A live webinar from

Leafcutter

AI for nonprofits

Leafcutter
Your Digital Partner

My introduction to AI

What we will cover

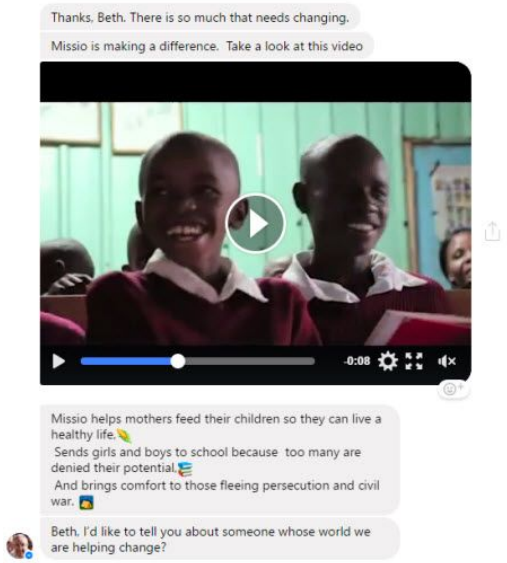

1. What is AI?

2. NFPs and AI

3. How to get started

What is AI?

Defining AI today

AI stands for **Artificial Intelligence**
Also known as **Machine Learning**

My definition:
a broad set of mathematical techniques employed at scale that simulate or surpass aspects of human intelligence (mainly pattern matching)

Why is AI so popular now?

Figure 1. Hype Cycle for Artificial Intelligence, 2018

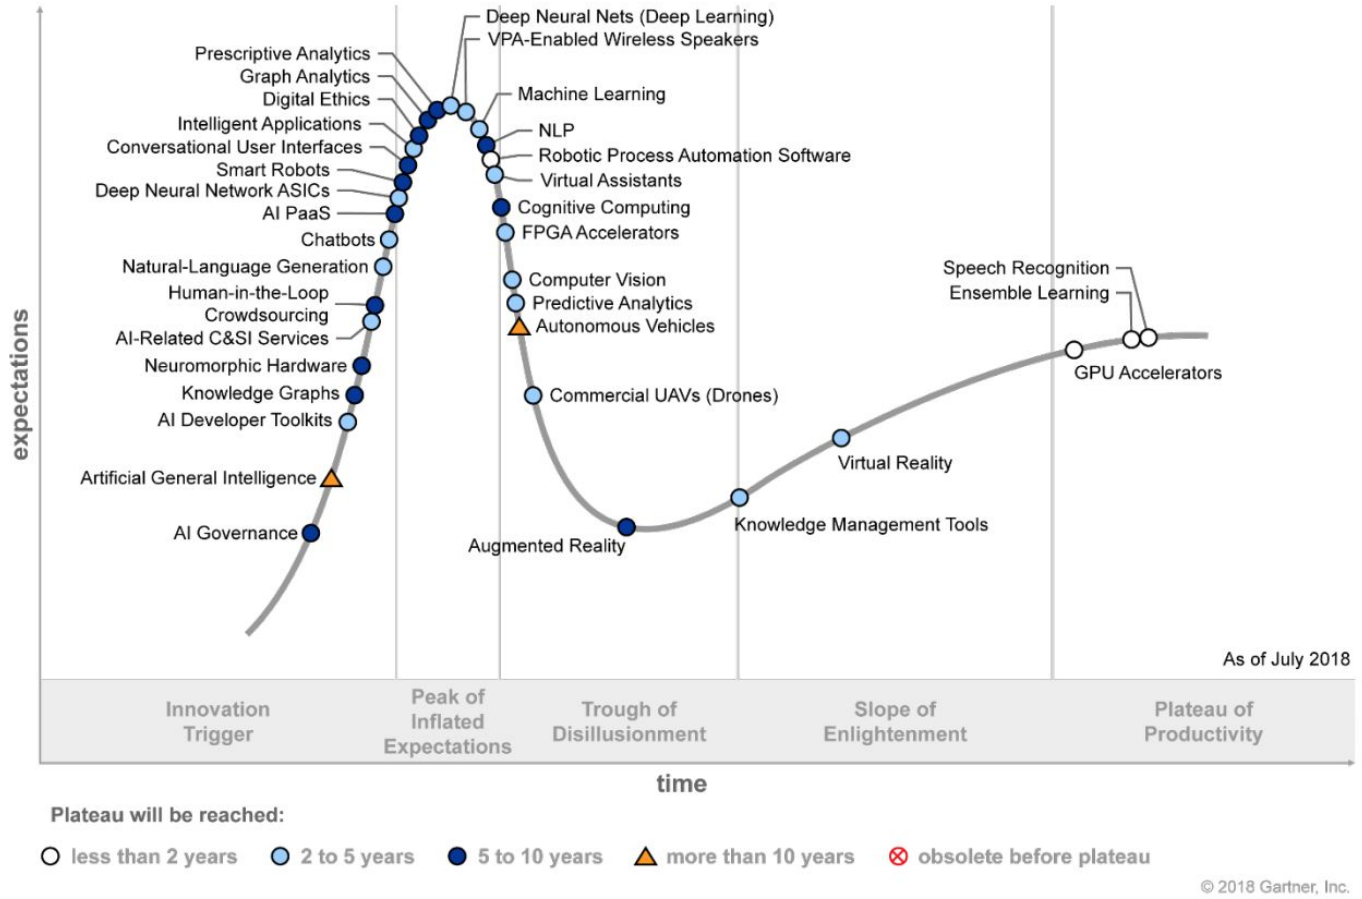

Source: Gartner (July 2018)

How it works

I have not failed.
I've just found 10,000 ways that won't work.

