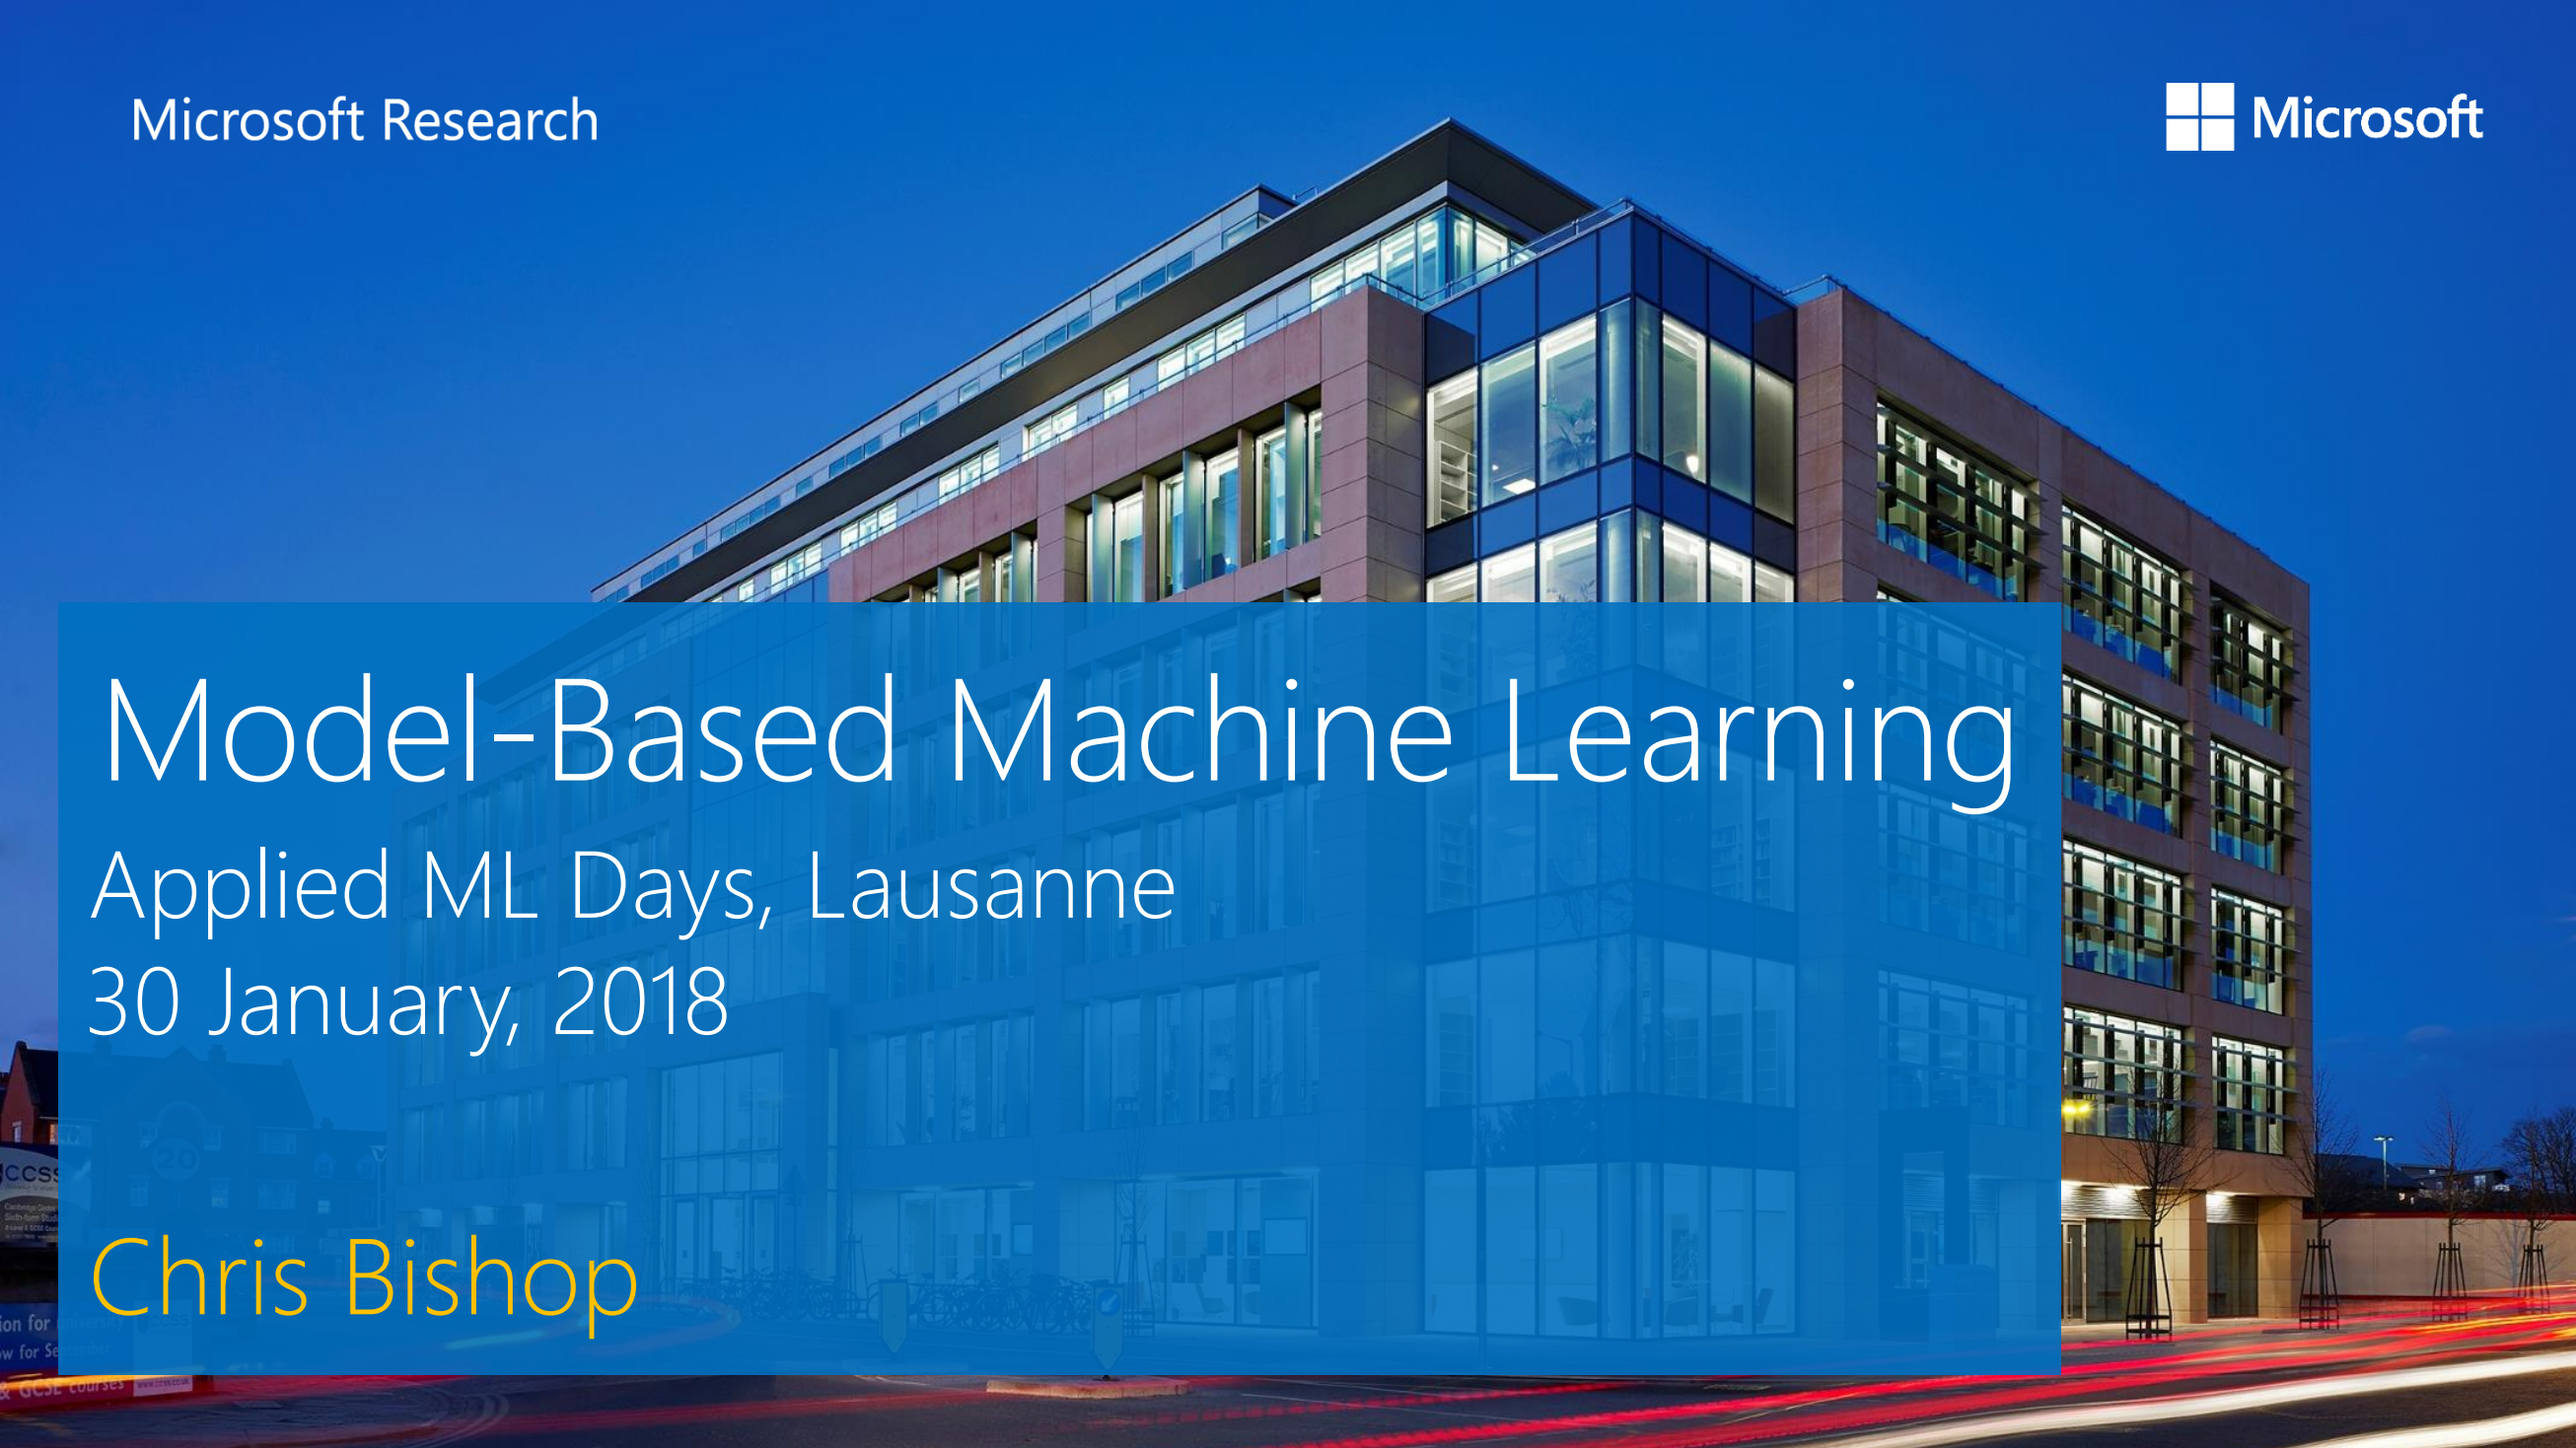

Microsoft Research

Microsoft

Model-Based Machine Learning

Applied ML Days, Lausanne
30 January, 2018

Chris Bishop

Machine learning algorithms

generative adversarial network

Markov random field

K-means clustering

Radial basis functions

decision trees

logistic regression

kernel PCA

Kalman filter

random forest

deep networks

principal components

Hidden Markov model

convolutional networks

support vector machines

Gaussian mixture

linear regression

Independent component analysis

neural networks

Gaussian process

factor analysis

Boltzmann machines

The ‘No Free Lunch’ Theorem

Averaged over all possible data-generating distributions, every classification algorithm has the same error rate when classifying previously unobserved points.

