

#TampaCC

Tampa Community Connect

Democratization of AI with Microsoft Cognitive Services

Adnan Masood, PhD.

October 2018

Slides Courtesy of Microsoft Corp. All product names, logos, and brands are property of their respective owners.





Adnan Masood, PhD.

Dr. Adnan Masood is an Artificial Intelligence and Machine Learning researcher, visiting scholar at Stanford AI Lab, software architect, and Microsoft MVP (Most Valuable Professional) for Artificial Intelligence. As Chief Architect of AI and Machine Learning, at UST Global, he collaborates with Stanford Artificial Intelligence Lab, and MIT AI Lab for building enterprise solutions

Author of Amazon bestseller in programming languages, "**Functional Programming with F#**", Dr. Masood teaches Data Science at Park University, and has taught Windows Communication Foundation (WCF) courses at the University of California, San Diego. He is a regular speaker to various academic and technology conferences (WICT, DevIntersection, IEEE-HST, IASA, and DevConnections), local code camps, and user groups. He also volunteers as STEM (Science Technology, Engineering and Math) robotics coach for elementary and middle school students.



Artificial Intelligence

Microsoft Practice Development Playbook

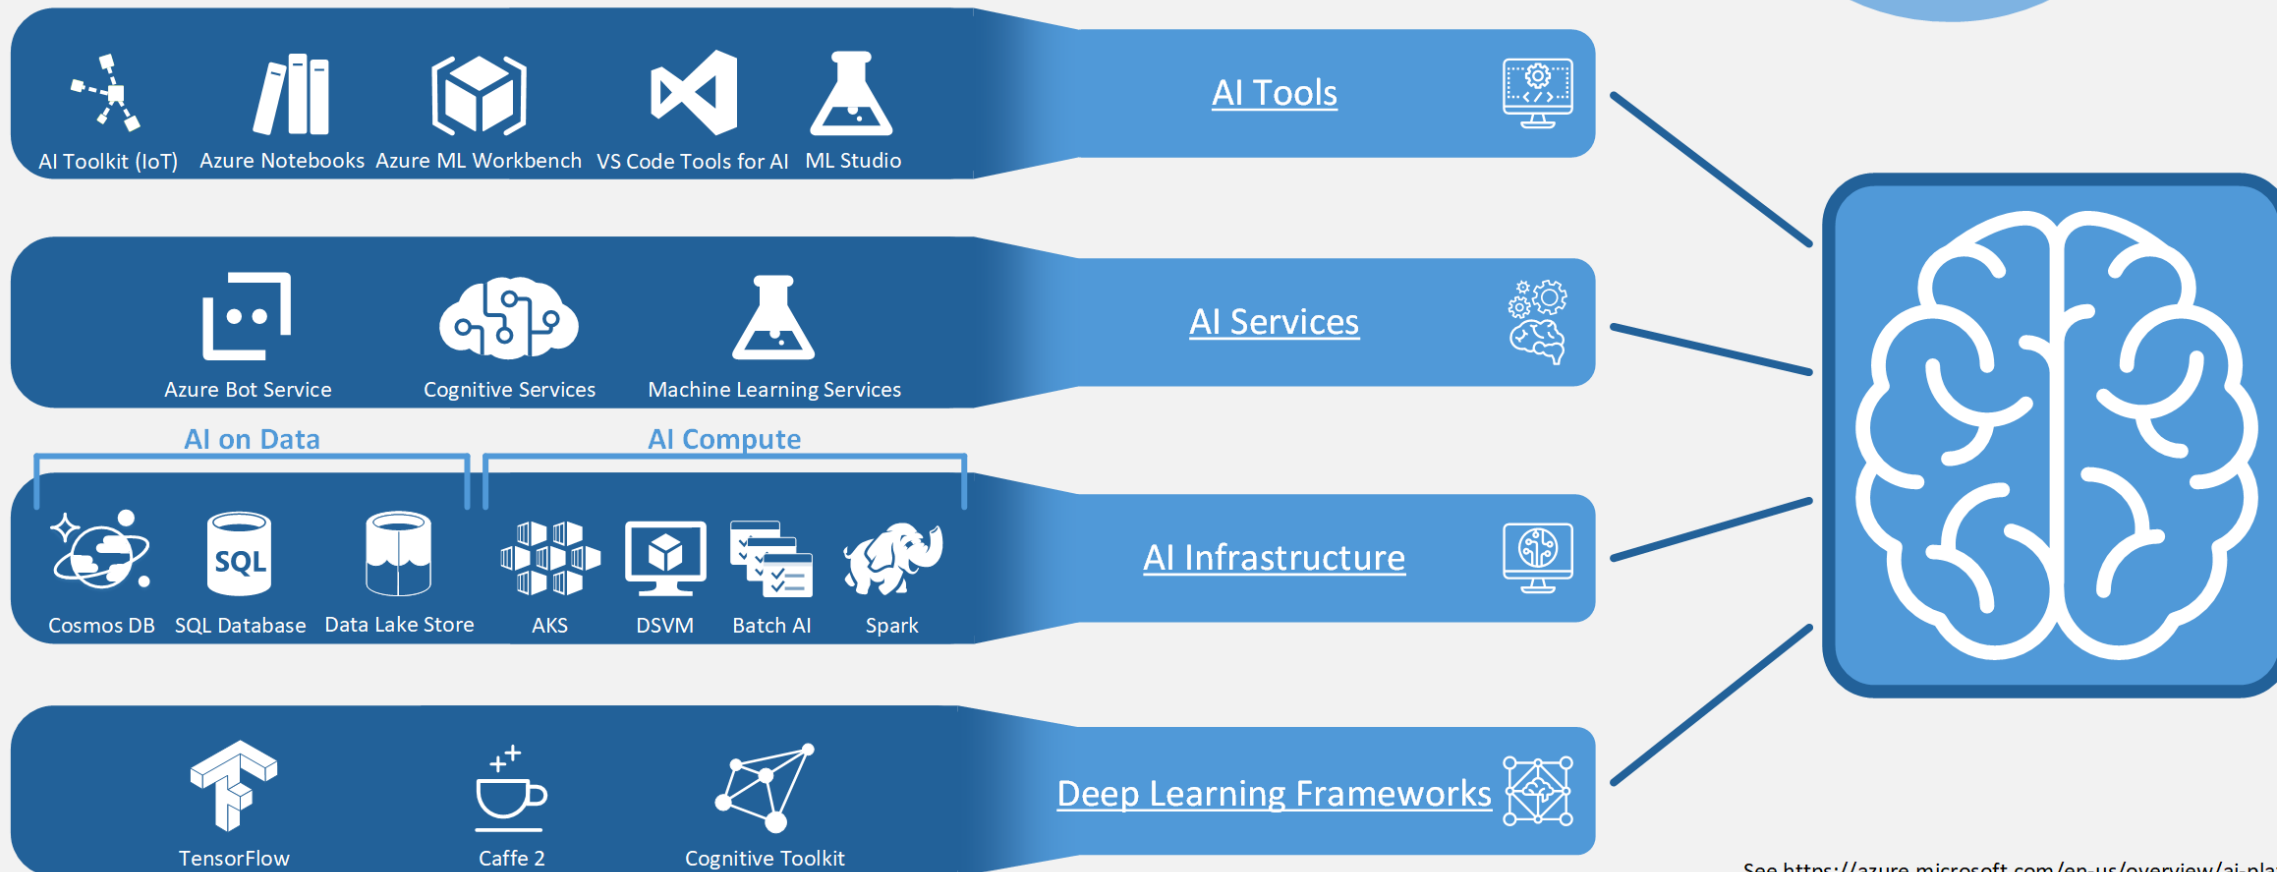


aka.ms/practiceplaybooks

Table of Contents


About this Playbook	2	Hire & Train	80
Partner Practice Development Framework	5	Executive Summary	81
What is Artificial Intelligence?.....	6	Hire, build, and train your team	82
AI Opportunity	8	Job Descriptions for your Technical Team	88
Industry Opportunities	9	Recruiting Resources.....	98
Define Your Strategy	17	Training & Readiness.....	99
Executive Summary.....	18	Operationalize	113
Define Your Practice Focus	19	Executive Summary	114
Understanding the AI Practice.....	20	Implement a Process.....	115
The Microsoft Approach to AI.....	23	Claim Your Internal Use Benefits	119
Pre-Built AI using Cognitive Services.....	28	Define Customer Support Program and Process	124
Building Custom AI	35	Manage and Support an AI solution in Azure.....	128
Microsoft AI Platform Summary.....	42	Support Ticket Setup and Tracking	130
Define and Design the Solution Offer.....	43	Implement Intellectual Property Offerings	131
Understanding Project Based Services.....	44	Setup Social Offerings	132
Understanding Managed Services	54	Create Engagement Checklists & Templates.....	133
Accelerate your Managed Service Model.....	60	Go to Market & Close Deals	134
Understanding Intellectual Property	61	Executive Summary	135
Define Industry Specific Offerings.....	65	Marketing to the AI Buyer.....	136
Define Your Pricing Strategy.....	66	Engage Technical Pre-Sales in Sales Conversations.....	138
Calculate Your Azure Practice Costs	69	Architecture Design Session (ADS)	140
Identify Partnership Opportunities	71	Go-to-Market and Close Deals Guide	142
Define Engagement Process.....	73	Optimize & Grow	143
Identify Potential Customers.....	74	Executive Summary	144
Join the Microsoft Partner Network	75	Understanding Customer Lifetime Value	145
Stay Informed on AI Matters	77	Guide: Optimize and Grow	147
Identify Solution Marketplaces.....	78	AI Playbook Summary	148

