BBM406

Fundamentals of Machine Learning

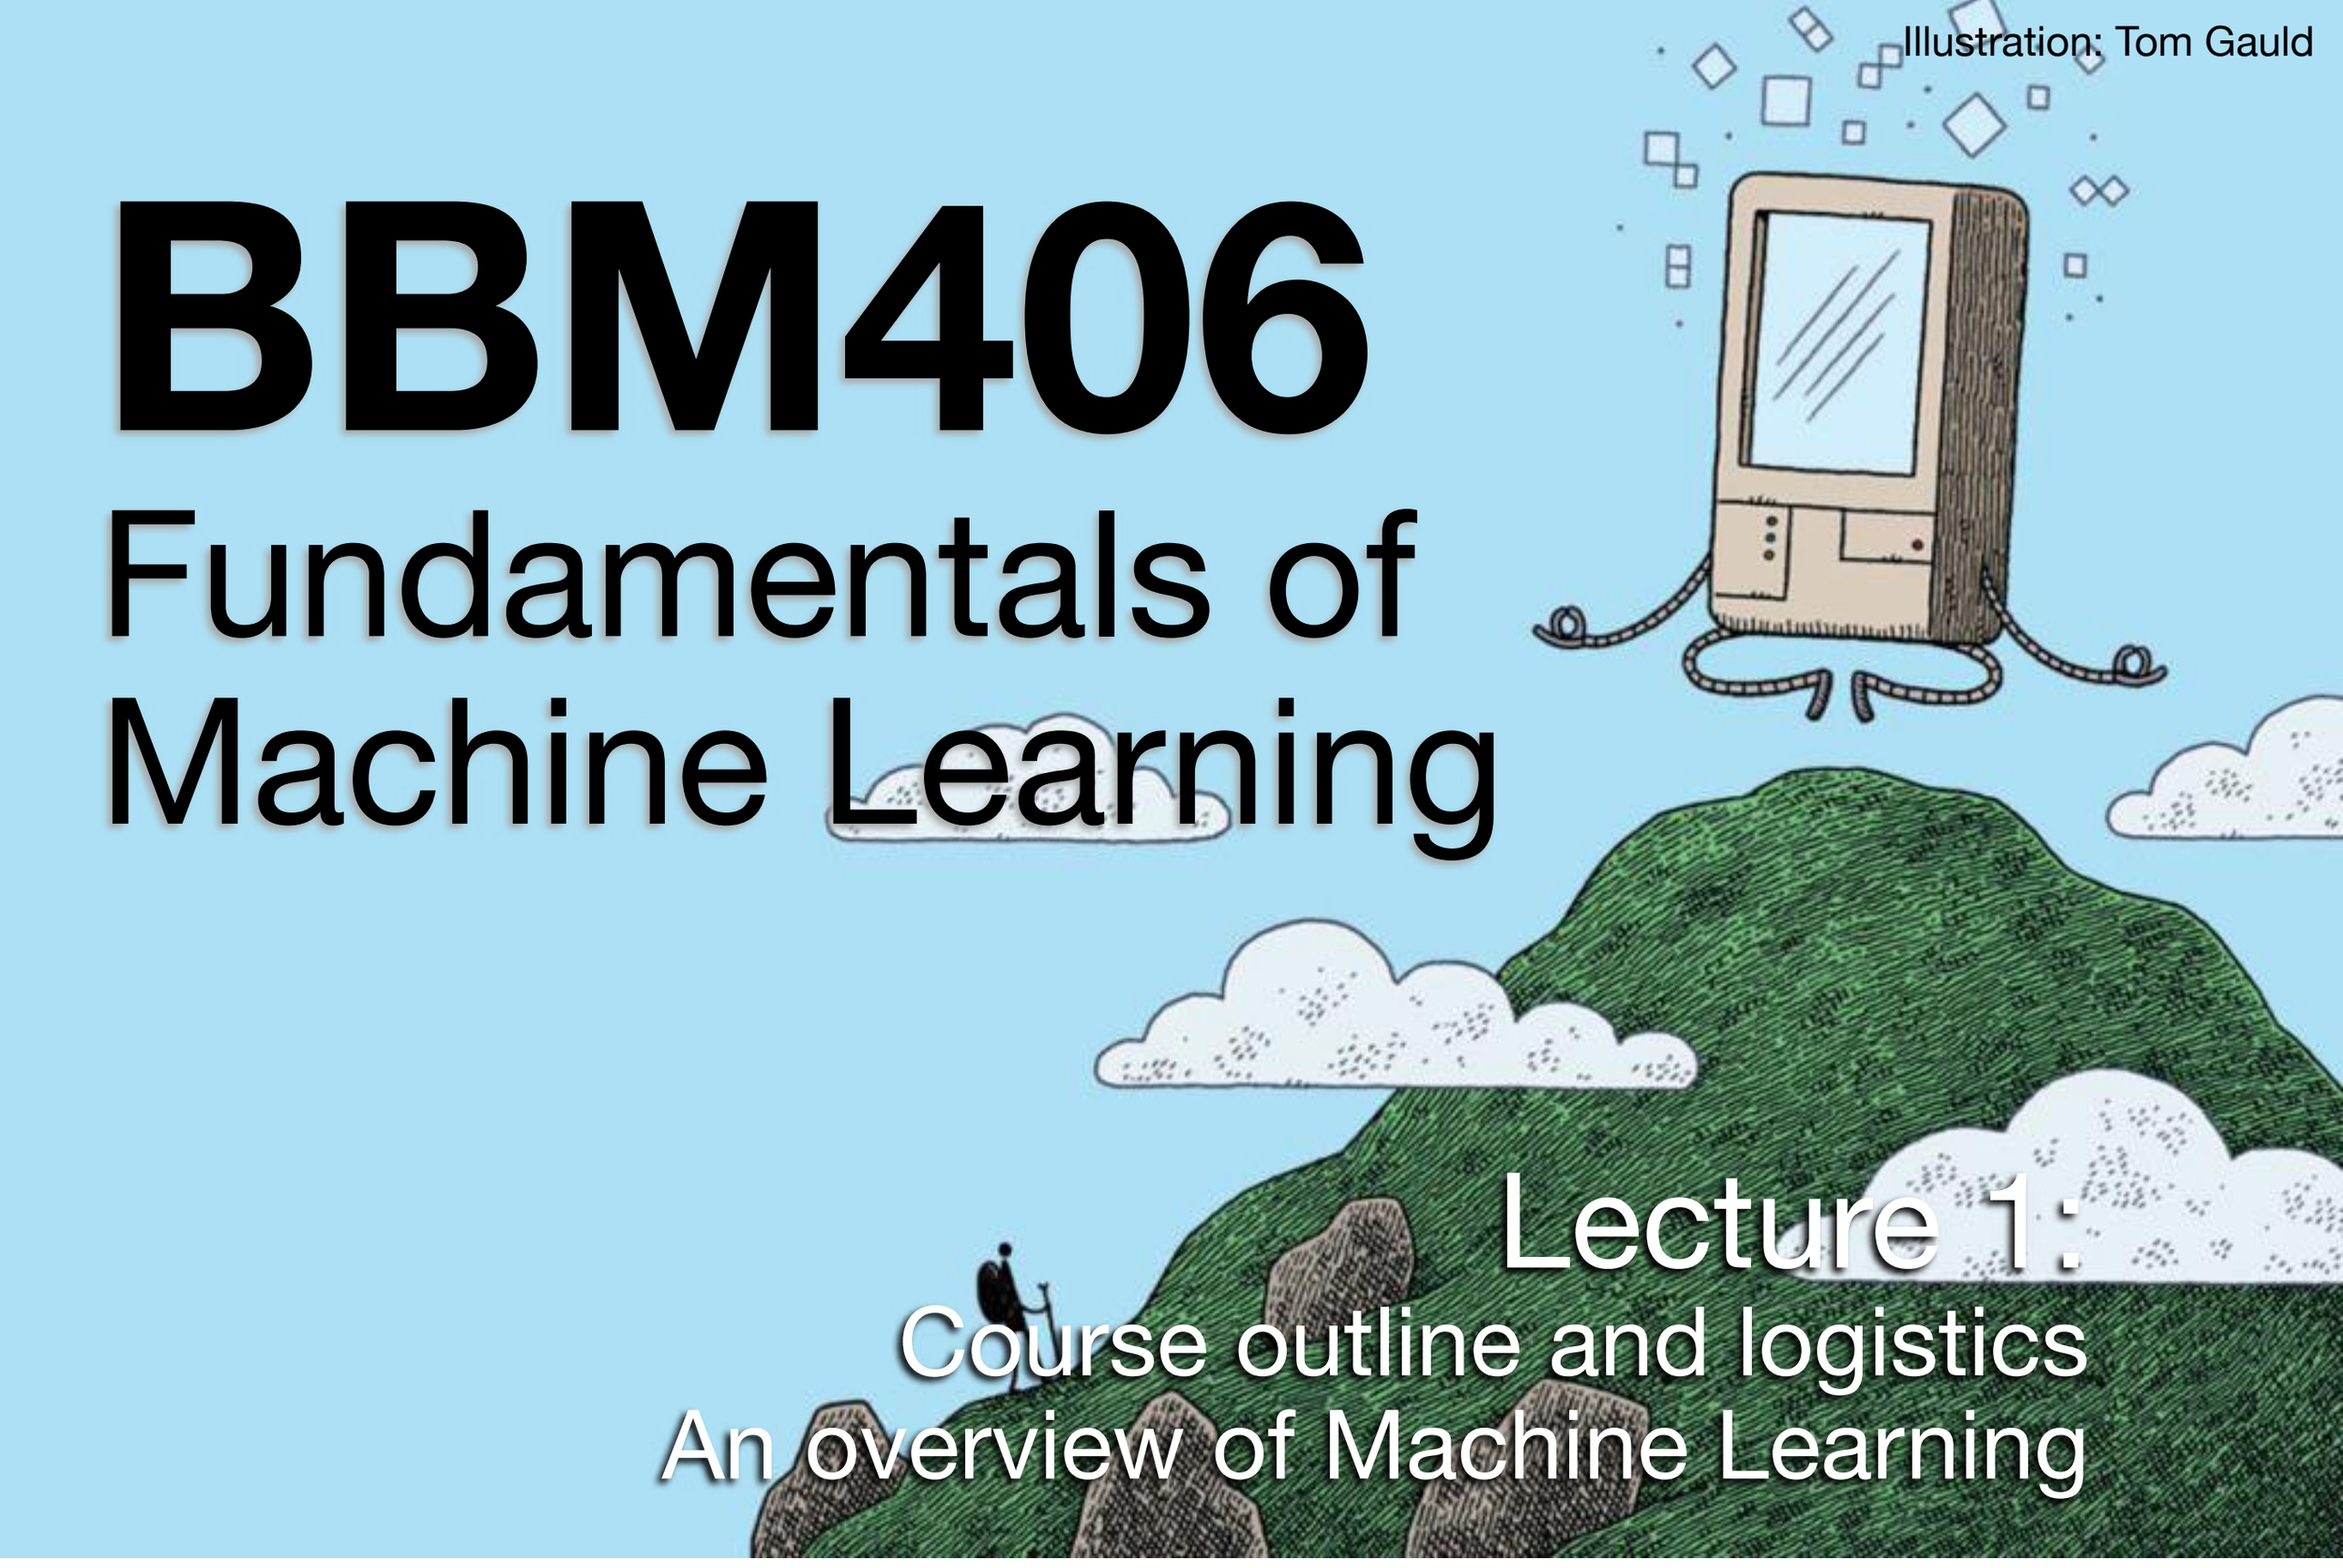

Illustration: Tom Gauld

Lecture 1:

Course outline and logistics
An overview of Machine Learning

HACETTEPE UNIVERSITY COMPUTER VISION LAB

Aykut Erdem // Hacettepe University // Fall 2019

Today's Schedule

- Course outline and logistics
- An overview of Machine Learning

Course outline and logistics

Logistics

- **Instructor:**

Aykut ERDEM
aykut@cs.hacettepe.edu.tr

- **Teaching Assistant:**

Burcak Asal
basal@cs.hacettepe.edu.tr

- **Lectures:** Wed 09:00 - 09:50_D4
  Fri 09:00 - 10:50_D4
- **Tutorials:** Mon 15:00 - 17:00_D8

About this course

- This is a undergraduate-level introductory course in machine learning (ML)
  - A broad overview of many concepts and algorithms in ML.
- **Requirements**
  - Basic algorithms, data structures.
  - Basic probability and statistics.
  - Basic linear algebra and calculus
  - Good programming skills

common distributions, Bayes rule, mean/median/model

vector/matrix manipulations, partial derivatives

- **BBM 409 Introduction to Machine Learning Practicum**
  - Students will gain skills to apply the concepts to real world problems.

Communication

- **Course webpage:** http://web.cs.hacettepe.edu.tr/~aykut/classes/fall2019/bbm406/

  - The course webpage will be updated regularly throughout the semester with lecture notes, programming and reading assignments and important deadlines.