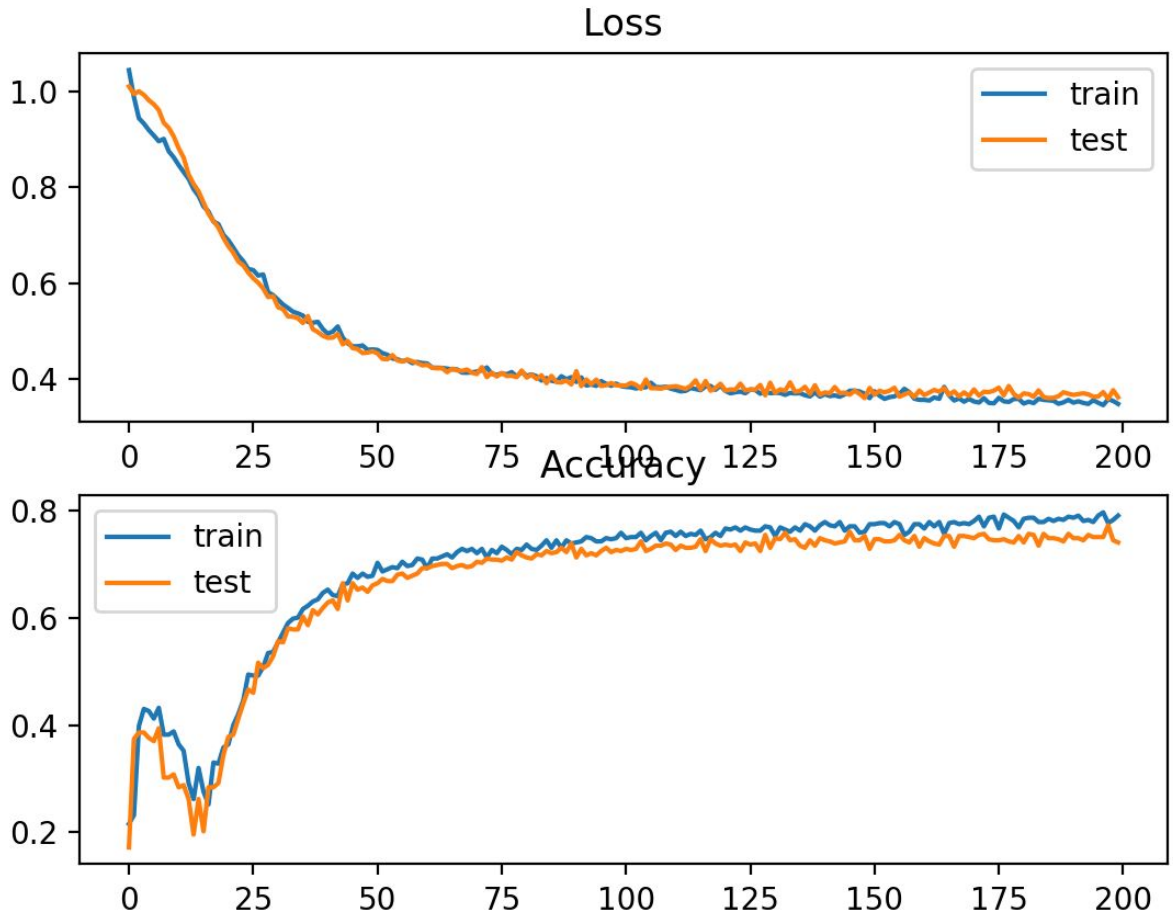

Probability vs determinism

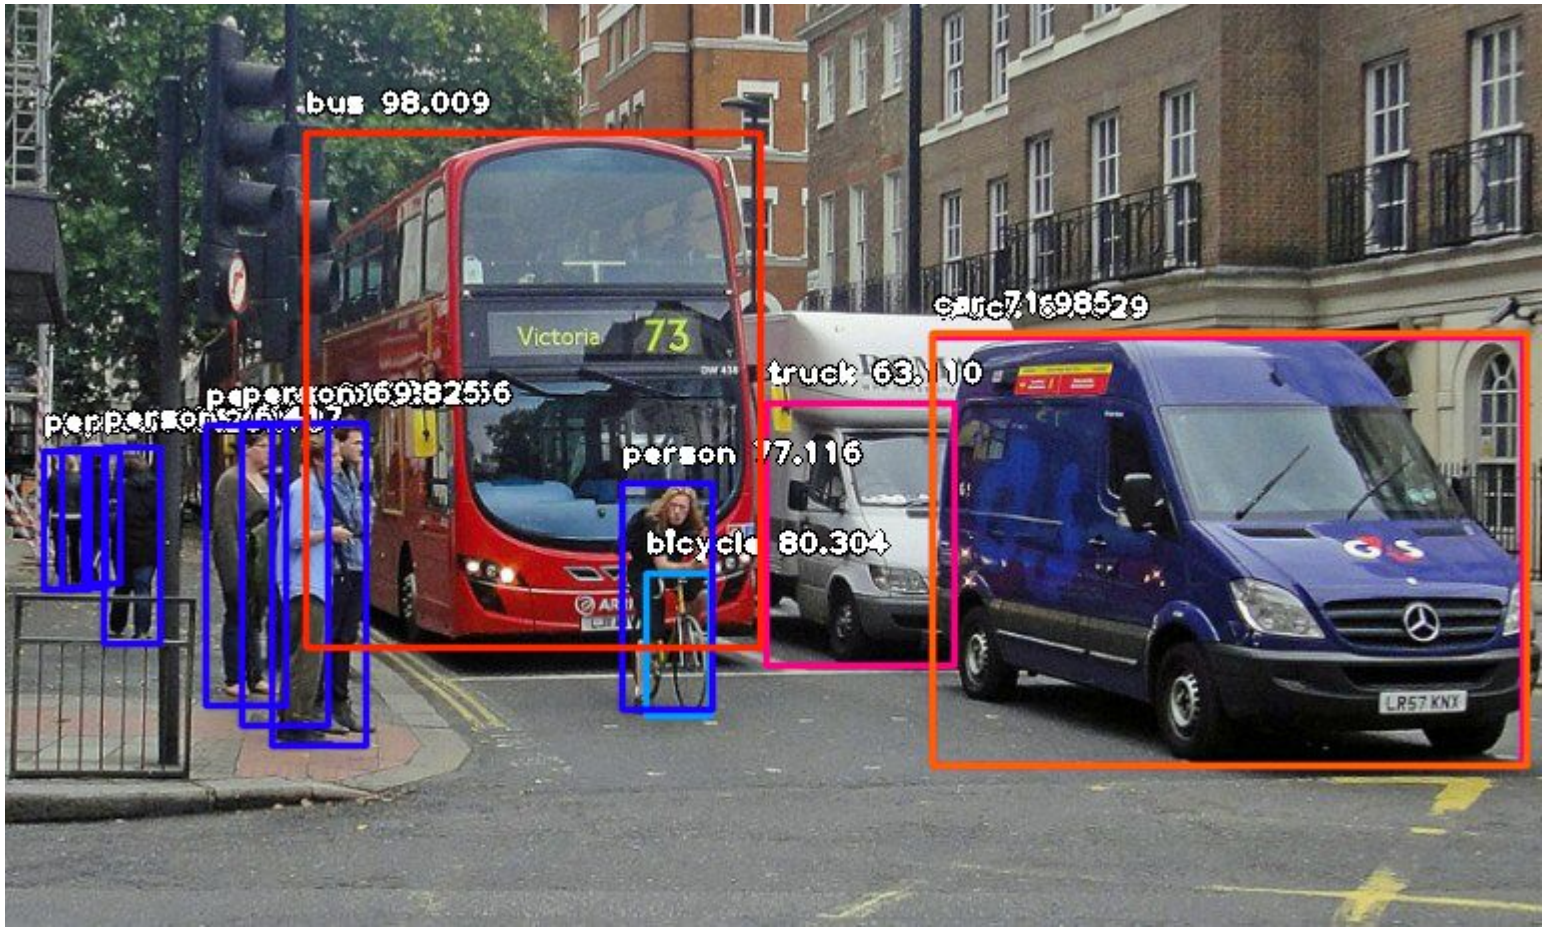

MACHINE LEARNING IN EMOJI

SUPERVISED UNSUPERVISED REINFORCEMENT

SUPERVISED — human builds model based on input / output

UNSUPERVISED — human input, machine output
human utilizes if satisfactory

REINFORCEMENT — human input, machine output
human reward/punish, cycle continues

BASIC REGRESSION

LINEAR — linear_model.LinearRegression()

Lots of numerical data

LOGISTIC — linear_model.LogisticRegression()

Target variable is categorical

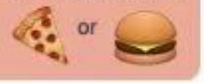

CLASSIFICATION

NEURAL NET — neural_network.MLPClassifier()

Complex relationships. Prone to overfitting
Basically magic.

K-NN — neighbors.KNeighborsClassifier()

Group membership based on proximity

DECISION TREE — tree.DecisionTreeClassifier()

If/then/else. Non-contiguous data
Can also be regression

RANDOM FOREST — ensemble.RandomForestClassifier()

Find best split randomly
Can also be regression

SVM — svm.SVC() svm.LinearSVC()

Maximum margin classifier. Fundamental
Data Science algorithm

NAIVE BAYES — GaussianNB() MultinomialNB() BernoulliNB()

Updating knowledge step by step with new info

CLUSTER ANALYSIS

K-MEANS — cluster.KMeans()

Similar datum into groups
based on centroids

ANOMALY DETECTION — covariance.EllipticalEnvelope()

Finding outliers
through grouping

FEATURE REDUCTION

T-DISTRIB STOCHASTIC NEIB EMBEDDING — manifold.TSNE()

Visualize high dimensional data. Convert
similarity to joint probabilities

PRINCIPLE COMPONENT ANALYSIS — decomposition.PCA()

Distill feature space into components that
describe greatest variance

CANONICAL CORRELATION ANALYSIS — decomposition.CCA()