Microsoft AI Platform

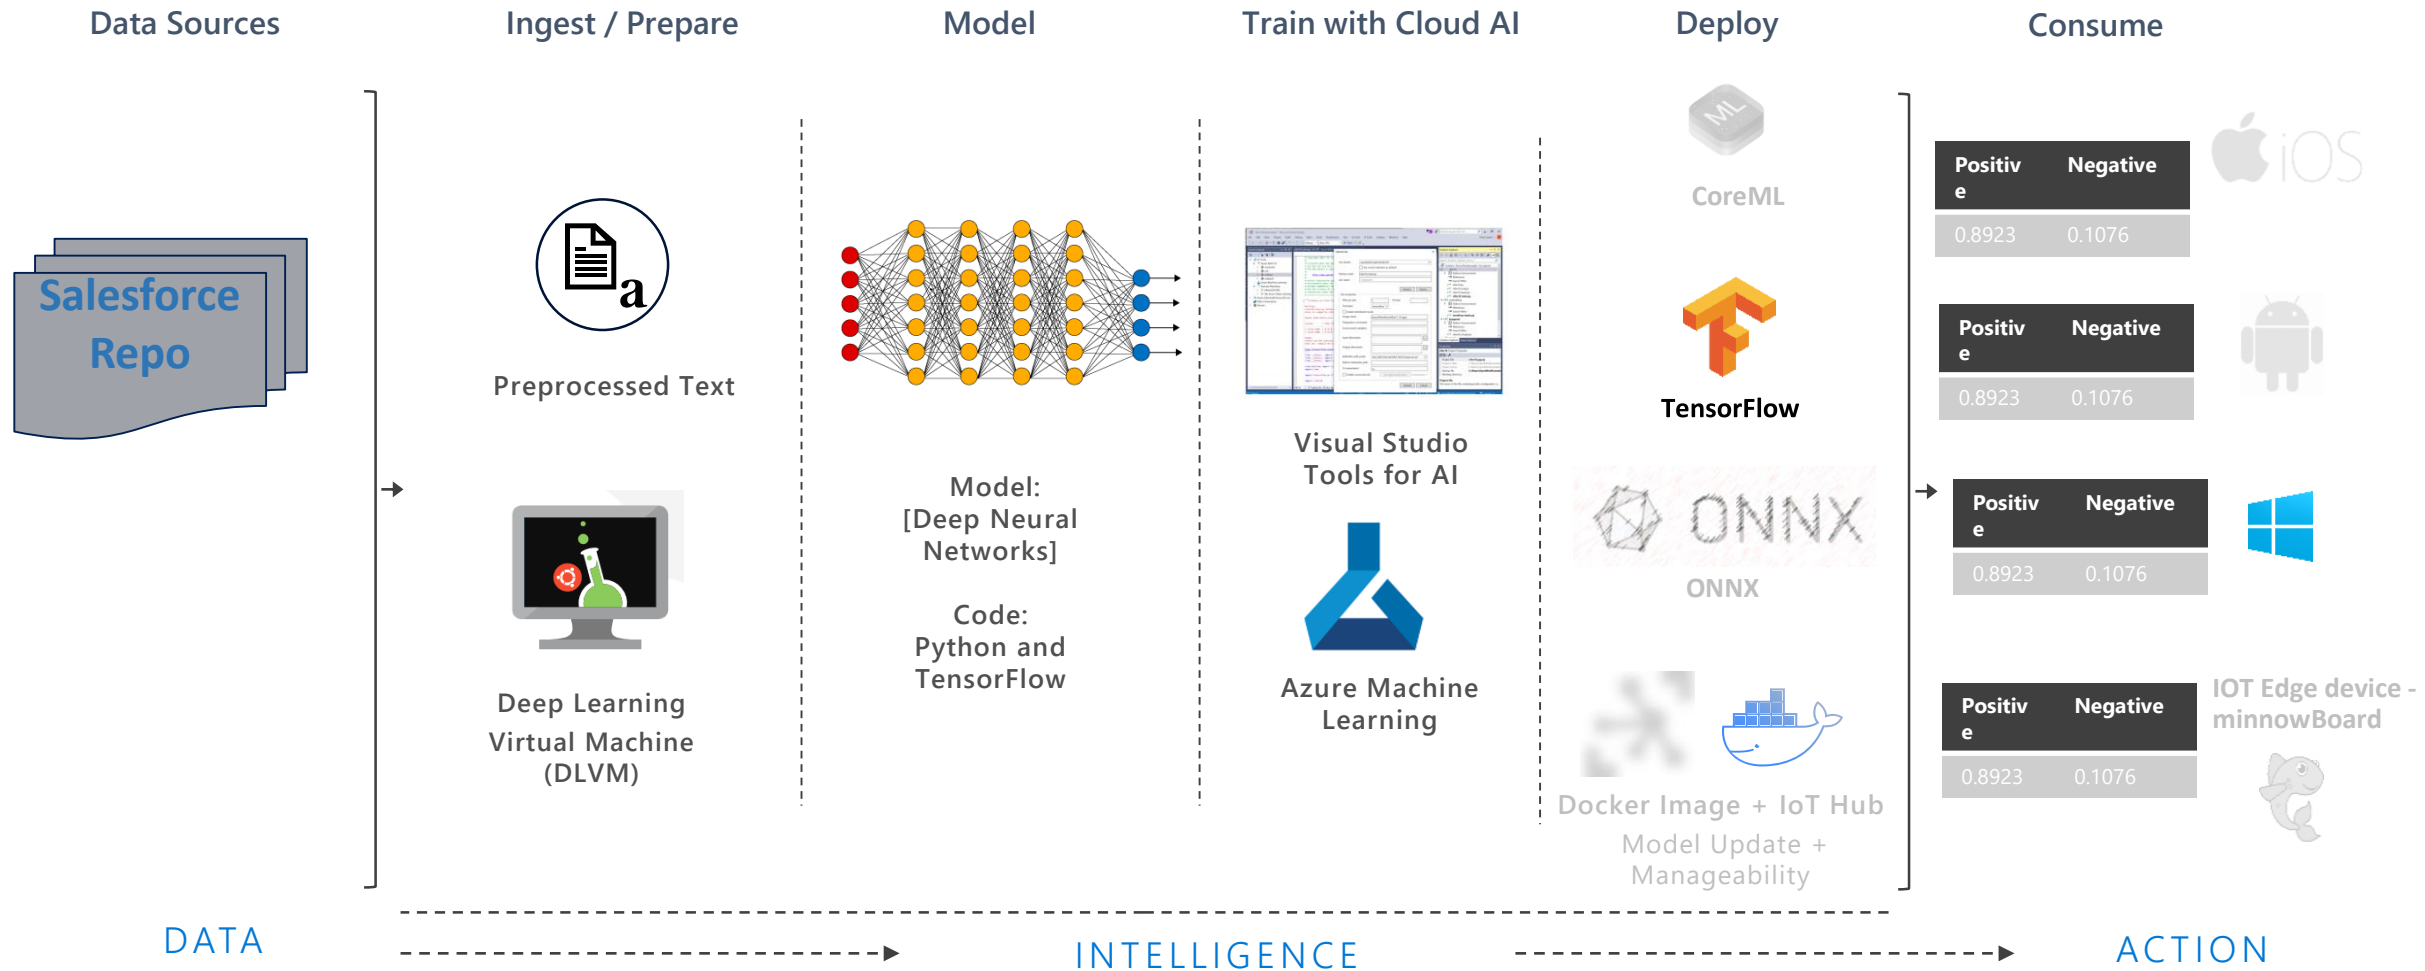



See <https://azure.microsoft.com/en-us/overview/ai-platform> for more information about the various services and features of the Microsoft AI Platform