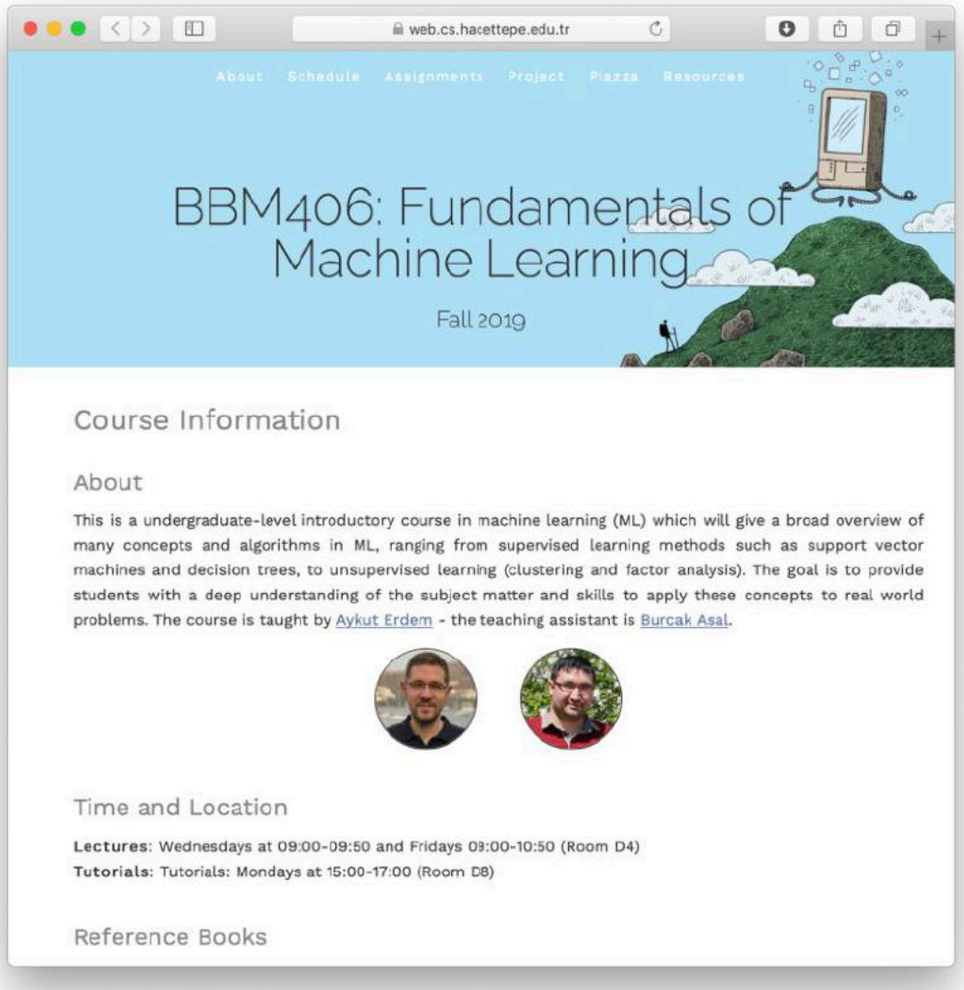

- We will be using Piazza for course related discussions and announcements. Please enroll the class on Piazza by following the link http://piazza.com/class#fall2019/bbm406

Reference Books

- A Course in Machine Learning, Hal Daumé III (**online version (v.0.99) available**), 2017
- Artificial Intelligence: A Modern Approach (3rd Edition), Russell and Norvig. Prentice Hall, 2009
- Bayesian Reasoning and Machine Learning, Barber, Cambridge University Press, 2012 (**online version available**)
- Introduction to Machine Learning (2nd Edition), Alpaydin, MIT Press, 2010
- Pattern Recognition and Machine Learning, Bishop, Springer, 2006
- Machine Learning: A Probabilistic Perspective, Murphy, MIT Press, 2012

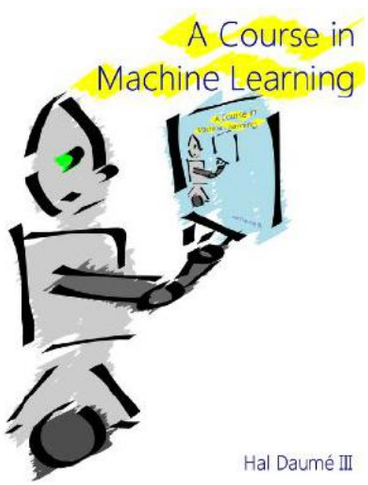

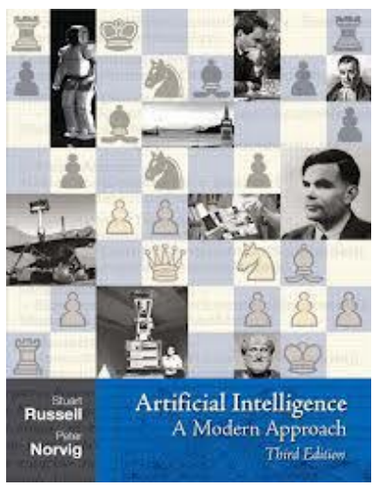

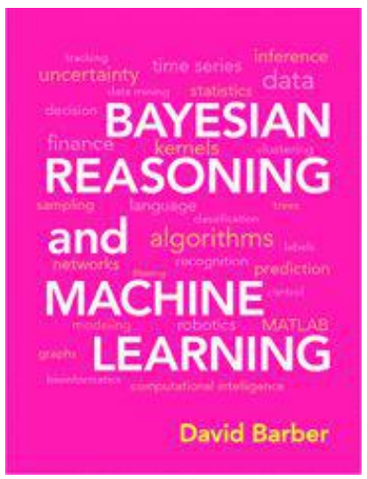

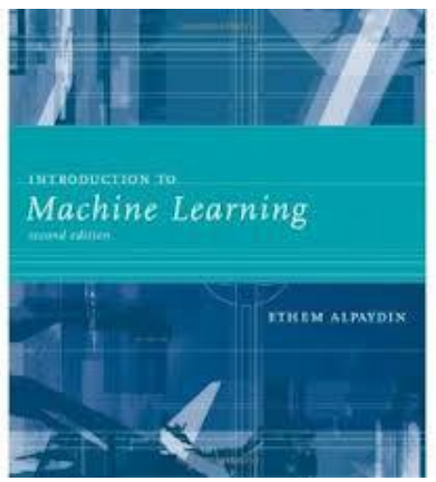

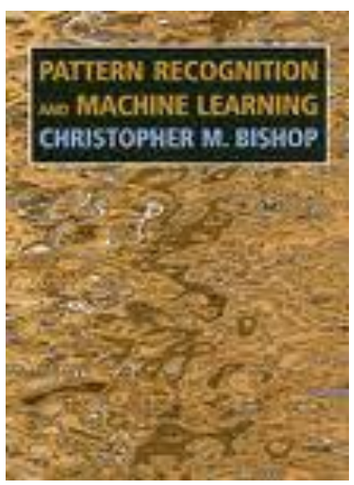

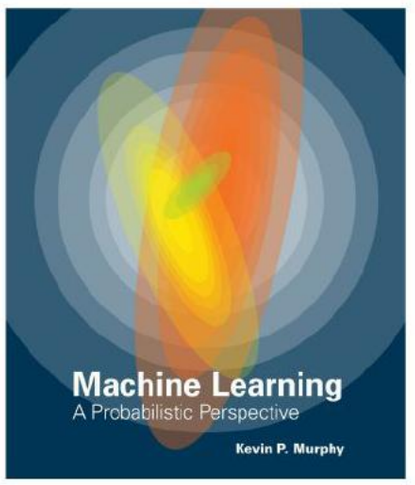

Grading Policy

- Grading for BBM 406 will be based on
  - course project (**done in groups of 2-3 students**) (30%),
  - midterm exam (30%),
  - final exam (35%), and
  - class participation (5%)

- In BBM 409, the grading will be based on
  - a set of quizzes (20%), and
  - 3 assignments (**done individually**)

Assignments

- 3 assignments
  - First one worth 20%, last two worth 30% each

- **Theoretical**: Pencil-and-paper derivations
- **Programming**: Implementing Python code to solve a given real-world problem

- A quick Python tutorial in this week's tutorial session.

KEEP CALM AND DO YOUR HOMEWORKS

Course Project

- Done in groups of 2 or 3 students.
- Choose your own topic (but focused on a specific theme) and explore ways to solve the problem

- **Proposal**: 1 page (Nov 15) (2%)
- **Project Blogs:** Regular blog posts (4%)
- **GitHub commits and meetings with TA:** (4%)
- **Progress Report**: 3-4 pages (Dec 20) (5%)
- **Project Presentation**: Classroom presentation and video presentation (7.5%) (Jan 8-10)
- **Final Report**: 6-8 pages (Jan 12) (30%)

Sample projects from 2016

BBM 406 Class Project - Final Report

Cem Güngör, Fatih Baltacı
Department of Computer Engineering
Hacettepe University
Ankara - TURKEY, Fall 2016
{b21328031, b21327689}@cs.hacettepe.edu.tr

Abstract

This paper is a final report of our project "What Am I Eating?" for BBM406 Introduction to Machine Learning lesson. "What Am I Eating?" is an image recognition project which predicts food labels from given images. Developments in the field of Machine Learning and increase of datasets in recent years encourage us to make an image recognition project. We are using deep learning. We performed transfer learning(from Inception v3 model [Szegedy et al. 2015]) and data augmentation. Our dataset is a combination of different datasets which has 113 classes. Each class has 1000 images.

Keywords: deep learning, image recognition, fine tuning

1 Introduction

In recent years there have been major developments in the field of machine learning. The datasets have grown up because of the increase in internet usage. Hardwares become stronger than before. Graphic cards become cheaper. Because of these conditions, researches have increased and new approaches such as deep learning has appeared. Open source libraries were developed.

Deep Learning is a new and very popular area of Machine Learning research. We decided to develop a project using deep learning to improve ourselves in this field. Deep learning is used in many areas such as image recognition, speech recognition, natural language processing and so on. We used deep learning for image recognition. So, What am I Eating? is a deep learning project that recognizes foods from images.

We saw that no dataset has any Turkish foods. We wanted our project to recognize Turkish foods too. Also we have some future thoughts about our project.

Figure 1: ***pizza (score = 0.84349)****, waffle (score = 0.04952), bruschetta (score = 0.02402), omelette (score = 0.01936), ...*

2 Related Work

There are three researches which are closely related to our research topic. All of them are new and made in 2016. One of them is [Liu et al. 2016]. The purpose of this research is to improve the accuracy of current measurements of dietary intake by analyzing the food images captured by mobile

Predicting Restaurant Ratings from Review Texts

Melis Mutlu & Raisha Abdillah
Department of Computer Engineering
Hacettepe University
Beytepe, Ankara 06800, Turkey
{b36160843,b21304258}@cs.hacettepe.edu.tr

ABSTRACT

Nowadays, with the growth of crowd-sourced review website, the quality of business is determined by its ratings and reviews. The costumer and the business owner will be able to see the trends, making decision, and getting recommendations based on their preferences just by looking at the reviews and ratings themselves. In this project, our goal is to predict the ratings which is given to a restaurant by looking at it its review text. We use Yelp Dataset for our training and testing. By applying machine learning and text mining principle, we analyzed the review text from the Yelp Dataset. We were researching for the best algorithm which would give us the best result. The algorithms which we used at this projects are Bayesian Ridge Regression, Suppert Vector Regression, and Random Forest Regression.

1 Introduction

The development of technology makes it easier for people to make the right decisions. In this matter, technology influences the field of business by delivering a more convenient way for people to evaluate their business. For example, nowadays costumer may look at the reviews and ratings which has been given and getting influenced by it, before deciding to go to a certain restaurant.

The goal of our project is to choose a supervised machine learning algorithm which will give us the best performance in predicting the restaurant ratings by looking at its review text that has been given in Yelp Dataset. Firstly we have to choose the most appropriate dataset to our problem. After that, in order to work with Machine Learning algorithm, we transform our raw data into vector or matrices form.

For our project we use Yelp Dataset, since it already provides the review and rating in an easily accessible format. Then, we did feature extraction from our dataset. We combined several feature extracting process in order to get the better result. For this, we use Bag of Werds and Werd2Vec model. We have tested these model and it gave us a satisfying result. For the better result, we also removed words which we considered unimportant. After we made our model, we use machine learning algorithm to test our model. We then choose the algorithm which gave us the best performance after we tested it. We treated this problem as regression problem, therefore we used regression algorithm. We made use of Yelp Dataset as our training set and testing set.

In this report, firstly, we will present you the dataset. Secondly, we will tell you about out feature extraction method (Bag of Words, Word2Vec). The next part is that we will explain about the algorithm which we use for this projects, which consists of Bayesian Ridge Regression, Support Vector Regression and Random Forest Regression. Then, by using Explained Variance Score (R^2 score) and Mean Square Error we calculate the accuracy of our model. We will share the result and the conclusion of our project by the last part of this report.

Finding The Ingredients of Pizza Using Deep Learning

Mümin Can Yılmaz — can.yilmaz12@hacettepe.edu.tr
Alim Giray Aytar — giray.aytar12@hacettepe.edu.tr
Hayati İbiş — hayati.ibis12@hacettepe.edu.tr

Abstract

Extracting ingredients from a dish can be a powerful tool for combatting obesity and making food inspection processes easier. For this purpose, we tried to create a program which extracts ingredients from a pizza, using convolutional neural networks. We also created a dataset which has 7405 images and 20 different labels as ingredients. Our experiments show us our model can predict small numbers of ingredients successfully (80 percent for one label), however as the number of ingredients increased, accuracy rate drops significantly (22 percent for 2 labels).

1. Introduction

Our aim is to create a model which can identify ingredients in the pizza. Our program should output a list of ingredients as output when feed with an image of a pizza.

First of all, we started with creating a new dataset from the scratch, because we couldn't find any ready-to-use dataset. To do this, we collected about twenty five thousand images from web and labeled all of them by hand with a little software we created for this purpose.

Secondly, we decided to use a Convolutional Neural Network, because they show much better performance in image recognition problems compared to other approaches. Also when using Convolutional Neural Networks, we don't need to extract any features because CNN's operates directly on images. There is also some downsides of using Convolutional Neural Networks as they need more data and require more computing power than other solutions.

Finally, we evaluated our project with the result that we get after the process of training our classifier model which we present in the results section.

Hardest part of this problem is, because food shapes are deformed after cooking, it might not be possible to predict them correctly for our model. Color information also isn't very helpful, because some different ingredients exactly have the same colour or same ingredients might have different colours.

hamsi: 0.58653
baklava: 0.30801
carrot cake: 0.05741
humus: 0.01253

yelp

```
{
 'type': 'business',
 'business_id': (encrypted business id),
 'name': (business name),
 'neighborhoods': [(hood names)],
 'full_address': (localized address),
 'city': (city),
 'state': (state),
 'latitude': latitude,
 'longitude': longitude,
 'review_count': review count,
 'categories': [(localized category names)]
 'open': True / False (corresponds to permanently closed, not business hours),
}
```

Green Pepper
Olive
Onion
Salami
Corn
Chicken
....

