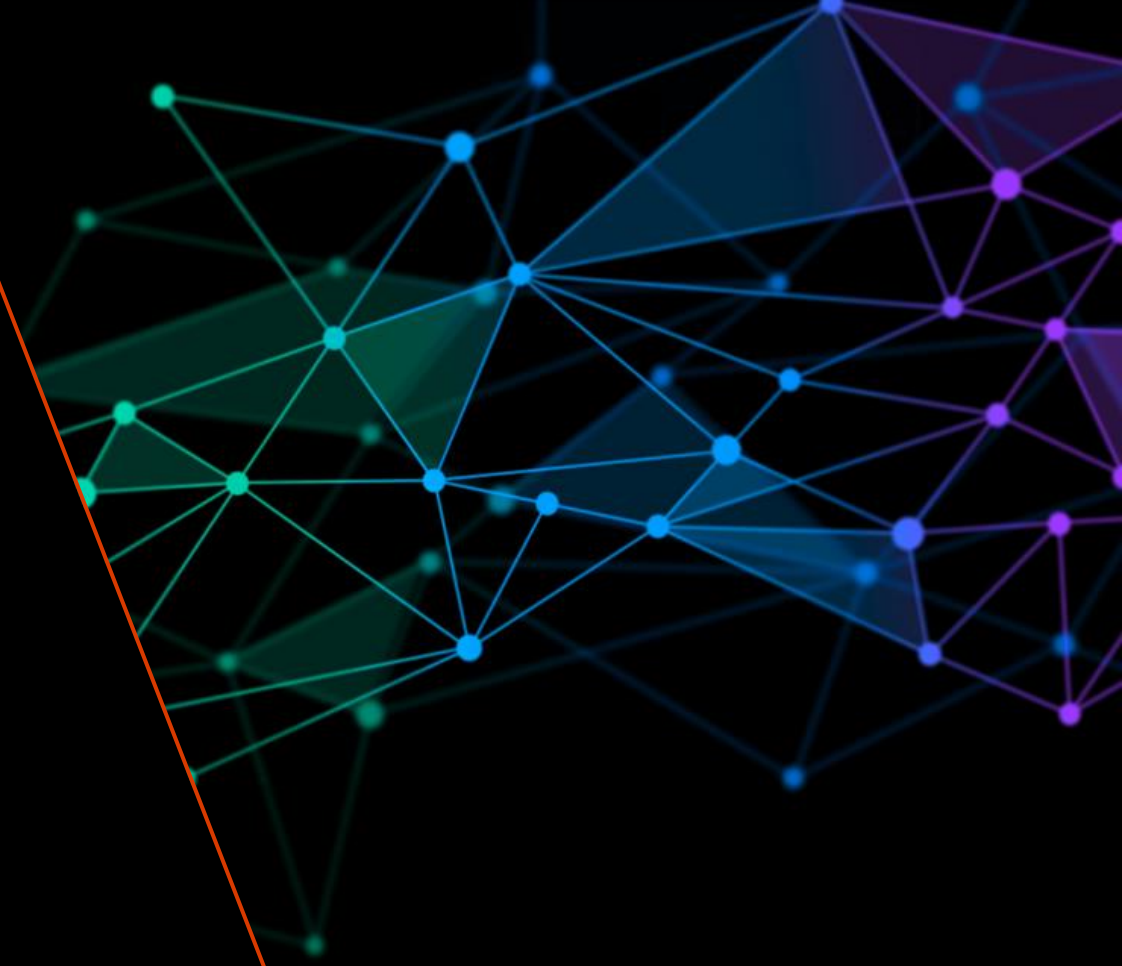

Understanding the latent value in all content

AI is here to help

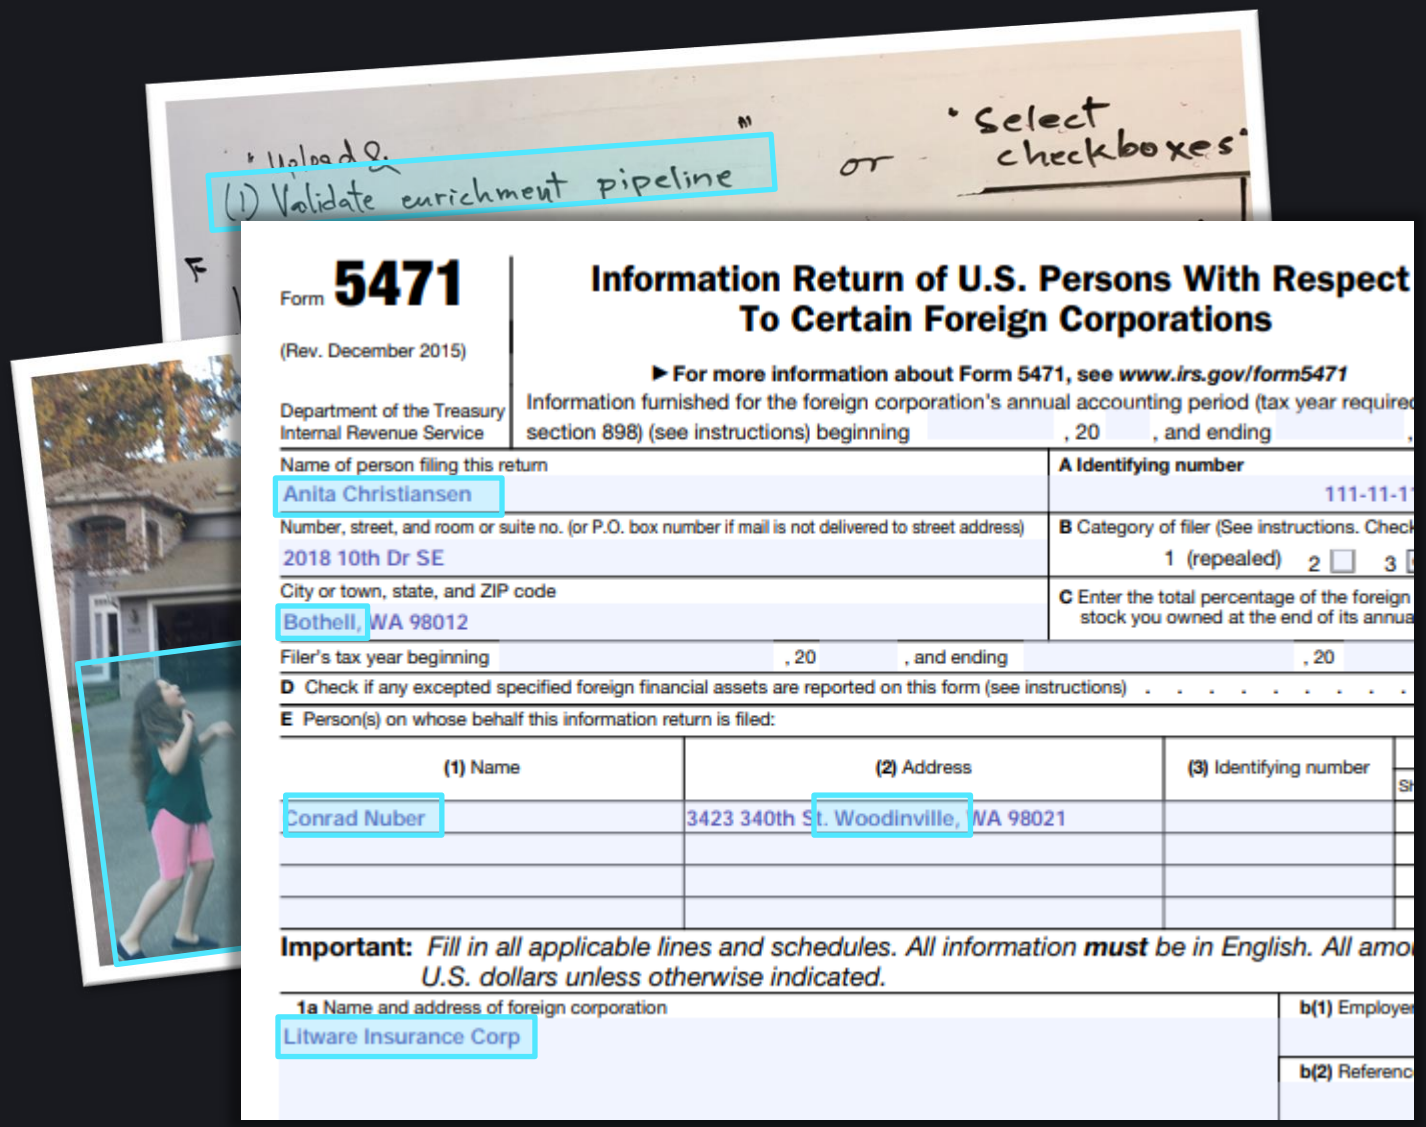

Text

(1) Validate enrichment pipeline

Tags

"throwing", "ball", "girl", "grass", "basketball"

Caption

"A girl throwing a ball"

Entities:

Person(s)

"Anita Christiansen", "Conrad Nuber",

Location(s)

"Bothell", "Woodinville"

Organization(s)

"Litware Insurance Corp."

