



Democratization of AI using Microsoft Cognitive Services



Adnan Masood, PhD

Microsoft MVP

Adnan.Masood@owasp.org

@adnanmasood

<https://github.com/adnanmasood>

Slides Courtesy of Microsoft Corporation

Tampa Bay Data Science Group

- ✓ Business Intelligence
- ✓ Statistical Computing
- ✓ Hadoop
- ✓ Big Data
- ✓ Natural Language Processing
- ✓ Machine Learning
- ✓ Text Analytics
- ✓ Data Analytics
- ✓ Data Visualization
- ✓ Predictive Analytics
- ✓ Data Mining
- ✓ Analytics
- ✓ Text Mining
- ✓ Data Science
- ✓ Big Data Analytics



About the Speaker



Adnan Masood, Ph.D. is a software architect, machine learning researcher, and Microsoft MVP for Data Platform. Before joining UST Global as Chief Architect of AI and Machine Learning, Dr. Masood worked at Green Dot Corporation, a leading prepaid financial technology institution as a Sr. Systems Architect. In the past life he has also served as principal engineer for an ecommerce start-up, and as a solutions architect for a leading British nonprofit organization.

A strong believer in the development community, Adnan is an active member of the Open Web Application Security Project (OWASP), an organization dedicated to software security. In the .NET community, he is a cofounder and president of the Pasadena .NET Developers group, co-organizer of Tampa Bay Data Science Group, and Irvine Programmer meetup. A certified ScrumMaster, Dr. Masood also hold certifications in big data, machine learning, and systems architecture from Massachusetts Institute of Technology; Application Security certification from Stanford University, and SOA Smarts certification from Carnegie Mellon University. he is a Microsoft Certified Solutions Developer, and Sun Certified Java Developer.

Dr. Masood teaches Data Science course 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 (, IEEE-HST, IASA, and DevConnections), local code camps, and user groups. He is also a volunteer STEM FLL robotics coach for middle school students.

For more details, visit Adnan's blog (<http://blog.adnanmasood.com>), GitHub repository (<http://github.com/adnanmasood>), and Twitter (@adnanmasood). Adnan can be reached at adnan.masood@owasp.org.

Drive Intelligence from Text in Smart Apps

The session covers how to use cognitive services to drive insights and intelligence in your applications. The session covers how to work with unstructured text and turn unstructured text into meaningful insights into mobile, web and line of business applications.

The session will be showing how to use a few lines of code to easily analyze sentiment, extract key phrases, detect topics, and detect language for any kind of text.

The session will provide an overview on Microsoft Cognitive Services and all related text analysis services including:

- Sentiment Analysis
- Key Phrase extraction
- Topic Detection
- Language detection

The session is code driven & will provide samples on how to build smart apps with cognitive services from Microsoft.