Making sense of cross-correlation
matrices

LINEAR DISCRIMINANT ANALYSIS — lda.LDA()

Linear combination of features that
separates classes

OTHER IMPORTANT CONCEPTS

BIAS VARIANCE TRADEOFF

UNDERFITTING / OVERFITTING

INERTIA

ACCURACY FUNCTION — $(TP + TN) / (P + N)$

PRECISION FUNCTION — $TP / (TP + FP)$

SPECIFICITY FUNCTION — $TN / (FP + TN)$

SENSITIVITY FUNCTION — $TP / (TP + FN)$

@emilyinamillion made this

Deep learning

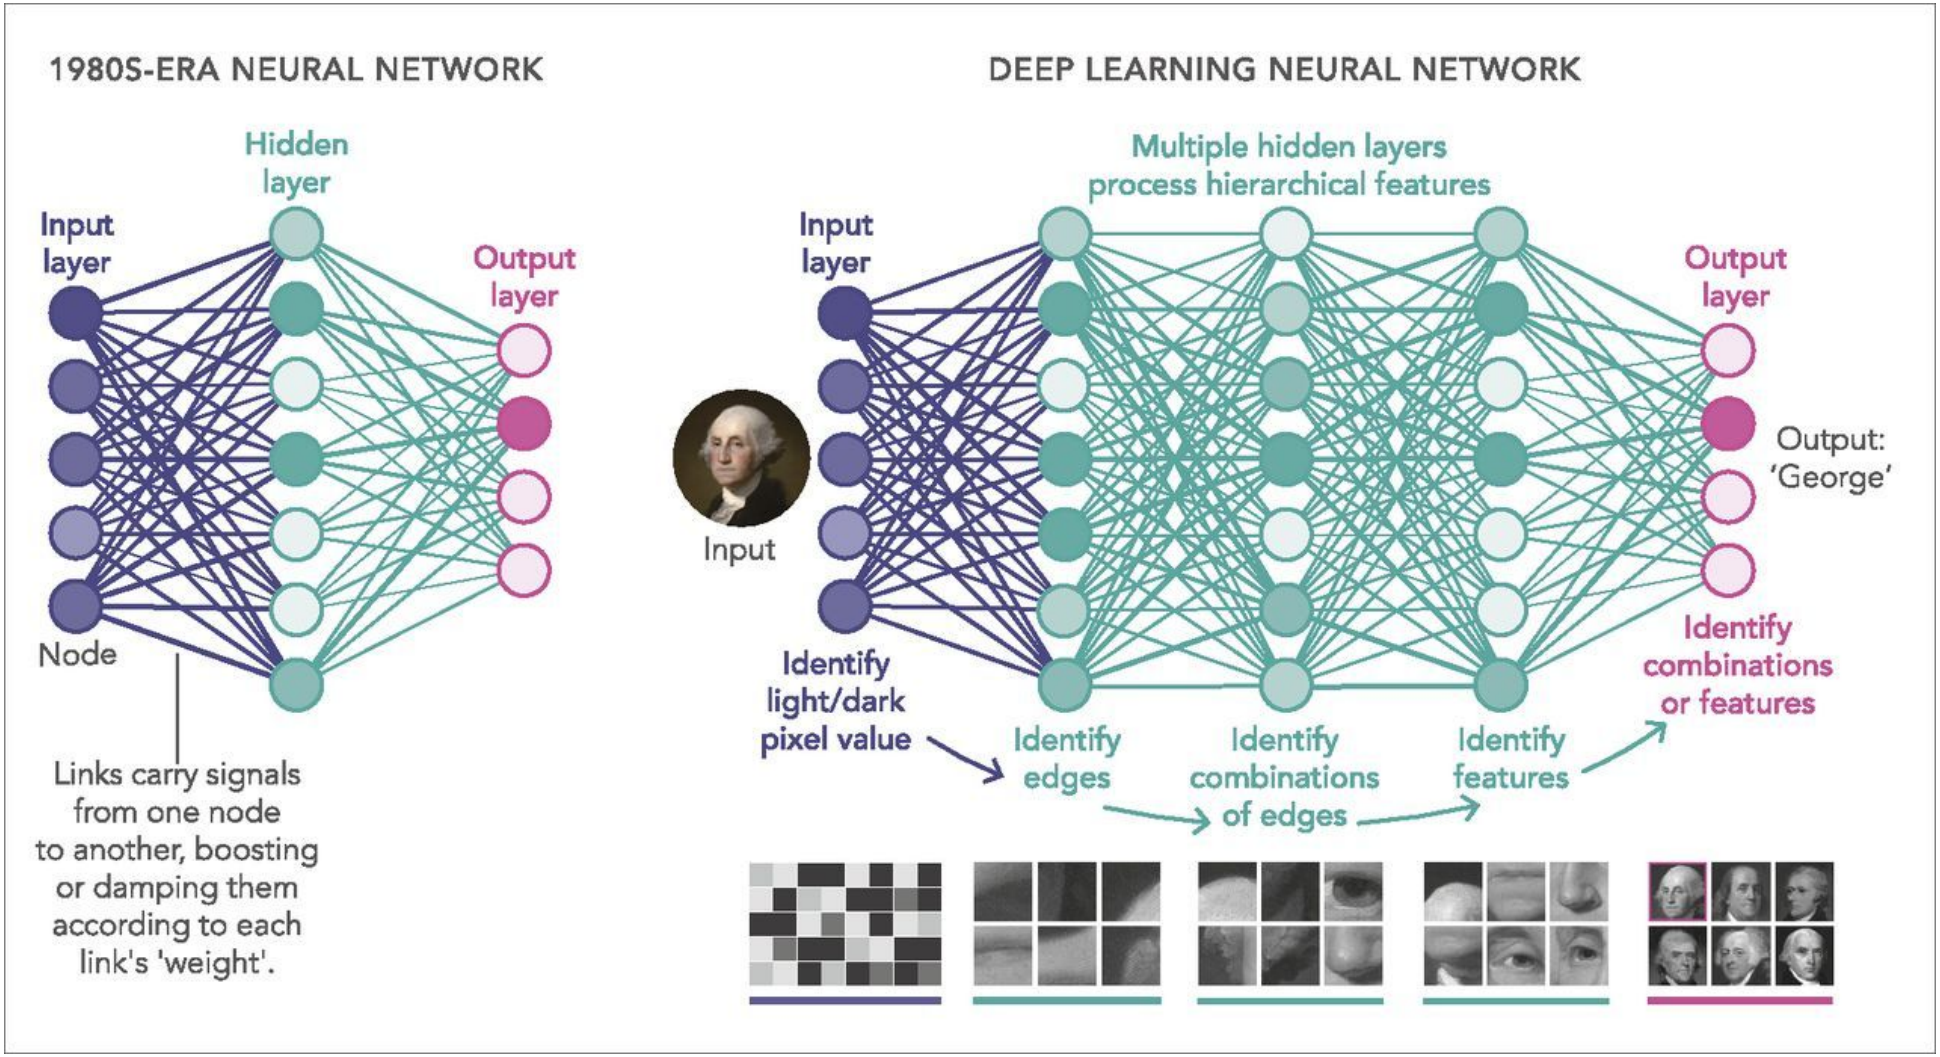

What today's AI can do

- Can recognise voice and text
- Can classify images (even in video)
- Can recommend movies, choices etc
- Can simulate various scenarios and outcomes

The pinnacle of AI achievement

What AI can’t do (today)

Today's AI is really great at solving narrow problems at scale

Less

More

Solving nonprofit challenges

Putting challenges in 2 baskets

Common challenges

- Automating customer service
- Image & video recognition
- Transcribing audio/video
- Fundraising campaigns
- Improving operations

