Microsoft

Artificial Intelligence

Current State and Future Considerations

Greg Milligan
Microsoft Canada

Envision

Artificial Intelligence

deep learning

machine learning

neural networks

search

probabilistic reasoning

etc.

A set of technologies that enable machine intelligence to simulate or augment elements of human thinking

vision

speech

language

knowledge

problem solving

etc.

AI, Machine Learning, Deep Learning

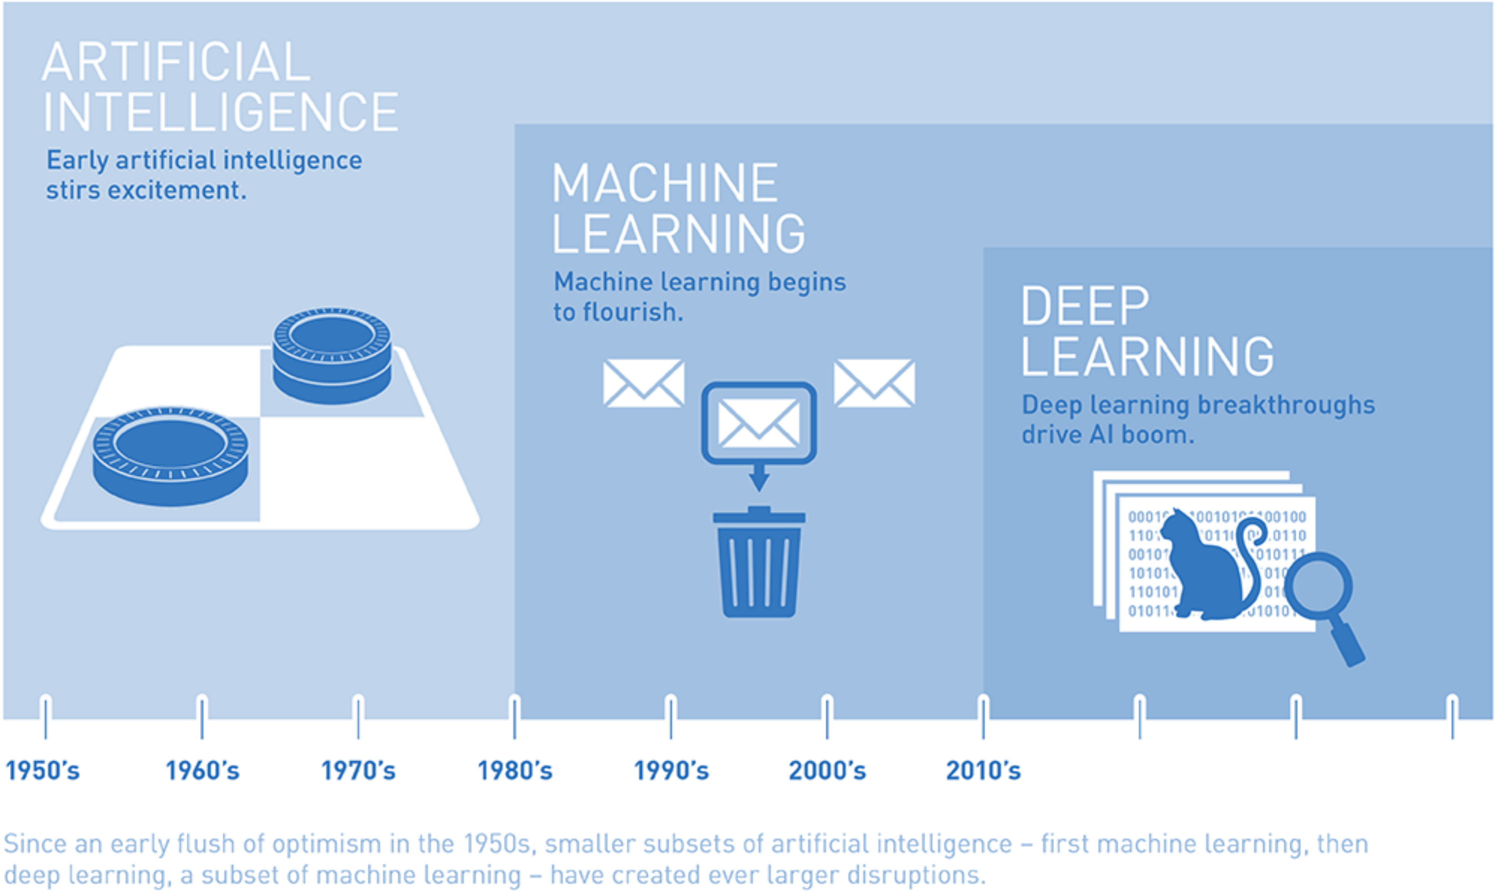

"What's the Difference Between Artificial Intelligence, Machine Learning, and Deep Learning?", Michael Copeland, 2016

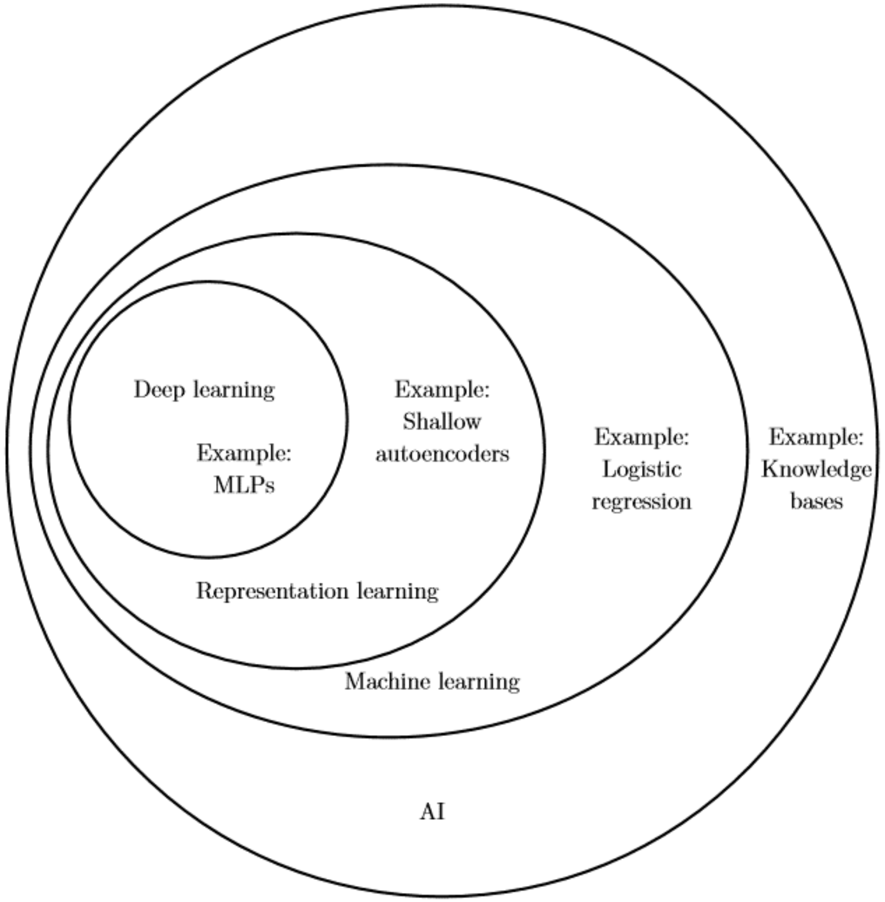

"Deep Learning", Ian Goodfellow, 2016

How is it different?