Microsoft Cognitive Services

Give your apps
a human side

Slides Courtesy of Microsoft Corporation



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

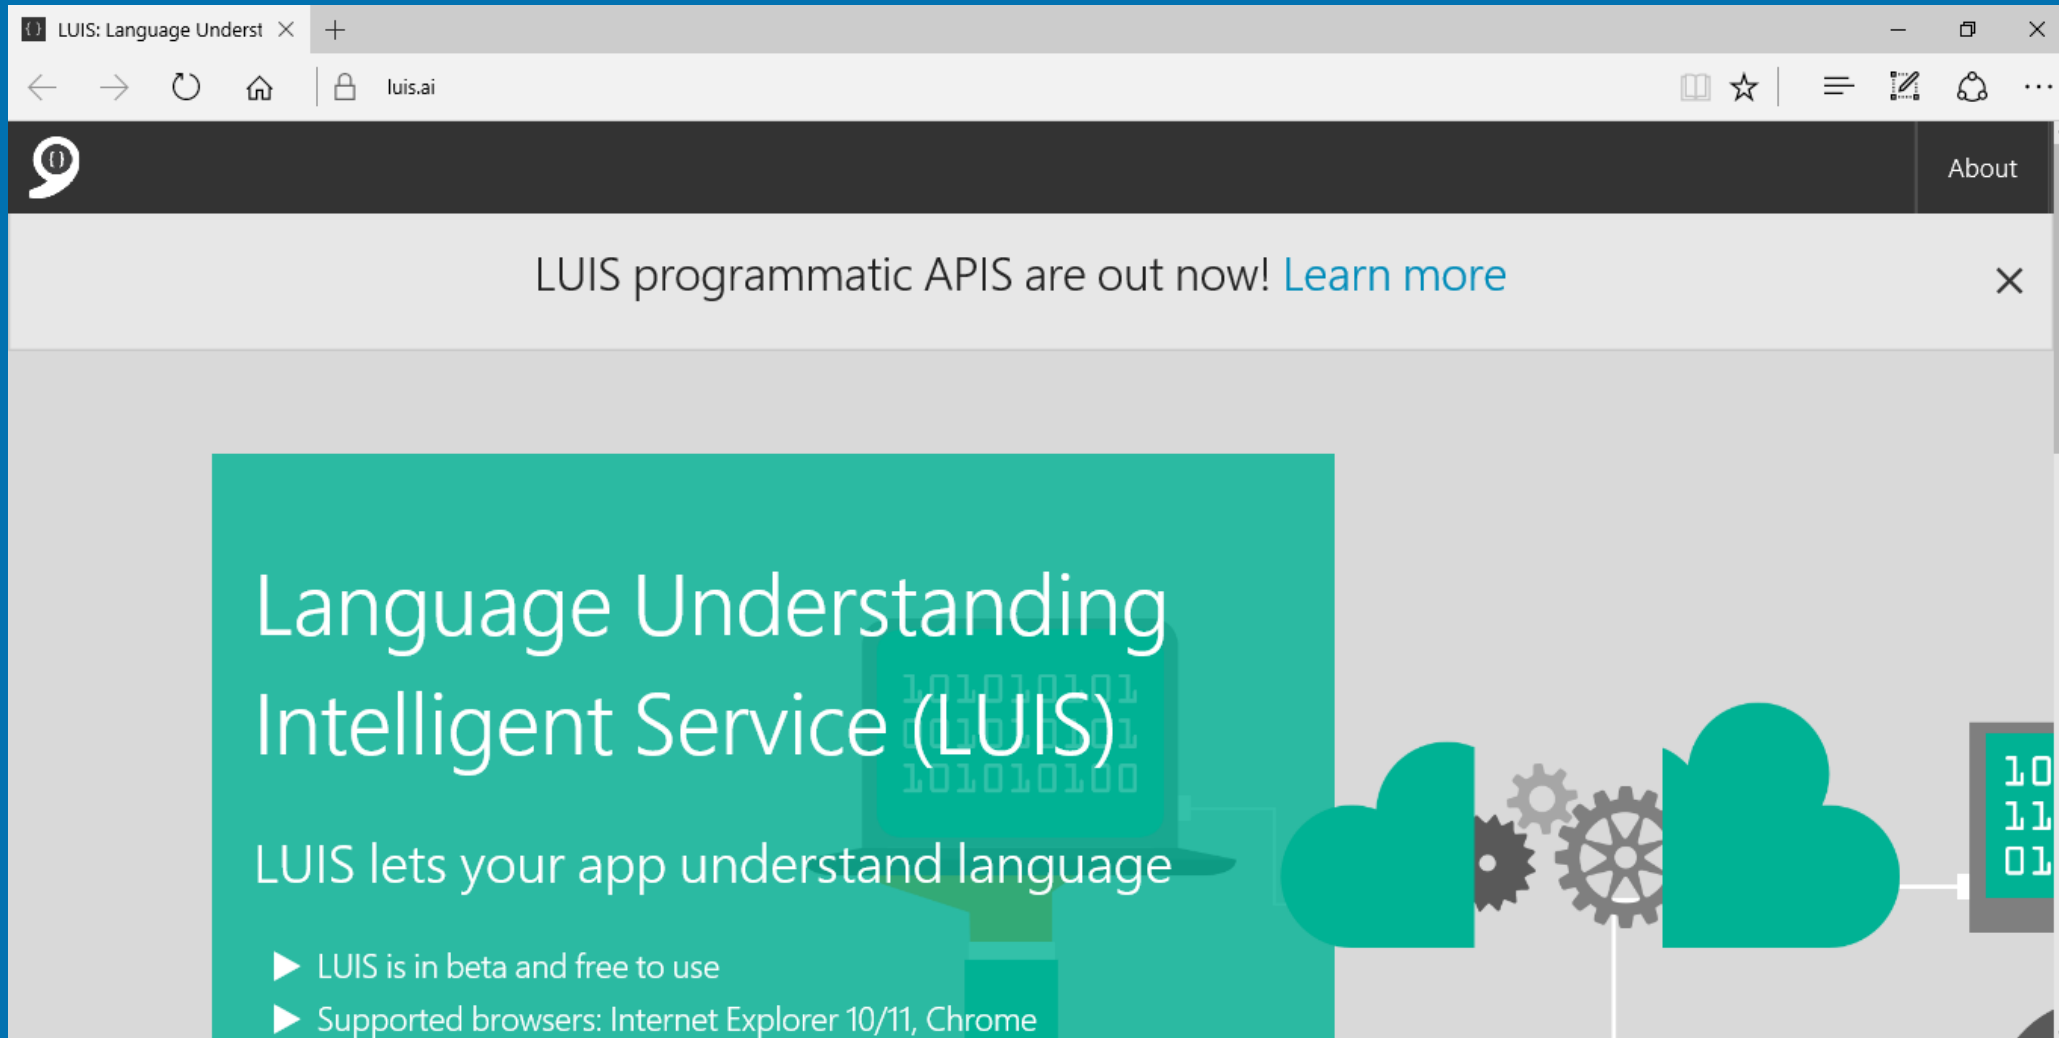
Tap into rich knowledge amassed from the web, academia, or your own data