Wolpert (1996)

There is no universal machine learning algorithm

The goal of machine learning is to find an algorithm that is well matched to the problem being solved

Machine Learning

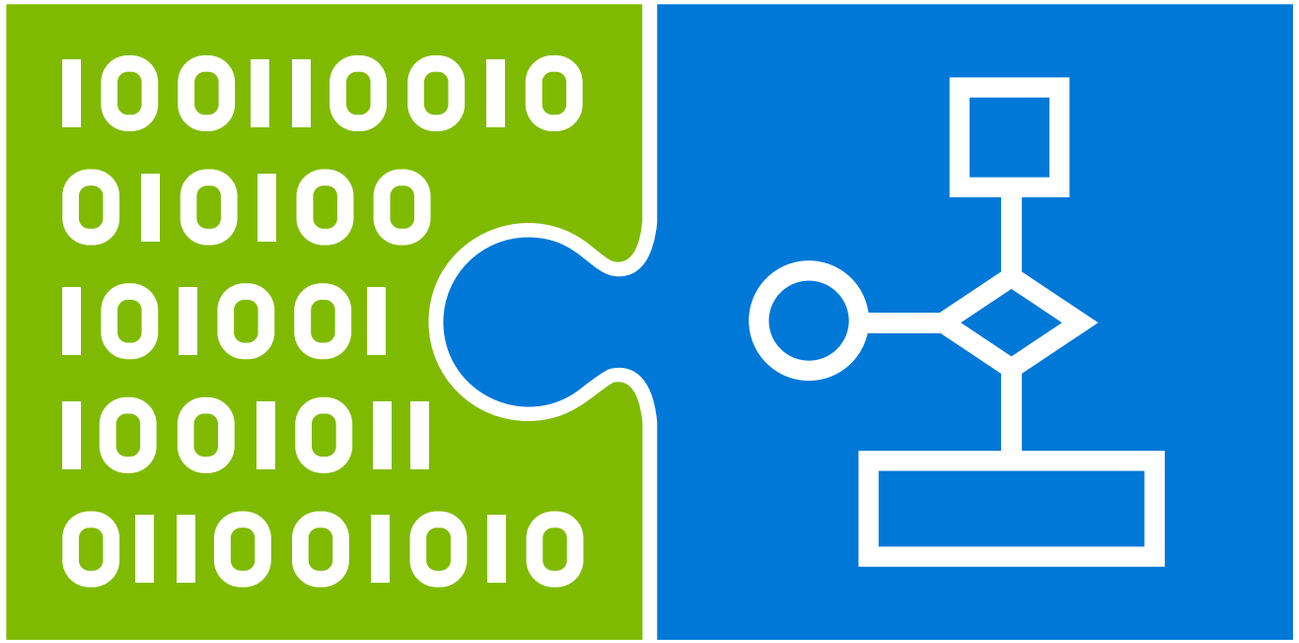

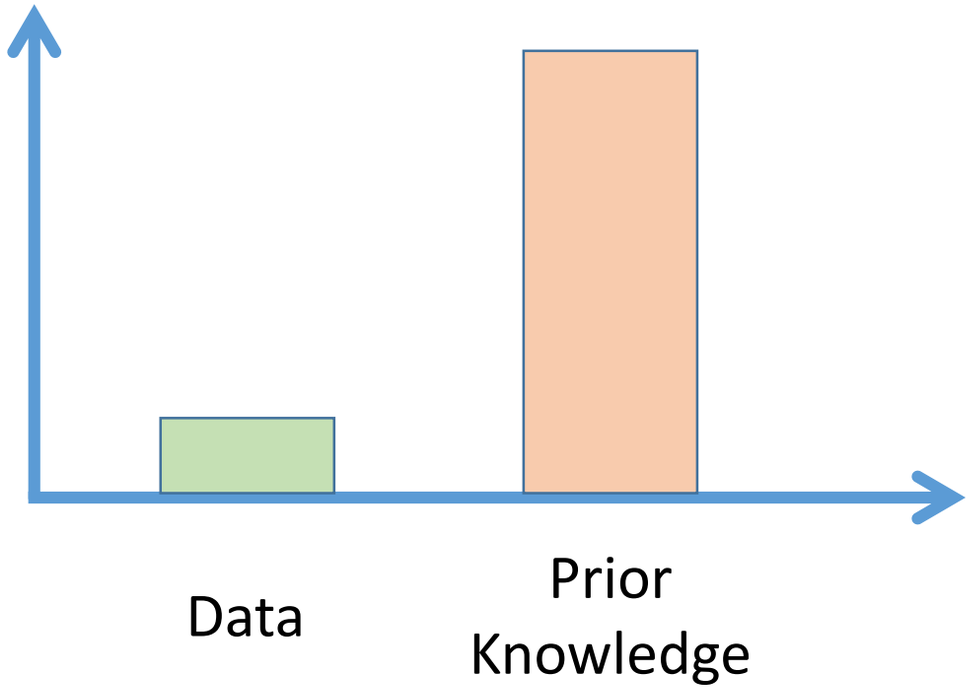
Data
Prior
Knowledge

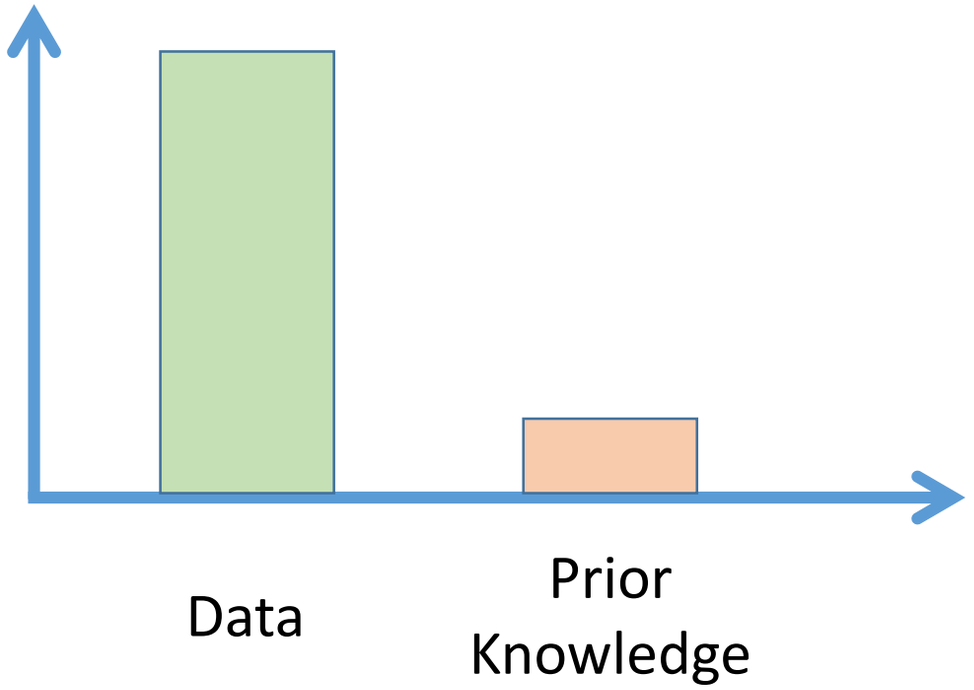
Data
Prior
Knowledge