Teach machines, not instruct

+

Increasing compute power

Massive growth in data

Advanced statistical methods

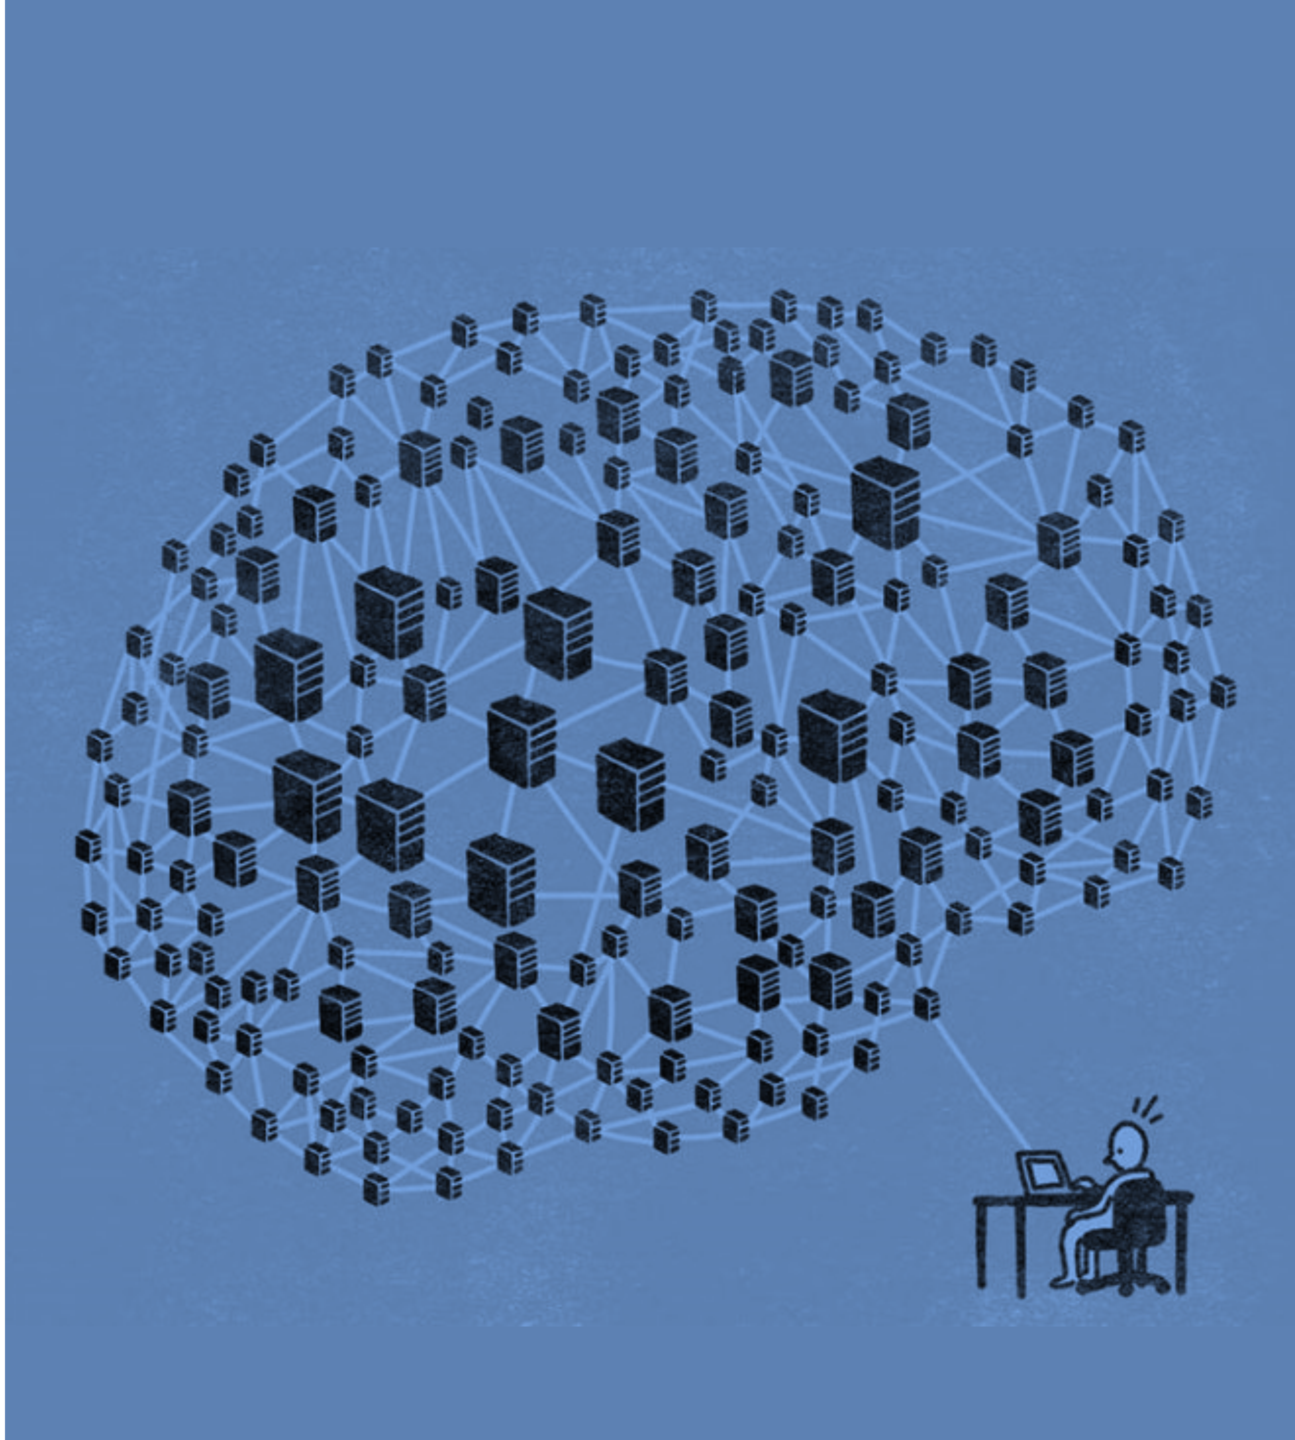

Where is it?

Alexa

amazon

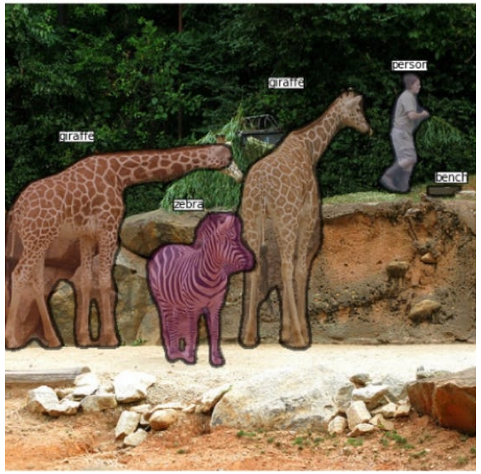

Facebook AI Research

facebook

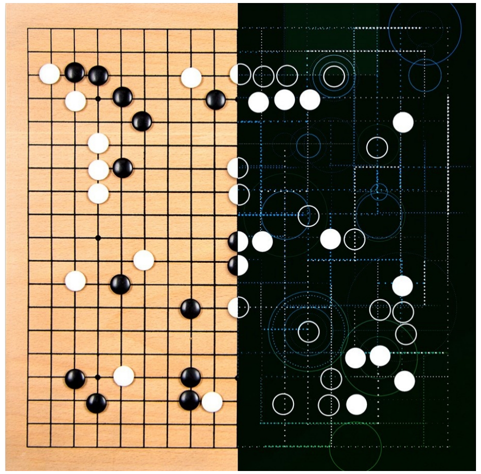

AlphaGo

Google

skype Translator

Microsoft

How you can use it?

process automation

task/decision augmentation

bring intelligence to data

customer service

etc.