Data Science Lifecycle



Example Real World ML Pipeline Architecture





Common AI/ML Problems:

- Most libraries provide state-of-the-art algorithms but little pertinent training data
- For many conversational domains, training data may be difficult or impossible to collect
- Pre-built domains streamline development but are largely irrelevant for most apps
- Tools for building custom domains can only handle narrow models and trivial apps
- ML capabilities only scratch the surface of what is typically required for production apps

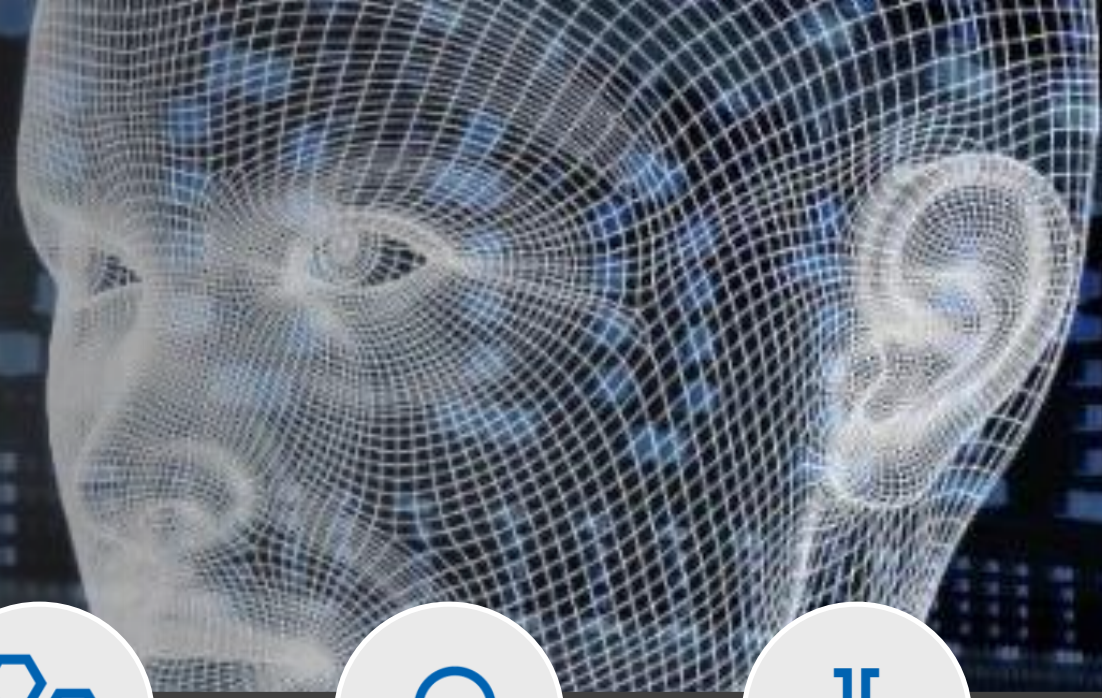
Machine Learning Development Lifecycle provides customized end to end solution from formal problem definition, domain modeling, creating training and test data, training models, evaluation of model, execution, deployment, and visualization.

Key Value Proposition:

- Not just offer an NLP library but provide expertise to work with bot framework for multiple modalities, commerce engine integration, and deployment infrastructure and expertise.

Microsoft Cognitive Services

Give your apps a human side



Vision

From faces to feelings, allow your apps to understand images and video



Speech

Hear and speak to your users by filtering noise, identifying speakers, and understanding intent



Language

Process text and learn how to recognize what users want



Knowledge

Tap into rich knowledge amassed from the web, academia, or your own data



Search

Access billions of web pages, images, videos, and news with the power of Bing APIs

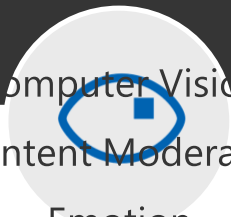



Labs

An early look at emerging Cognitive Services technologies: discover, try and give feedback on new technologies before general availability

Microsoft Cognitive Services

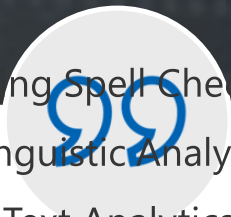
Give your apps a human side



Computer Vision
Content Moderator
Emotion
Face Vision
Video Indexer
Custom Vision Service



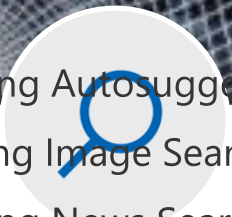
Bing Speech
Speaker Recognition
Custom Speech
Speech



Bing Spell Checker
Linguistic Analysis
Text Analytics
Language
Translator Text & Speech
Web Language Model
Language Understanding



Academic Knowledge
Entity Linking
Knowledge Exploration
Knowledge
Recommendations
QnA Maker
Custom Decision Service



Bing Autosuggest
Bing Image Search
Bing News Search
Search
Bing Video Search
Bing Web Search
Bing Entity Search
Bing Custom Search



Project Prague (gesture)
Project Cuzco (events)
Project Johannesburg (r**labs**)
Project Nanjing (isochrones)
Project Abu Dhabi (distance matrix)
Project Wollongong (location)

Microsoft Cognitive Services

Give your apps a human side



Vision

Computer Vision
Content Moderator
Emotion
Face
Video Indexer



Speech

Bing Speech
Speaker Recognition




Language

Bing Spell Check
Linguistic Analysis
Text Analytics
Translator Text & Speech
Web Language Model



Knowledge

Academic Knowledge
Entity Linking
Knowledge Exploration
Recommendations
QnA Maker



Search

Bing Autosuggest
Bing Image Search
Bing News Search
Bing Video Search
Bing Web Search
Bing Entity Search



Labs

Project Prague (gesture)
Project Cuzco (events)
Project Johannesburg (routing)
Project Nanjing (isochrones)
Project Abu Dhabi (distance matrix)
Project Wollongong (location)

CUSTOMIZATION

Custom Vision Service

Custom Speech Service

Language Understanding

Custom Decision Service

Bing Custom Search

Microsoft Cognitive Services Customization


The set of custom services available for your business, allowing customers to use their own data to train models.



Vision




Speech



Language



Knowledge



Search



Labs

Custom Vision
Service

Custom Speech
Service

Language
Understanding

Custom Decision
Service

Bing Custom
Search

Why Microsoft Cognitive Services?

Easy

Roll your own with REST APIs
Simple to add: just a few lines of code required



Flexible

Integrate into the language and platform of your choice
Breadth of offerings helps you find the right API for your app
Bring your own data for your custom experience




Tested

Built by experts in their field from Microsoft Research, Bing, and Azure Machine Learning
Quality documentation, sample code, and community support




A variety of real-world applications

Vision


 What is in the image?

Computer Vision




Category	People; 5 faces
Adult/Racy?	False/False
Dominant colors	
Accent color	

Speech


 Give me directions to the nearest local branch.

Bing Speech



- Convert spoken audio to text
- Convert text to spoken audio
- Extract intent of user


Language

 Play today's customer call recording.

Language Understanding

Natural Language Processing

Intent: PlayCall
Content: Customer#
DateTime.date: today



Now Playing

11/29/2016 Customer Call

Knowledge


 Top publications in customer lifecycle trends?

Knowledge Exploration

Here are the top results:




- [Customer Relationship Management – 5 Key Trends for 2014 CRM](#)
 Oct 28, 2015 – Here are FIVE key trends in 2014 that would help marketers in rolling ... Of late, marketers are looking at customer lifecycle management (CLM)
- [Predictive Customer Lifecycle Management \(CLM\)](#)
 The purpose of Customer Life-cycle Management (CLM) is to maximize both customer retention and Predictive trend analysis provides business visibility.
- [Trends 2016: The Future of Customer Service](#)
 Jan 5, 2016 – The top 10 customer service trends for 2016 that ... North American Consumer
- [Language Around Customer Lifecycles in the Banking Industry](#)
[View PDF](#)

Search

 Search for 'fraud prevention'

Bing News Search

Here is what I found:

-  [Information Communications Media Market News](#)
 It also investigates the top three expected Fraud Detection and Prevention programs, in terms of demand in key markets...
-  [The Big Question: In-House or Outsourced Fraud Protection?](#)
 First, let's point out that there is not one absolute answer—there are "pros" and "cons" to each. Those who favor in-house...
-  [How to Protect Your Business from Online Fraud this Holiday Season](#)
 Michael heads fraud prevention tool. Online and mobile shopping are expected to continue growing apace...