Sample projects from 2017

Predicting the Location of a Photograph

Ali Yunus Emre ÖZKÖSE
Hacettepe University
ANKARA, TURKEY
aliozkose@hacettepe.edu.tr

Tarık Ayberk YILIKOĞLU
Hacettepe University
ANKARA, TURKEY
tarikyilikoglu@hacettepe.edu.tr

Abstract

In this paper, we addressed to prediction of an image location problem. It is still a hard problem because of several kinds of other problems. We use convolutional neural networks (CNNs) to tackle this problem. We collect data from Flickr[13], create a dataset which we call Turkey15 and test with basic algorithms. After testing the dataset, we train AlexNet and ResNet-18 with Turkey15 from scratch. Since Turkey15 is very small, we use transfer learning to improve results. We use feature extracting and fine-tuning[14]. We also freeze some layers to get better accuracy.

(a) Ankara (b) İstanbul (c) Bursa

Figure 1: Images from Turkey15

1. Introduction

Although there are a lot of works on this issue and it is very popular research topic in recent years, predicting the location of an image is still a hard problem. There are various problems such that constructing features [3], viewpoint problem[4], illumination and structural modification[12] etc.. It can be used for many areas such as estimation people's perception [5]. But how can we predict the location of given image? In this work, We focus on exactly the problem of city classification.

With the development of technology and the increase of applications, people are taking photos and upload to internet much more than ever. The significant point of sharing is that a huge data has existed and it can be used for creating artificial machines as an experience. At this point, we collected images from Flickr where are taken in Turkey, created a dataset which we called Turkey15 and predict image locations where is limited to Turkey.

Figure 2: Images from Turkey15

First of all, we tested our dataset with hand-crafted features which are Tiny images, GIST features, and Hog features, because we should know that our dataset is convenient enough to use as a dataset or not. Details in this process are explained in section 3.1.

After testing the dataset, we trained existed models which are AlexNet and ResNet-18 models with our dataset. We trained from scratch in this step and get some results and compare with training with hand-crafted features. Details and result are written in section 4.1.

Thirdly, we used transfer learning, in particular, fine-tuning and feature extracting. We trained pre-trained models which are trained with places365 and imageNet datasets. Models are AlexNet and ResNet-18 again. Details are written in section 4.2.

Finally, we froze some layers of models and trained AlexNet and ResNet-18 again. Details are written in section 4.3.

2. Related Work

Because of the popularity of this challenge, there are many kinds of proposed methods and works for predicting location. Li et al. propose to represent features with SIFT and match query image features to database image features mutually[11], but matching is only among the prioritized features. They keep informative points. In this way, they reduce computational cost. We also used hand-crafted features for testing dataset, but we use convolutional neural networks for training.

Sound of The City

Buğrahan Akbulut
Department of CS Hacettepe University
akbulutbugrahan@gmail.com

Mustafa Çağdaş Çaylı
Department of CS Hacettepe University
cagdas_cayli@hotmail.com

Abstract

In this paper we will introduce our project that is detects and classify leading sounds on urban sounds. We focused on audio because it was more attractive then working on image or some numerical data and also because sound is a very important tool for understanding the world.Also another reason is working with sound is very challenging because it is hard to find only one pure sound on outside world there are lots of sound sources and we generally hear the mixture of these sounds,so our data sets that we used in this project have real field records - has lots of mixed sounds-.We worked on UrbanSound8K and UrbanSound data sets containing 27 hours of audio with 18.5 hours of annotated sound event occurrences across 10 sound classes(air conditioner,car horn, children playing, dog bark, drilling, engine idling, gun shot, jackhammer, siren, and street music).Our goal was extract leading sounds with a correct shape by using Shogun and classify them correctly.

1. Introduction

Since new audio technologies developed rapidly recently, audio processing and classification are growing research fields and it contains many challenges. Especially separating audio into its components is a very tough problem.However working on an analysis of urban sounds instead of working on the analysis of speech, music, bioacoustics is relatively easy and relaxing. Furthermore we worked on extraction of the leading sounds with correct shape.

One of the main challenges in this project was lack of labeled mixed sounds.Previous work focused on classification of single labeled audio data. We needed lots of audio data to get our final results correct. With this purpose we created our own multi-labeled audios by using shogun.Actually we first wanted to separate a given any kind of mixed sound into its components by using ICA (independent component analysis) but we could not find any working library or implementation of this algorithm and due to the restricted time we could not achieve this goal.But we wanted to make it so we have done some more researches and find a new library named shogun which provide some tools for mixing and separating sounds not like ICA but it works for us to get some results by making tests on mixed sounds.

After all these things we also want to improve our results getting from tests, we decided to combine two different machine learning approach to get higher results and it was another challenge for us to increase our results by using neural networks and support vector machines combination.The aproach we use to combine these two algorithms will be explained in more detail at "The Approach" section.

Here you can see wave-plot form of single and mixed sound sources we worked on :

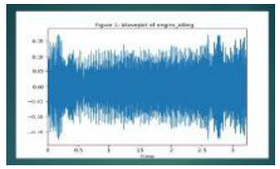

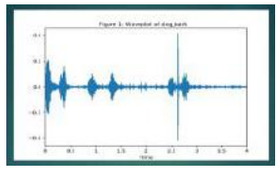

Prediction Of Life Quality

Tark Ramazan BASOGLU
Department of Computer Engineering
Hacettepe University, Ankara, TURKEY
tarik.basoglu@hacettepe.edu.tr

Emre DOGAN
Department of Computer Engineering
Hacettepe University, Ankara, TURKEY
emre.dogan@hacettepe.edu.tr

Abstract

In this study, we mention about the usage of using a machine learning approach to specify life qualities of cities instead of public research. We create an assorted dataset that contains statistical and physical features. To do that, we utilize from MAPZEN. We expect to predict the scores on MOVEHUB with high accuracy.

1. Introduction

Nowadays, we can easily see that cities differ considerably from each other in terms of their physical and social characteristics and that difference is highly influential in human life. We are making great efforts to determine the effects of these differences on human life and to make cities more livable and to change this imbalance positively.

In this situation, we are faced with a notion named quality of life.

"Quality of life (QOL) is the general well-being of individuals and societies, outlining negative and positive features of life. It observes life satisfaction, including everything from physical health, family, education, employment, wealth, religious beliefs, finance and the environment."[2]

By this definition, there are various social and physical criteria that influence the quality of life. The number of researchs and studies carried out in this area is increasing day by day. While life quality information for large cities is easily accessible, it is not possible to find reliable results for cities that are not big enough.

In this project, we purpose to achieve higher efficiency in shorter time and reduce the burden on a human in such researches. Rather the laborious and time-consuming processes of public researches we also aim to provide a new, flexible and developable method by making use of machine learning experiences. Thus, we get a chance to detect the life qualities for any cities in the world. At the same time, we are expecting to be able to observe which physical factors effects the life quality with which rates.

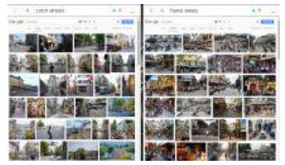

Figure 1. The reflection of the crowd difference between the Hanoi and Zrich on the street photos

MoVEHUB
There is a platform named MOVEHUB that helps you make informed decisions about where to move to around the world. And it has a city ranking list consists of over 200 cities. We utilized this list as the main target in the estimation results.

MAPZEN
Mapzen is an open and accessible mapping platform that is focused on the core components of geo platforms, including search, rendering, navigation, and data.

2. Related Work

There are numberless researchs done to measure life quality in cities every year. In this researches generally, lots of criteria are considered to obtain correct results. Such researches have been carried out in the form of public opinion polls up to now.

MOVEHUB: MOVEHUB is similar research that includes

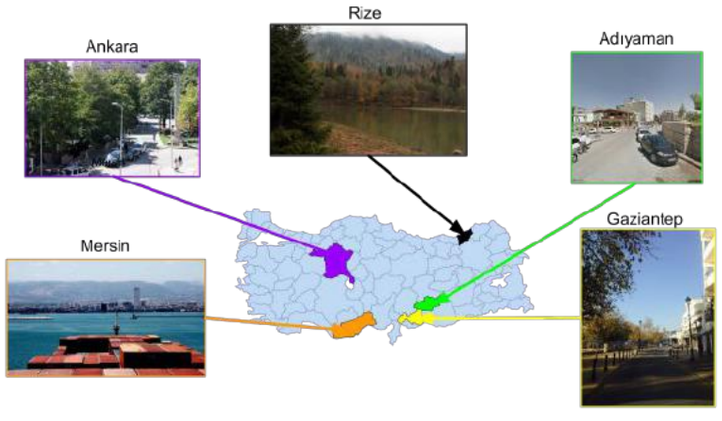

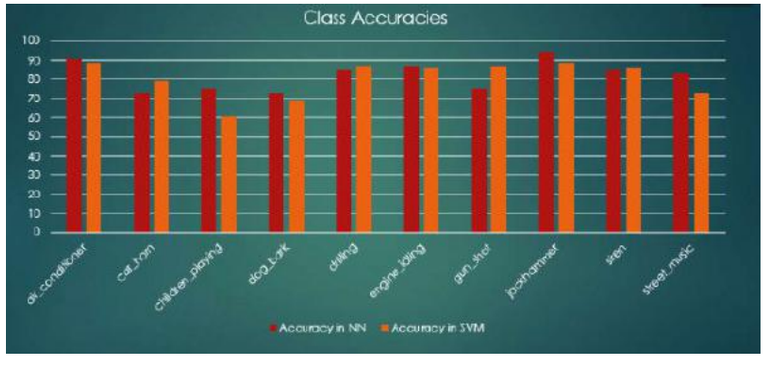

Sample projects from 2018

Wi-Fi Based Indoor Positioning Systems

Burak Emre Ozer
Hacettepe University
Ankara, Turkey
burak.ozer@hacettepe.edu.tr

Furkan Caglar Gulmez
Hacettepe University
Ankara, Turkey
furkancaglargulmez@gmail.com

Huzeyfe Kocabas
Hacettepe University
Ankara, Turkey
kocabashuzeyfe@gmail.com

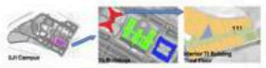

Figure 1: Source [12] Left: map of UJI campus and Tx buildings. Middle: red indicates ESTCE - Tx building. Right: example of a reference point.

Abstract

Wi-Fi Fingerprint-based positioning approach that detects the position of user or device is widely used in the indoor positioning systems instead of Global Positioning System (GPS). In this approach, Received Signal Strength (RSS) values that are known as Wi-Fi fingerprints used. Received Signal Strength values are the measurement of the power present in a received radio signals. We use UJIIndoorLoc dataset with 19937 training records and 1111 test records. This dataset of RSS values are collected by using previously placed wireless access points (WPAs) in Tx Buildings of University of Jaume I campus. We aim to predict location points with respect to floor IDs, building IDs, longitude and latitude values with supervised machine learning algorithms such as K-Nearest Neighbor Algorithm, Random Forest Algorithm, Support Vector Machine and Decision Tree Algorithm. Then we use the model with the highest accuracy in the rest of the progress. Classification techniques are used for building and floor classification and regression techniques are used for detection of location points.

***Keywords*—** indoor positioning, received signal strength indication (RSSI), machine learning algorithms, classification, random forest

1. Introduction

Global Positioning System (GPS), which uses satellites,is the most popular outdoor positioning system,however its signals can be easily blocked by various structures and factors then it becomes useless for indoor environment because of signal loss. Unlike the GPS, Indoor Positioning Systems aims to detect the position of user or device by using Access Points signal also called Wi-Fi ngerprint. With the advancing technology and spread of wireless networks, Indoor Positioning Systems become even more important place in the fields of augmented reality, social networking, personal tracking, guiding blind people, tracking small children or elderly individuals and location-based advertising etc.

Wi-Fi-based fingerprint methods have some problems when positioning phase in indoor. These problems can be caused by the fact that the devices in which the radio signals are collected during the training stage and the devices in the test phase are different.

Another reason is that the number of access points in the environment varies greatly. Inevitably, these problems negatively affect positioning success. However, we will try to determine the position with regression algorithms using the real latitude and longitude values of the collected locations.

Then we will turn our problem into classification problem by using the building and floor features in the data set. In the test section, we will try to estimate which building is located or which floor of the building.

Different machine learning algorithms will be tried and we will decide the most suitable algorithm for indoor positioning.

We will use the UJIIndoorLoc database throughout the entire project. Classification and regression problems will be solved using the RSS values from 520 wireless access points (WAPs). In the classification part, since the data set contains 3 buildings, we will divide the data into three and try to estimate which floor is located.

The rest of the paper is organized as follows: Presentation of related studies, explanation of the data set used in the experiment, explanation of solution approaches, experimental results

Country Classification Using House Photos

Meltem TOKGOZ
Hacettepe University
21527381
meltemtokgoz@hacettepe.edu.tr

Enes Furkan CIGDEM
Hacettepe University
21526877
enescigdem@hacettepe.edu.tr

Asma AIOUEZ
Hacettepe University
21504074
asma.aiouvez@hacettepe.edu.tr

Abstract

Home designs vary from country to country and when we talk about housing, we should refer to both modern and traditional styles. You can come across a picture of a house taken by someone anywhere in the world and you may wonder where it has been taken from. In this project, we tried to find out which country the photo of a house was taken from. In short, we worked on the problem of classification according to where the photographs were taken.

We used our own World dataset for this project. This dataset contains over 4000 pictures for 15 different countries. In our project, we collected our data from the Flickr [1] , Pinterest [3], and Google Photos [2]. We first tested our data with a single layer neural network and then with convolutional neural networks (CNN). We used ResNet18 and AlexNet models when implementing CNN in our project. In accordance with the results, we applied some methods to increase the accuracy and we got the best accuracy with ResNet18.

1. Introduction

Recognizing home photos and classifying them by country is a quite difficult problem. Because the houses in many countries in the modern world are similar to each other. Beside that, there are some features to distinguish these houses. For example, each country's climate, people's lifestyle and culture are different. This gives us some hints on the architecture of the houses in that country. From this point of view, especially the design of traditionally styled houses begins to change from a country to another. The main problem here is that the houses in the same continent are very similar to each other. For example as shown in *Figure 1*, in the Asian continent, traditionally styled houses of some countries such as South Korea, Japan, Indonesia and Malaysia are very similar. This factor complicates the solution of the problem. In addition, many factors such as the shooting angle, light, shadow and seasonal differences affect the solution of this problem.

Figure 1. Example of similar data

Since this is an image classification problem, there are many algorithms and methods used in its solution. K-nearest neighbors, logistic regression, support vector machine and convolutional neural networks are some of these solutions. Especially in recent years, CNN is a successful algorithm preferred to solving problems in this area.