but – **start with the problem**

Machine Learning

Designing a solution using Machine Learning

Sample device telemetry analytics solution architecture

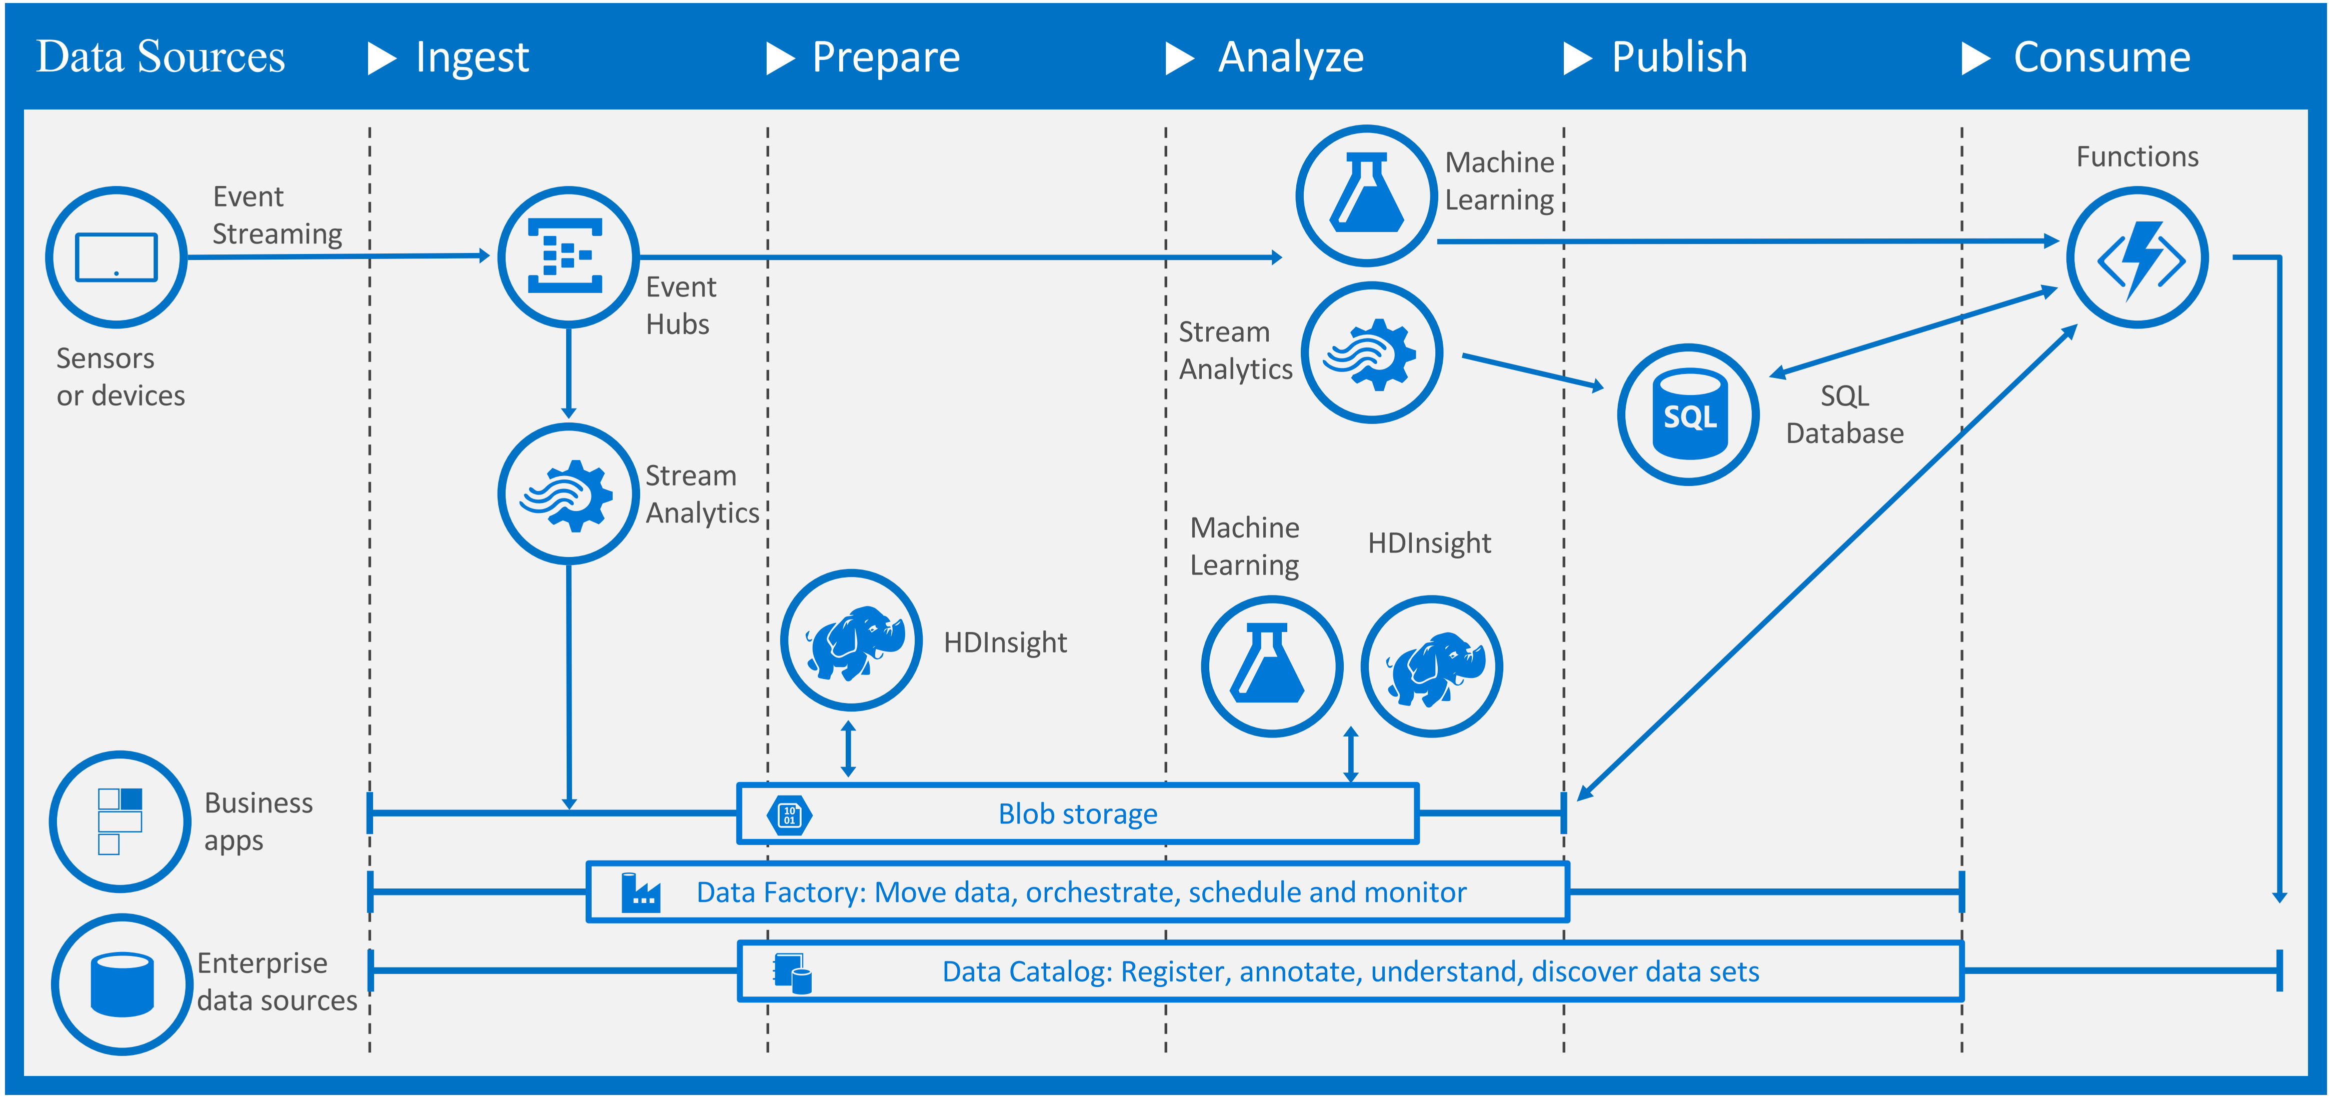

Deep Learning

Deep Learning

Based on advancements in artificial neural networks

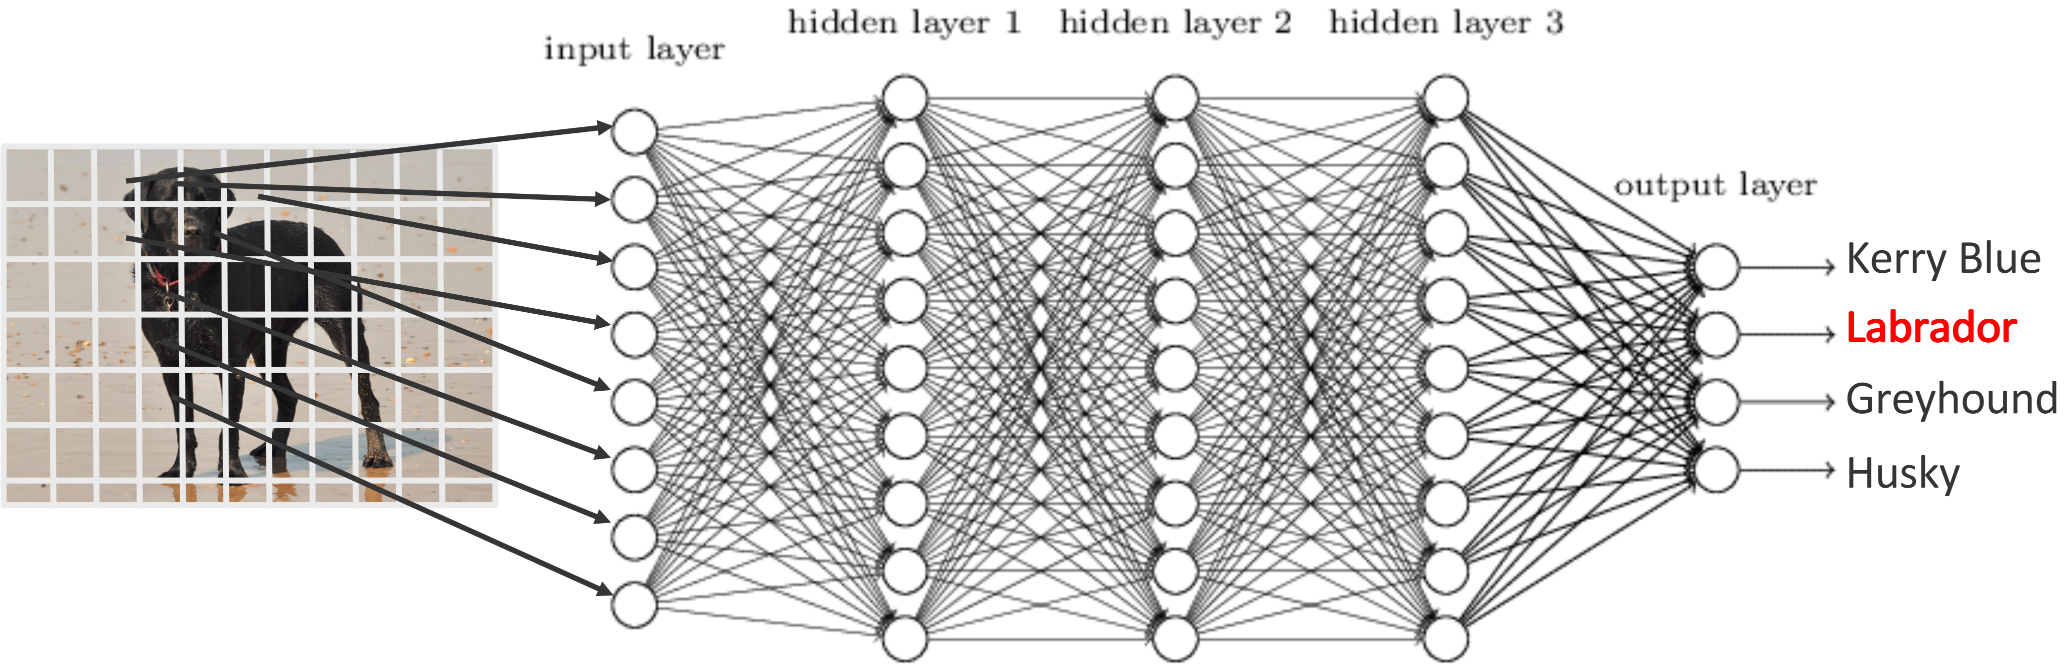

Larger and deeper networks

- Many layers; some up to 150 layers
- Billions of learnable parameters
- Feed Forward, Recurrent, Convolutional, Sparse, etc.

Training on big data sets

- 10,000+ hours of speech
- Millions of images
- Years of click data

Highly parallelized computation

- Long-running training jobs (days, weeks, months)
- Acceleration with GPU
- Recent advances in more computer power and big data