Unique challenges

- Creating a unique supporter experience
- Research
- Improving service delivery

Automating customer service

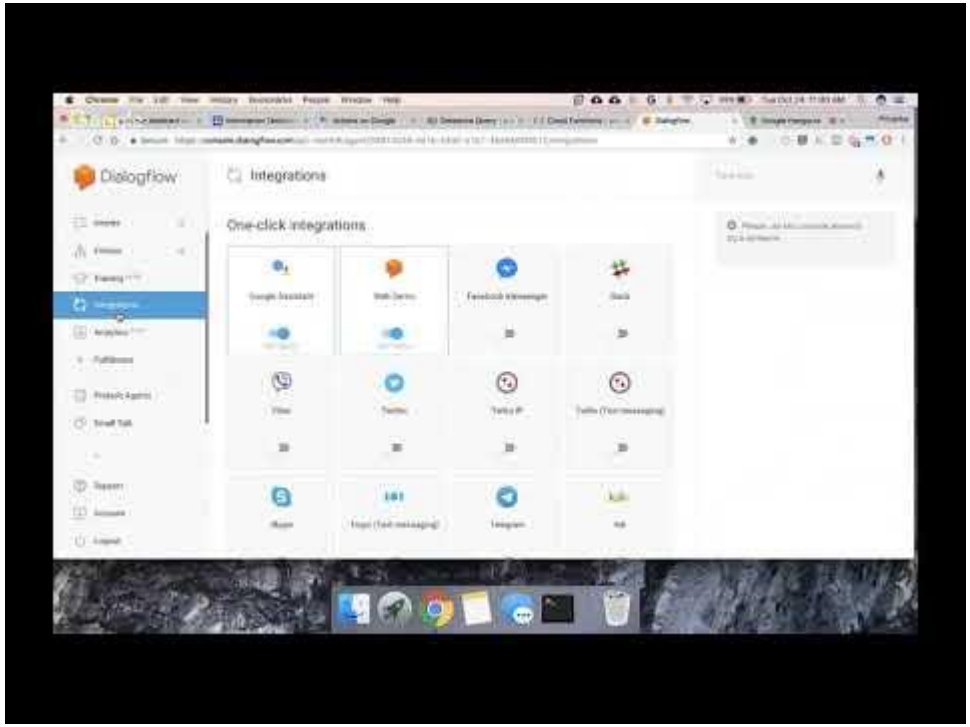

Image & video recognition