The Pre-Built and Custom AI Spectrum

Pre-Built AI

Custom AI



Microsoft Cognitive Services

Azure Machine Learning

Models are pre-trained using
Microsoft supplied data

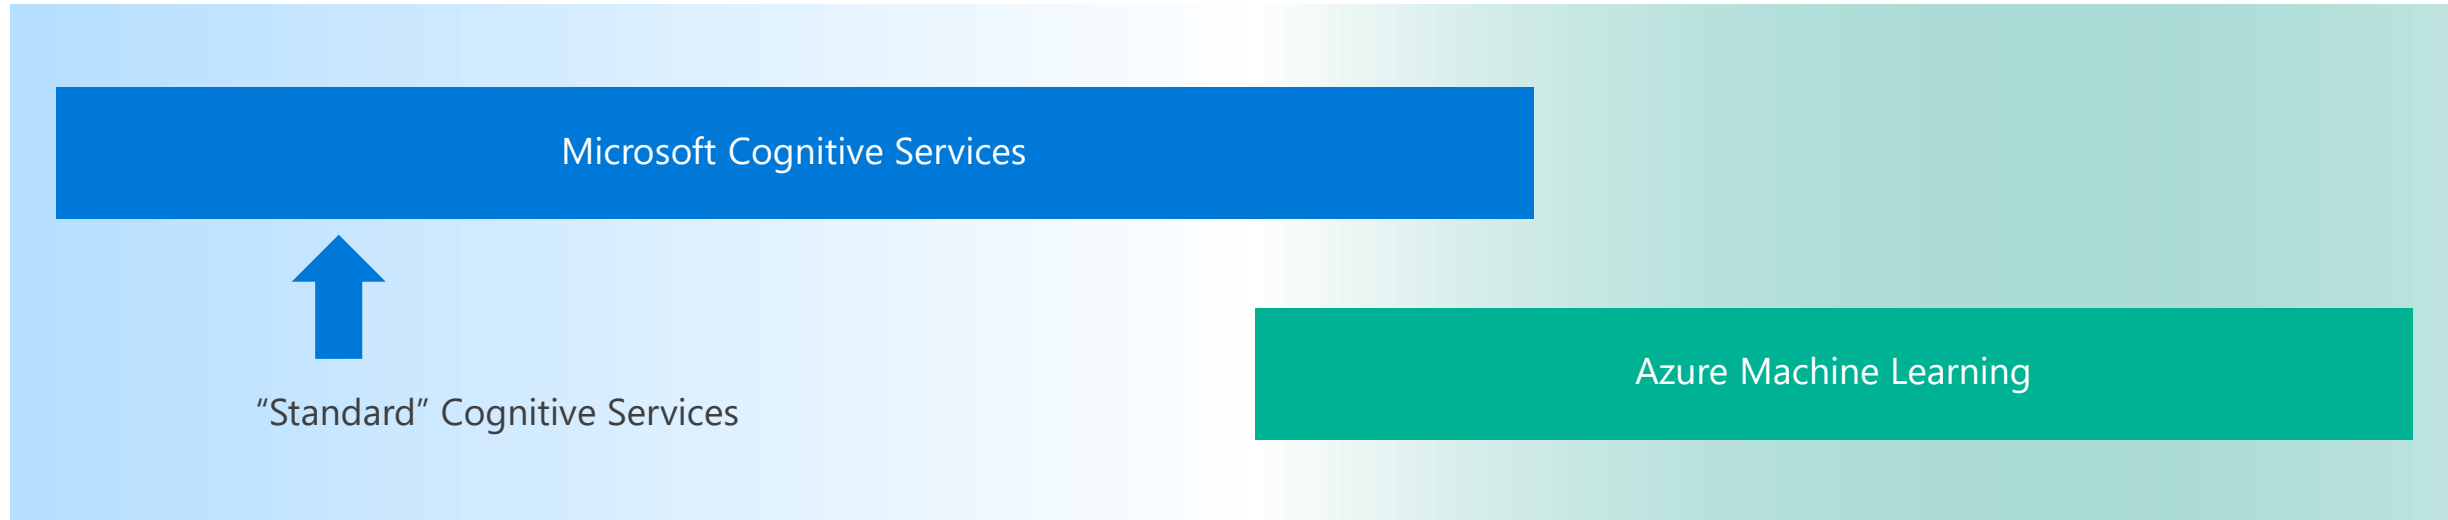
Models "customized" with your data

Models tailored to your
scenario and your data

The Pre-Built and Custom AI Spectrum

Pre-Built AI

Custom AI



Models are pre-trained using
Microsoft supplied data

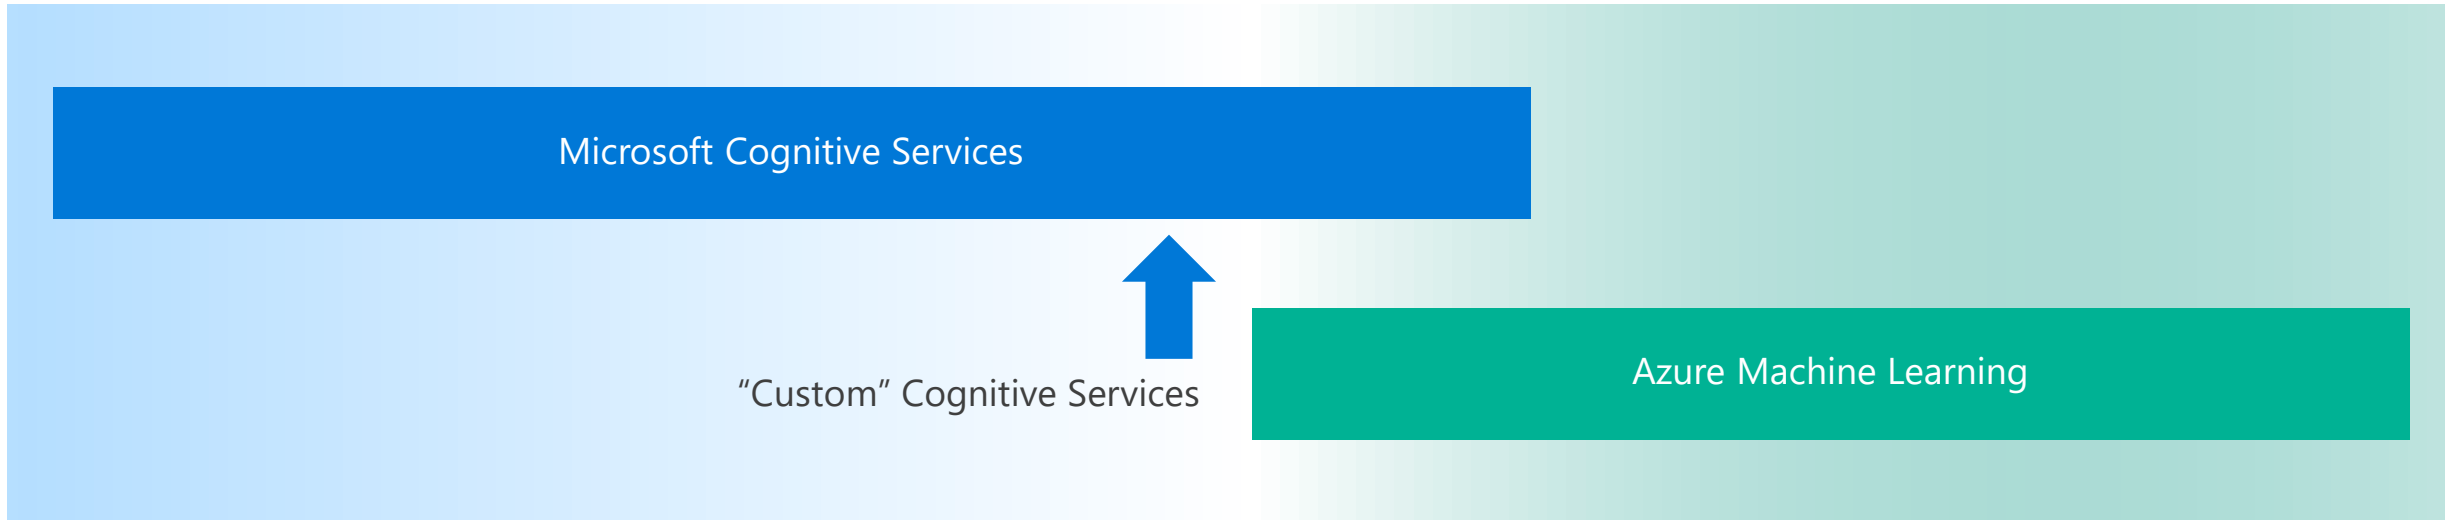
Models "customized" with your data

Models tailored to your
scenario and your data

The Pre-Built and Custom AI Spectrum

Pre-Built AI

Custom AI



Models are pre-trained using Microsoft supplied data

Models "customized" with your data

Models tailored to your scenario and your data



Vision



Computer Vision API

Distill actionable information from images



Face API

Detect, identify, analyze, organize, and tag faces in photos



Emotion API

Personalize experiences with emotion recognition




Content Moderator

Machine-assisted moderation of text and images, augmented with human review tools



Custom Vision Service

Customizable web service that learns to recognize specific content in imagery



Video Indexer

Process and extract smart insights from videos



Speech



Bing Speech API

Convert speech to text and back again, and understand its intent



Speaker Recognition API

Give your app the ability to know who's talking



Custom Speech Service

Fine-tune speech recognition for anyone, anywhere



Language




Bing Spell Check API

Detect and correct spelling mistakes within your app



Text Analytics API

Detect sentiment, key phrases, topics, and language from your text



Language Understanding Intelligent Service

Teach your apps to understand commands from your users



Translator

Easily perform speech and text translation




Knowledge



Recommendations API

Provide personalized product recommendations for your customers



Custom Decision Service

Create custom web content experiences with adaptive, contextual decision-making



QnA Maker

Distill information into conversational, easy-to-navigate answers




Search



Bing Web Search API

Connect powerful search to your apps



Bing Image Search API

Bring advanced image and metadata search to your app



Bing Video Search API

Trending videos, detailed metadata, and rich results



Bing News Search API

Link your users to robust and timely news searches



Bing Autosuggest API

Give your app intelligent autosuggest options for searches



Bing Entity Search

Enrich user experiences with contextual entity search results



Bing Custom Search

Create a highly-customized web search experience



Microsoft Cognitive Services Labs

An early look at emerging Cognitive Services technologies:
CAUTION these are experimental and are not guaranteed to release.



Project Academic Knowledge