Grocery item object detection and recognition

- Automated grocery inventory management in connected refrigerators
- Implemented Fast R-CNN object detection in CNTK. REST API published using Python Flask
- Annotated 311 images, split into 71 test and 240 training images. In total 2578 annotated objects, i.e. on average 123 examples per class
- Prototype classifier has a precision of 98% at a recall of 80%, and 93% precision at recall of 90%

LIEBHERR

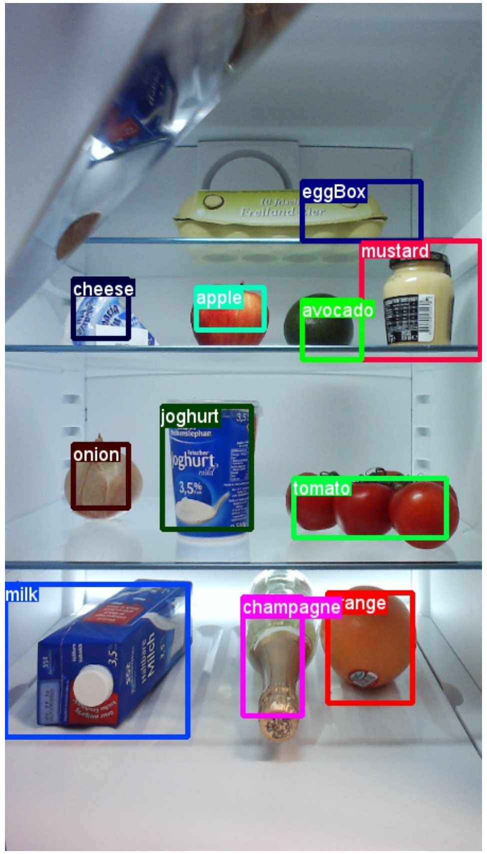

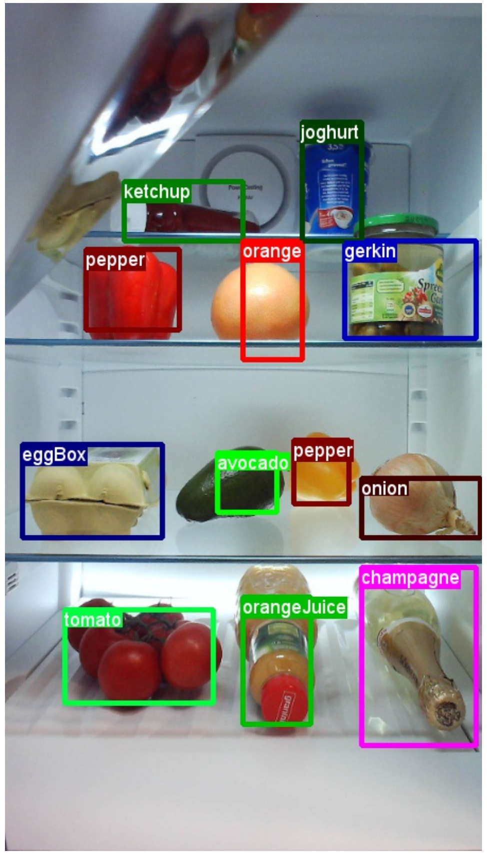

- https://blogs.technet.microsoft.com/machinelearning/2016/09/02/microsoft-and-liebherr-collaborating-on-new-generation-of-smart-refrigerators/