Figure 2. Example of similar data

In our study, we deal with the problem of classification according to the country where the house pictures were

Rock or Not?

Defne Tuncer [1] **Kutay Barcin** [1]

Abstract

In the era of technology, millions of songs are brought to people everyday. The dramatic increase in the size of music collections has made the music genre recognition (MGR) an important task on machine learning. The goal of this paper is to give machines a chance to predict music genres given input features from music tracks. To do that, we applied various techniques based on machine learning on the dataset called Free Music Archive (FMA), and we have reached an accuracy score of 67.80% as our highest.

1. Introduction

When there is people, there is music. As people, living in today's world, music is always at our reach through technology. The ease of it has brought the demand of automatically generated playlists and customized music recommendations. The task in both those challenges is to be able to group songs in semantic categories. In this work, we aim to model and classify music genres with the assumption of different music genres are also different at the bit level.

In this paper, we will put forward the efforts we made concerning the classification models that allow us to recognize the genre of a given song from its audio features. As for the beginning, we introduced studies on the subject music genre recognition. Then we made a brief introduction to the dataset we bring into use, and explained how we handled our data. Thereafter, we implemented various baseline classification models, and discussed towards advancing the models to solve the problem of music genre recognition. These methods include: 4.1.1 Nearest Neighbor Classifier with/without dimensionality reduction through Principal Component Analysis (PCA) and weighting hyperparameter. 4.1.2 Logistic Regression through one-vs-one scheme, multinomial approach and one-vs-rest scheme with variety of solvers and regularization. 4.1.3 Support Vector Machines with linear and radial basis function (RBF) kernels. 4.1.4 Deep Learning method Neural Network also known as Multi-Layer Perceptron through various optimizers. To represent the audio tracks in building our baseline models we planned to use the combination of all the features, which have been shown to be effective in the task of predicting genres. We improved our methods with model and feature selection by using k-fold cross validation afterwards. Based on the results obtained from the algorithms, we performed experimental analysis. Finally, ended our work with a detailed conclusion, and proposed our feature work.

2. Related Work

For the music genre recognition task, the most common datasets are GTZAN (Tzanetakis & Cook, 2002), Million Song Dataset (MSD) (Bertin-mahieux et al., 2011) and FMA: A Dataset For Music Analysis (Defferrard et al., 2017). While FMA, which consists of 161 sub-genres among 106,574 tracks and published in 2017, is the most up-to-date dataset, and is especially suited for MGR as it features fine genre information. A challenge took place as one of challenges of Web Conference (WWW2018) by the publishers of FMA Dataset on the subject predicting genres of the music (Defferrard et al., 2018). The winner succeeded by examining through artist-related information and scored an accuracy of 66.29% on predicting 16 genres (Kim et al., 2018).

In Music Information Retrieval (MIR), there have been various number of studies on building effective models to predict genre of music using audio features. Mel-Frequency Cepstral Coefficients (MFCCs), one of the audio features, are generally used in music genre classification as the perceptual scale of pitches of a human hearing are represented by the Mel-scale. A Hidden Markov model with MFCCs is used to classify pop, country, jazz and classical genres in (Shao et al., 2004). On the other hand, another study focuses on a new feature called Renyi Entropy Cepstral Coefficients (RECCs) (Tsai & Bao, 2010). The highest achieved accuracy scores reported on the datasets ISMIR2004 which is from the contest (Cano et al., 2006) and GTZAN are accomplished by representing the auditory human perception with a proposed spectrogram (Panagakis et al., 2009). Most of their studies are done through researching the timbre texture,

[1]Department of Computer Engineering, Hacettepe University, Ankara, Turkey. Correspondence to: Defne Tuncer <defnetuncer@hacettepe.edu.tr>, Kutay Barcin <kutaybarcin@hacettepe.edu.tr>.

BBM406 Fundamentals of Machine Learning, Fall 2018.

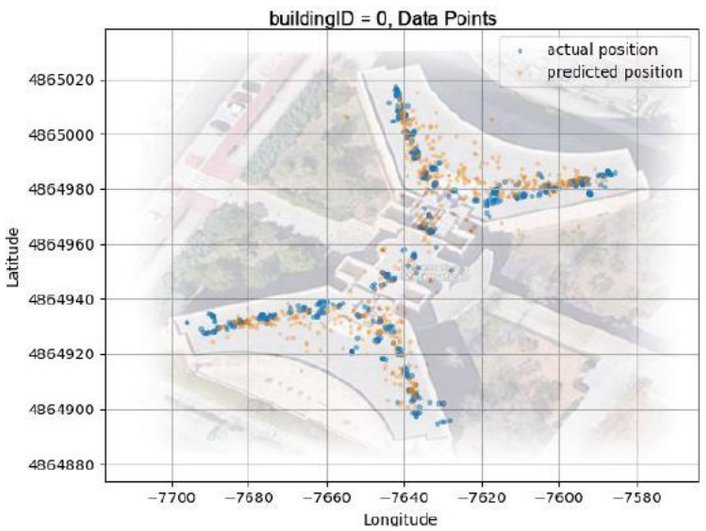

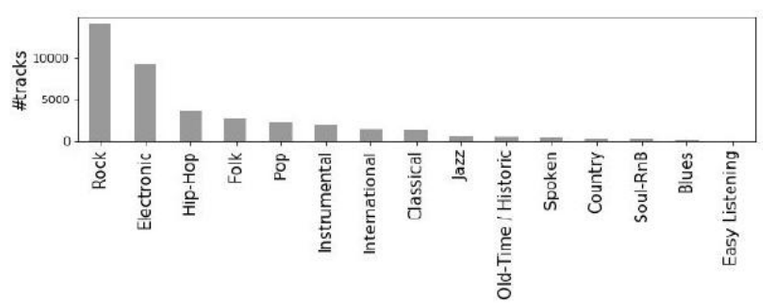

Collaboration Policy

- All work on assignments have to be done **individually**. The course project, however, can be done **in groups of 2-3**.

- You are encouraged to discuss with your classmates about the given assignments, but these discussions should be carried out in an abstract way.

- **In short, turning in someone else's work, in whole or in part, as your own will be considered as a violation of academic integrity.**

- Please note that the former condition also holds for the material found on the web as everything on the web has been written by someone else.

http://www.plagiarism.org/plagiarism-101/prevention/

Course Outline

- **Week1** Overview of Machine Learning, Nearest Neighbor Classifier
- **Week2** Linear Regression, Least Squares

Assg1 out

- **Week3** Machine Learning Methodology
- **Week4** Statistical Estimation: MLE, MAP, Naïve Bayes Classifier

Assg1 due

- **Week5** Linear Classification Models: Logistic Regression, Linear Discriminant Functions, Perceptron

Assg2 out

- **Week6** Neural Networks

Course project proposal due

- **Week7** Deep Learning

Assg2 due

Course Outline (cont'd.)

- **Week8** *Midterm Exam*

Assg3 out

- **Week9** Support Vector Machines (SVMs)

- **Week10** Multi-class SVM, Kernels, Support Vector Regression

Assg3 due

- **Week11** Decision Tree Learning, Ensemble Methods: Bagging, Random Forests, Boosting

Project progress report due

- **Week12** Clustering: K-Means Clustering, Spectral Clustering, Agglomerative Clustering

- **Week13** Dimensionality Reduction: PCA, SVD, ICA, Autoencoders

- **Week14** Course Wrap-up, Project Presentations

Final project report due

Machine Learning:

An Overview

Quotes

- *"If you were a current computer science student what area would you start studying heavily?*
  - *–Answer: Machine Learning.*
  - *–"The ultimate is computers that learn"*

–Bill Gates, Reddit AMA

- *"Machine learning is today's discontinuity"*

–Jerry Yang,
Co-founder, Yahoo

- *"AI is the new electricity! Electricity transformed countless industries; AI will now do the same."*