Azure Cognitive Search
Search-as-a-Service

management free
keyword search
faceting
language analyzers
geospatial support
suggestions/auto-complete
customizable scoring
proximity search
synonyms
complex types
etc.

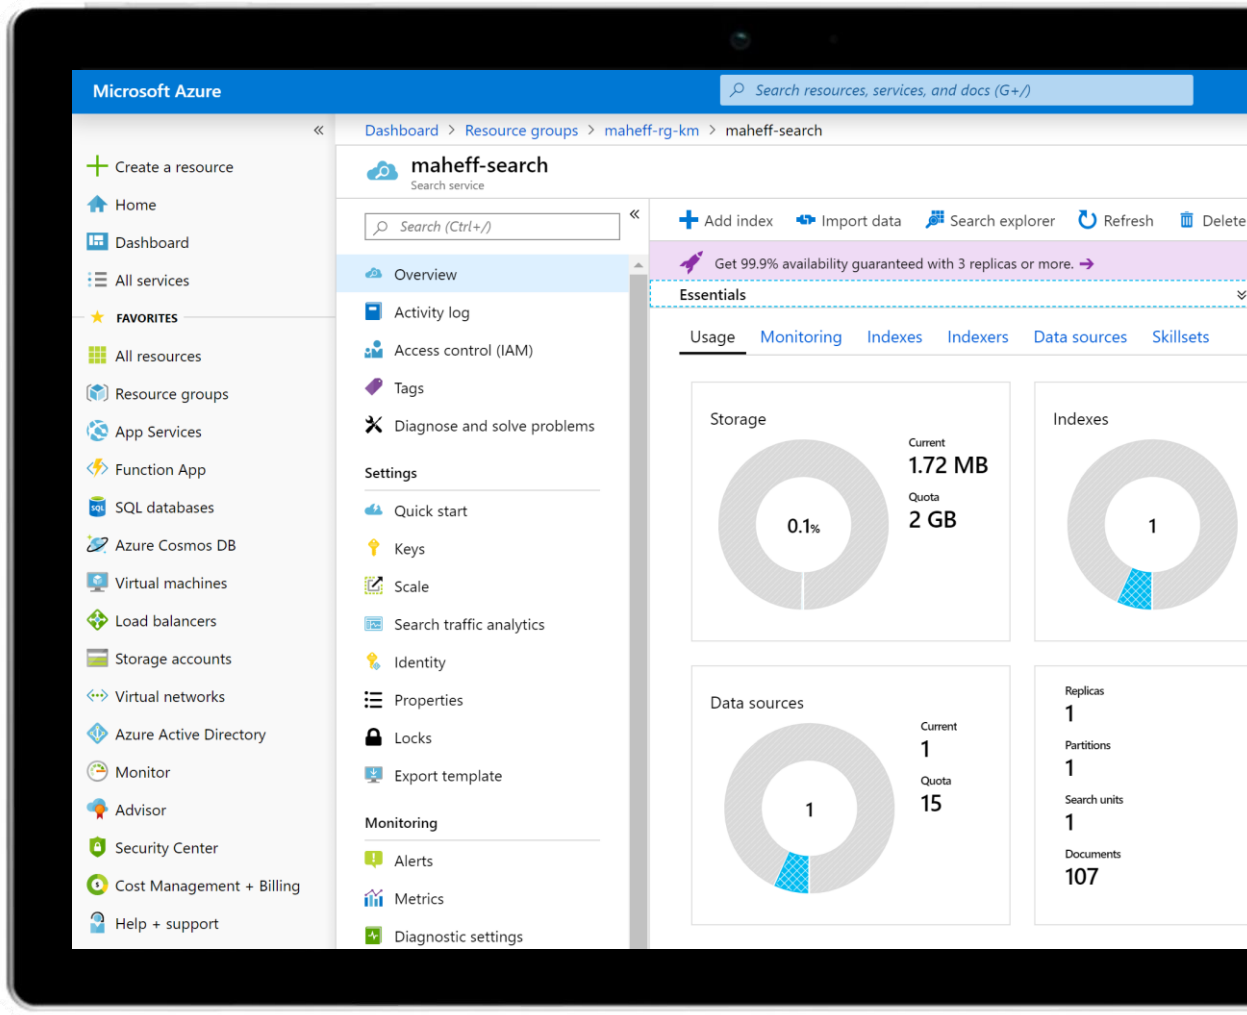

Azure Cognitive Search

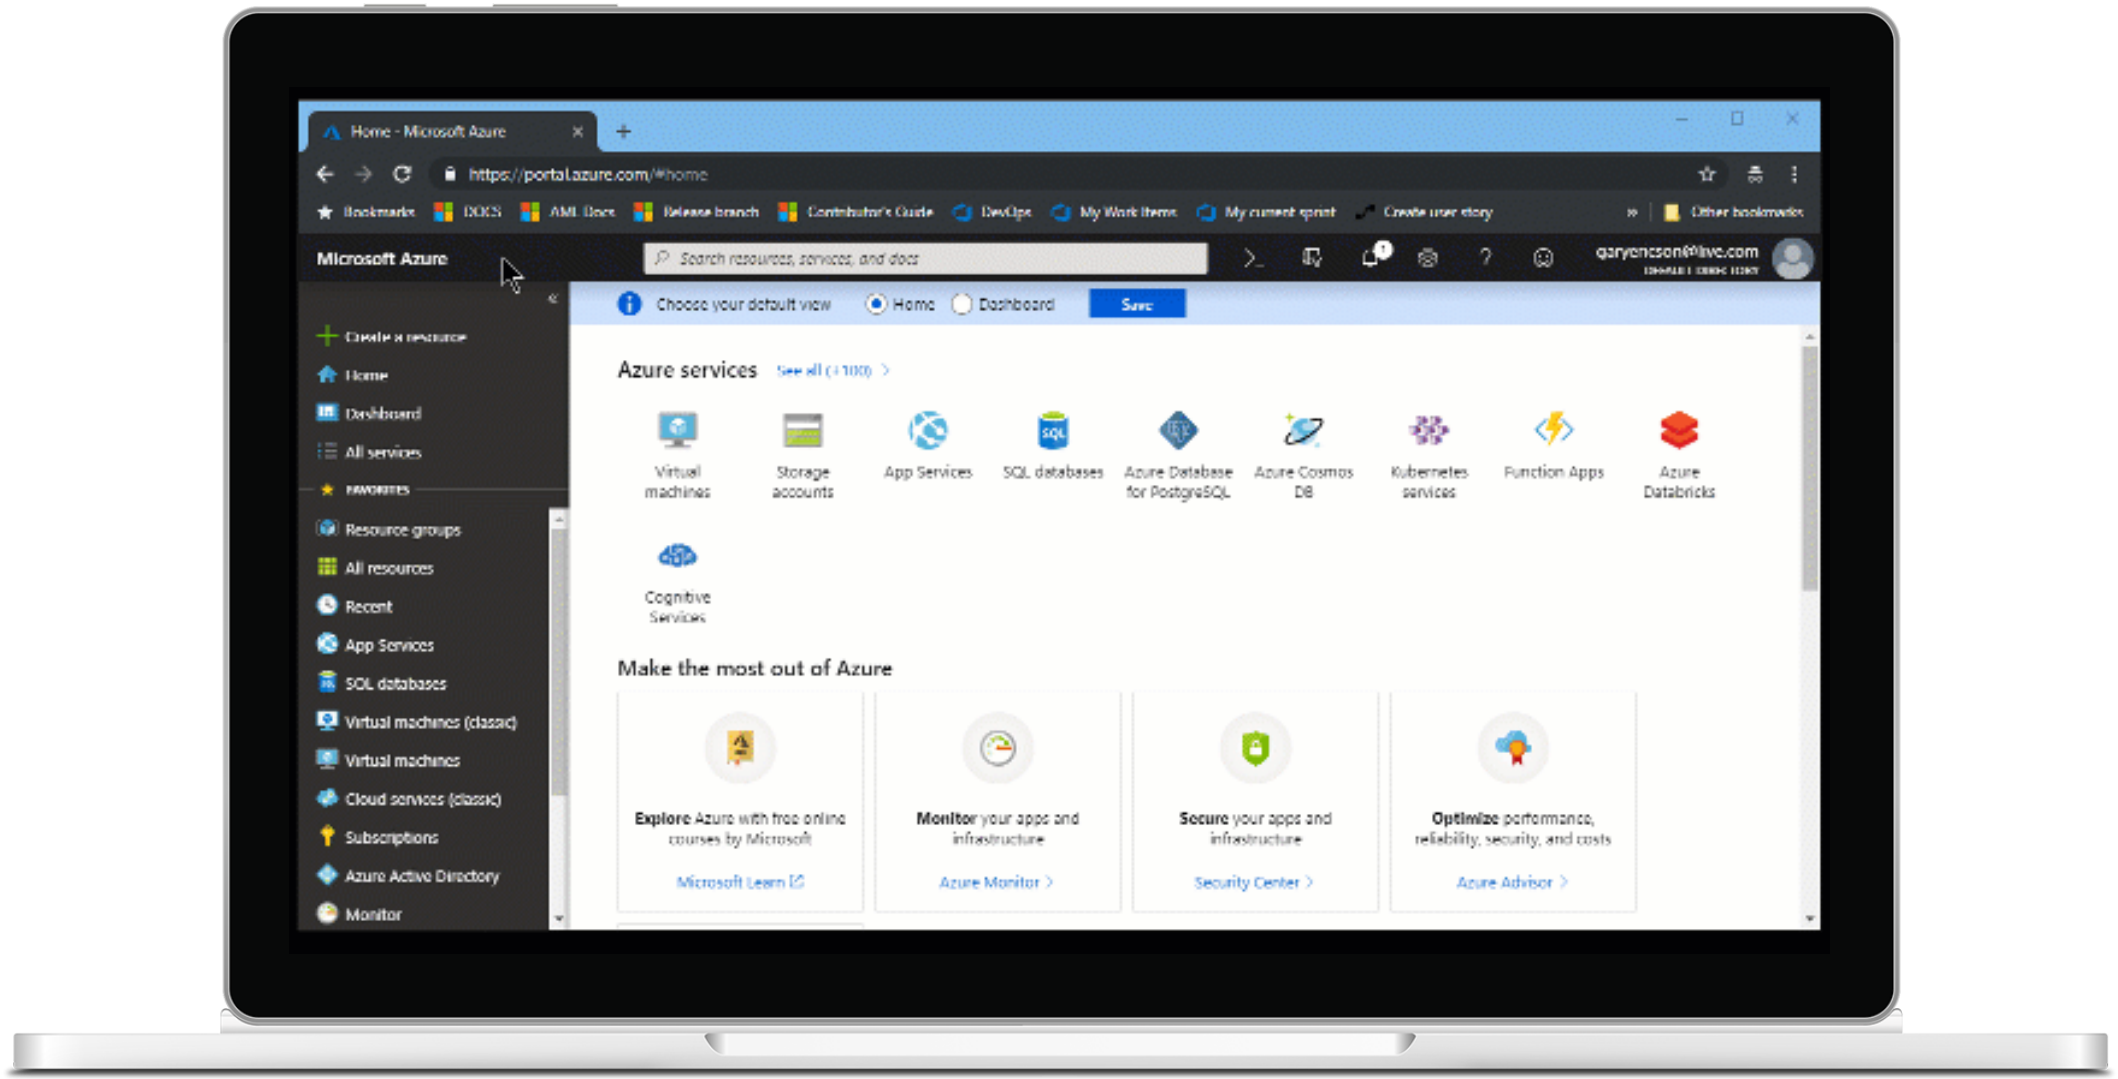

At a high level...