Search

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

Language Understanding Intelligent Service



The screenshot shows a web browser window with the URL [luis.ai](https://www.luis.ai). The page features a dark navigation bar with a speech bubble icon and an 'About' link. A white notification banner at the top reads 'LUIS programmatic APIs are out now! [Learn more](#)'. The main content area has a light gray background with a large green graphic on the left containing the text 'Language Understanding Intelligent Service (LUIS)' and 'LUIS lets your app understand language'. To the right of this graphic is a decorative illustration of teal clouds, gray gears, and a small screen displaying binary code. Below the main text, there are two bullet points: '▶ LUIS is in beta and free to use' and '▶ Supported browsers: Internet Explorer 10/11, Chrome'.

LUIS programmatic APIs are out now! [Learn more](#)

Language Understanding Intelligent Service (LUIS)

LUIS lets your app understand language

- ▶ LUIS is in beta and free to use
- ▶ Supported browsers: Internet Explorer 10/11, Chrome

<https://www.luis.ai/>

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

Tested

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



GitHub



Scenarios

Emotion detection
at retail displays

Facial identification to
find missing children

Sentiment analysis
to learn how
customers feel


Facial detection
to calculate the
male/female ratio
at a nightclub

Language
understanding to allow
automated support
bots to understand
natural language

Object recognition to
enable a blind person
to read a menu



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



Video API

Analyze, edit, and process videos within your app

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 Recognition Intelligent Service

Fine-tune speech recognition for anyone, anywhere

Language



Bing Spell Check API

Detect and correct spelling mistakes within your app



Web Language Model API

Leverage the power of language models trained on web-scale data



Linguistic Analysis API

Easily parse complex text with language analysis



Language Understanding Intelligent Service

Teach your apps to understand commands from your users



Text Analytics API

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



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

Topic detection

Use clustering techniques to identify the trending topics on a large set of text records

Language detection

Identify the language, 120 supported languages



Text analytics

Sentiment analysis **English, Spanish, French, and Portuguese**
Understand if a record has positive or negative sentiment

Key phrase extraction **English, Spanish, German, and Japanese**
Extract key phrases from a piece of text, and retrieve topics

Topic detection **English**
Use clustering techniques to identify the trending topics on a large set of text records

Language detection
Identify the language, 120 supported languages

Demo

Text analytics

<http://text-analytics-demo.azurewebsites.net>

Language understanding (LUIS)

Define entities and intents

Entities–DepartureCity, ArrivalCity, DepartureDate, ReturnDate

Intent–book a flight

Map some utterances to an intent

Examples: “I want to go to Paris from Sept 25 to Sept 29, 2016”, “Book me a flight from DTW to CDG leaving on 9/25/2016 and returning 9/28/2016”, etc.

Help your model improve over time based on real feedback

See what real users are sending to your model, and map those utterances to intents (or create new intents based on what your users are asking).



Knowledge



Academic Knowledge API

Explore relationships among academic papers, journals, and authors




Knowledge Exploration Service

Add interactive search over structured data to your project



Entity Linking Service

Contextually extend knowledge of people, locations, and events



Recommendations API

Provide personalized product recommendations for your customers

Apps Powered by MS Cognitive Services

I think it's a person sitting in front of a computer and he seems 😊. I am 99% sure that's **Bill Gates**



CaptionBot.ai

CelebsLike.Me
Which Oscar nominee do you look like?

Your Oscar nominee match is...