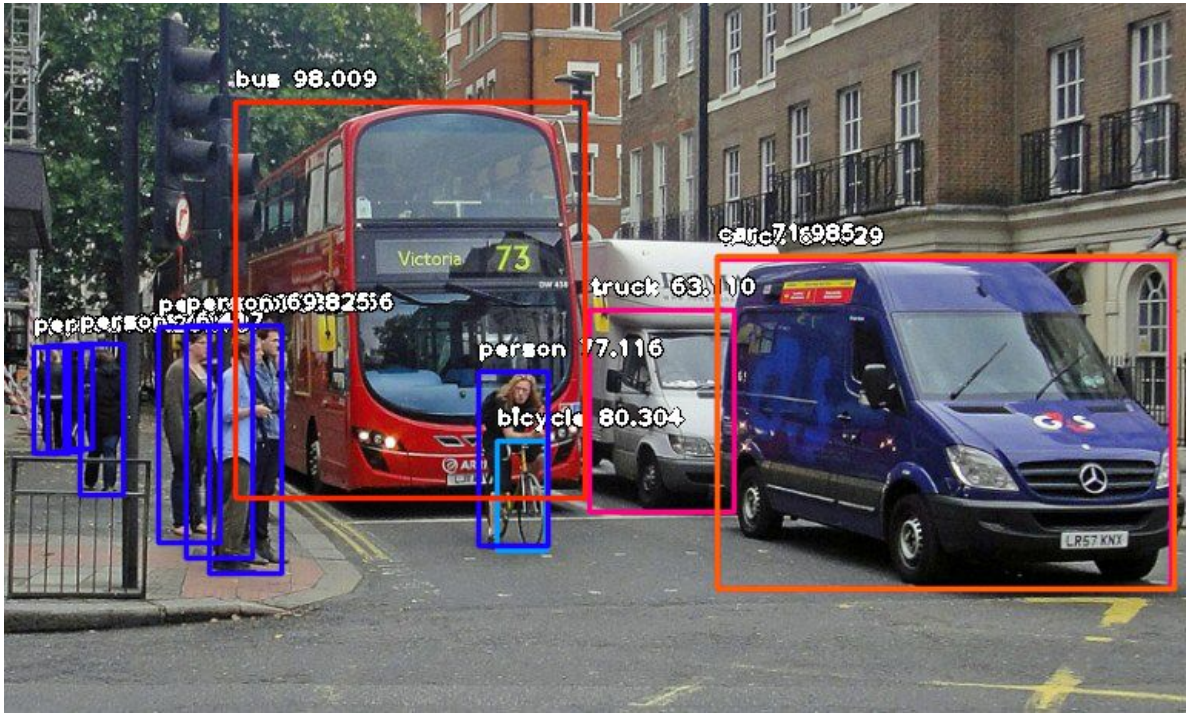

Transcribing audio and video

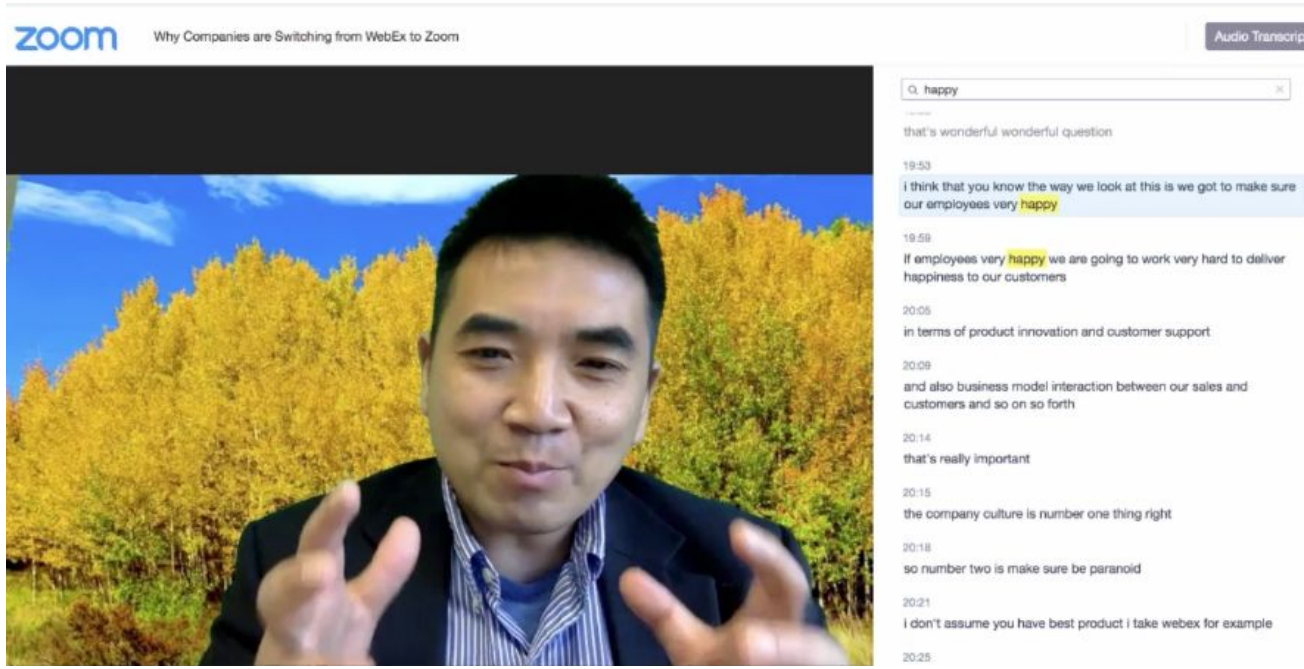

Fundraising campaigns

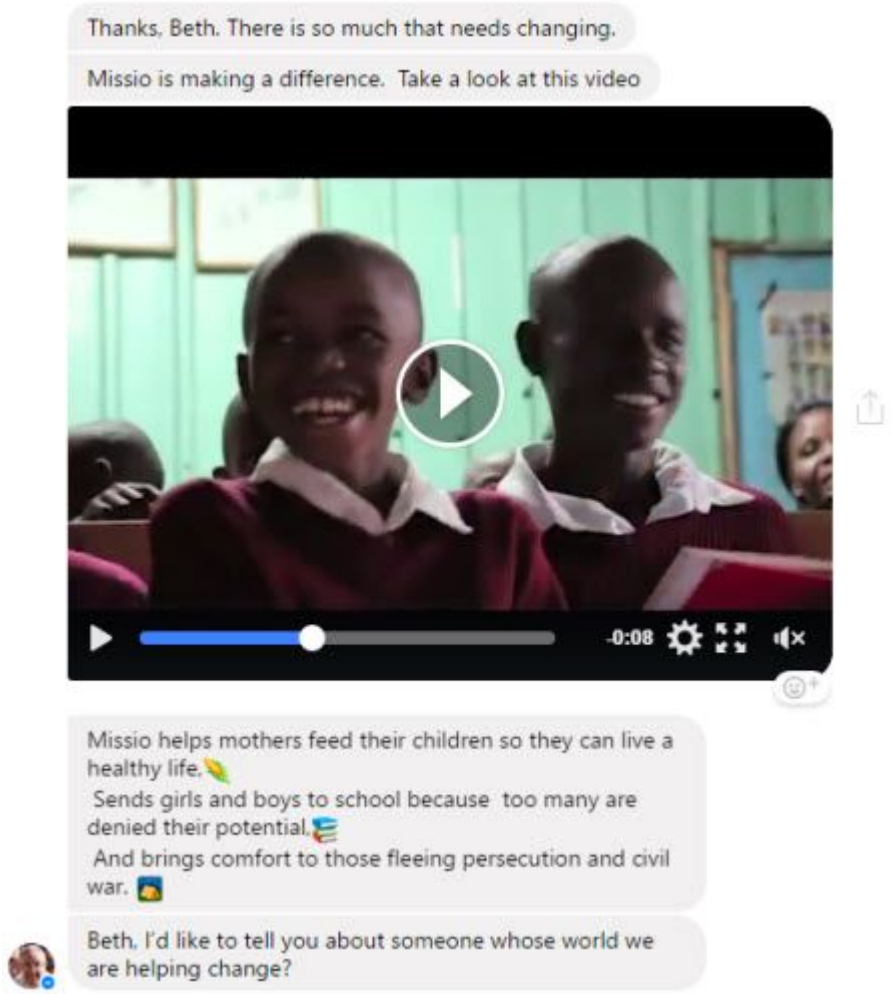

Improving operations (CRM, Finance, etc)