‘Big data’

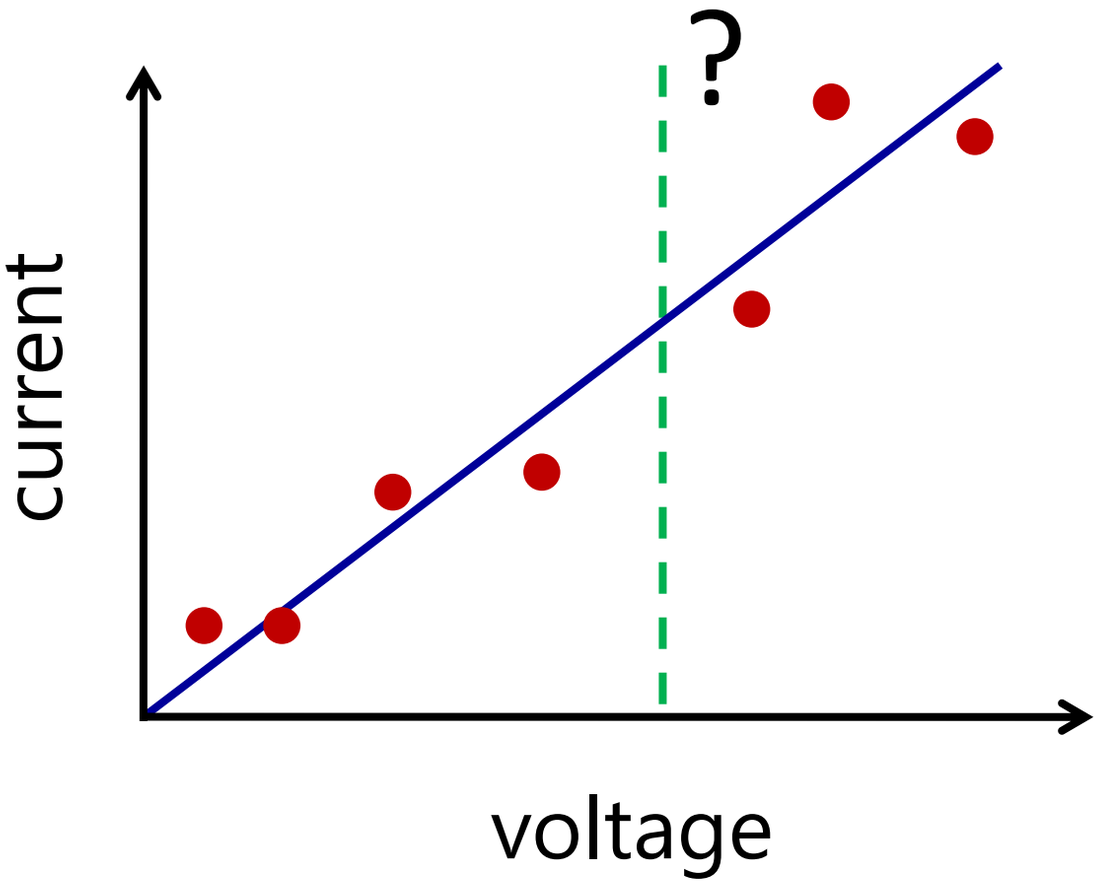
?
current
voltage


CitySightseeing Cambridge
Hop On
Hop Off
7

Model-based machine learning

Derive the appropriate ML algorithm by making modelling assumptions explicit

Traditional:

"how do I map my problem onto standard algorithms"?

Model-based:

"what is the *model* that represents my problem"?