Anthony Hopkins






Celebslike.me



ProjectMurphy.net

Cognitive Services

microsoft.com/cognitive

 Vision	 Speech	 Language	 Knowledge	 Search
Computer Vision	Custom Recognition	Bing Spell Check	Academic Knowledge	Bing Web Search
Emotion	Speaker Recognition	Linguistic Analysis	Entity Linking	Bing Image Search
Face	Speech	Language Understanding	Knowledge Exploration	Bing Video Search
Video	Translator	Text Analytics	Recommendations	Bing News Search
		WebLM		Bing Autosuggest

Cognitive Services

microsoft.com/cognitive

 Vision

 Speech

 Language

 Knowledge

 Search

Computer Vision

Custom
Recognition

Bing Spell Check

Academic
Knowledge

Bing
Web Search

Emotion

Speaker
Recognition

Linguistic Analysis

Entity Linking

Bing
Image Search

Face

Speech

Language
Understanding

Knowledge
Exploration

Bing
Video Search

Video

Translator

Text Analytics

Recommendations

Bing
News Search

WebLM

Bing
Autosuggest

Vertical Search APIs

Get more results, features and metadata tailored to each search vertical

Image Search API



<https://bingapis.azure-api.net/v5/images/search?q=shuttle+launch>

- Enhanced metadata and filters (size, license, style, freshness, color)
- Image insights (entity recognition, visually similar)


Video Search API



<https://bingapis.azure-api.net/v5/videos/search?q=viral+videos>

- Enhanced metadata and filters (price, resolution, length, freshness)
- Motion thumbnails (video preview)

News Search API



<https://bingapis.azure-api.net/v5/news/search?q=cuba>

- News by category/market, and trending news
- Rich article metadata (featured entities)

*screenshots show actual search results in bing.com

Accessing the APIs

1. Obtain API subscription key from microsoft.com/cognitive
2. Call REST endpoint, and pass API key via special header

```
GET https://bingapis.azure-api.net/v5/search?q=nasa HTTP/1.1  
OCP-Apim-Subscription-Key: <API KEY>
```

LUIS + Computer Vision

Language Understanding Models



```
{
  "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
    }
  ]
}
```

Language Understanding Models

Reduce labeling effort with interactive featurizing

Seamless integration to Speech API

Deploy using just a few examples with active learning

Supports 5 languages (English, Chinese, Italian, French, Spanish)



Demo

Updated Computer Vision API



Content of Image:

Categories

```
v0: [{ "name": "animal", "score": 0.9765625 }]  
V1: [{ "name": "grass", "confidence": 0.9999992847442627 },  
  { "name": "outdoor", "confidence": 0.9999072551727295 },  
  { "name": "cow", "confidence": 0.99954754114151 },  
  { "name": "field", "confidence": 0.9976195693016052 },  
  { "name": "brown", "confidence": 0.988935649394989 },  
  { "name": "animal", "confidence": 0.97904372215271 },  
  { "name": "standing", "confidence": 0.9632768630981445 },  
  { "name": "mammal", "confidence": 0.9366017580032349,  
    "hint": "animal" },  
  { "name": "wire", "confidence": 0.8946959376335144 },  
  { "name": "green", "confidence": 0.8844101428985596 },  
  { "name": "pasture", "confidence": 0.8332059383392334 },  
  { "name": "bovine", "confidence": 0.5618471503257751,  
    "hint": "animal" },  
  { "name": "grassy", "confidence": 0.48627158999443054 },  
  { "name": "lush", "confidence": 0.1874018907546997 },  
  { "name": "staring", "confidence": 0.165890634059906 }]
```

Describe

```
0.975 "a brown cow standing on top of a lush green field"  
0.974 "a cow standing on top of a lush green field"  
0.965 "a large brown cow standing on top of a lush green field"
```

Translator API

NEW:

- Translate speech

Not NEW - but still very useful:

- Translate text between 50 languages, any to any
- Highly customizable translation
 - Collaborative methods for engaging the community to improve translation
 - Self-service custom training, using your previously translated documents
- AJAX, REST and SOAP interface
- Methods:
 - Translate, Detect, Speak, AddTranslation, GetTranslations, BreakSentences
 - Array variants of the above

Developer Call to Action

- Sign up and get started today for free at

www.microsoft.com/cognitive

Developer Resources

Preview Pricing

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

Documentation

<https://www.microsoft.com/cognitive-services/en-us/computer-vision-api/documentation>

Client SDKs and Samples

<https://www.microsoft.com/cognitive-services/en-us/sdk-sample>

Join Our Community

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

<https://social.msdn.microsoft.com/forums/azure/en-US/home?forum=mlapi>

<https://cognitive.uservoice.com/>

Q & A