Explore relationships among academic papers, journals, and authors



Project Entity Linking

Contextually extend knowledge of people, locations, and events



Project Gesture

Gesture based controls



Project Event Tracking

Event associated with Wikipedia



Project Knowledge Exploration

Add interactive search over structured data to your project



Project Local Insights

Score location attractiveness



VISION

From faces to feelings, allow your apps to understand images and video

Computer Vision | Content Moderator | Emotion | Face | Video Indexer | Custom Vision Service

Computer Vision API

Analyze an image

Understand content within an image

OCR

Detect and recognize words within an image

Generate thumbnail

Scale and crop images, while retaining key content

Recognize celebrities

Thanks to domain specific models, ability to recognize 200K celebrities from business, politics, sports and entertainment around the world



Analyze image

Type of image

Clip Art Type	0 Non-clipart
Line Drawing Type	0 Non-Line Drawing
Black & White Image	False

Content of image

Categories	[{ "name": "people_swimming", "score": 0.099609375 }]
Adult Content	False
Adult Score	0.18533889949321747
Faces	[{ "age": 27, "gender": "Male", "faceRectangle": { "left": 472, "top": 258, "width": 199, "height": 199 } }]

Image colors

Dominant Color Background	White
Dominant Color Foreground	Grey
Dominant Colors	White
Accent Color	




Age: 27
Gender: Male

Is Adult Content: False
Categories: people_swimming

OCR

Life is like riding
a bicycle

To keep your
balance you must
keep moving




OCR

JSON:

```
{  
  "language": "en",  
  "orientation": "Up",  
  "regions": [  
    {  
      "boundingBox": "41,77,918,440",  
      "lines": [  
        {  
          "boundingBox": "41,77,723,89",  
          "words": [  
            {  
              "boundingBox": "41,102,225,64",  
              "text": "LIFE"  
            },  
            {  
              "boundingBox": "356,89,94,62",  
              "text": "IS"  
            },  
            {  
              "boundingBox": "539,77,225,64",  
              "text": "LIKE"  
            }  
          ]  
        }  
      ]  
    }  
  ]  
}
```

....




OCR

Good at

Scanned documents

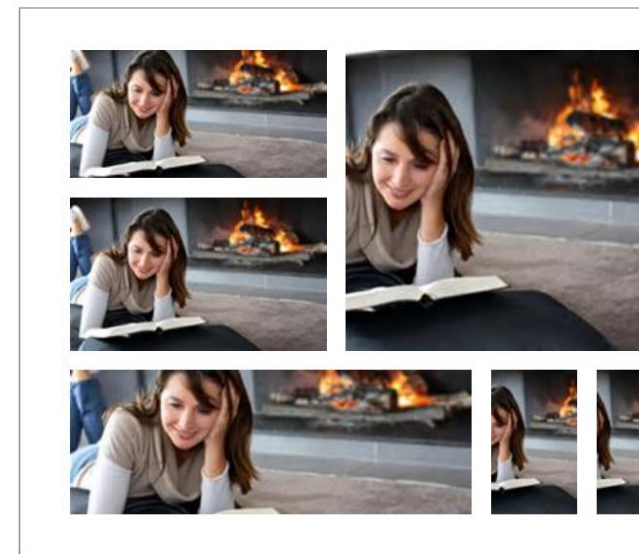
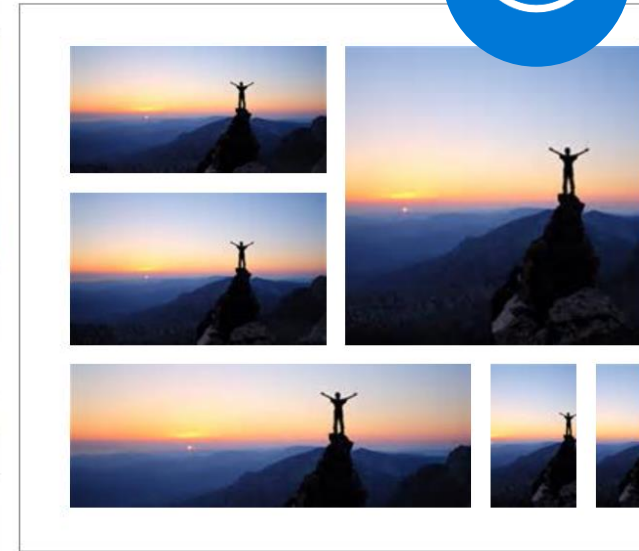

Photos with text

Fine grained
location information



Smart thumbnail

Smart cropping off



Face API

Face detection

Detect faces and their attributes within an image

Face verification

Check if two faces belong to the same person

Similar face searching

Find similar faces within a set of images

Face grouping

Organize many faces into groups

Face identification

Search which person a face belongs to



Face API



Detection

```
"faceRectangle": {"width": 193, "height": 193,  
"left": 326, "top": 204}
```

...