Machine learning algorithm

PCA as an algorithm

$$\overline{\mathbf{x}} = \frac{1}{N}\sum_{n=1}^{N}\mathbf{x}_n$$

$$\mathbf{S} = \frac{1}{N}\sum_{n=1}^{N}(\mathbf{x}_n - \overline{\mathbf{x}})(\mathbf{x}_n - \overline{\mathbf{x}})^{\mathrm{T}}$$

$$\mathbf{S}\mathbf{u}_i = \lambda_i \mathbf{u}_i$$

retain $M < D$ eigenvectors

PCA as a model

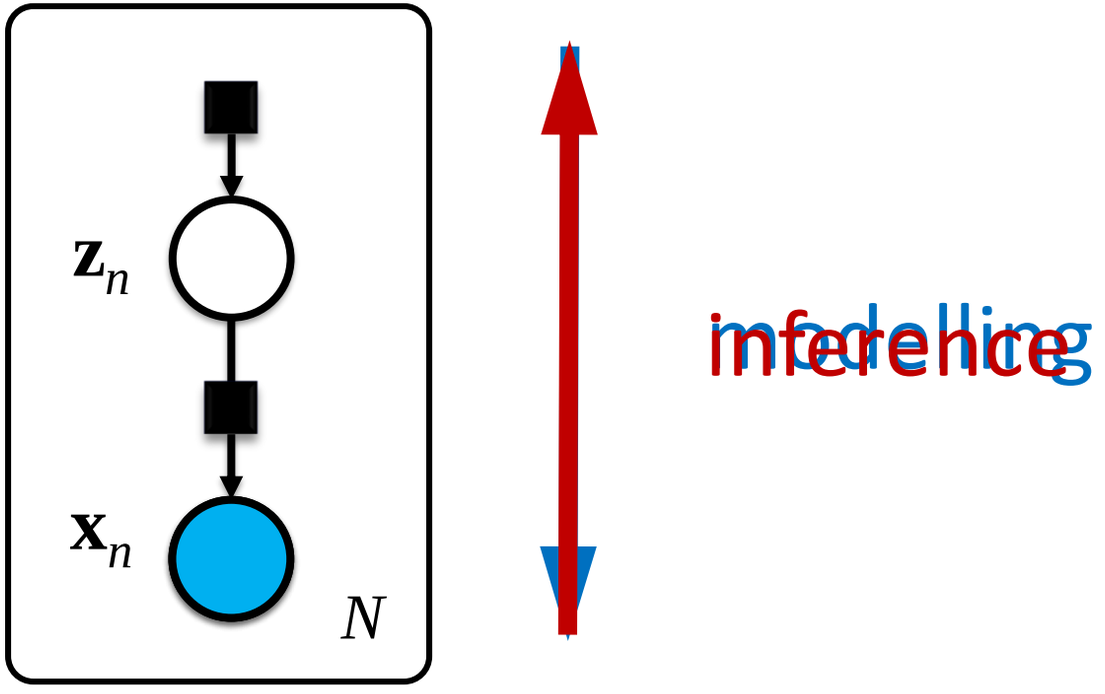

M. E. Tipping and C. M. Bishop (1997)

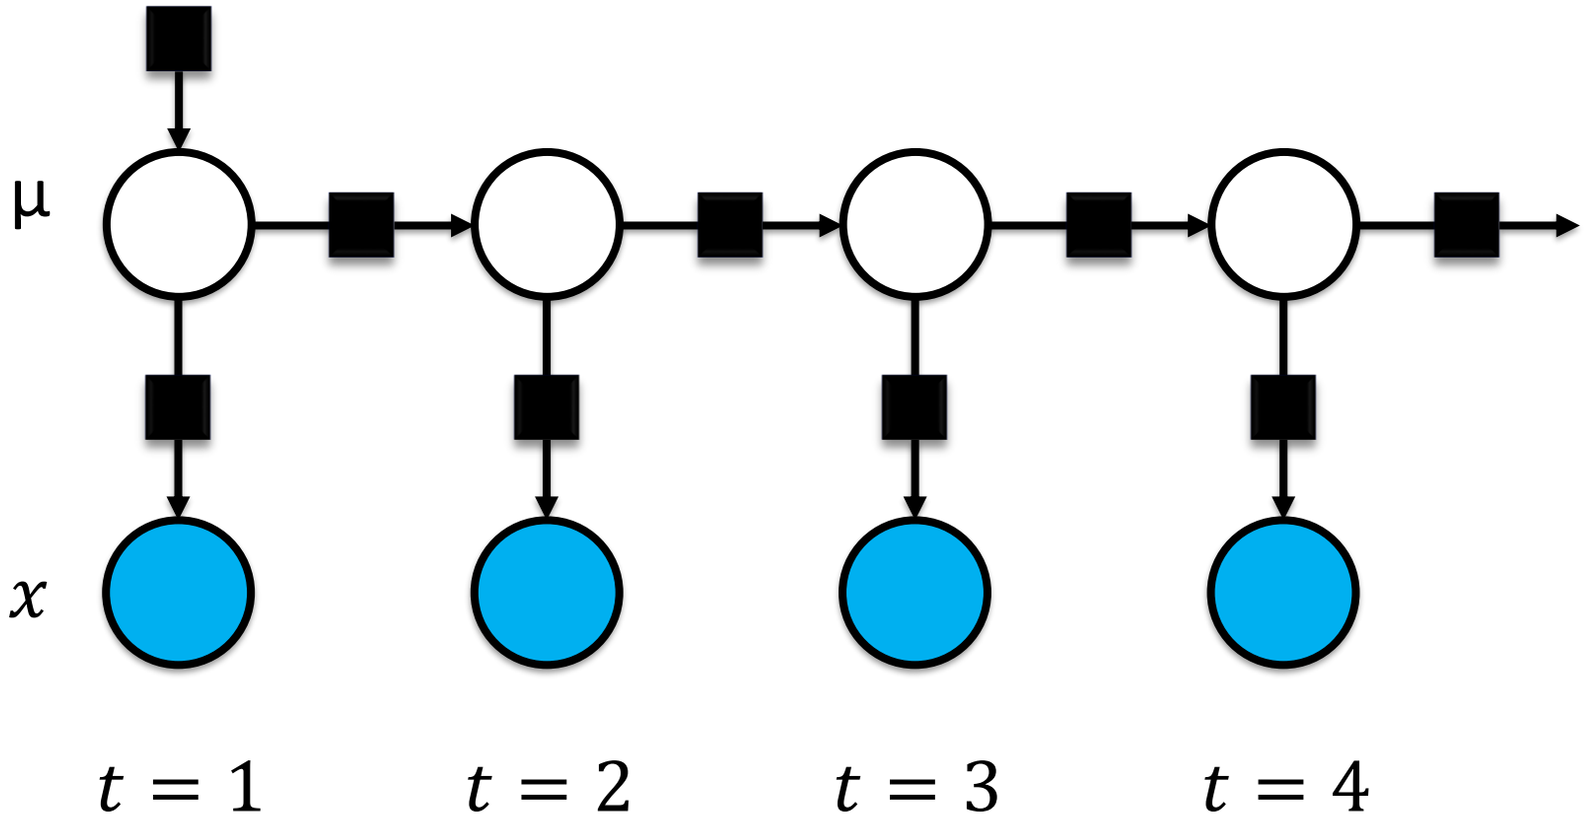

The Kalman filter

The hidden Markov model

Logistic Regression

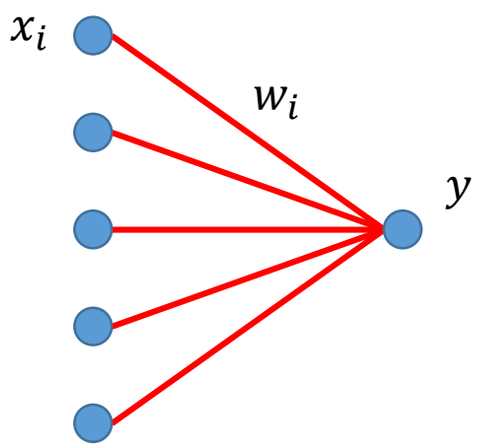

$$y = \sigma\left(\sum_i w_i x_i\right)$$

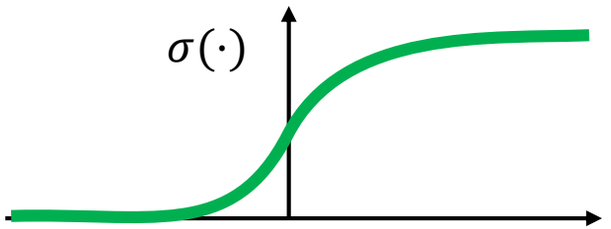