–Andrew Ng

slide by David Sontag

Google Trends

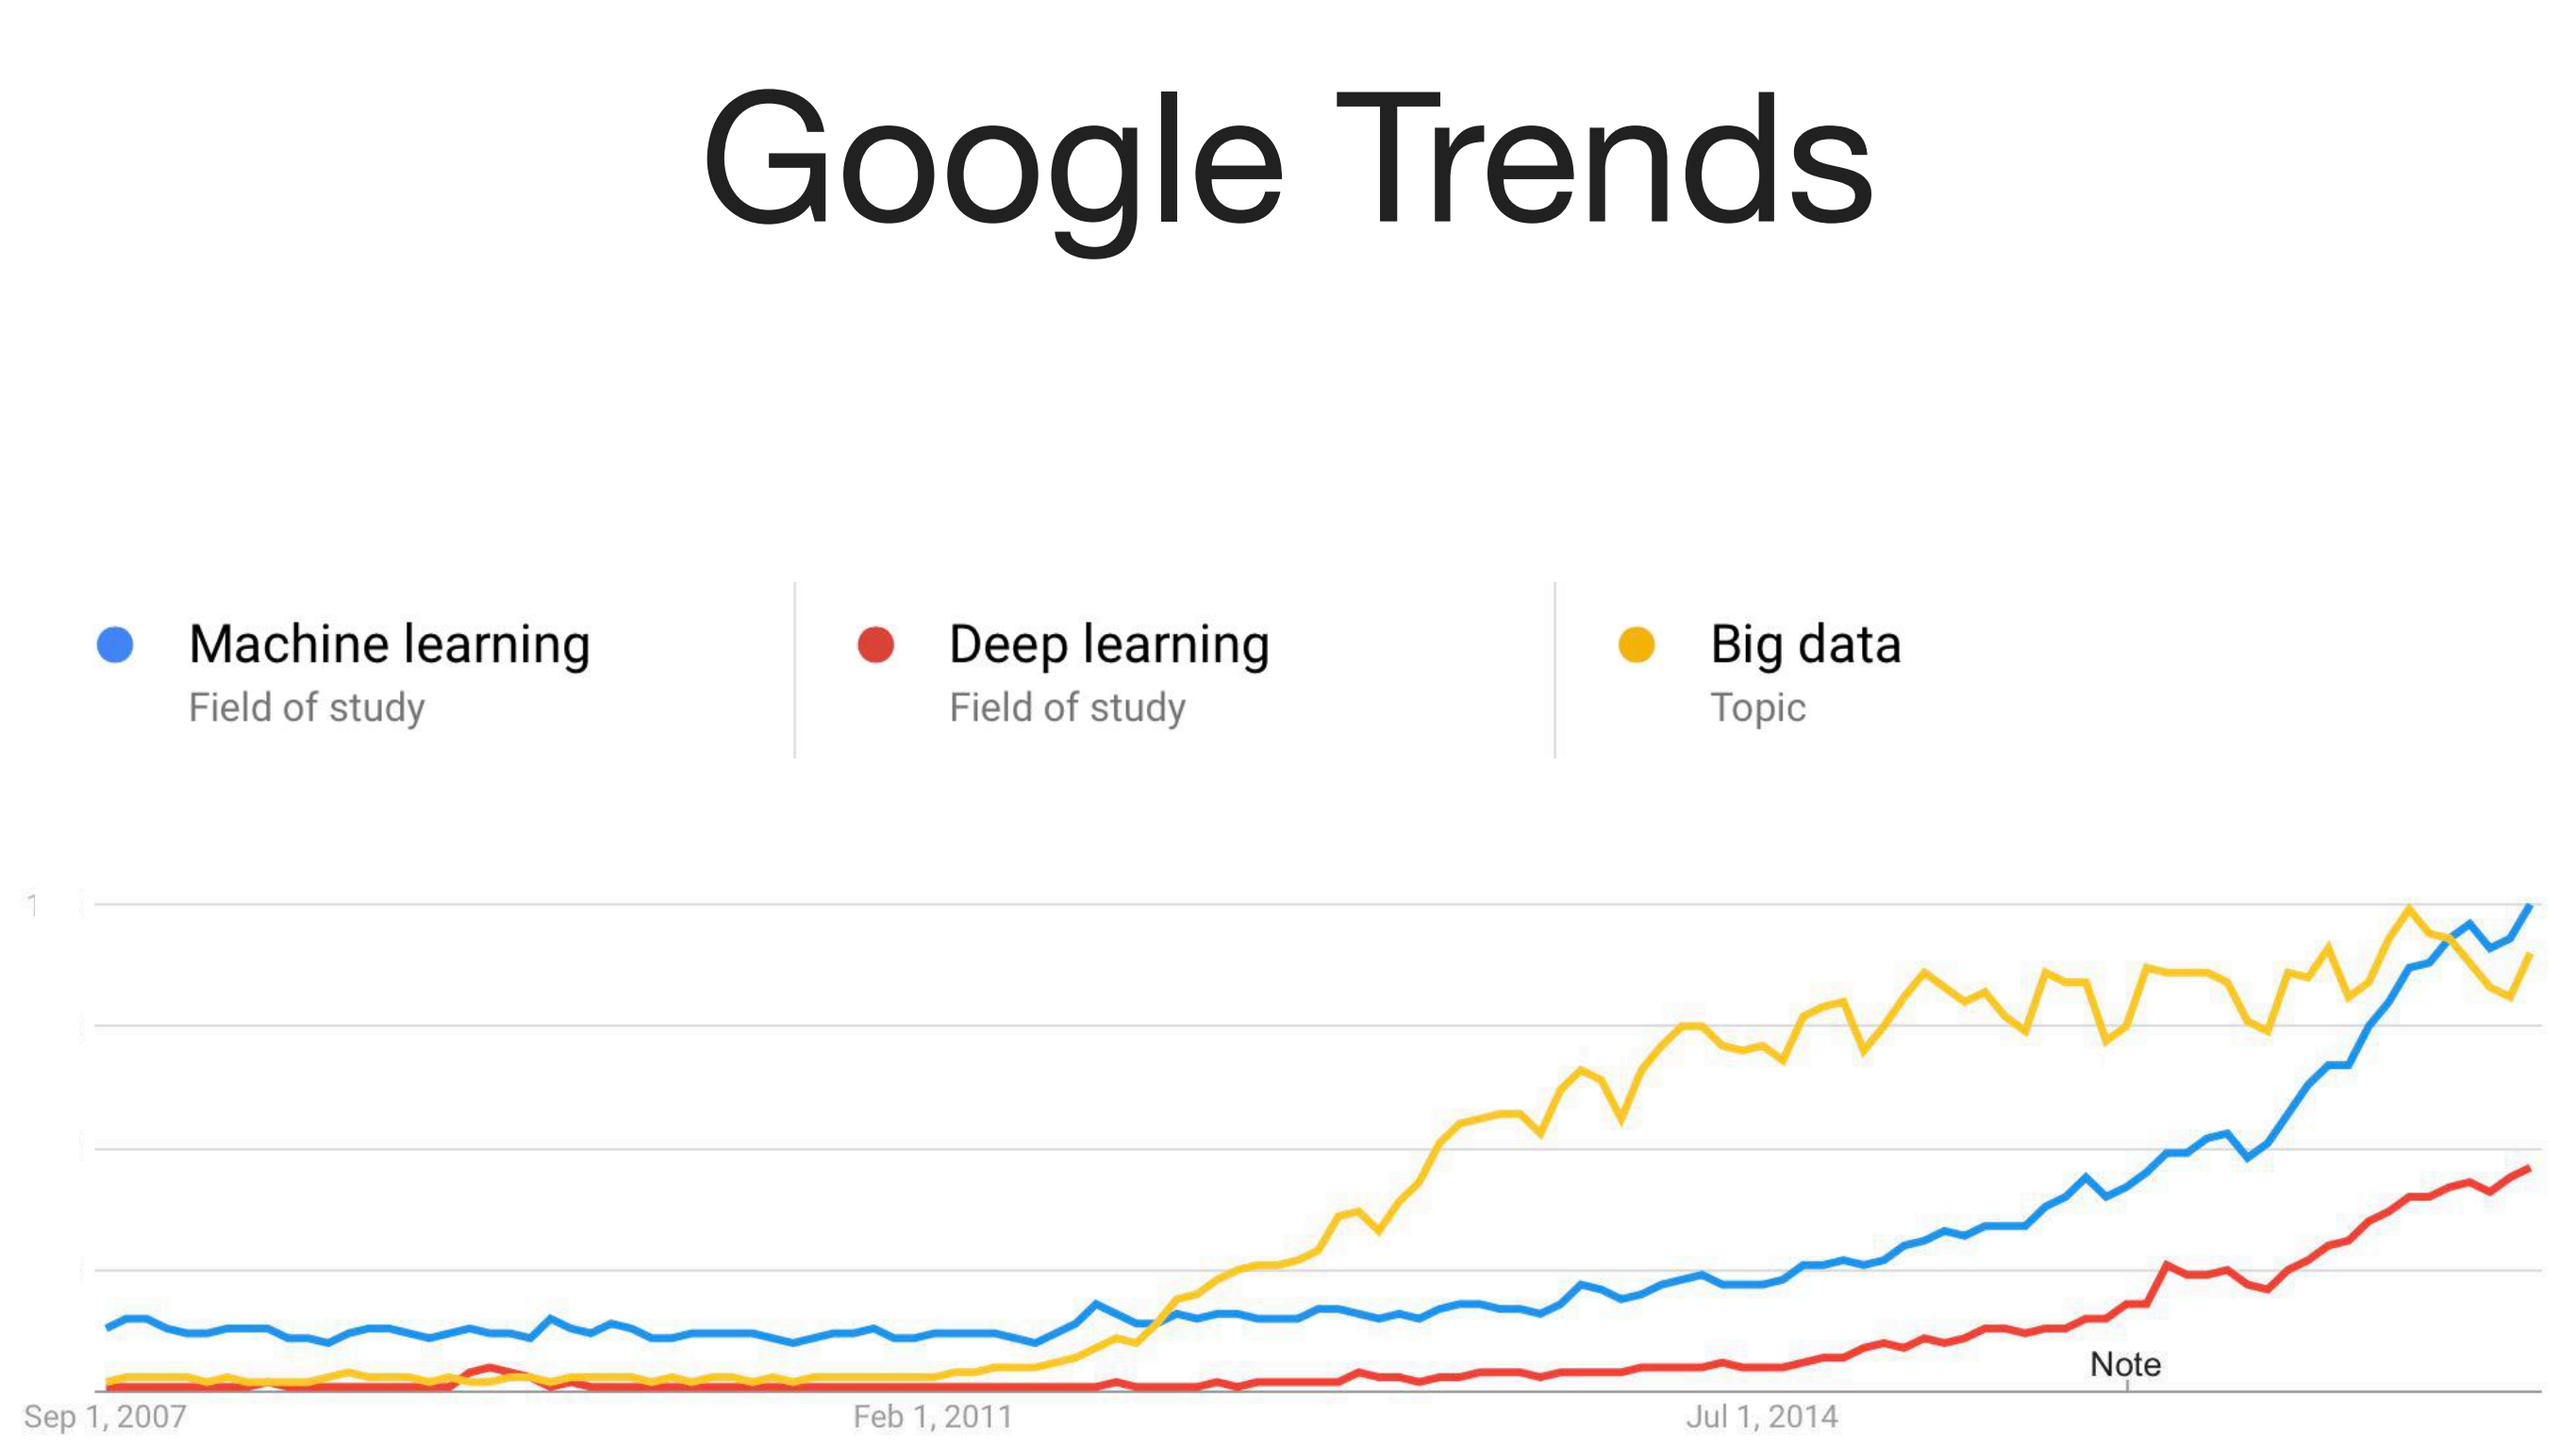

Machine Intelligence LANDSCAPE

2015 Edition

CORE TECHNOLOGIES

ARTIFICIAL INTELLIGENCE

IBM WATSON · MetaMind · Numenta · ai-one · Cycorp · Microsoft Research · nara · Reactor · SI SCALED INFERENCE

DEEP LEARNING

vicarious · Vision Factory · facebook · Baidu IDL · Google · ersatz · SKYMIND · SignalSense

MACHINE LEARNING

rapidminer · context relevant · Oxdata H2O · DATARPM · LiftIgniter · SPARKBEYOND · Azure ML · yhat · Wise.io · Sense · GraphLab · Alpine · tonian

NLP PLATFORMS

cortical.io · idibon · LUMINOSO · wit.ai · MaluubA

PREDICTIVE APIS

AlchemyAPI · MINDOPS · Google · bigml · indico · ALGORITHMIA · Expect Labs · PredictionIO

IMAGE RECOGNITION

clarifai · MADBITS · DNNresearch · DEXTRO · ViSENZE · lookflow

SPEECH RECOGNITION

GRIDSPACE · popup archive · NUANCE

RETHINKING ENTERPRISE

SALES

Preact · AVISO · RelateIQ · NGDATA · CLARABRIDGE · FRAMED · infer · ATTENSITY · causata

SECURITY / AUTHENTICATION

CROSSMATCH · conjur · EYEVERIFY · BITSIGHT · CYLANCE · AREA 1 SECURITY · bionym

FRAUD DETECTION

sift science · SOCURE · ThreatMetrix. · feedzai · Brighterion. · VERAFIN

HR / RECRUITING

TalentBin · entelo · predikt · Connectifier · gild · hiQ · CONCEPT NODE

MARKETING

brightfumel · bloomreach · CommandIQ · AIRPR · RADIUS · TellApart · people pattern · Freshplum

PERSONAL ASSISTANT

Siri · Google now · Cortana · cleversense · tempo · Robinlabs · KASISTO · fuse|machines · VIV · CLARA LABS

INTELLIGENCE TOOLS

ADATAD · Palantir · Quid · Digital Reasoning · FirstRain

RETHINKING INDUSTRIES

ADTECH

METAMARKETS · dstillery · rocketfuel · YieldMo · ADBRAIN

AGRICULTURE

BLUE RIVER · TerrAvion · ceresimaging · HONEYCOMB · THE CLIMATE CORPORATION · tule · mavrx

EDUCATION

declara · coursera · KNEWTON · kidaptive

FINANCE

Bloomberg · FinGenius · alphasense · KENSHO · Dataminr · minettabrook · Orbital Insight · BINATIX

LEGAL

Lex Machina · brightleaf · COUNSELYTICS · RAVEL · JUDICATA · Brevia · DiligenceEngine

MANUFACTURING

SIGHT MACHINE · MICROSCAN · BOULDER IMAGING · IVISYS

MEDICAL

Parzival · transcriptic · ZEPHYR HEALTH · Genescient · grand round table · bina TECHNOLOGIES · TUTE GENOMICS

OIL AND GAS

kaggle · AYASDI · TACHYUS · biota · Flutura

MEDIA / CONTENT

outbrain · newsle · ARRIA · SAILTHRU · wavii · Owlin · NarrativeScience · yseop · Summly · Prismatic · ai AUTOMATED INSIGHTS

CONSUMER FINANCE

Affirm · inVenture · zestfinance · BILL GUARD · LendUp · LendingClub · Kabbage

PHILANTHROPIES

DataKind · thorn · THE DATA GUILD

AUTOMOTIVE

Google · Continental · TESLA · MOBILEYE · CRUISE

DIAGNOSTICS

enlitic · 3SCAN · lumiata · ENTOPSIS

RETAIL

BAY SENSORS · PRISM SKYLABS · celect · euclid

RETHINKING HUMANS / HCI

AUGMENTED REALITY

wearable intelligence · AUGMATE · APX · blippAR · META · layar

GESTURAL COMPUTING

THALMICLABS · omek · FLUTTER · LEAP MOTION · eyeSight · 3Gear Systems · GestureTek · nod

ROBOTICS

intel · LIQUID ROBOTICS · SoftBank · iRobot · Boston Dynamics · jibo · evolution robotics · anki

EMOTIONAL RECOGNITION

affectiva · BEYONDVERBAL · EMOTIENT · BRSLABS · cogito

SUPPORTING TECHNOLOGIES

HARDWARE

NVIDIA · XILINX · QUALCOMM · NERVANA SYSTEMS · TERADEEP · Artificial Learning · rigetti

DATA PREP

TRIFACTA · Paxata · tamr · Alation

DATA COLLECTION

diffbot · kimono · CrowdFlower · Connotate · WorkFusion · import.io

www.shivonzilis.com/machineintelligence

Bloomberg BETA

MACHINE INTELLIGENCE 2.0

2016 Edition

AGENTS

PROFESSIONAL

Howdy! · x.ai · clara · KASISTO · DigitalGenius · OVERLAP.CC · meekan · fuse|machines · PRIMER

PERSONAL

facebook · XIAOICE · large · assistant.ai · nestor · @wesome · Magic

OS INTERFACES

Siri · Cortana · VIV · maluuba · api.ai · CogNEA · Google Now

AUTONOMOUS SYSTEMS

AIR

3DR · dji · PROJECT LOON Google · VERTICAL · DroneDeploy · AIRDOG · SKYCATCH · SKYDIO · Airware · LILY

GROUND

Google · UBER · TESLA · CRUISE · MOBILEYE · COMMA · AdasWorks

SEA

LIQUID ROBOTICS · bluefin data · OPENROV · BluHaptics

INDUSTRIAL

KIVA Systems · fetch · HARVEST AUTOMATION · CLEARPATH · AVIDBOTS · ENERGID · rethink robotics · GREYORANGE · OSARO

ENTERPRISE

SECURITY / FRAUD

Sentinel · graphistry · BITSIGHT · feedzai · AREA 1 · drawbridge · sift science · CYLANCE · Brighterion

HR / RECRUITING

textio · hiQ · gild · SpringRole · entelo · unifive · GIGSTER

SALES

6sense · infer · people pattern · Preact · Prism · AVISO · Vidora · sentient · salespredict · Gainsight

MARKETING

LiftIgniter · RADIUS · brightfunnel · retention SCIENCE · AIRPR

CUSTOMER SUPPORT

CLARABRIDGE · QUANTIFIND · Wiseio · ACTIONIQ · FRAMED · DigitalGenius

INTERNAL INTEL

Alation · ADATAO · Palantir · sapho · lucid · Rainbird · SKIPFLAG · agolo · Digital Reasoning · Narrative Science

MARKET INTEL

Quid · mattermark · DataFox · bottlenose · PREMISE · enigma · CB INSIGHTS

PLATFORMS

RESEARCH / AGI

OpenAI · vicarious · Google DeepMind · Numenta · Cycorp · nnaisense · SCALED INFERENCE · CURIOUS · GEOMETRIC INTELLIGENCE

FULL STACK

context relevant · CognitiveScale · NVIDIA · TERADEEP · QUALCOMM · nervana SYSTEMS

MACHINE LEARNING

Dato · rapidminer · cortical.io · AYASDI · amazon web services · Azure Machine Learning · narrative logics · PredictionIO · SKYTREE · bigml · blueyonder

INDUSTRIAL IOT

ThingWorx · UPTAKE · IMUBIT · Preferred Networks · Alluvium · xively · PLANET OS

AUDIO

Gridspace · TalkIQ · nexidia · vocaliq · NUANCE · Expect Labs · popup archive

VISION

ORBITAL INSIGHT · Descartes Labs · DEXTRO · cortica · clarifai · PLANET LABS · MetaMind

DATA ENRICHMENT

diffbot · Paxata · TRIFACTA · iDIBON · WorkFusion · loop AI Labs · CrowdFlower

INDUSTRIES

ADTECH

ADTHEORENT · dstillery · BEYONDVERBAL · METAMARKETS · TAPAD · rocketfuel · affectiva

AGRICULTURE

BLUE RIVER · tule · TerrAvion · mavrx · THE CLIMATE CORPORATION · CERES TECHNOLOGIES · HONEYCOMB

FOR GOOD

Conservation Metrics · DataKind · DATA POP · thorn · BAYES IMPACT

RETAIL FINANCE

inVenture · Affirm · earnest · MIRADOR · Lendo · zestfinance · LendUp

LEGAL

Everlaw · RAVEL · LEGAL ROBOT · Sseal · BEAGLE · ROSS · Lex Machina

MATERIALS & MFG

zymergen · AUGMATE · GINKGO BIOWORKS · ITRINE TECHNOLOGIES · SIGHT MACHINE · CALCULARIO · Eigen Innovations

HEALTHCARE

deep genomics · 3SCAN · enlitic · Calico · BUTTERFLY Network, Inc · Atomwise · Recombine · color · METABIOTA · GRAND ROUNDS · Google Life Sciences · IBM Watson Health

INDUSTRIES (CONT'D)