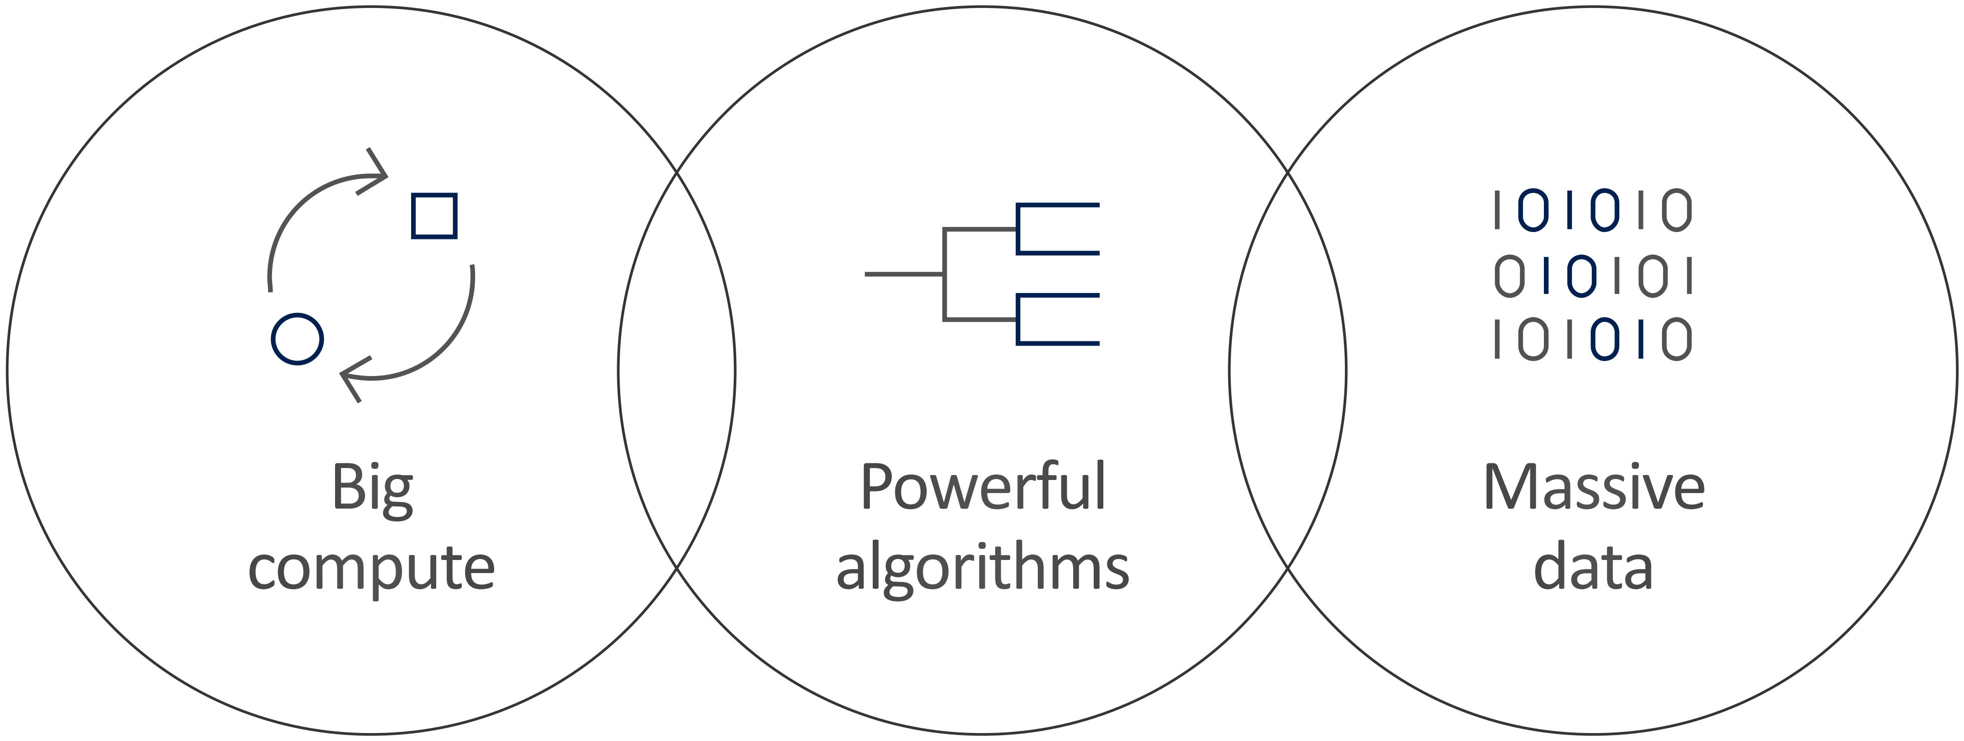
1010101
0101010
1010101
Big compute
Powerful algorithms
Massive data

Examples 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.

Speech

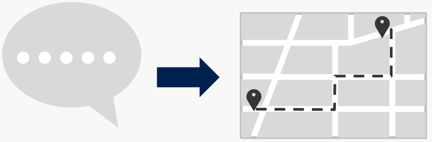

- Convert spoken audio to text
- Convert text to spoken audio
- Extract intent of user

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'

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...

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	

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"
      }
. . .
```


Emotion 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}
```