Deep Neural Networks

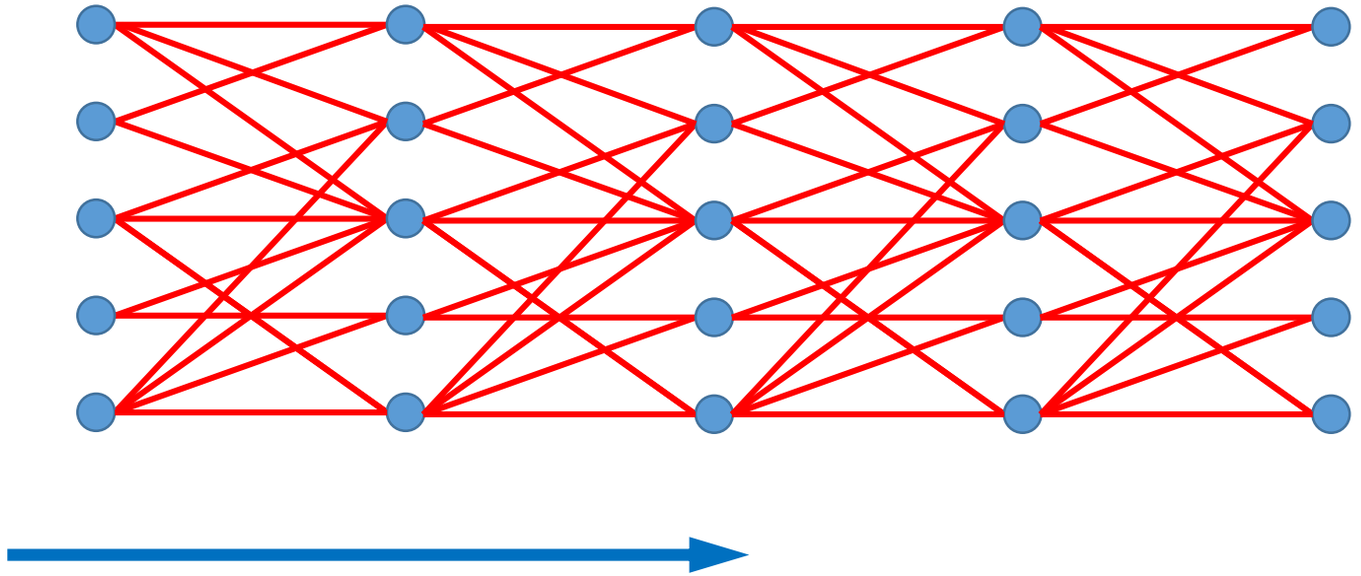

Data and prior knowledge

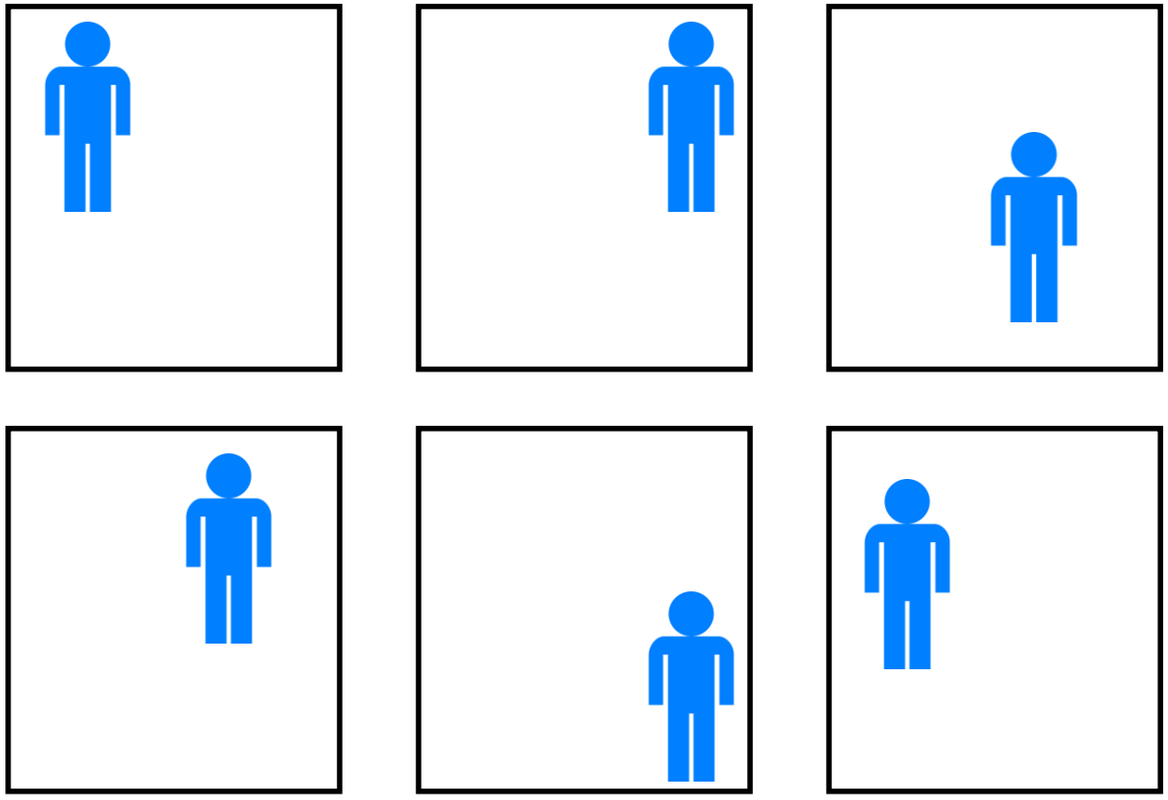

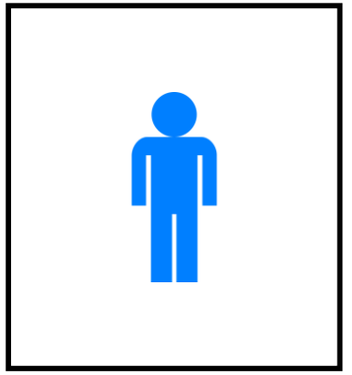

Translation invariance

Convolutional Neural Networks

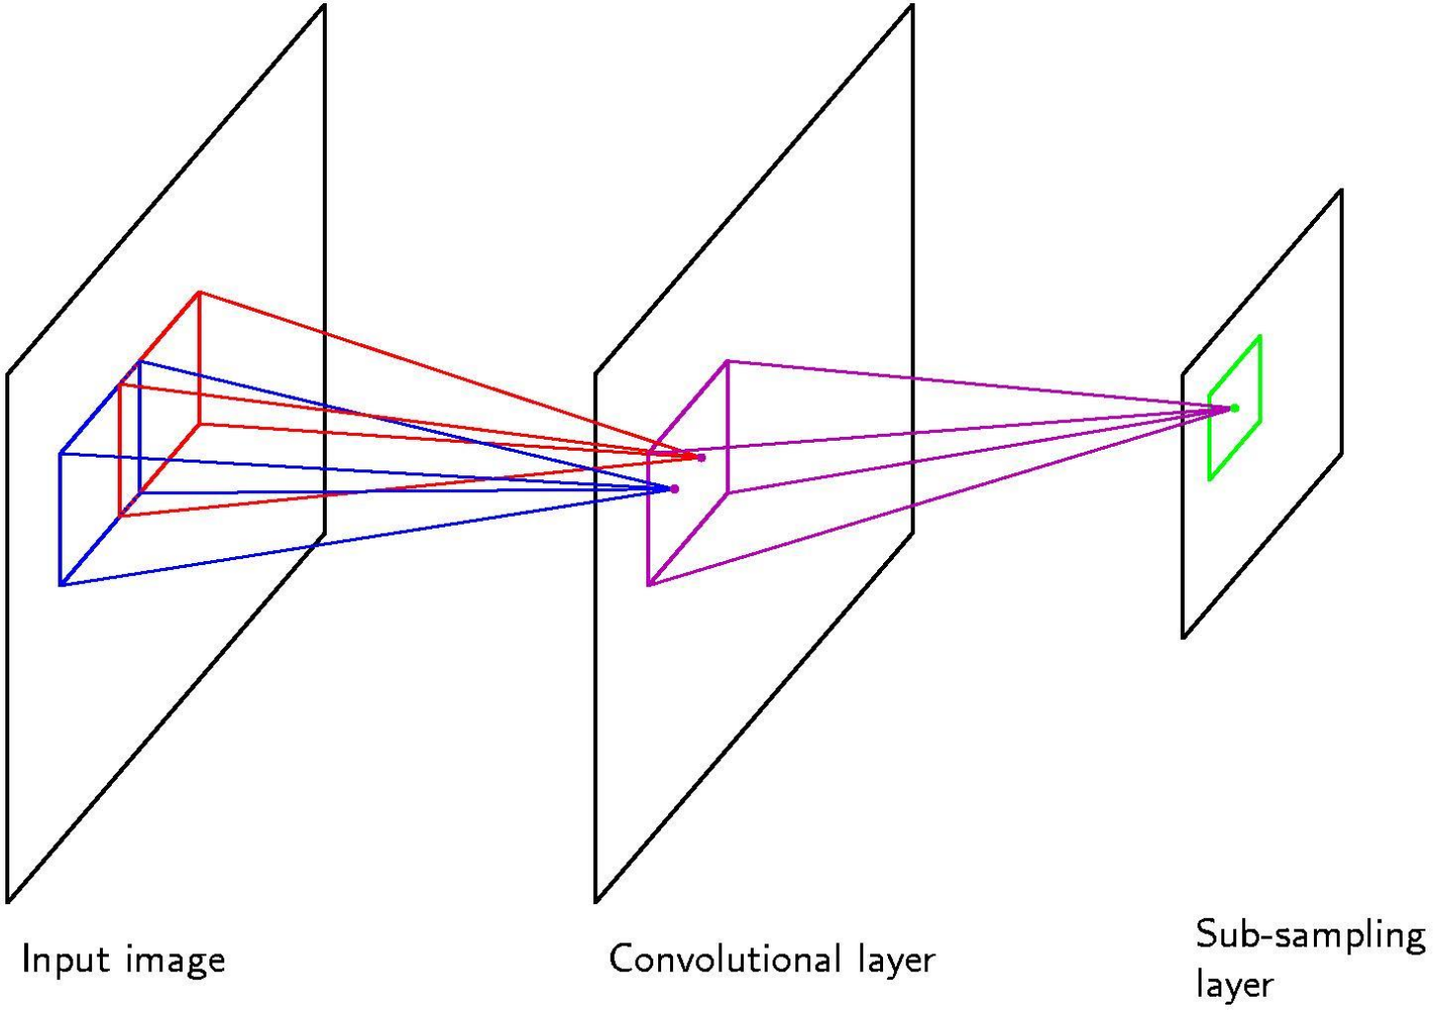

mbmlbook.com

Feedback please:

mbmlbook@microsoft.com

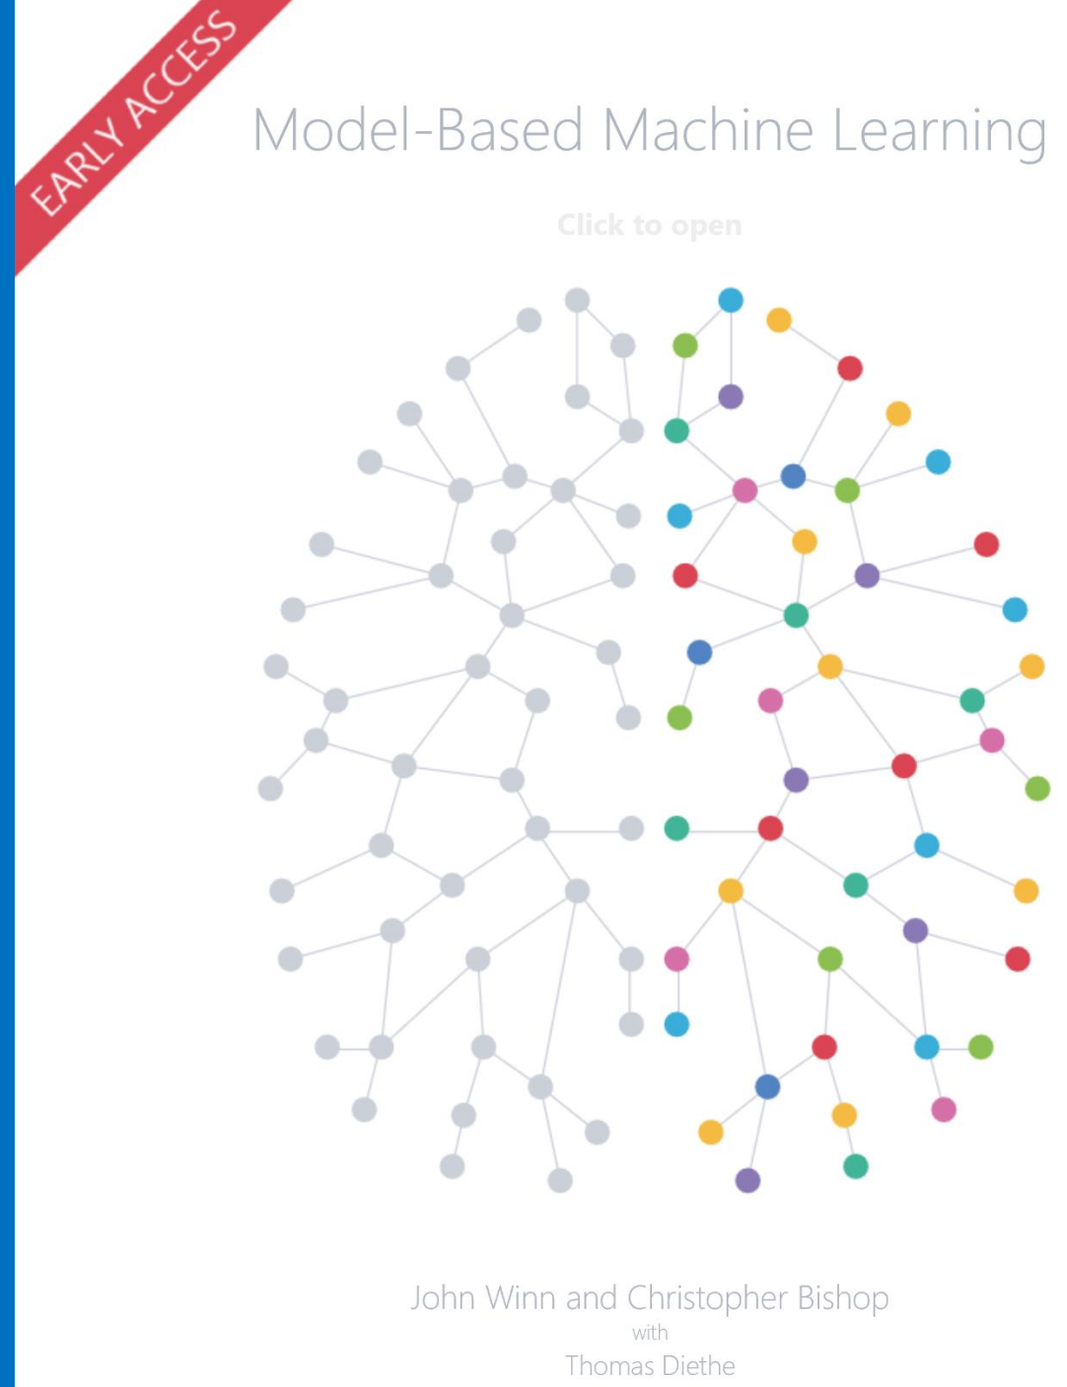

AI Residency Programme

1 year, full-time

Real-world projects with potential to empower millions

Develop AI solutions for healthcare, entertainment, productivity and more

Training and mentorship

Applications close March 30, 2018
Residency starts September 2018

www.aka.ms/AIResidencyCam

Thank you!