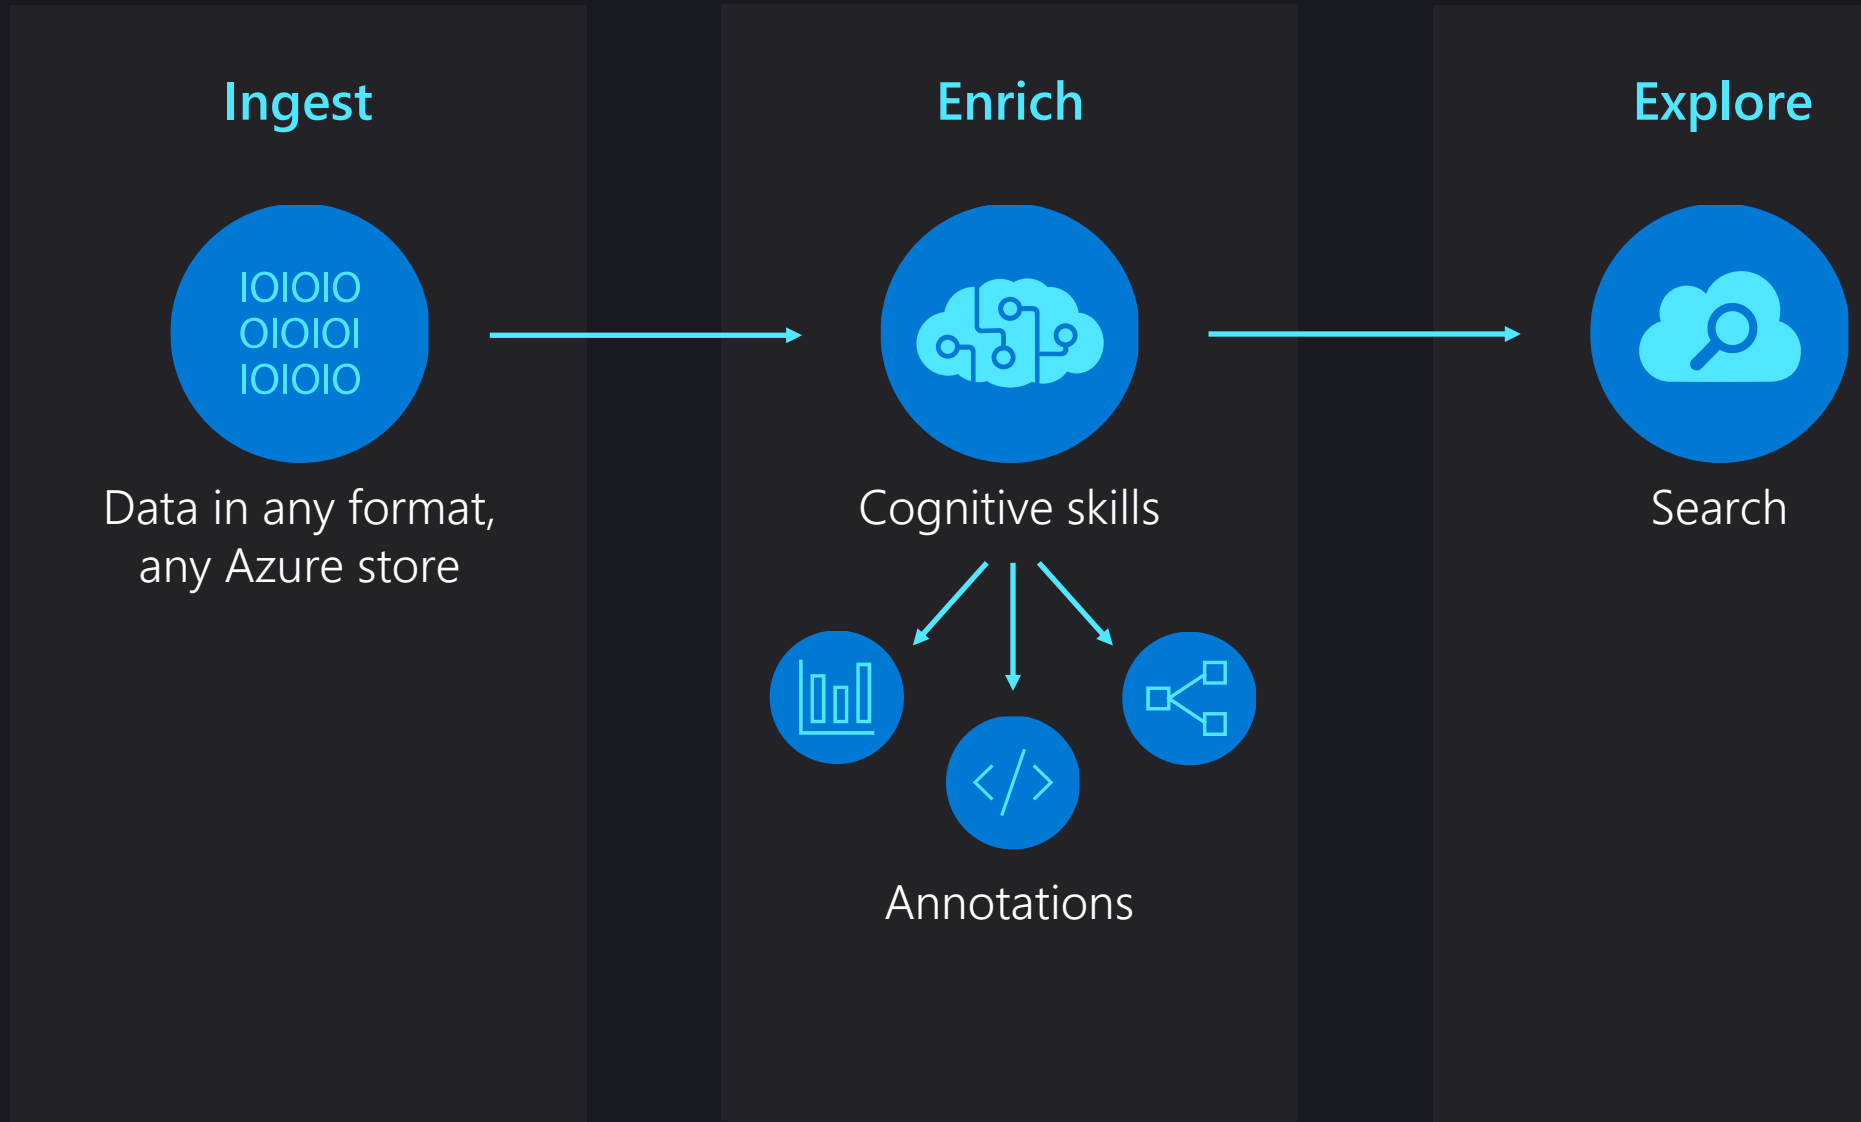

Cognitive Search Architecture

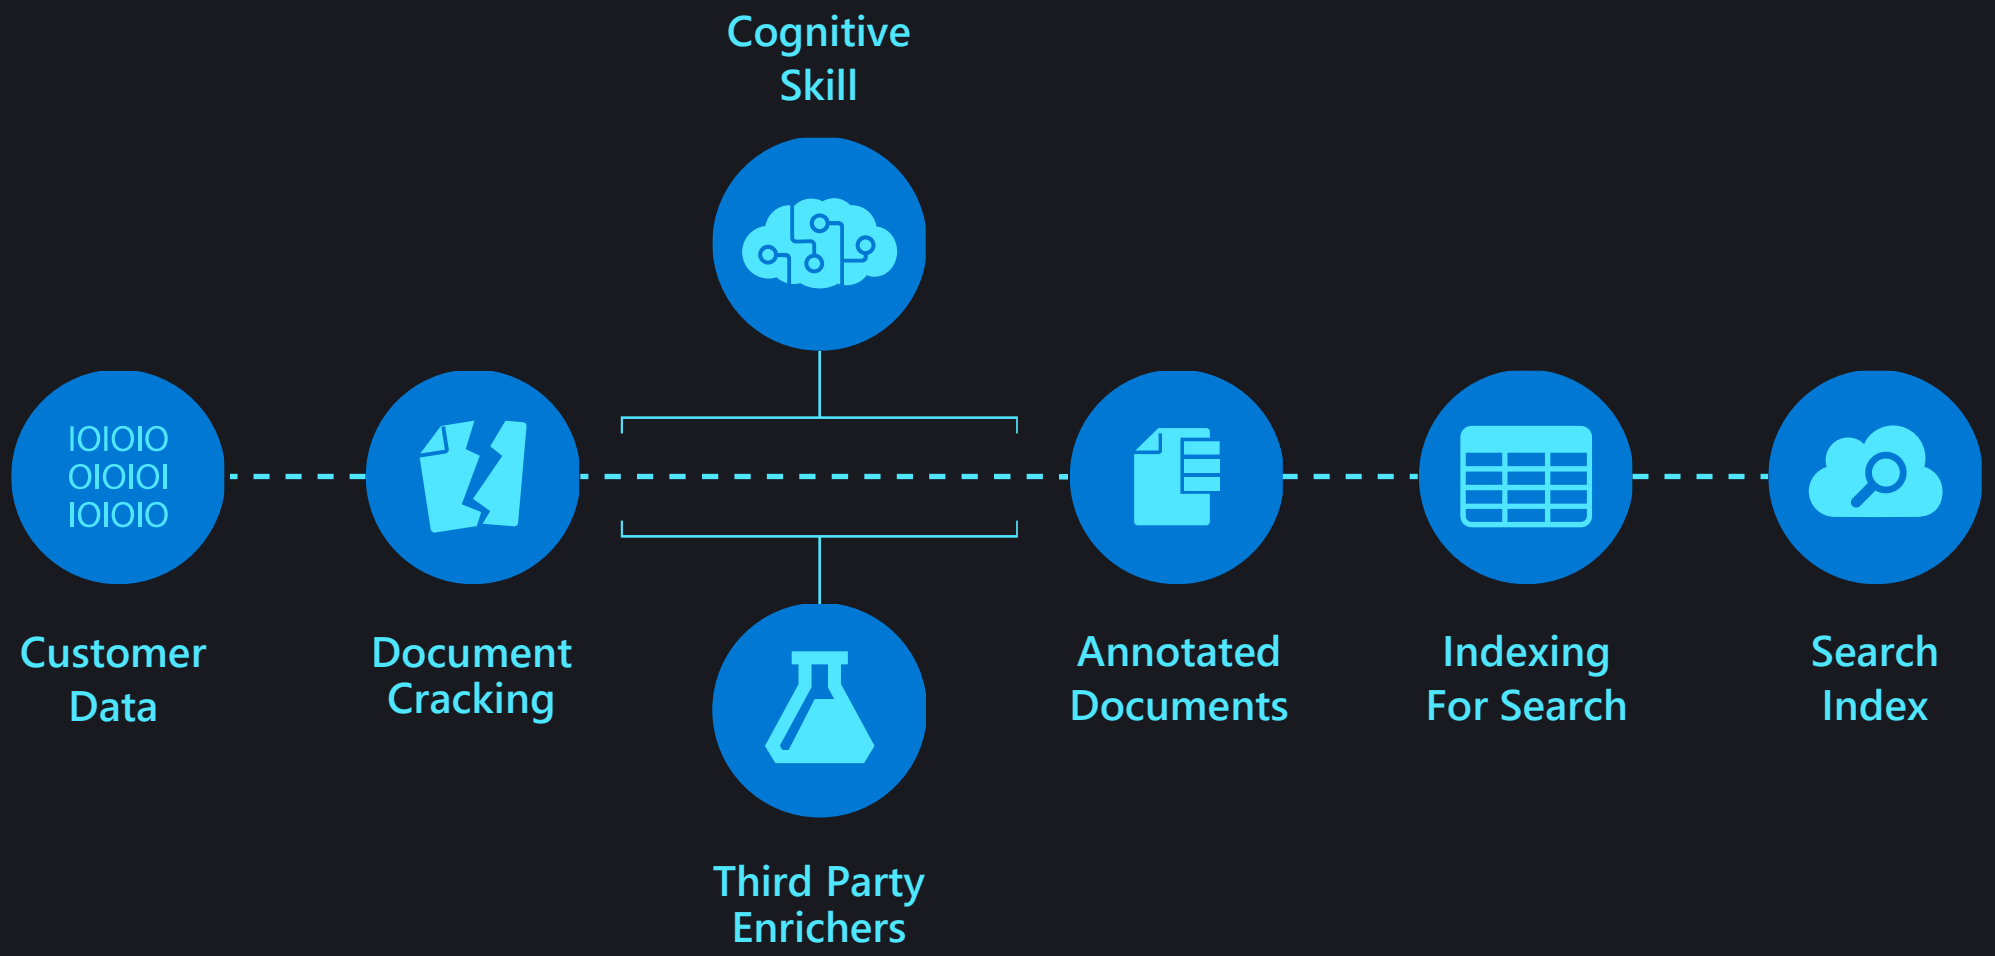

Cognitive Search Architecture

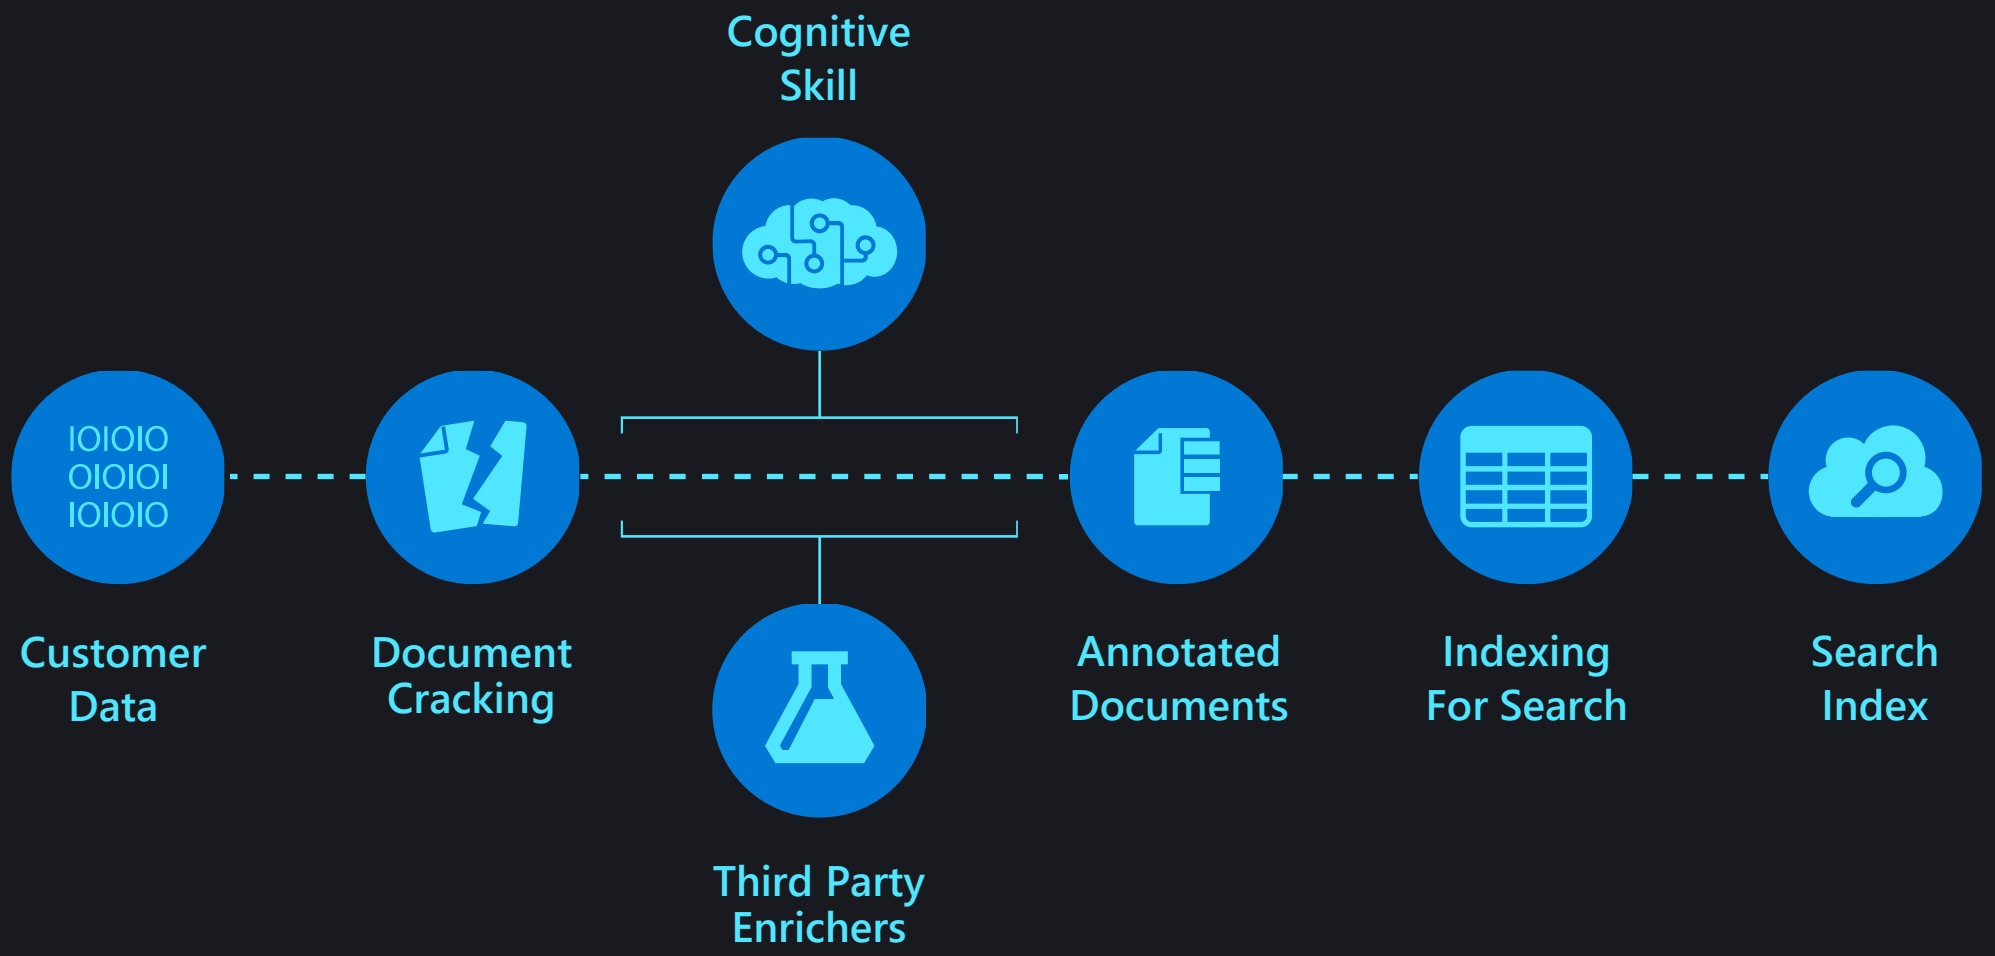

Document Cracking

Different types of data sources

- Azure Blob Storage
- Azure SQL
- Azure Cosmos DB
- Azure Table Storage
- ADLS Gen2

File formats supported in blob storage

- PDF
- Microsoft Office formats: DOCX/DOC, XLSX/XLS, PPTX/PPT, MSG (Outlook emails)
- HTML
- XML
- ZIP
- EML
- RTF
- Plain text files (see also Indexing plain text)
- JSON (see Indexing JSON blobs)
- CSV (see Indexing CSV blobs preview feature)