Creating a unique supporter experience

Research

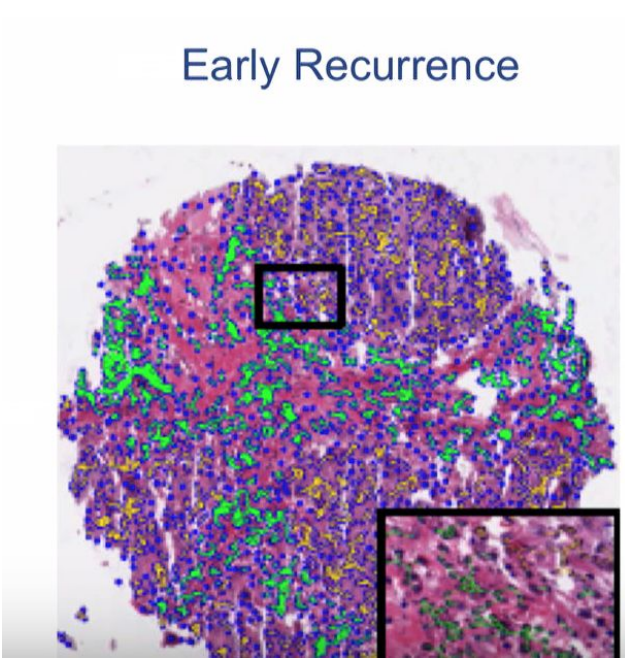

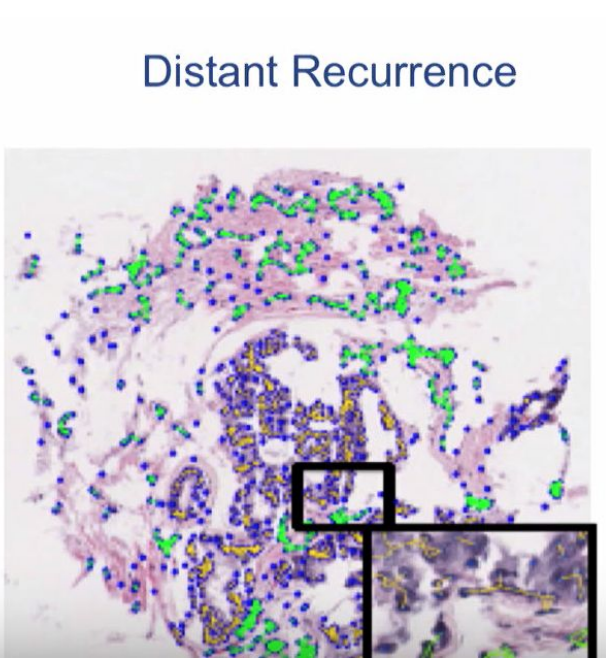

Improving service delivery

How to get started

You still need...