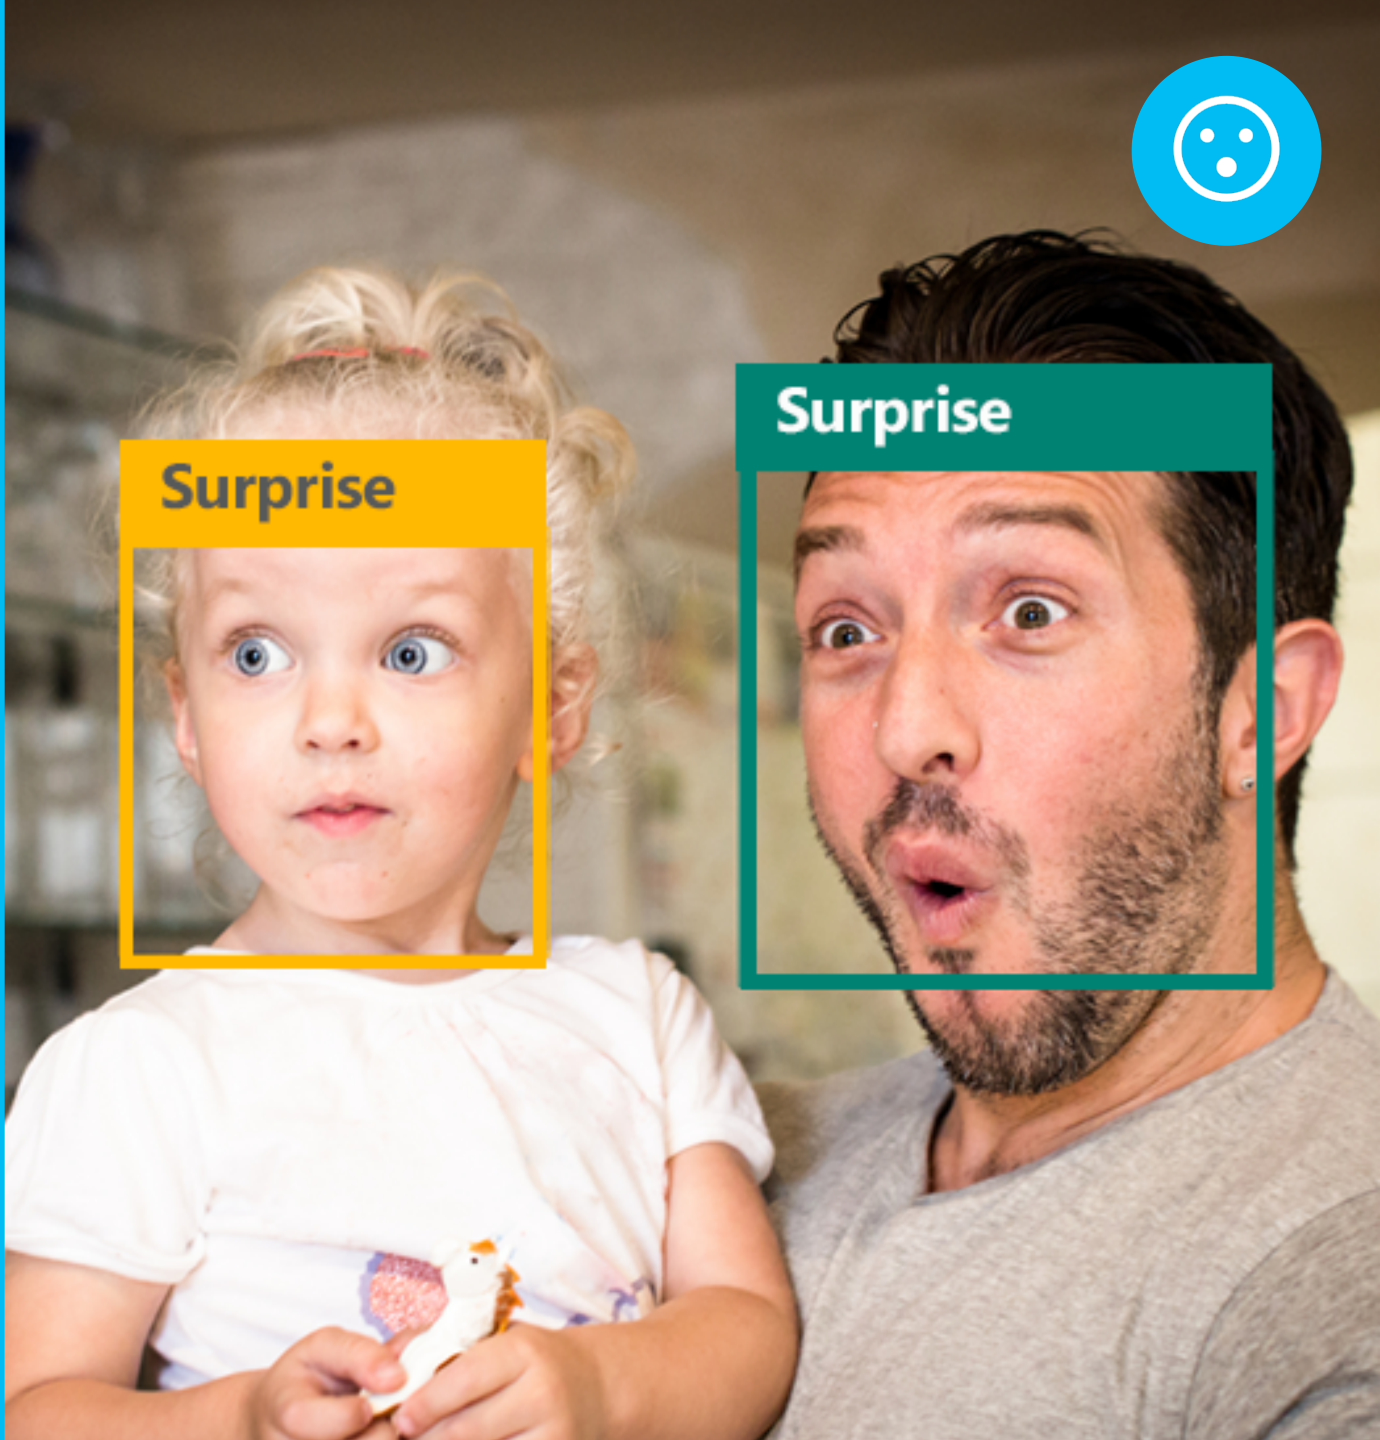

Content Moderator

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

Enable the scoring of possible adult content in videos.