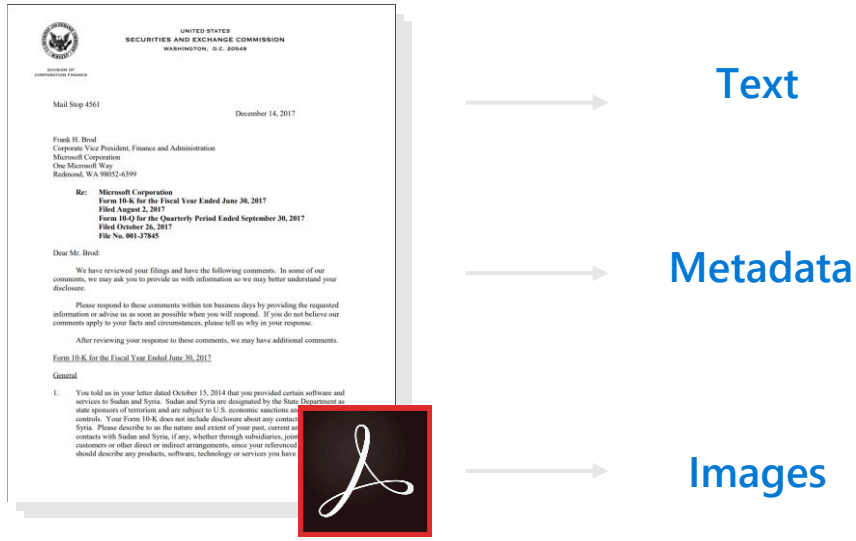

Cognitive Search Architecture

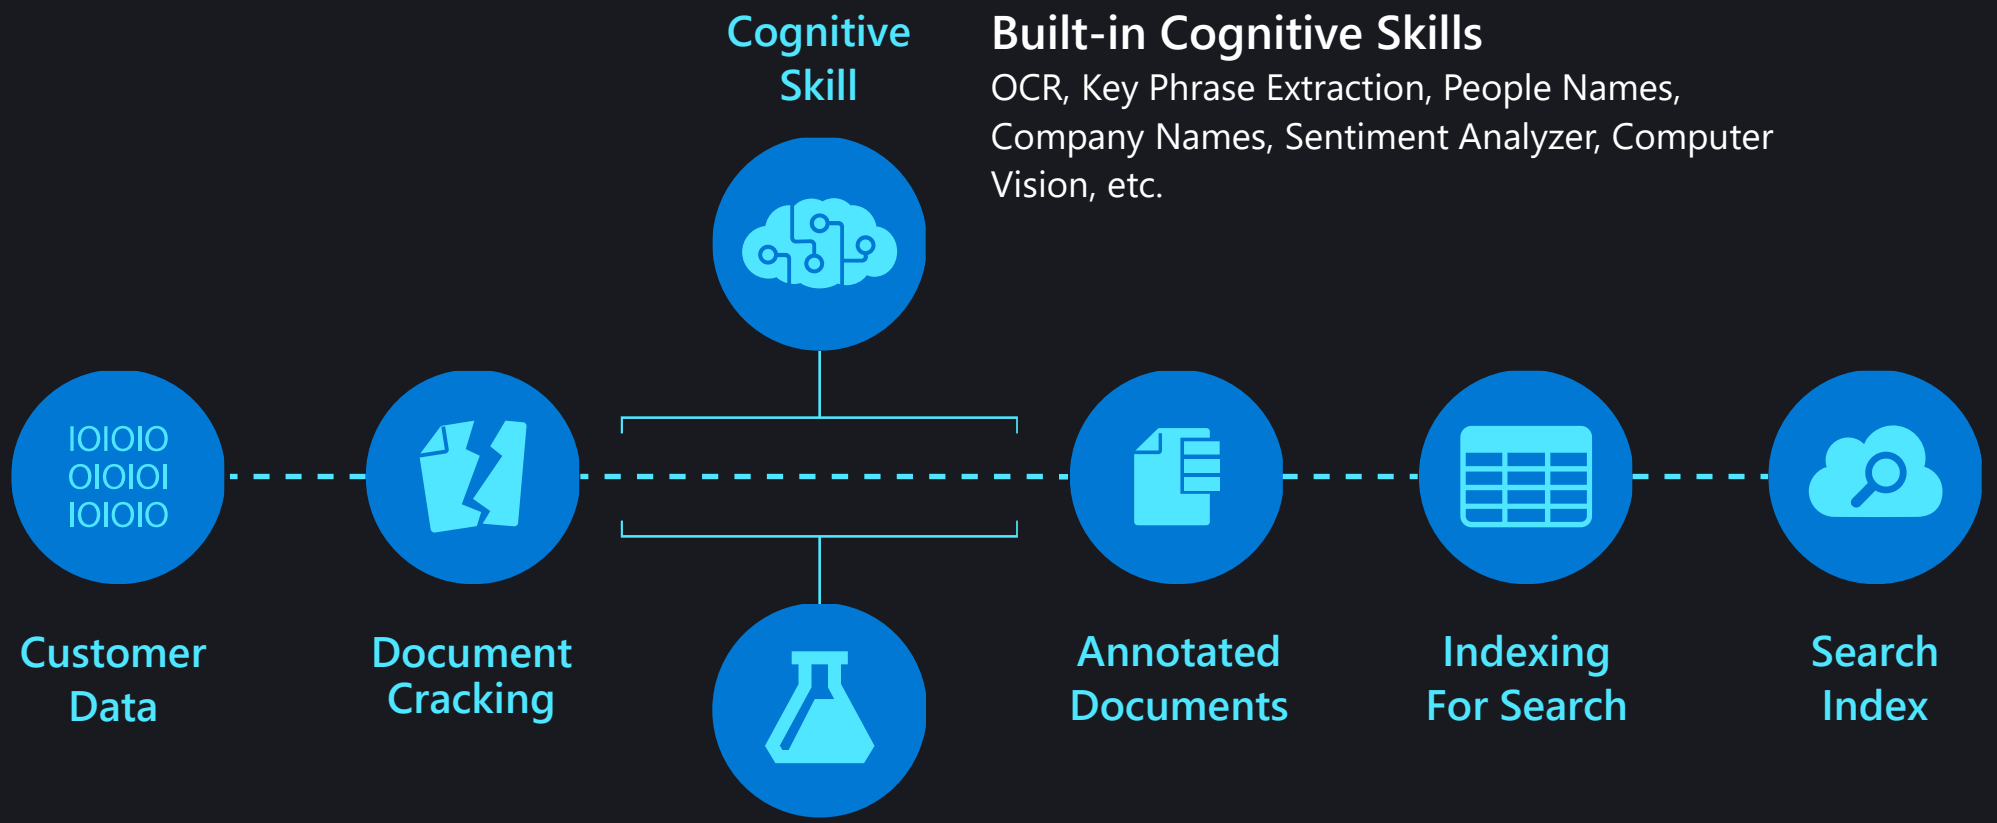

Built-in Cognitive Skills
OCR, Key Phrase Extraction, People Names, Company Names, Sentiment Analyzer, Computer Vision, etc.

Third Party Enrichers
Custom classification models, Custom entity extraction, etc.

Built-in Skills in Multiple Languages

Natural Language Processing

Key Phrase Extraction
Sentiment Analysis
Organization Entity Extraction
Location Entity Extraction
Persons Entity Extraction
Language Detection
Text Translation

Image Processing

Face Detection
Tag Extraction
Celebrity Recognition
Landmark Detection
Handwriting Recognition
Printed Text Recognition