Feature attributes


```
"attributes": { "age": 42, "gender": "male",  
"headPose": { "roll": "8.2", "yaw": "-37.8",  
"pitch": "0.0" }}
```

Grouping



Identification

Jasper Williams



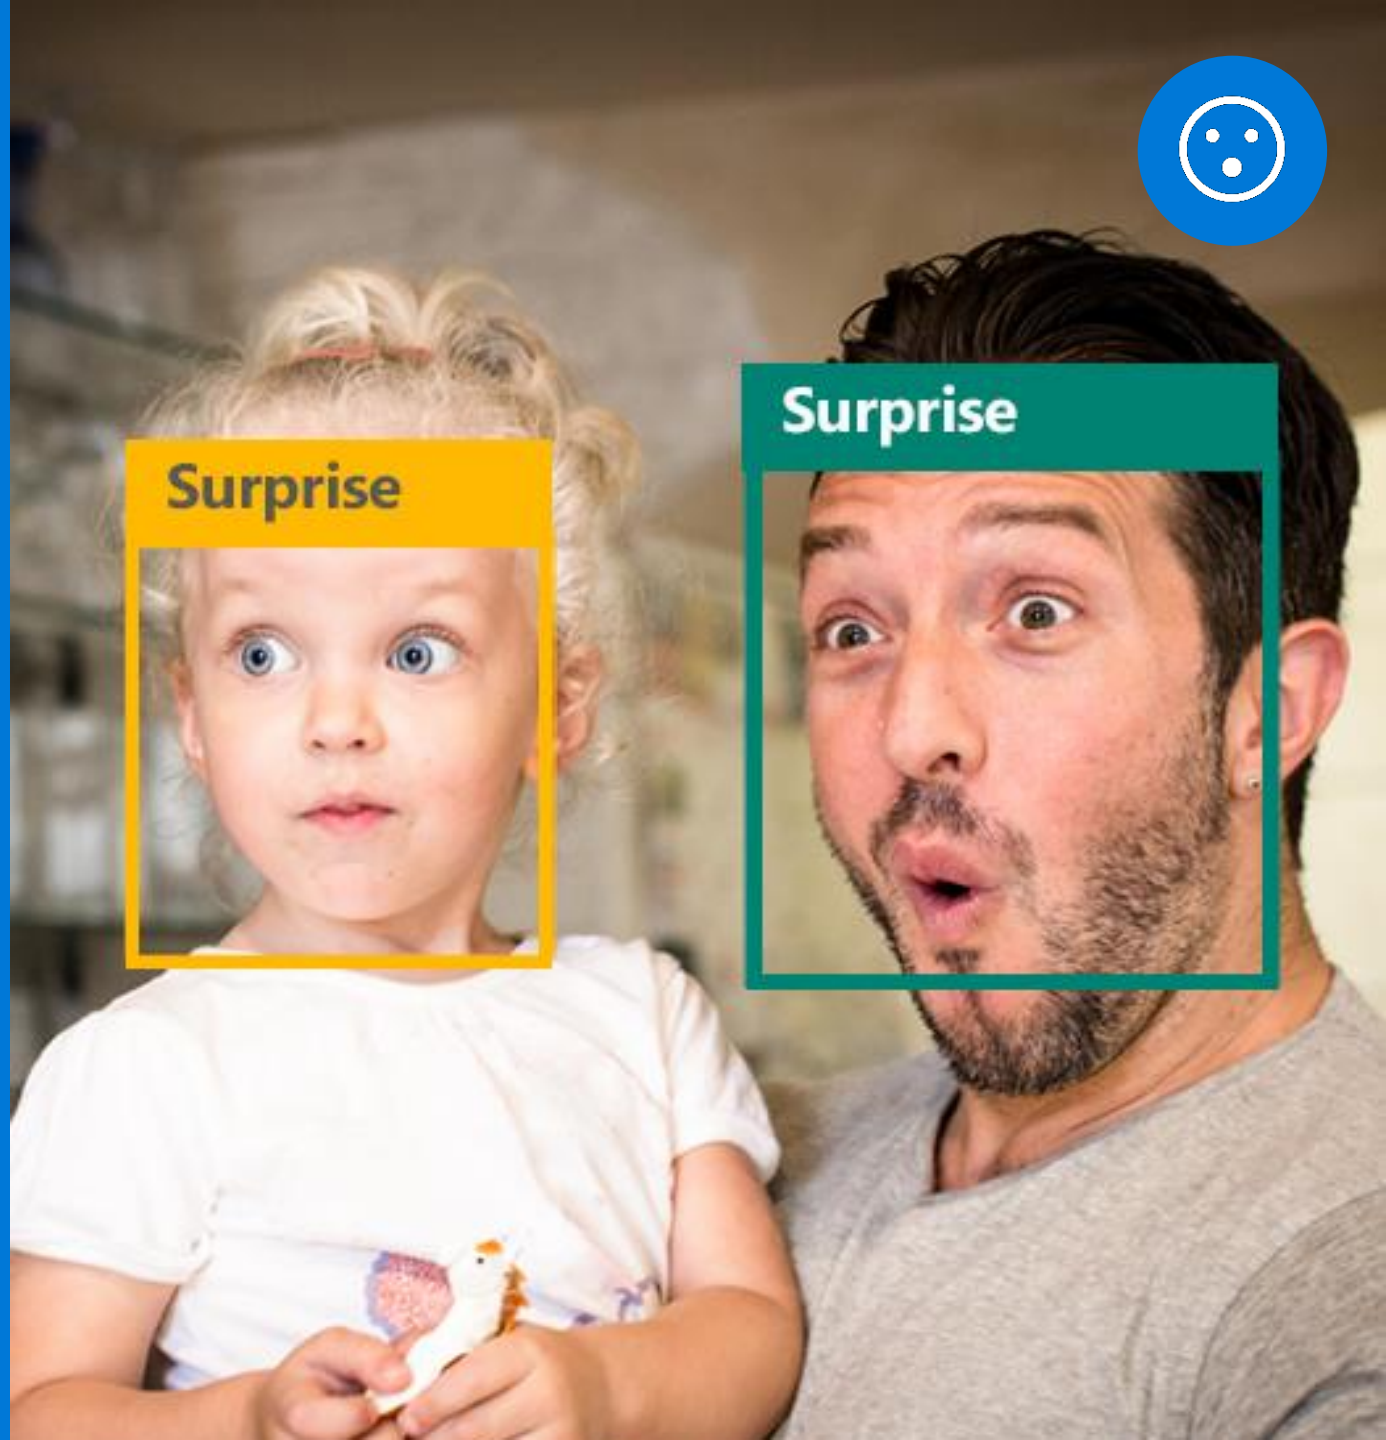
Face API

Face detection

```
"faceRectangle": {"width": 193,  
  "height": 193,  
  "left": 326,  
  "top": 204} ...
```

Emotion scores

```
"scores": { "anger": 5.182241e-8,  
  "contempt": 0.0000242813,  
  "disgust": 5.621025e-7,  
  "fear": 0.00115027453,  
  "happiness": 1.06114619e-8,  
  "neutral": 0.003540177,  
  "sadness": 9.30888746e-7,  
  "surprise": 0.9952837}
```



Content Moderator

Machine-assisted moderation of text and images, augmented with human review tools

Image moderation


Enhance your ability to detect potentially offensive or unwanted images through machine-learning based classifiers, custom blacklists, and Optical Character Recognition (OCR)

Text moderation

Helps you detect potential profanity in more than 100 languages and match text against your custom lists automatically. Content Moderator also checks for possible Personally Identifiable Information (PII)

Video moderation (in Azure Media Services)

Enable the scoring of possible adult content in videos. Video moderation is currently deployed in preview on Azure Media Services



Custom Vision Service

A customizable web service that learns to recognize specific content in imagery

Upload images

Upload your own labeled images, or use Custom Vision Service to quickly tag any unlabeled images

Train


Use your labeled images to teach Custom Vision Service the concepts you want it to learn

Evaluate

Use simple REST API calls to quickly tag images with your new custom computer vision model

Active learning

Images evaluated through your custom vision model become part of a feedback loop you can use to keep improving your classifier



Video Indexer

Unlock video insights

Upload your video and go

Start turning your video into insights right away. No more tedious and error-prone manual indexing. And no need for specialized expertise. With Video Indexer, just upload your video, and start finding insights right away, without writing a single line of code

Make your content more discoverable

Quickly and easily extract insights from videos using artificial intelligence. Enhance content discovery experiences such as search results by detecting spoken words, faces, characters, and emotions

Improve engagement with your video

Metadata extracted by Video Indexer can be used to build powerful engagement experiences with recommendations, highlight clips, and interactive videos





SPEECH

Hear and speak to your users by filtering noise, identifying speakers, and understanding intent

Custom Recognition | Speaker Recognition | Speech

Bing Speech API

Voice recognition (speech to text)

Converts spoken audio to text

Voice output (text to speech)

Synthesize audio from text

Speech intent recognition

Convert spoken audio to intent



Custom Speech Service

Customize both language and acoustic models

Tailor speech recognition to your app and environment



Custom Speech Service

Create custom language models for the vocabulary of the application

Adapt acoustic models to better match the expected environment of the application's users

Deploy to a custom endpoint and access from any device



Record audio



Transcribe



Adapt

Deploy



Speaker Recognition API

Speaker verification

Check if two voices are the same

Speaker identification

Identify who is speaking



Speaker Recognition API

Enrollment


Create a unique voiceprint for a profile

Recognition

After enrolling one or more voices, identify who is speaking from an audio clip

Verification

Confirm if a voice belongs to a previously enrolled profile





LANGUAGE

Process text and learn how to recognize what users want

Bing Spell Check | Language Understanding |
Linguistic Analysis | Text Analytics | Web Language Model |
Translator Text and Speech

Bing spell check API

State-of-the-art cloud-based spelling algorithms

Recognizes a wide variety of spelling errors

Recognize name errors and homonyms in context

Difficult to spot errors that use the context
of the words around them

Updates over time

Support for new brands and coined
expressions as they emerge



Bing spell check API

Check a single word or a whole sentence

"Our engineers developed this **four** you!"