LANGUAGE

Process text and learn how to recognize what users want

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

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
    }
  ]
}
```

Linguistic analysis

Analysis tolls for natural language processing

Access to part-of-speech tagging and parsing, identifying concepts, and actions

Linguistic analysis

Enter a sentence

The Linguistic Analysis API simplifies complex languages to help you easily parse text.

POS tags

[["DT","NNP","NNP","NNP","VBZ","JJ","NNS","TO","VB","PRP","RB","VBP","NN","."]]

Constituency tree

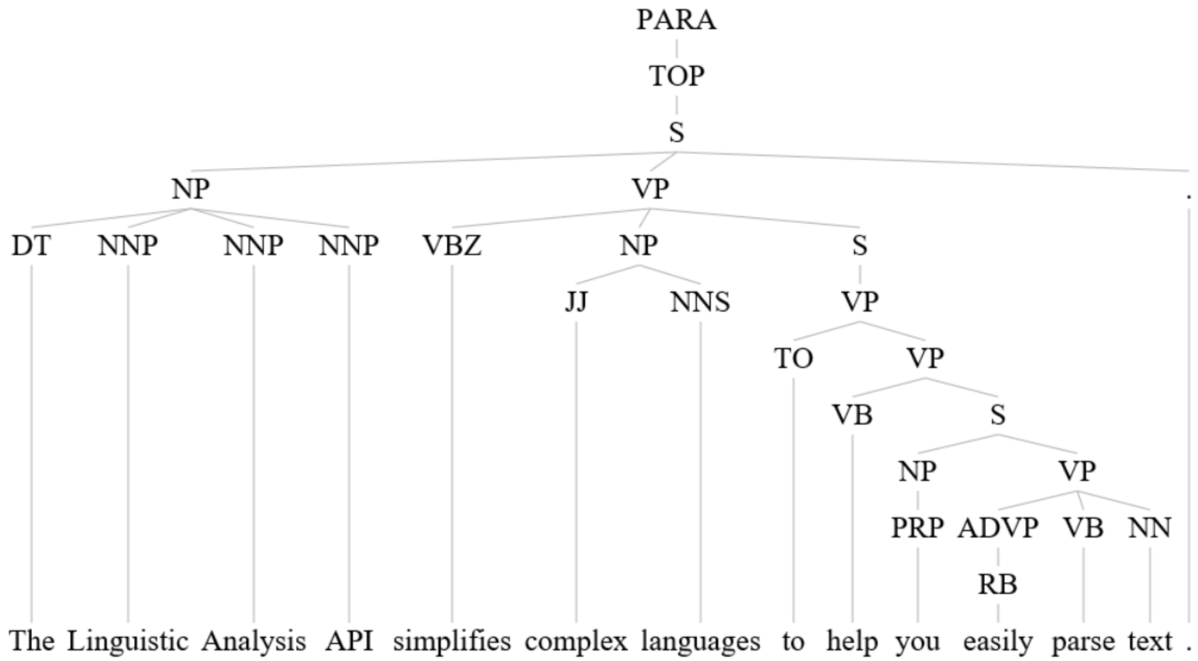

Text analytics

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

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

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

Recommendations

Increase catalog discovery
Help customers easily discover items that they may be interested in

Personalize your experience
Show suggestions that are targeted to each specific user

Increase the bottom line
Increase your conversion rate by offering the right products at the right time

Progress And Predictions

Artificial Intelligence: Progress and Predictions

Ece Kamar

Researcher, Microsoft Research Redmond

http://www.ecekamar.com/

Exciting Times

Exciting Times

Public Perception

Stephen Hawking warns artificial intelligence could end mankind

By Rory Cellan-Jones
Technology correspondent

2 December 2014 | Technology

Stephen Hawking: "Humans, who are limited by slow biological evolution, couldn't compete and would be superseded"

Intelligent Robots Will Overtake Humans by 2100, Experts Say

by Tia Ghose, Senior Writer | May 07, 2013 01:03pm ET

622
Share
80
Tweet
21
Submit
23

Bill Gates claims 'AI dream is finally arriving' - and says machines will outsmart humans in some areas within a decade

- Claims in next 10 years there will be robots to do tasks like driving and warehouse for humans
- AI will also outpace humans in certain areas of knowledge

What Happens in the Real World

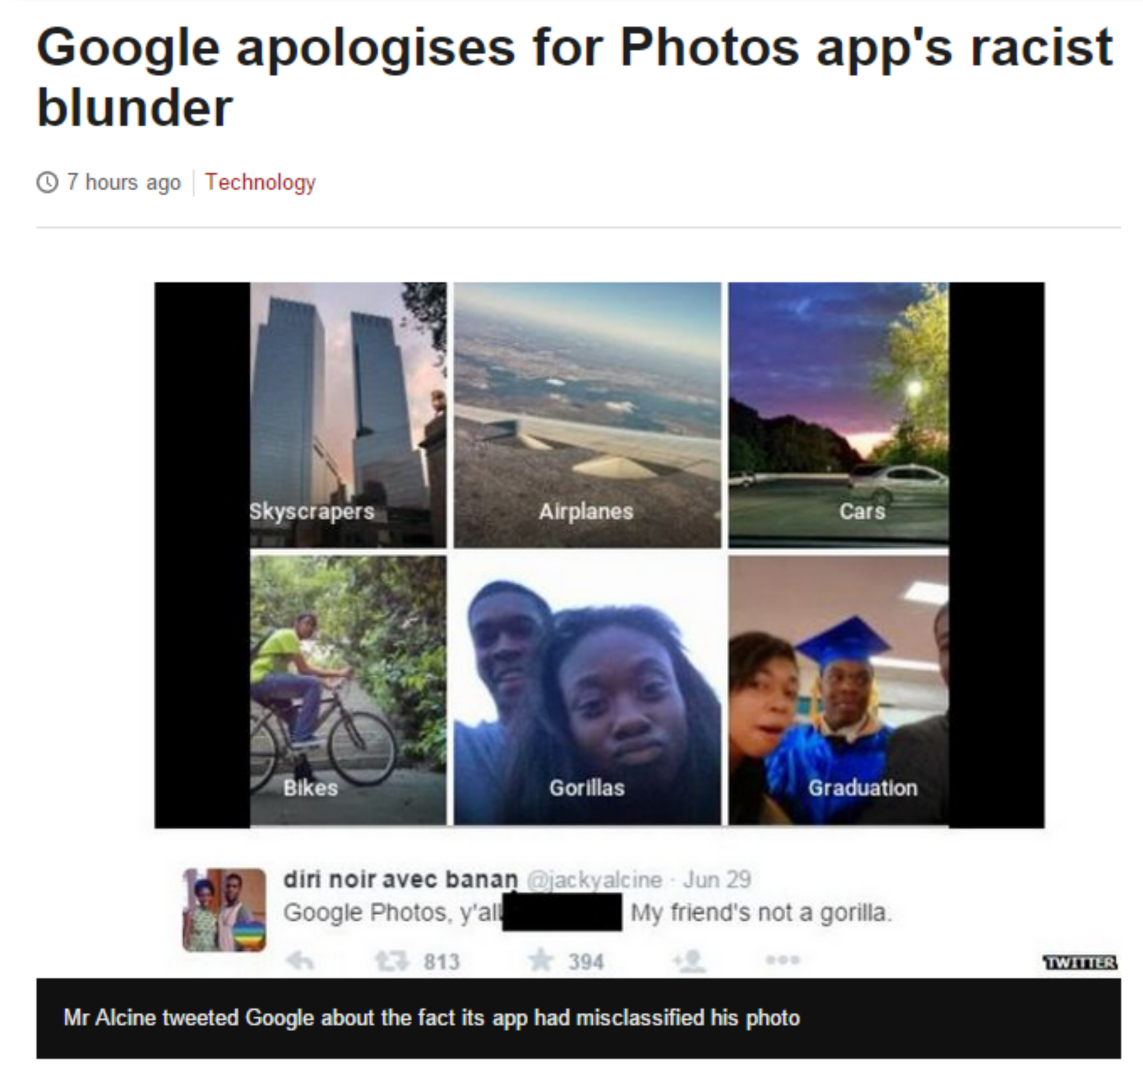

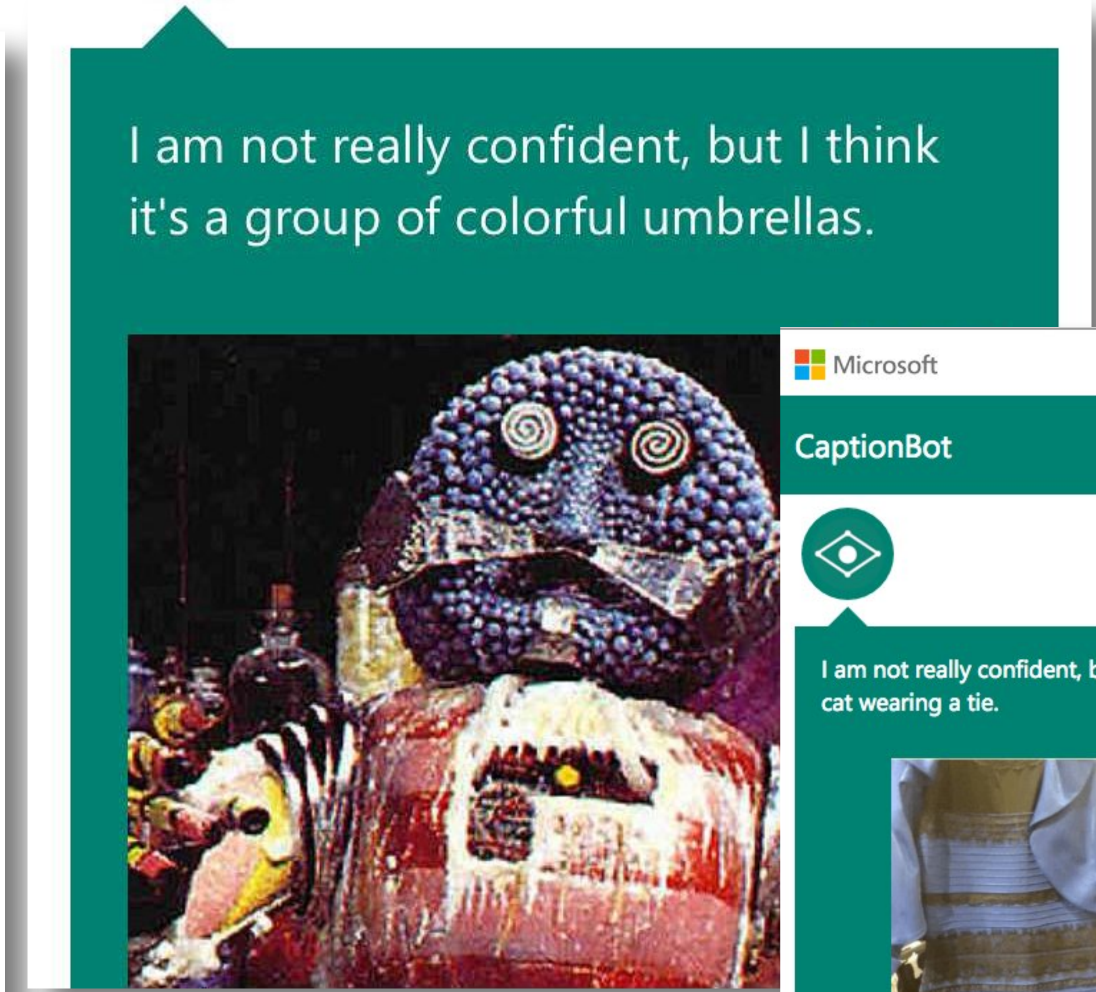

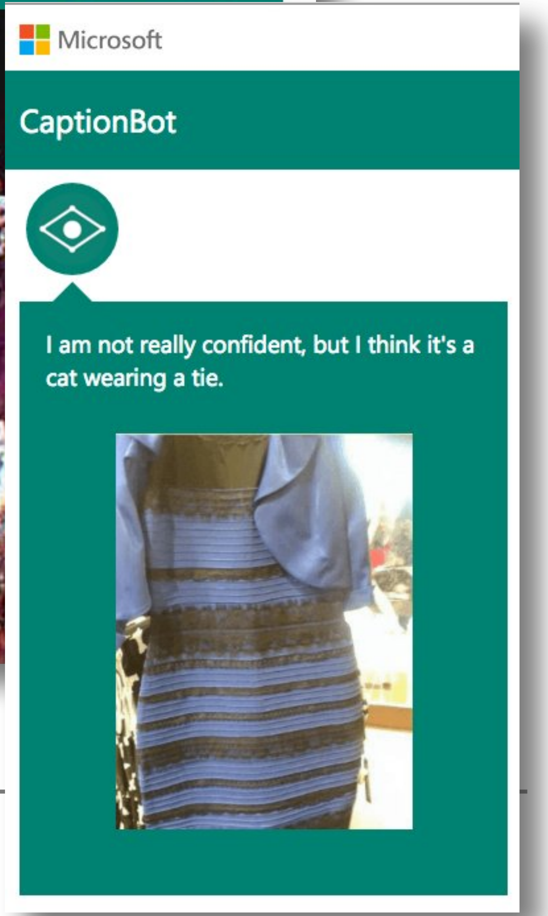

What Happens in the Real World

3 Questions

- What is AI?
- Where are we?
- Where are we headed?

AI 100: One Hundred Year Study of AI

- A longitudinal study of influences of AI advances on people and society:
  - Analysis of trends, developments, potential disruptions
  - Formulating recommendations and guidance
- Target audience:
  - AI researchers
  - Industry
  - General public
  - Policy makers

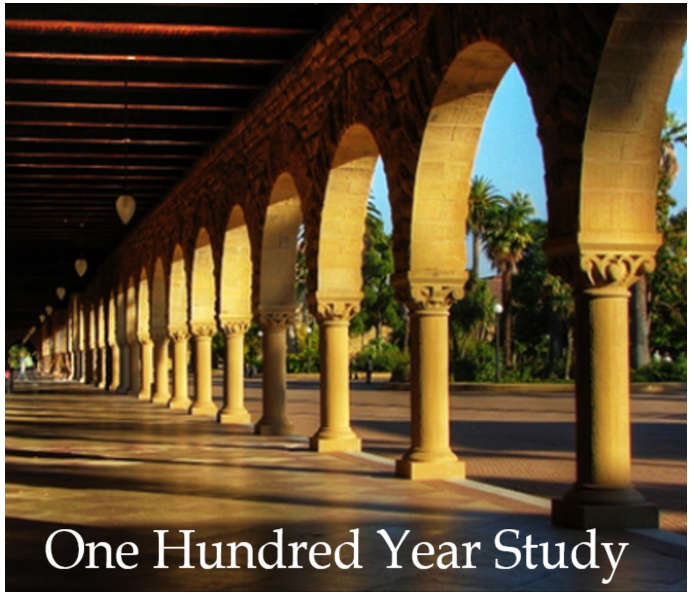

https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf

Charge for the 1st Study

- Focus: Large urban areas (typical North American city)
  - Identify possible **advances** in AI over 15 years and influences on daily life
  - Specify **scientific, engineering and legal** efforts needed
  - Consider actions needed to shape outcomes for **societal good**, deliberating **design, ethical and policy** challenges

Study Panel

Chair: Peter Stone, UT Austin

- Rodney Brooks, Rethink Robotics
- Erik Brynjolfsson, MIT
- Ryan Calo, University of Washington
- Oren Etzioni, Allen Institute for AI
- Greg Hager, Johns Hopkins
- Julia Hirschberg, Columbia
- Shivaram Kalyanakrishnan, IIT Bombay
- Ece Kamar, Microsoft
- Sarit Kraus, Bar Ilan
- Kevin Leyton-Brown, UBC
- David Parkes, Harvard
- William Press, UT Austin
- Julie Shah, MIT
- Astro Teller, X
- Milind Tambe, USC
- AnnaLee Saxenian, Berkeley

Report Structure

- Preface for context
- Executive Summary (1 page)
- Overview (5 pages)
- Introduction
  - Defining AI; Current research trends
- AI by domain
  - 8 areas with likely urban impact by 2030
  - Look backwards 15 years and forward 15 years
  - Opportunities, barriers and realistic risks
- Policy and legal issues
  - Current status, recommendations

Quick History Lesson

1950 In his famous paper Computing Machinery and Intelligence, Alan Turing posits that computer programs could think like humans and proposes a test to ascertain whether a computer's behavior is "intelligent."

1956 Stanford computer scientist John McCarthy, above, convenes the Dartmouth conference on "artificial intelligence," a term he defined. At this conference Herbert Simon and Allen Newell demonstrate a program that uses artificial intelligence to prove theorems in Principia Mathematica, by Bertram Russell and Alfred North Whitehead about logical foundations of mathematics. Simon and Newell also start work on computerized chess.