Utilities

Complex Type Shaping
Text Merging
Text Splitting
Conditional Skill

Sample skillset

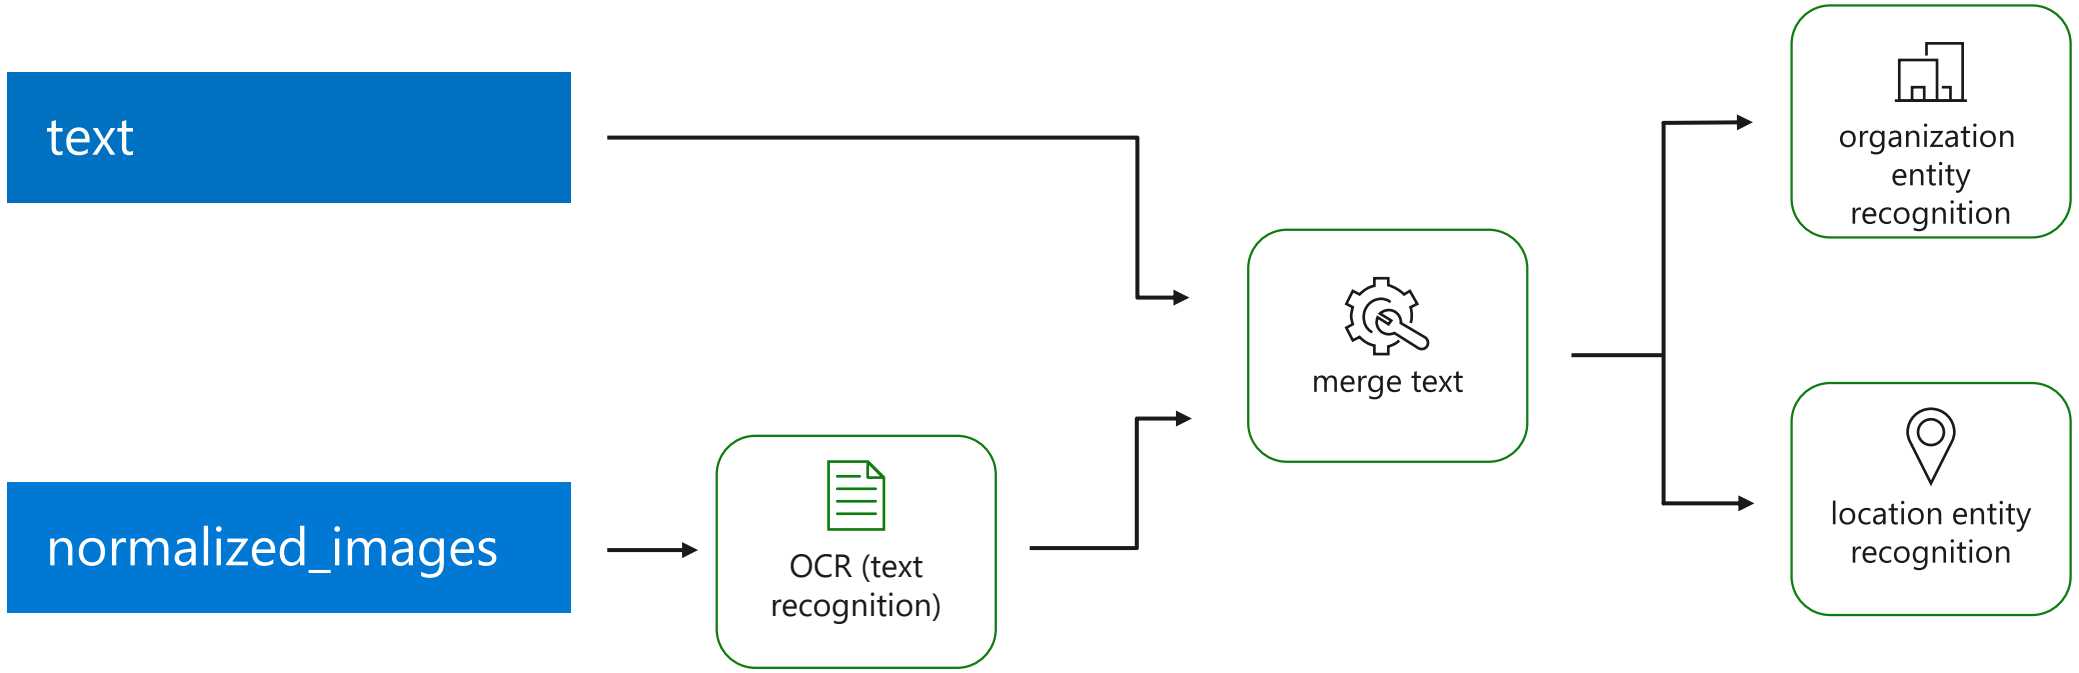

Extend with your own custom skills…

```
…,
   {
     "@odata.type": "#Microsoft.Skills.Custom.WebApiSkill",
     "uri": "https://myskill.azurewebsites.net/api/OrgId"
     "httpHeaders": {"Api-Key": "mySecret" },
     "context": "/document/organizations/*" ,
     "inputs":
     [
       {  "name": “organizationName", "source": "/document/organizations/*" },
     ],
     "outputs":
     [
       {  "name": "organizationId", "targetName": "organizationId" }
     ]
   },
```

Extend with your own skills…

Well defined input/output schema

```
{
    "values": [
      {
        "recordId": "7cad2",
        "data":
           {
             "myInput1": "fox",
             "myInput2": "cat",
           }
      },
      {
        "recordId": "7cad3",
        "data":
           {
             "myInput1": "blue",
             "myInput2": "red",
           }
      },
      …
    ]
}
```

Custom Cognitive Skill

```
{
    "values": [
      {
        "recordId": "7cad2",
        "data":
           {
             "myOuput1": “animals"
           }
      },
      {
        "recordId": "7cad3",
        "data":
           {
             "myOutput1": “colors"
           }
      },
      …
    ]
}
```

Custom Skills

Azure Function Framework

Custom Cognitive Skill

Azure Container Service

Custom Cognitive Skill

Azure Machine Learning

Custom Cognitive Skill

Review: Your options with Cognitive skills

Built-in skills

- Key Phrase extraction
- Location entity extraction
- Sentiment analysis
- Organization entity extraction
- Persons entity extraction
- Language detection
- Face detection
- Celebrity recognition
- Image tag extraction
- Text Utilities
- Landmark detection
- Printed text recognition

Custom skills

Your custom skill goes here!