Corrected Text: "four" → "for"

Identify errors & get suggestions

```
"spellingErrors": [  
  {  
    "offset": 5,  
    "token": "gona",  
    "type": "UnknownToken",  
    "suggestions": [  
      { "token": "gonna" }  
    ]  
  }  
]
```



A new service from **microso ft!**

Microsoft



Director **stephen** Spielberg should use it in the next AI movie!

Steven



Our service is like **lyft** for word processing!

Lyft

Language Understanding Intelligent Service

Understand what your users are saying

Use pre-built Bing and Cortana
models or create your own




Language Understanding Intelligent Service

Reduce labeling effort with interactive featuring

Use visualizations to gauge performance and improvements

Leverage speech recognition with seamless integration

Deploy using just a few examples with active learning



Language understanding models

"News about flight delays"



```
{
  "entities": [
    {
      "entity": "flight_delays",
      "type": "Topic"
    }
  ],
  "intents": [
    {
      "intent": "FindNews",
      "score": 0.99853384
    },
    {
      "intent": "None",
      "score": 0.07289317
    },
    {
      "intent": "ReadNews",
      "score": 0.0167122427
    },
    {
      "intent": "ShareNews",
      "score": 1.0919299E-06
    }
  ]
}
```



[My apps](#)[My keys](#)[Docs](#)[Pricing](#)[Support](#)[About](#)

Pizza Bot

Version: 0.1

[Settings](#)[Dashboard](#)[Intents](#)[Entities](#)[Prebuilt domains ^{PREVIEW}](#)[Features](#)[Train & Test](#)[Publish App](#)[← Back to App list](#)

Test your application

Use this tool to test the current and published versions of your application, to check if you are progressing on the right track ... [Learn more](#)

[Train Application](#)

Last train: Jan 18, 2017 2:41:01 PM | Last publish: Jan 18, 2017 2:41:25 PM

Interactive Testing [Batch Testing](#)

 Enable published model

Labels view (Ctrl+E)

Entities

[Reset console](#)

Type a test utterance & press Enter



i ' d like a [\$Size] pie with [\$Toppings]

i need a [\$Size] pizza with [\$Toppings] and tomatoes

i need a [\$Size] pizza

Current version results

Top scoring intent

Order-Pizza (0.88)

Other intents

None (0.06) Cancel (0)

Text analytics

Sentiment analysis

Understand if a record has positive or negative sentiment

Key phrase extraction

Extract key phrases from a piece of text, and retrieve topics

Language detection

Identify the language,
120 supported languages



Microsoft Translator

Translator Text API

Automatically detect language and easily power translation to and from 60 supported text languages

Translator Speech API

Easily translate real-time speech conversations in 9 support languages





KNOWLEDGE

Tap into rich knowledge amassed from the web, academia, or your own data

Academic Knowledge | Entity Linking |
Knowledge Exploration | Recommendations |
QnA Maker | Custom Decision Service

QnA Maker

Create a FAQ service from existing content

Extract questions and answers

Extract all possible pairs of questions and answers from user provided content – FAQ URLs, documents and editorial content

Test, train and publish

Edit, remove, or add pair before testing and training the knowledge base and publishing your knowledge base as an API endpoint

Integrates with other APIs and solutions

Use QnA Maker with Cognitive Services such as LUIS & create something as elegantly simple as a chat bot that answers FAQs, or as complex as an interactive virtual guide



Custom Decision Service

A cloud-based, contextual decision-making API that sharpens with experience.

Contextual

Understanding context from information you provide, Custom Decision Service ranks the options and makes a decision

Rapid learning

Custom Decision Service automatically optimizes based on your feedback. It even experiments with new options to see if the best decision has changed, enabling it to adjust to emerging trends

Easy to use

Custom Decision Service is cloud-based, so it's easy to run, able to plug into your application and help to make decisions in real time





SEARCH

Access billions of web pages, images, videos,
and news with the power of Bing APIs

[Bing Web Search](#) | [Bing Image Search](#) | [Bing News Search](#) |
[Bing Video Search](#) | [Bing Auto Suggest](#) | [Bing Custom Search](#) |
[Bing Entity Search](#)

Bing web search

**Intelligent search to your apps
with the ability to comb billions
of webpages, images, videos,
and news with a single API call**

Retrieve web documents indexed by Bing and
narrow the results down with filters such as
by answer type and freshness



Bing image search

A variety of image search options, from trending images to detailed insights


Scour the web for images and get results that include thumbnail and full image URLs, publishing website, image metadata, related images, and more



Bing video search


A variety of advanced video search features, including trending videos, price, and other useful metadata

Find videos from across the web and get responses that provide useful metadata including creator, encoding format, video size and quality, and source view count




Bing news search

Turn any app into a news desk with world news grouped and filtered by topic, local news, and metadata you can mine



Bing autosuggest

Query completion suggestions capabilities, so users can type less and get to what they want faster



Bing custom search

An easy-to-use, ad-free, commercial-grade search tool lets you deliver results you want

Customize your search

Quickly and reliably define the slices of the web that you want to draw from. Change the parameters of the sites you want and don't want at any time

Easy-to-use

Custom Search features a straightforward UI that enables you to create your web search without a line of code

Commercial-grade

Usage is ad-free – no matter how much or how little of the service you use



Bing entity search

Enrich user experiences with contextual entity search results

Tap into the power of the web to search for the most relevant entities (such as movies, books, famous people, U.S. local businesses...) and easily provide primary details and information sources about them

Augment your content with entity search results

Showcase of local businesses nearby



Academic knowledge

Interpret

Interprets a natural language user query string. Returns annotated interpretations which can enable rich search-box auto-completion experiences that anticipate what the user is typing

Evaluate

Evaluates a query expression and returns academic knowledge entity results

Calchistogram

Calculates a histogram of the distribution of attribute values for the academic entities returned by a query expression, such as the distribution of citations by year for a given author



Entity linking

Power your app's data links with named entity recognition and disambiguation

A word might be used as a named entity, a verb, or another word form within a given sentence

The Entity Linking Intelligence Service will recognize and identify each separate entity based on the context



Knowledge exploration

Enable interactive search experiences over structured data via natural language inputs

Attribute histograms

To enable rich visualization and interactive faceted experience

Structured query evaluation

To efficiently retrieve detailed information about matching objects

Query auto-completion

To reduce user effort and help with discovery of rich capabilities

Natural language understanding

To interpret natural language queries into structured query expressions



Deeper Dive

Invoking Computer Vision

```
subscription_key = "3c0062ef12644cc99e6c4f857972f2a2"
```

```
vision_base_url =  
"https://westus.api.cognitive.microsoft.com/vision/v1.0/"  
vision_analyze_url = vision_base_url + "analyze"
```

```
image_url = "https://demo-site.com/images/swimmer.png"
```

```
import requests  
headers = {'Ocp-Apim-Subscription-Key': subscription_key }  
params = {'visualFeatures':  
          'Categories,Description,Tags,Color,Adult'}  
data = {'url': image_url}
```

```
response = requests.post(vision_analyze_url, headers=headers,  
                          params=params, json=data)
```

```
analysis = response.json()
```



```
{'adult': {'adultScore': 0.023512152954936028,  
          'isAdultContent': False,  
          'isRacyContent': False,  
          'racyScore': 0.04207553341984749},  
 'categories': [{'name': 'others_', 'score': 0.39453125},  
                {'name': 'trans_car', 'score': 0.44140625}],  
 'color': {'accentColor': '895D42',  
          'dominantColorBackground': 'White',  
          'dominantColorForeground': 'White',  
          'dominantColors': ['White'],  
          'isBwImg': False},  
 'description': {'captions': [{'confidence': 0.9485308427051494,  
                               'text': 'a truck is parked on the side of a road'}],  
 'tags': ['outdoor',  
          'road',  
          'truck',  
          'car',  
          'parked',  
          'street',  
          'large',  
          'traffic']},  
 'metadata': {'format': 'Jpeg', 'height': 1080, 'width': 1920},  
 'requestId': '046593b3-2313-4867-bc83-6b7eec88fad7',  
 'tags': [{'confidence': 0.9950141310691833, 'name': 'outdoor'},  
          {'confidence': 0.9936342239379883, 'name': 'road'},  
          {'confidence': 0.981715738773346, 'name': 'truck'},  
          {'confidence': 0.749627411365509, 'name': 'transport'},  
          {'confidence': 0.16133838891983032, 'name': 'trailer'}]}
```