EDUCATION

KNEWTON · coursera · turnitin · gradescope · UDACITY · KHANACADEMY

TRANSPORT & LOGISTICS

NAUTO · taleris · PRETECKT · clearmetal

INVESTMENT FINANCE

Bloomberg · Quantopian · Dataminr · KENSHO · iSENTIUM · NEURENSIC · alphasense

TECH USER TOOLS

DATA SCIENCE

DOMINO · kaggle · Sentenai · sense · yseop · Outlier · ŷhat · DataRobot

MACHINE LEARNING

Cortana Analytics · AlchemyAPI · glowfish · IBM Watson Platform · Anodot · MonkeyLearn · (h [s]) HyperScience · fuzzy.io · SIGOPT · 0xdata H2O · SPARKBEYOND · indico

OPEN SOURCE

SKYMIND · TensorFlow · seldon · Caffe · theano · Spark MLlib · Microsoft DMTK · spaCy · DL4J · SciKit · CGT

SHIVONZILIS.COM/MACHINEINTELLIGENCE

MACHINE INTELLIGENCE 3.0

2017 Edition

ENTERPRISE INTELLIGENCE

VISUAL
Orbital Insight, planet, clarifai, DEEPVISION, cortica, Sigocian, SPACE_KNOW, Captricity, netra, deepomatic

AUDIO
Gridspace, TalkIQ, nexidia, twilio, CAPIO, Expect Labs, Clover, Mobvoi, Qurious.AI, popUP archive

SENSOR
PREDIX, C3IoT, MAANA, Sentenai, PLANET OS, UPTAKE, IMUBIT, Preferred Networks, thingworx, KONUX, Alluvium

INTERNAL DATA
PRIMER, IBM WATSON, Cycorp, Palantir, ARIMO, Alation, sapho, Outlier, Digital Reasoning

MARKET
mattermark, Quid, DataFox, PREMISE, Bottlenose, MOTIVA, enigma, CB INSIGHTS, Tracxn, predata

ENTERPRISE FUNCTIONS

CUSTOMER SUPPORT
DigitalGenius, Kasisto, ELOQUENT, Wise.io, ACTIONIQ, zendesk, Preact, CLARABRIDGE

SALES
collective[i], 6sense, fuse|machines, AVISO, salesforce, INSIDE SALES.COM, clari, Zensight

MARKETING
MINTIGO, Lattice, RADIUS, Liftigniter, [PERSADO], brightfunnel, retention SCIENCE, COGNICOR, AIRPR, msg.ai

SECURITY
CYLANCE, DARKTRACE, ZIMPERIUM, deepinstinct, Sentinel, DEMISTO, graphistry, drawbridge, SignalSense, AppZen

RECRUITING
textio, entelo, Wade & Wendy, hi, unitive, SpringRole, GIGSTER, HireVue

AUTONOMOUS SYSTEMS

GROUND NAVIGATION
drive.ai, AdasWorks, ZOOX, MOBILEYE, UBER, Google, TESLA, nuTonomy, Auro Robotics

AERIAL
SKYDIO, SHIELD AI, Airware, dji, LILY, DroneDeploy, pilot.ai, SKYCATCH

INDUSTRIAL
JAYBRIDGE, OSARO, CLEARPATH ROBOTICS, fetch robotics, KINDRED, HARVEST AUTOMATION, rethink robotics

AGENTS

PERSONAL
amazon alexa, Cortana, Allo, facebook, Siri, Replika

PROFESSIONAL
butter.ai, pogo, SKIPFLAG, clara, x.ai, slack, talla, Zoom.ai, sudo

INDUSTRIES

AGRICULTURE
BLUE RIVER, mavrx, tule, TRACE GENOMICS, Pivot Bio, TerrAvion, AGRI-DATA, Descartes Labs, udio, abundant robotics

EDUCATION
KNEWTON, volley, gradescope, CTI, coursera, UDACITY, altschool

INVESTMENT
Bloomberg, sentient, SENTIUM, KENSHO, alphasense, Dataminr, CEREBELLUM CAPITAL, Quandl

LEGAL
blue J, BEAGLE, Everlaw, RAVEL, seal, ROSS, LEGAL ROBOT

LOGISTICS
NAUTO, Acerta, PRETECKT, clearmetal, Routific, MARBLE, PITSTOP

INDUSTRIES CONT'D

MATERIALS
zymergen, Citrine, Eigen Innovations, SIGHT MACHINE, GINKGO BIOWORKS, nanotronics, CALCULARIO

RETAIL FINANCE
TALA, zestfinance, Lendo, earnest, affirm, MIRADOR, wealthfront, Betterment

HEALTHCARE

PATIENT
PULSE, CareSkore, ZEPHYR HEALTH, IBM Watson Health, Oncora, SENTRIAN, Atomwise, Numerate

IMAGE
BUTTERFLY, 3SCAN, ARTERYS, enlitic, BAYLABS, imagia, Google DeepMind

BIOLOGICAL
iCarbonX, color, GRAIL, deep genomics, RECURSION, LUMINIST, Numerate, Atomwise, verily, WHOLE BIOME

TECHNOLOGY STACK

AGENT ENABLERS
OCTANE.AI, howdy, MaluubA, KITT.AI, OpenAI Gym, Kasisto, AUTOMAT, semanticmachines

DATA SCIENCE
DOMINO, SPARKBEYOND, rapidminer, kaggle, DataRobot, ŷhat, AYASDI, data iku, seldon, yseop, bigml

MACHINE LEARNING
CognitiveScale, GoogleML, context relevant, Cycorp, HyperScience, narrative logics, minds.ai, H2O.ai, SCALED INFERENCE, sparkcognition, loop, GEOMETRIC INTELLIGENCE, deepsense.io, reactive, skymind, bonsai

NATURAL LANGUAGE
agolo, AYLIEN, LEXALYTICS, Narrative Science, loop.ai Labs, spaCy, LUMINOSO, cortical.io, MonkeyLearn

DEVELOPMENT
SIGOPT, HyperOpt, fuzzy.io, okite, rainforest, lobe, Anodot, Signifai, LAYER 6 AI, bonsai

DATA CAPTURE
CrowdFlower, diffbot, CrowdAI, import.io, Paxata, DATASIFT, amazon mechanical turk, enigma, WorkFusion, DATALOGUE, TRIFACTA, parsehub

OPEN SOURCE LIBRARIES
Keras, Chainer, CNTK, TensorFlow, Caffe, H2O, DEEPLEARNING4J, theano, torch, DSSTNE, Scikit-learn, AzureML, neon, MXNet, DMTK, Spark, PaddlePaddle, WEKA

HARDWARE
KNUPATH, TENSTORRENT, Cirrascale, NVIDIA, intel nervana, Movidius, tensilica, GoogleTPU, 10^{26} Labs, Qualcomm, Cerebras, Isosemi

RESEARCH
OpenAI, nnaisense, ELEMENT AI, vicarious, KNOGGIN, Numenta, Kimera Systems, Cogitai

shivonzilis.com/MACHINEINTELLIGENCE · BloombergBETA

Two definitions of learning

(1) Learning is the acquisition of knowledge about the world.

Kupfermann (1985)

(2) Learning is an adaptive change in behavior caused by experience.

Shepherd (1988)

Empirical Inference

- Drawing conclusions from empirical data (observations, measurements)

Empirical Inference

- Drawing conclusions from empirical data (observations, measurements)
- Example 1: scientific inference

slide by Bernhard Schölkopf

Empirical Inference

- Drawing conclusions from empirical data (observations, measurements)
- Example 1: scientific inference

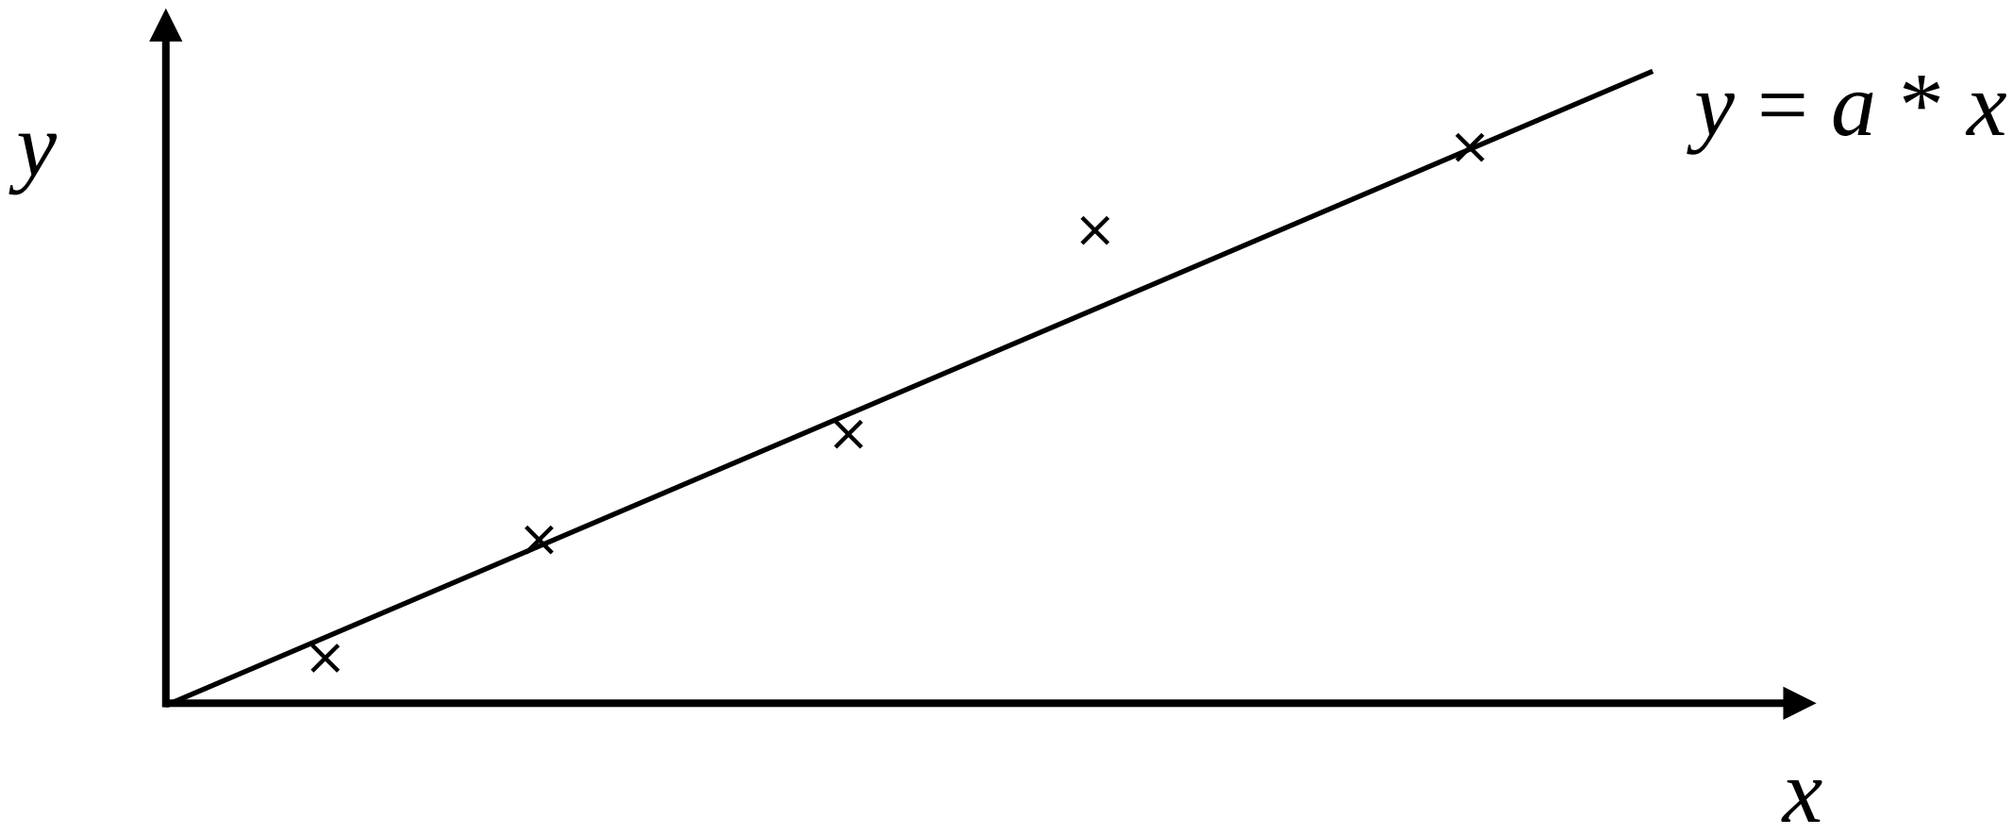

slide by Bernhard Schölkopf

Empirical Inference

- Drawing conclusions from empirical data (observations, measurements)
- Example 1: scientific inference

Leibniz, Weyl, Chaitin

slide by Bernhard Schölkopf

Empirical Inference

- Drawing conclusions from empirical data (observations, measurements)
- Example 1: scientific inference