Azure Databricks

Azure Machine Learning

Machine Learning VMs

Sample skillset with custom skills

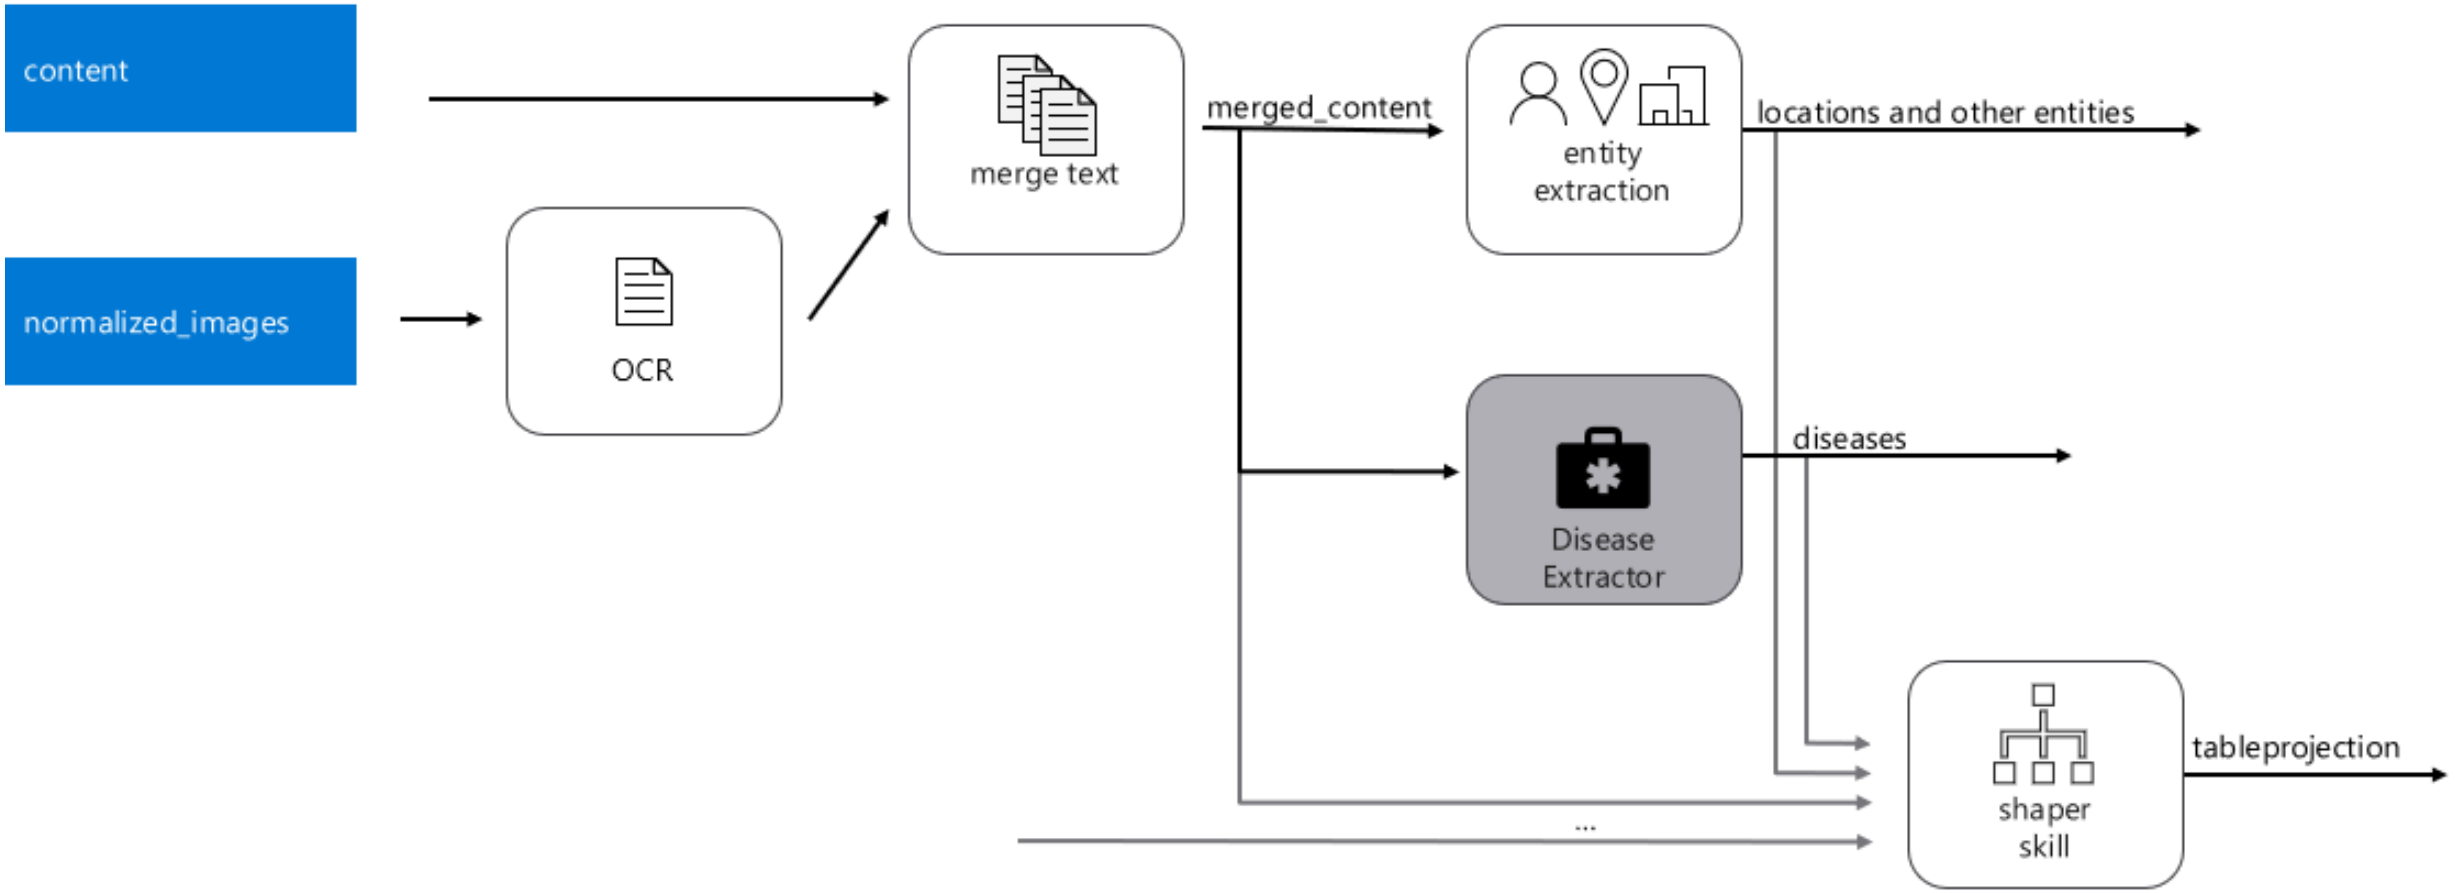

Challenges / Use Cases

Challenge 1

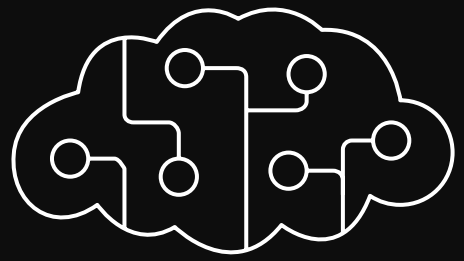

Smart Attendance System

Challenge 2

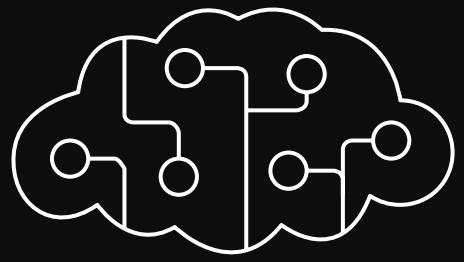

Intelligent Minutes of Meeting

Challenge 3

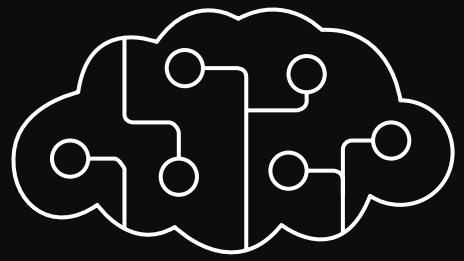

Addressing customers concerns in real-time

Challenge 4

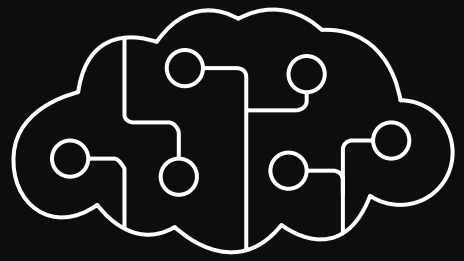

Waste Management

Challenge 5