Good people

Good data

Choosing challenges to solve

1. List your organisation challenges and priorities
2. Identify areas for automation (low hanging fruit)
3. Identify areas where you have data but don't have insight (unique challenges)
4. Identify areas where you could offer an interactive experience (unique challenges)
5. Rank them according to problem scope, understanding of the problem, data volume & quality, and human resources available

The structure of an AI Project

1. Define the problem and hypothesis
2. Gather and clean the data
3. Train POC model from test data
4. Compare POC model to existing solutions
5. Refine and implement MVP
6. Continually monitor and retrain as required in production

NOTE: AI projects have probabilistic outcomes

Technology tools

Google Cloud Platform

Azure

amazon web services

IBM Watson

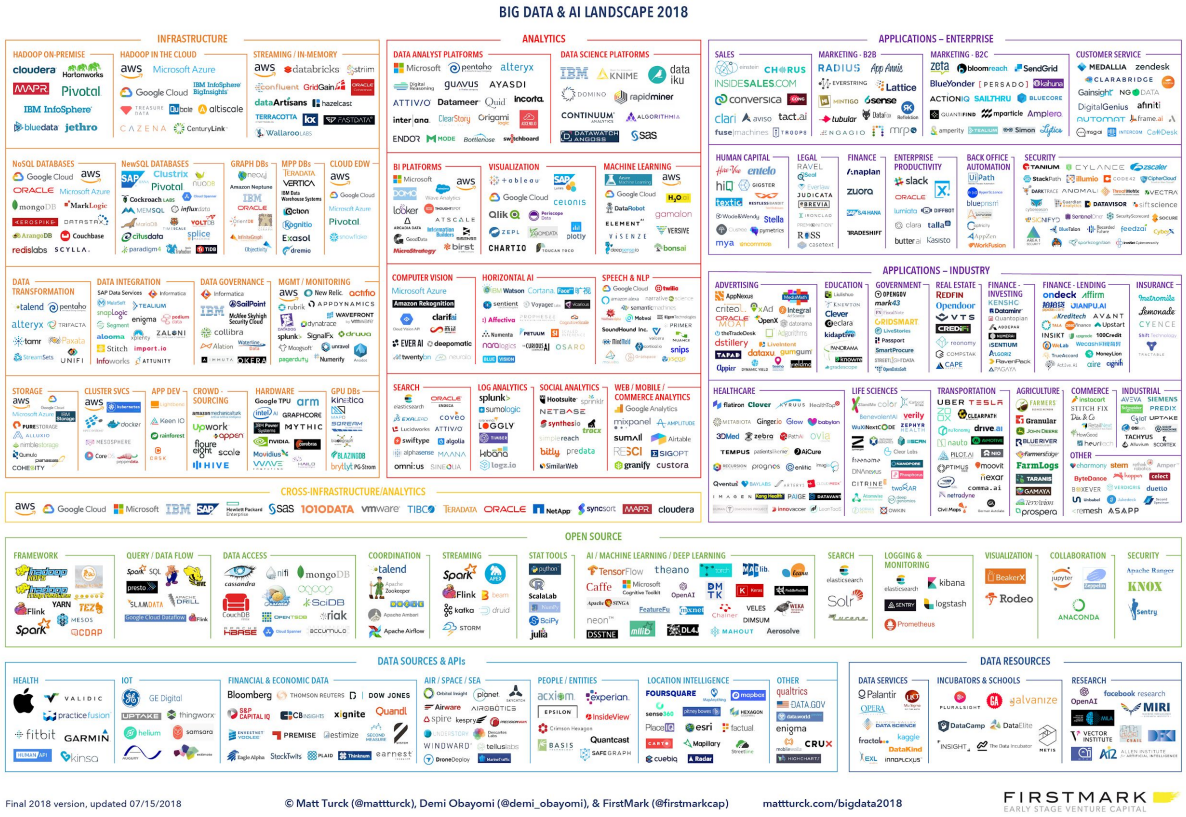

http://mattturck.com/wp-content/uploads/2018/07/Matt_Turck_FirstMark_Big_Data_Landscape_2018_Final.png

Wrapping up

Takeaways

1. AI is not a silver bullet - its good for solving narrow problems at scale
2. AI is happening - just not overnight
3. You need good process, good people and good data
4. People still need to work with the AI. You need to be able to tell the AI the problem to solve and ‘supervise’ its learning.

Thank you

Your Digital Partner

T +61 2 9438 1280
E hello@leafcutter.com.au
W leafcutter.com.au

Suite 401, 272 Pacific Hwy
Crows Nest NSW
Australia 2065

Leafcutter
Your Digital Partner