Computer Vision - Scenarios

Good For:

- Object detection - Automatic captioning and tagging of images
- Extracting printed text from images
- Reading handwritten text on whiteboards, paper and sticky notes

Not So Good For:

- Reading license plates
- Object identification

Computer Vision - Considerations

General:

Image can be supplied as uploaded binary or as publicly available URL (TLS supported)

Requires an Internet connection – no offline version (see Custom Vision Service)

Not currently HIPAA or PCI DSS compliant

see latest list: <https://gallery.technet.microsoft.com/Overview-of-Azure-c1be3942>

Categories:

Category taxonomy is fixed, see list at <https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/category-taxonomy>

Domain-specific details:

Currently limited to celebrities and landmarks

Faces:

Produces only the face rectangle and gender, use Face API for more details including emotions.

Computer Vision - Considerations

Tagging & Captioning:

Custom tags not currently supported

Yes, tags are repeated in the response. Once with confidence, one without. Both sorted by confidence. The tags provided in the description object are the exhaustive list, those in the tags object are just the top tags.

If you want to choose from multiple candidate captions, use the “describe” operation instead of “analyze”
Analyze operations supports English and simplified Chinese only. Describe operations supports English only.

Text Recognition:

OCR not in all languages – currently at 18 languages (see documentation for latest list)

ICR (handwriting recognition) supports only English

ICR results can take several seconds, so expect multi-second delay for longer texts

ICR can handle text that is rotated by 30 to 40 degrees, more than that you will need to pre-process image

Computer Vision - Considerations

Video:

Computer Vision does not support video as an input, instead you need to send frames as individual image files.

Regional Availability:

Available in many but not all regions (see documentation for latest list)

Limits:

Min image dimensions are 50x50 pixels

Max image size 4MB

Supported image formats are JPEG, PNG, GIF, BMP

Standard tier supports max of 10 calls per second, but this is a soft limit raised with a call to support

Easter Egg:

You can read the research papers behind Computer Vision [here](#)

Invoking Text Analytics (sentiment)

```
text_analytics_subscription_key =
"6fd91a9f199a49f29e748543fdc05ea4"
text_analytics_base_url =
"https://westus.api.cognitive.microsoft.com/text/analytics/
v2.0/"
sentiment_api_url = text_analytics_base_url + "sentiment"
claim_text = """"We are just happy the damage was minimal
and that everyone is safe. We are thankful for your
support.""""
documents = {'documents' : [
    {'id': '1', 'language': 'en', 'text': claim_text}
]}

headers = {"Ocp-Apim-Subscription-Key":
text_analytics_subscription_key}
response = requests.post(sentiment_api_url,
headers=headers, json=documents)
sentiments = response.json()

score = sentiments['documents'][0]['score']
score_interpretation = "neutral"
if (score < 0.45):
    score_interpretation = "negative"
elif (score >= 0.55):
    score_interpretation = "positive"
score_interpretation
```

```
{
  'documents': [
    {
      'id': '1',
      'score': 0.934346616268158
    }
  ],
  'errors': []
}
```

```
score_interpretation = 'positive'
```



Invoking Text Analytics (key phrases)

```
text_analytics_subscription_key =  
"6fd91a9f199a49f29e748543fdc05ea4"  
text_analytics_base_url =  
"https://westus.api.cognitive.microsoft.com/text/analytics/  
v2.0/"
```

```
keyphrase_api_url = text_analytics_base_url + "keyPhrases"
```

```
claim_text = "I was driving down El Camino and stopped at  
a red light.."
```

```
documents = {'documents' : [  
    {'id': '1', 'language': 'en', 'text': claim_text}  
]}
```

```
headers = {"Ocp-Apim-Subscription-Key":  
text_analytics_subscription_key}  
response = requests.post(keyphrase_api_url,  
headers=headers, json=documents)  
key_phrases = response.json()
```

```
key_phrases
```

I was driving down El **Camino** and stopped at a **red light**. It was about 3pm in the **afternoon**. The sun was bright and shining just behind the **stoplight**. This made it hard to see the **lights**. There was a **car** on my left in the **left turn lane**. A few **moments** later another car, a **black sedan** pulled up behind me. When the left turn light changed **green**, the **black sedan** hit me thinking that the light had changed for us, but I had not moved because the light was still red. After hitting my **car**, the **black sedan** backed up and then sped past me. I did manage to catch its **license plate**. The **license plate** of the **black sedan** was **ABC123**.

```
{'documents': [  
    {'id': '1',  
    'keyPhrases': ['black sedan',  
    'car',  
    'red light',  
    'left turn lane',  
    'license plate',  
    'ABC123',  
    'Camino',  
    'moments',  
    'afternoon',  
    'stoplight',  
    'lights'  
    ]  
    }  
],  
'errors': []  
}
```

Text Analytics - Scenarios

Good For:

- Scoring sentiment
- Highlighting key phrases in text
- Detecting the languages in a collection of documents

Not So Good For:

- Detecting mood
- Dealing with sarcasm
- Summarizing texts
- Detecting multiple languages in a single document

Text Analytics - Considerations

General:

Requires an Internet connection – no offline version

Model quietly updated over time, expect that same text may yield different results in time.

18 languages supported – but varies by operation (sentiment or key-phrases). See [list in docs](#).

Not currently HIPAA or PCI DSS compliant

see latest list: <https://gallery.technet.microsoft.com/Overview-of-Azure-c1be3942>

Language Detection:

Detects only the predominant language in mixed language documents and provides a confidence score (which will be less than 1 in this case)

Sentiment Scoring:

Most accurate for short texts, 1-3 sentences, not as accurate for longer texts.

Sentiment of exactly 0.5 may mean it was not processed for sentiment (expects English, but Spanish text provided or the longer text was deemed objective)

Text Analytics - Considerations

Key Phrase Extraction:

Performs better on large blocks of text.

A domain specific vocabulary is not supported, meaning some words may not be picked up (like scientific terms).

Output consists of nouns and objects of the sentence.

Output listed in descending order of importance.

Importance is measured by the number of times a particular concept is mentioned, or the relation of that element to other elements in the text.

Limits:

Max document size of 5k characters (about 2 pages of text)

Max size of entire request is 1 MB

Max of 1k documents per request

Max of 100 calls per minute

Regional Availability:

Available in many but not all regions (see documentation for latest list)



Get started for free at
<http://azure.com/Cognitive>



Learn more on the Cortana
Intelligence Suite [website](#) and
Cognitive Services [website](#)



Schedule a workshop to identify
areas in your business where
analytics and intelligence can
drive transformation



Talk with your Microsoft contact
about licensing options and partners



DEVELOPER RESOURCES

Pricing

<https://azure.microsoft.com/en-us/pricing/details/cognitive-services/>

Documentation

<https://docs.microsoft.com/en-us/azure/#pivot=products&panel=cognitive>

Client SDKs

<https://azure.microsoft.com/en-us/resources/samples/?sort=0&term=cognitive+services>
<https://github.com/southwood/project-oxford-python>

Example Code

<https://github.com/jsturtevant/happy-image-tester-django>
<https://github.com/Microsoft/Cognitive-Face-Android>
<https://github.com/Microsoft/Cognitive-Samples-IntelligentKiosk>

Join Our Community

<https://stackoverflow.com/questions/tagged/microsoft-cognitive>
<https://cognitive.uservoice.com/>

Q&A