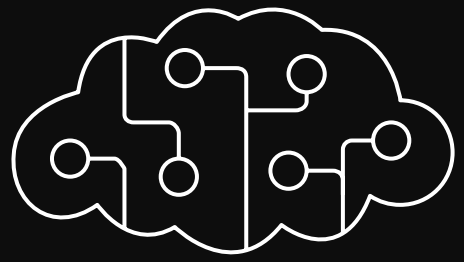

Build Smart, Safe & healthy Countries

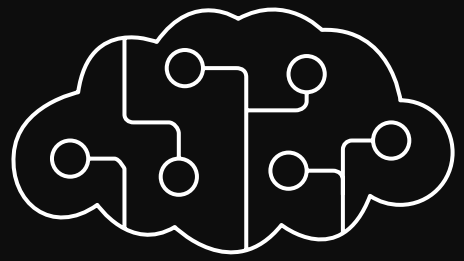

FOR GOOD (Others)

Where to start

Microsoft

AI School

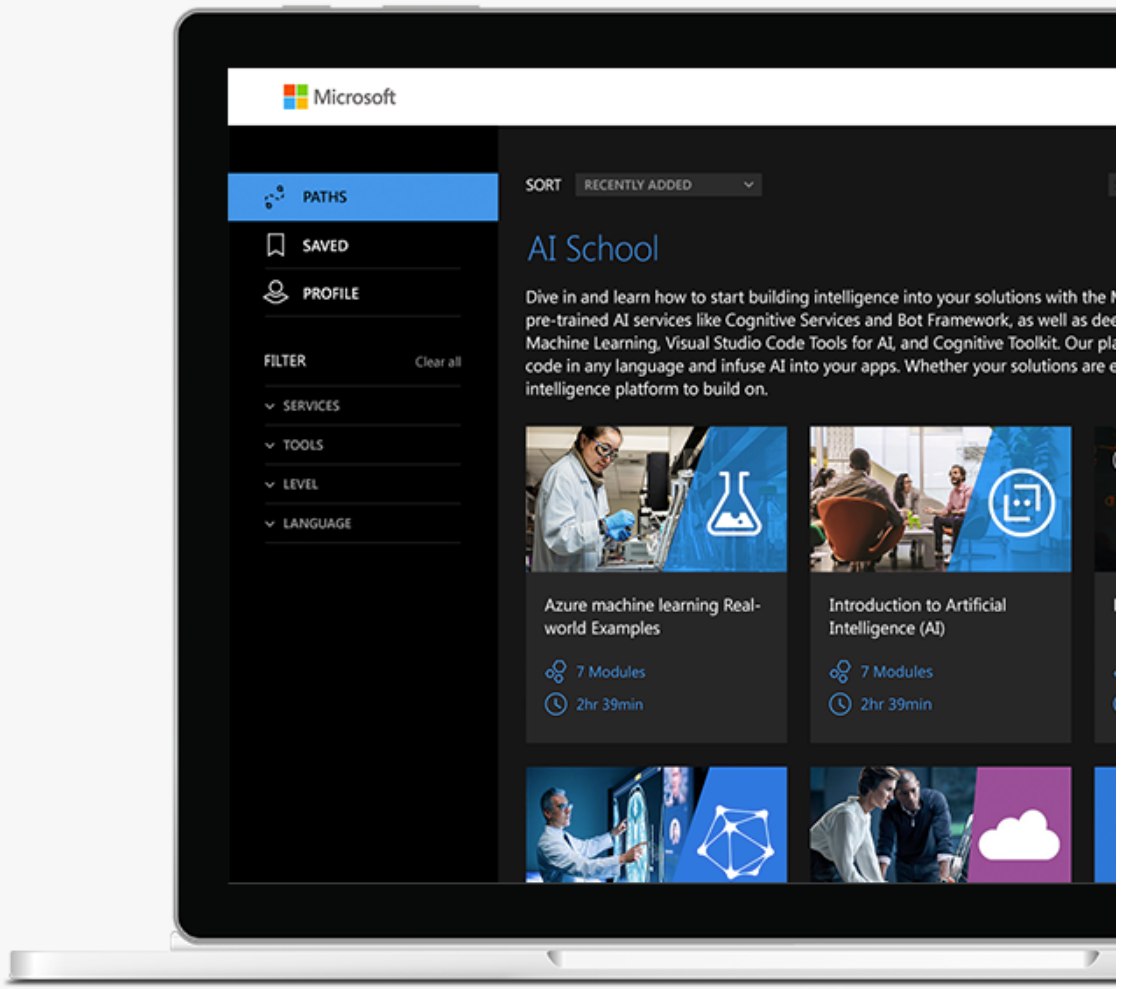

http://aka.ms/ai-school

Key Links

AutomatedML

Hyperdrive

Azure Machine Learning Documentation

AI on IoT Edge Devices

Azure Cognitive Services

Azure Cognitive Services in Containers

Azure Machine Learning Studio

Key Links

aka.ms/ACE-Blog

aka.ms/kma

aka.ms/AILabs

Question and Answer