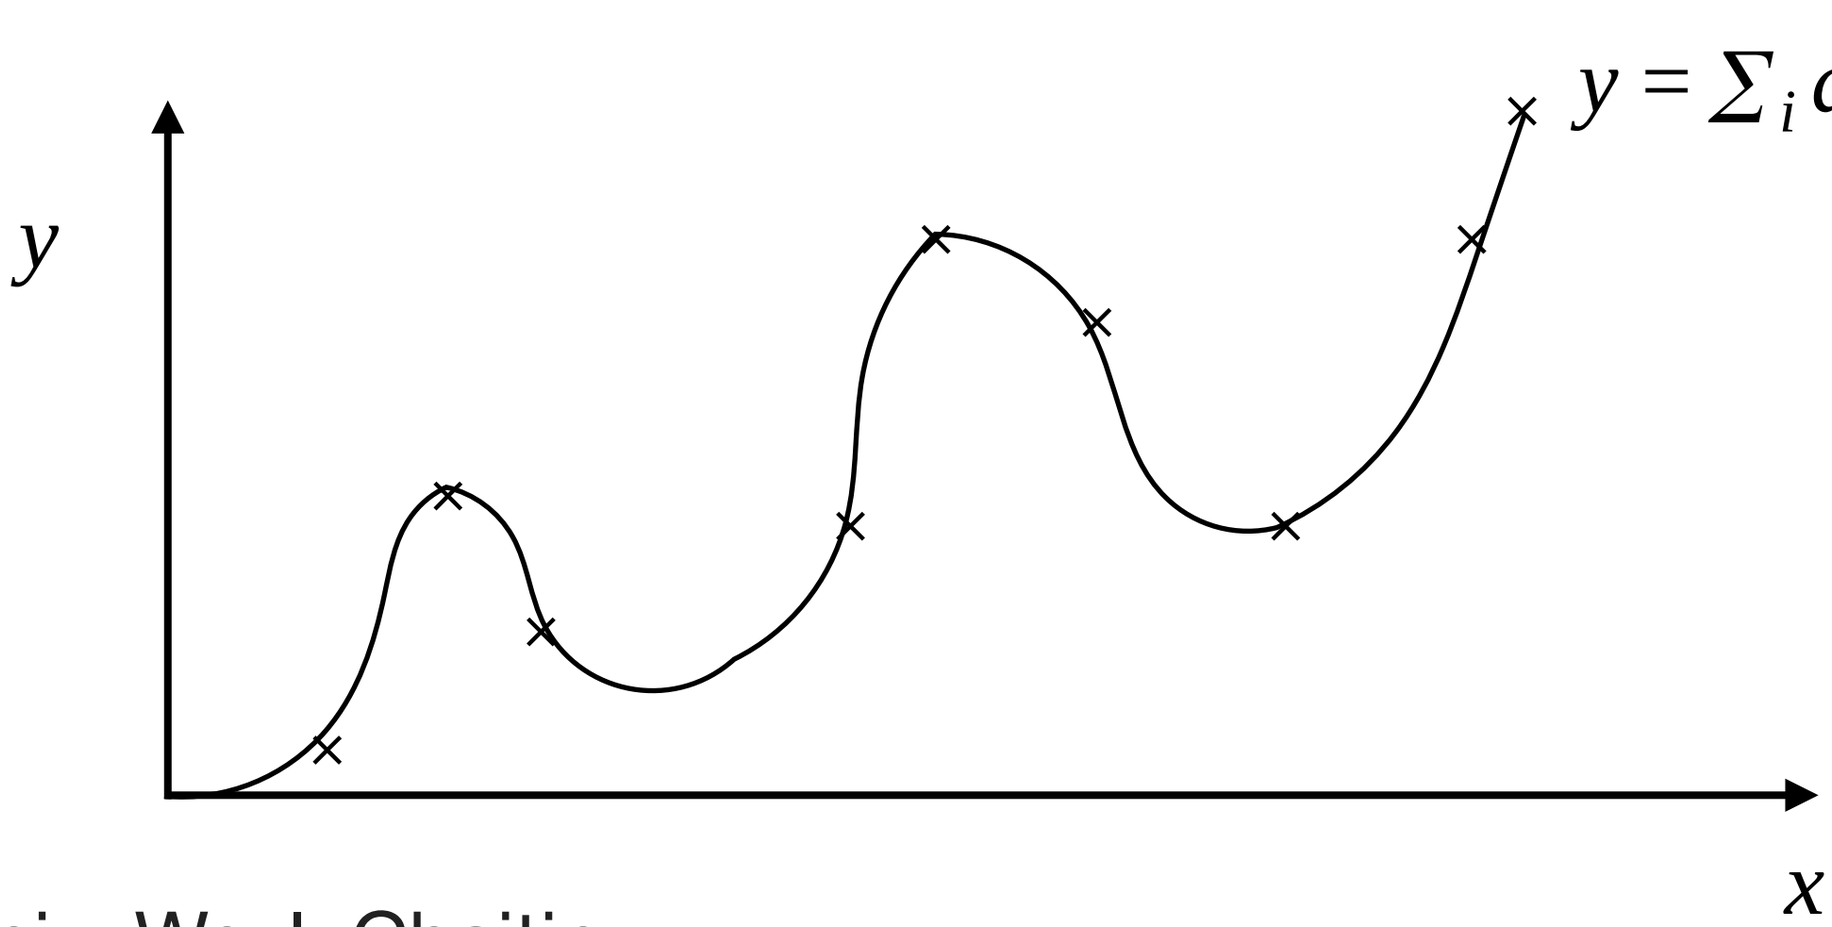

Leibniz, Weyl, Chaitin

slide by Bernhard Schölkopf

Empirical Inference

- Example 2: perception

9

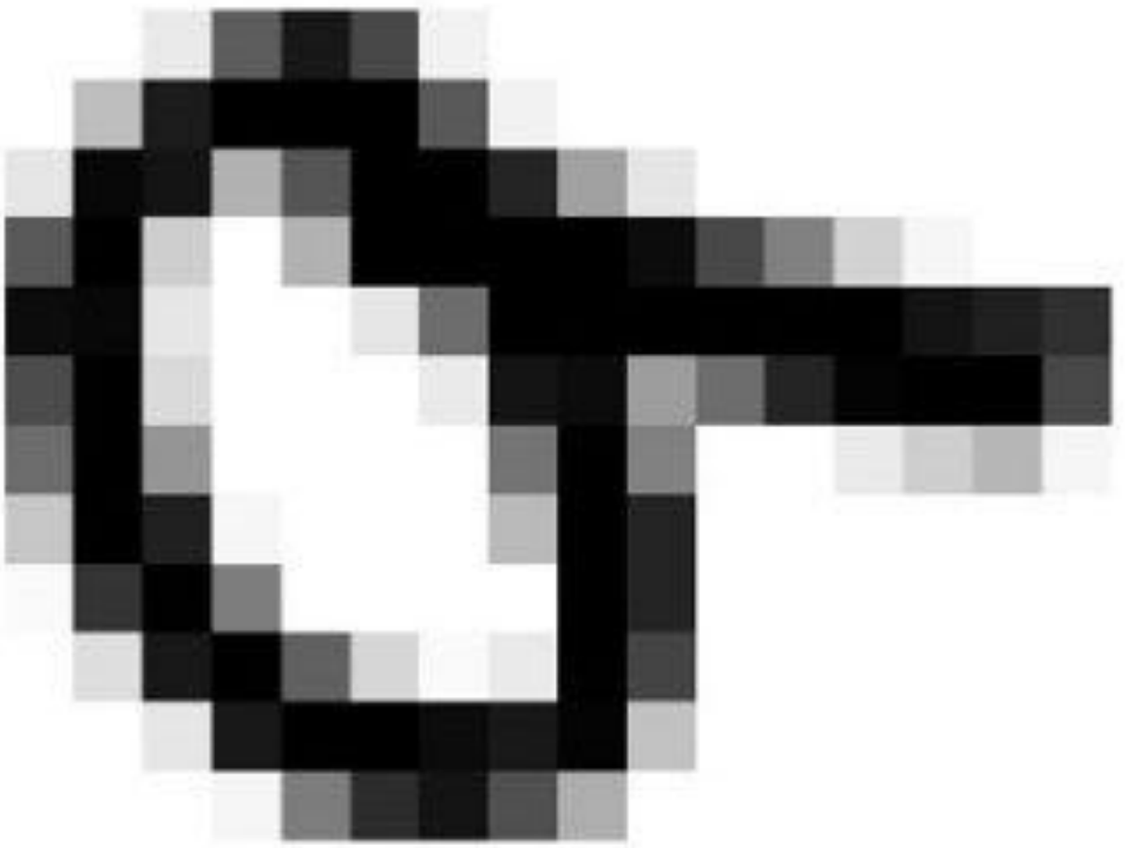

9

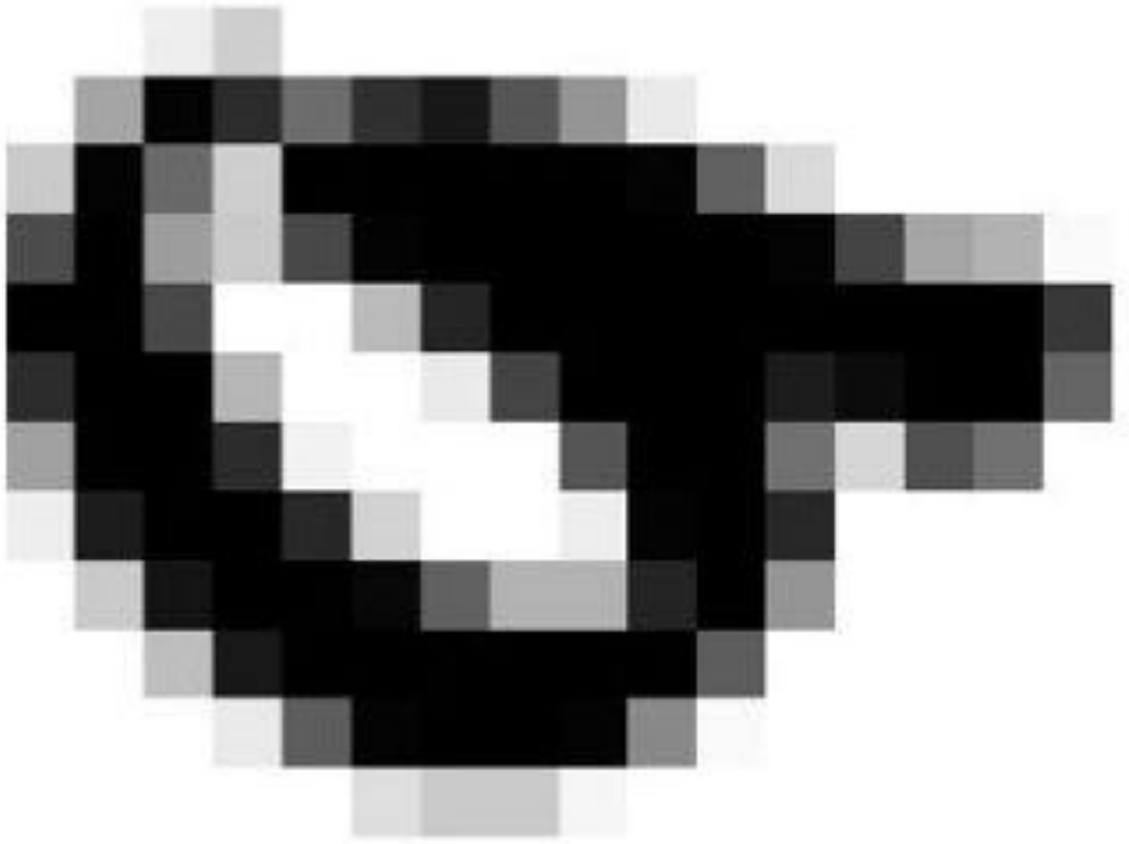

∞

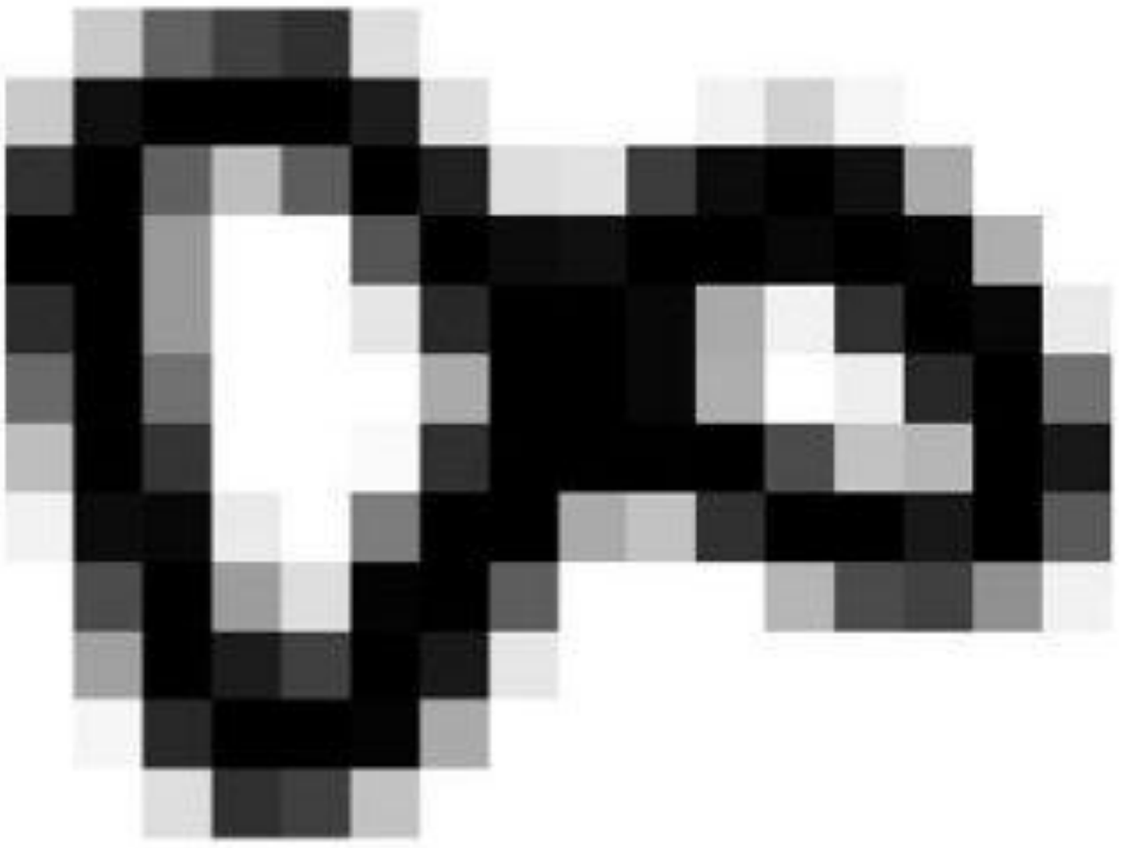