1962 Arthur Samuel, an IBM computer scientist who later became a Stanford professor, creates a self-learning program that proves capable of defeating one of America's top-ranked checkers champions.

1965-1970
Stanford researchers Ed Feigenbaum, seated above, Joshua Lederberg, Bruce Buchanan and Carl Djerassi create DENDRAL, the first "expert system." It creates scientific hypotheses about molecular structure using measured data.

1970-1980
Researchers develop more expert systems with applications to biology, medicine, engineering and the military.

1973 SRI's Artificial Intelligence Group creates Shakey the Robot, which crosses an obstacle-filled room autonomously using vision and locomotion systems. Shakey is the Computer History Museum's iconic exhibit for AI and Robotics.

1997 IBM's Deep Blue beats world chess champion Garry Kasparov in a six-game match, capping what Simon and Newell started four decades earlier.

2000 Statistical machine learning research that began in the 1980s achieves widespread practical use in major software services and mobile devices.

2005 Computer scientist Sebastian Thrun, above, and a team from the Stanford Artificial Intelligence Laboratory build a driverless car named Stanley. It becomes the first autonomous vehicle to complete a 132-mile course in the Mojave Desert, winning the DARPA Grand Challenge. Stanley is now on exhibit in the Smithsonian.

2009 Computer scientist Eric Horvitz assembles an AAAI study group on long-term AI futures, which holds its final meeting at Asilomar in California.

2011 IBM's Watson supercomputing system beats the two best human players of the TV game show Jeopardy, demonstrating an ability to understand and answer the types of nuanced questions that had previously bedeviled computer programs.

2014 Stanford accepts proposal to host One-Hundred-Year Study on Artificial Intelligence.

Dartmouth Artificial Intelligence (AI) Conference

> *We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire.*
>
> *The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.*
>
> \- Dartmouth AI Project Proposal; J. McCarthy et al.; Aug. 31, 1955.

Where are We? Where are We Headed?

- Transportation ← Short-term
- Home-service robots
- Healthcare
- Education
- Public safety and security
- Low-resource communities
- Employment and workspace ← Longer-term
- Entertainment

Employment and Workplace

- Near term: AI will replace more tasks than jobs
  - Will create jobs, but harder to predict what types
- Medium term: Path for lower costs for goods and services
  - Distribution of wealth will be an issue
- Long term: Fear of replacing all human jobs is overblown
  - Social safety nets may be needed

Key Takeaways

- “The Study Panel found no cause for concern that AI is a imminent threat to humankind.”
- “No machines with self-sustaining long-term goals and intent have been developed, nor are they likely to be developed in the near future.”

Key Takeaways continued…

- “Emerging technologies have potential to profoundly transform society and economy for the better by 2030.”
- “AI will not replace people, it will augment them.”
- “Need increased focus on building systems that can collaborate effectively with people.”
- “Towards intelligent systems that are human-aware and trustworthy”.

Interdisciplinary Issues

- Biases
- Transparency of AI
- Fair access
- Definition of responsibility
- Potential for good and bad
- Technical expertise at all levels of decision-making

AI Implications (and a few more examples)

- We overestimate the impact of technology in the short term and underestimate it in the long term
- Things go from being really bad to really good in a few short years (voice, translation, etc.)
- Kaltura – just an API call…

- *McDonalds – voice ordering*
- *Manulife – voice analysis – live suggestions by a bot on customer (and agent) emotions, etc.*

Implications in Higher Education

- LMS – student engagement tracking
- Collaboration Platform – usage, plagiarism, etc.
- Classrooms – student (faculty?) engagement, mood, tone (facial analysis and voice)
- CCTV – security, location, mood
- Student Residence – safety, waking hours, attendance
- Cafeteria/campus restaurants – wellness

Concerns

- Student Privacy & Buy-In
- Faculty Concerns
- Other Potentially Inappropriate Uses:
  - Prediction of Student Success
  - Confirmation Bias
  - Likely lots more…

Establishing an initial governance framework will be critical to maintaining stakeholder trust