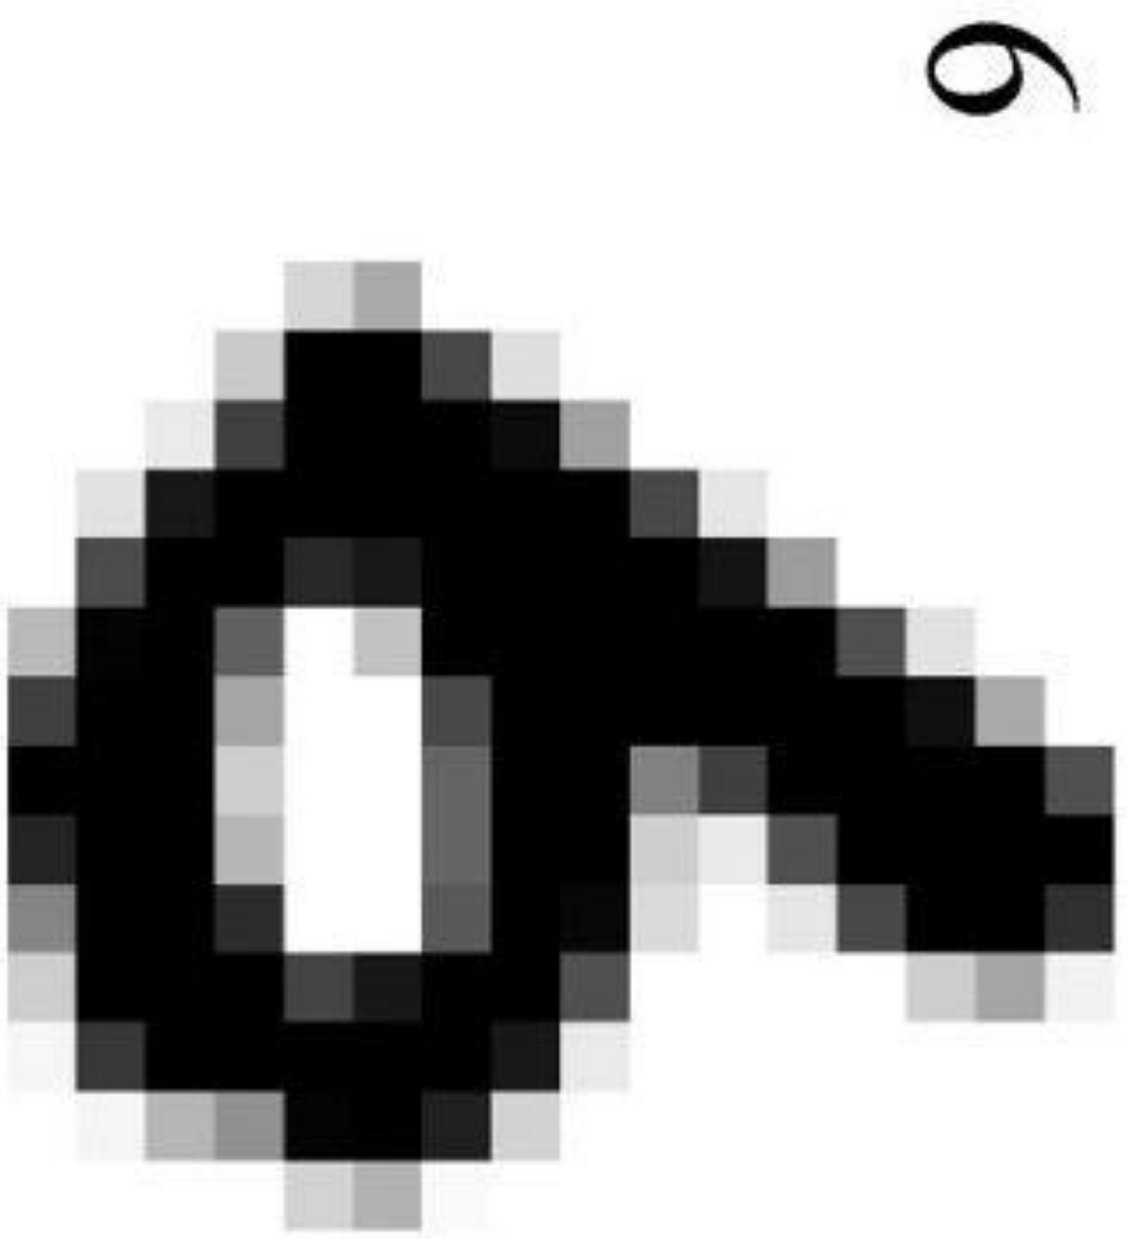

slide by Bernhard Schölkopf

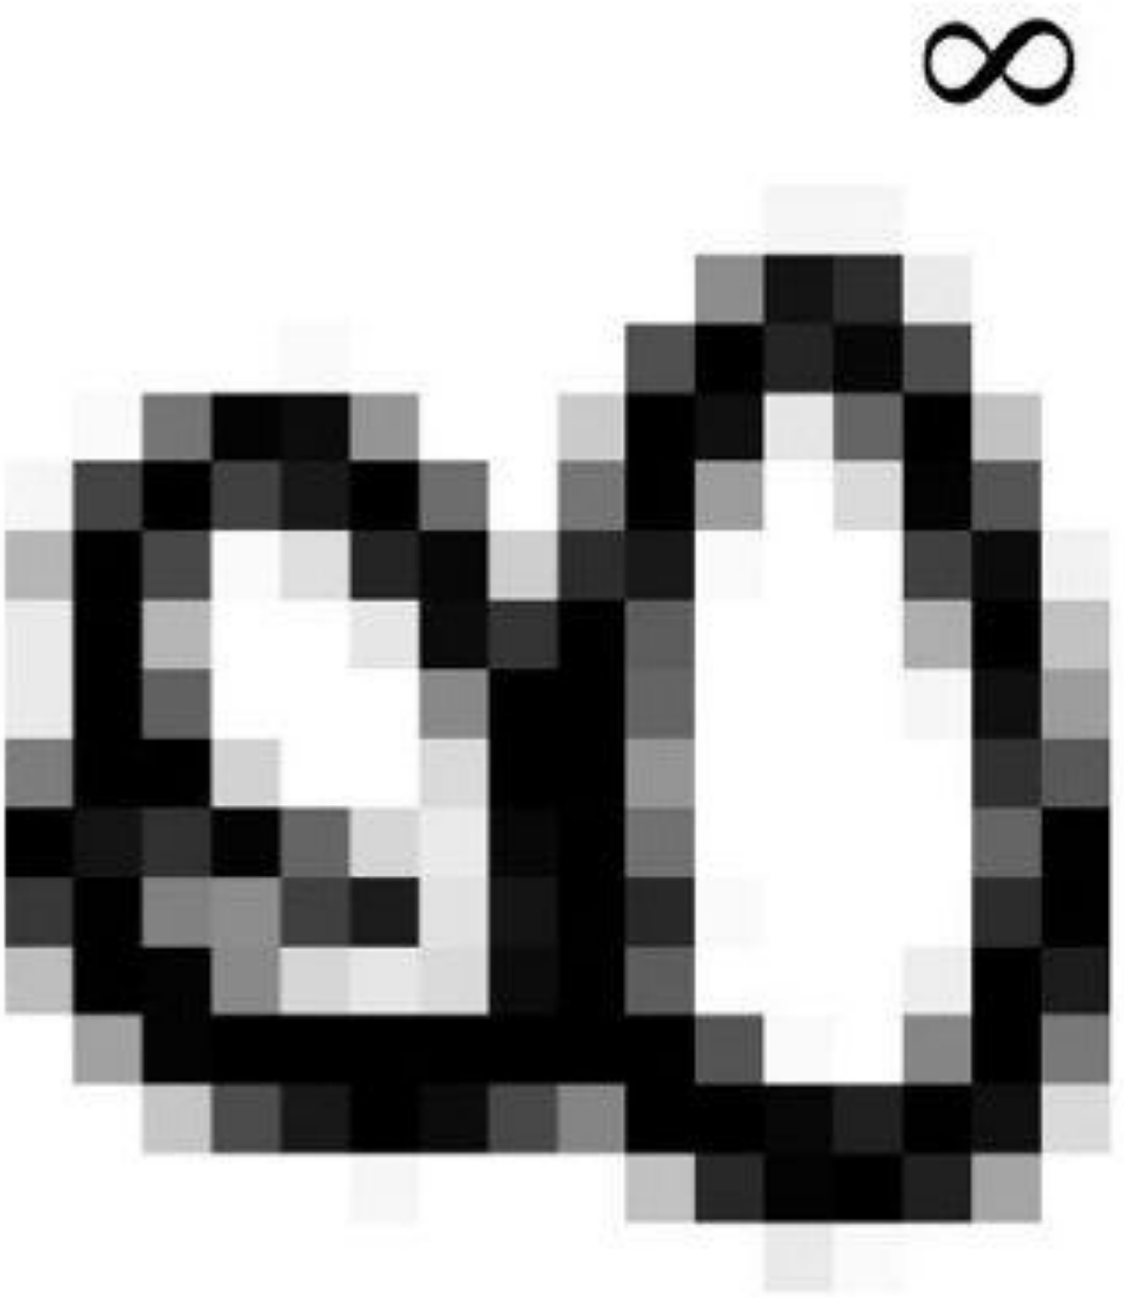

slide by Bernhard Schölkopf

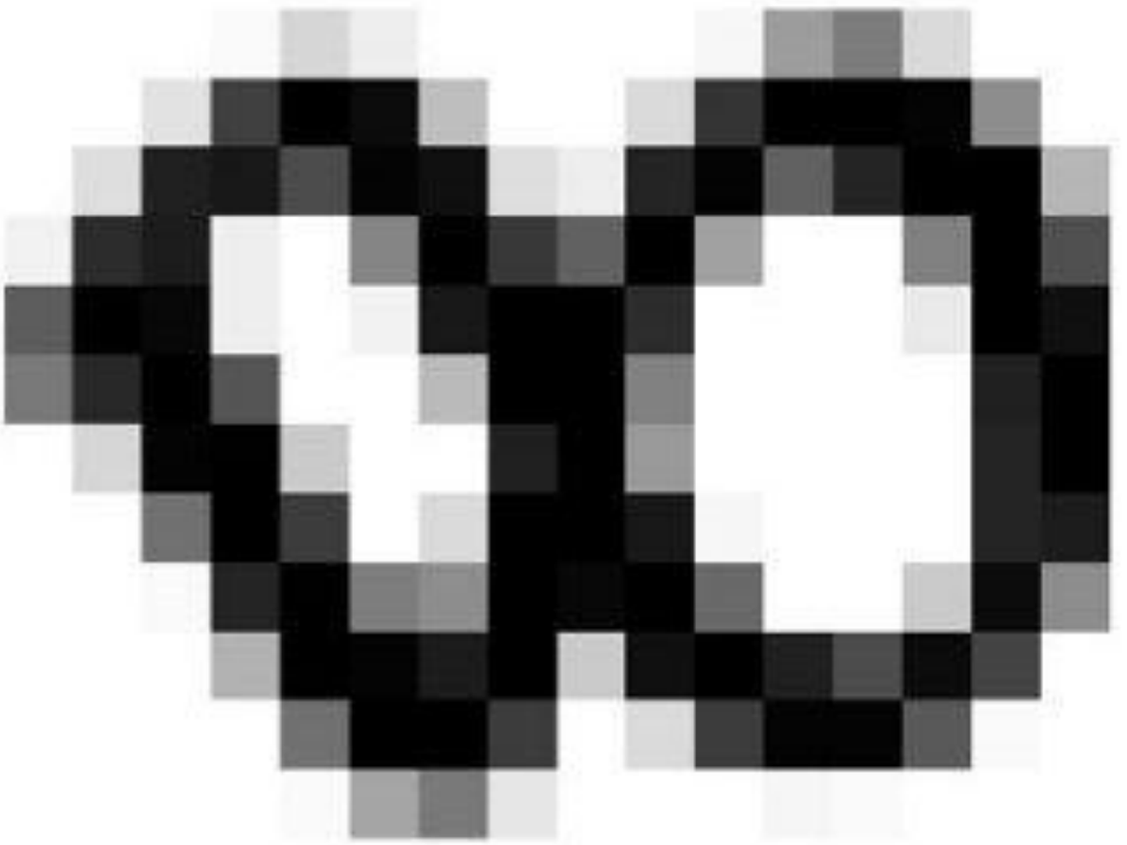

slide by Bernhard Schölkopf

9

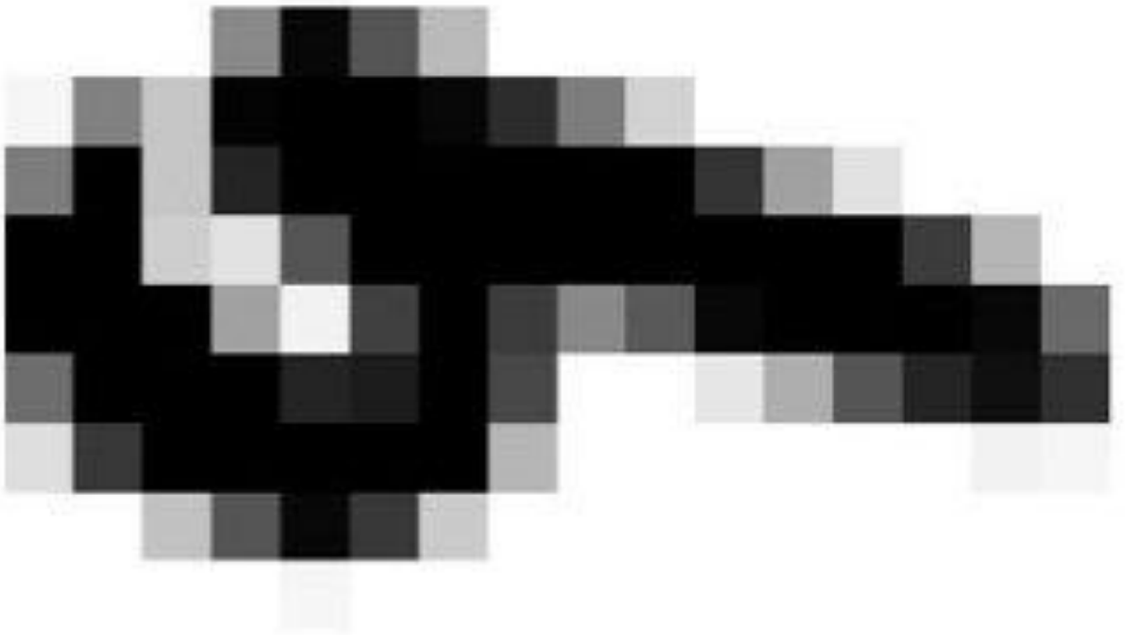

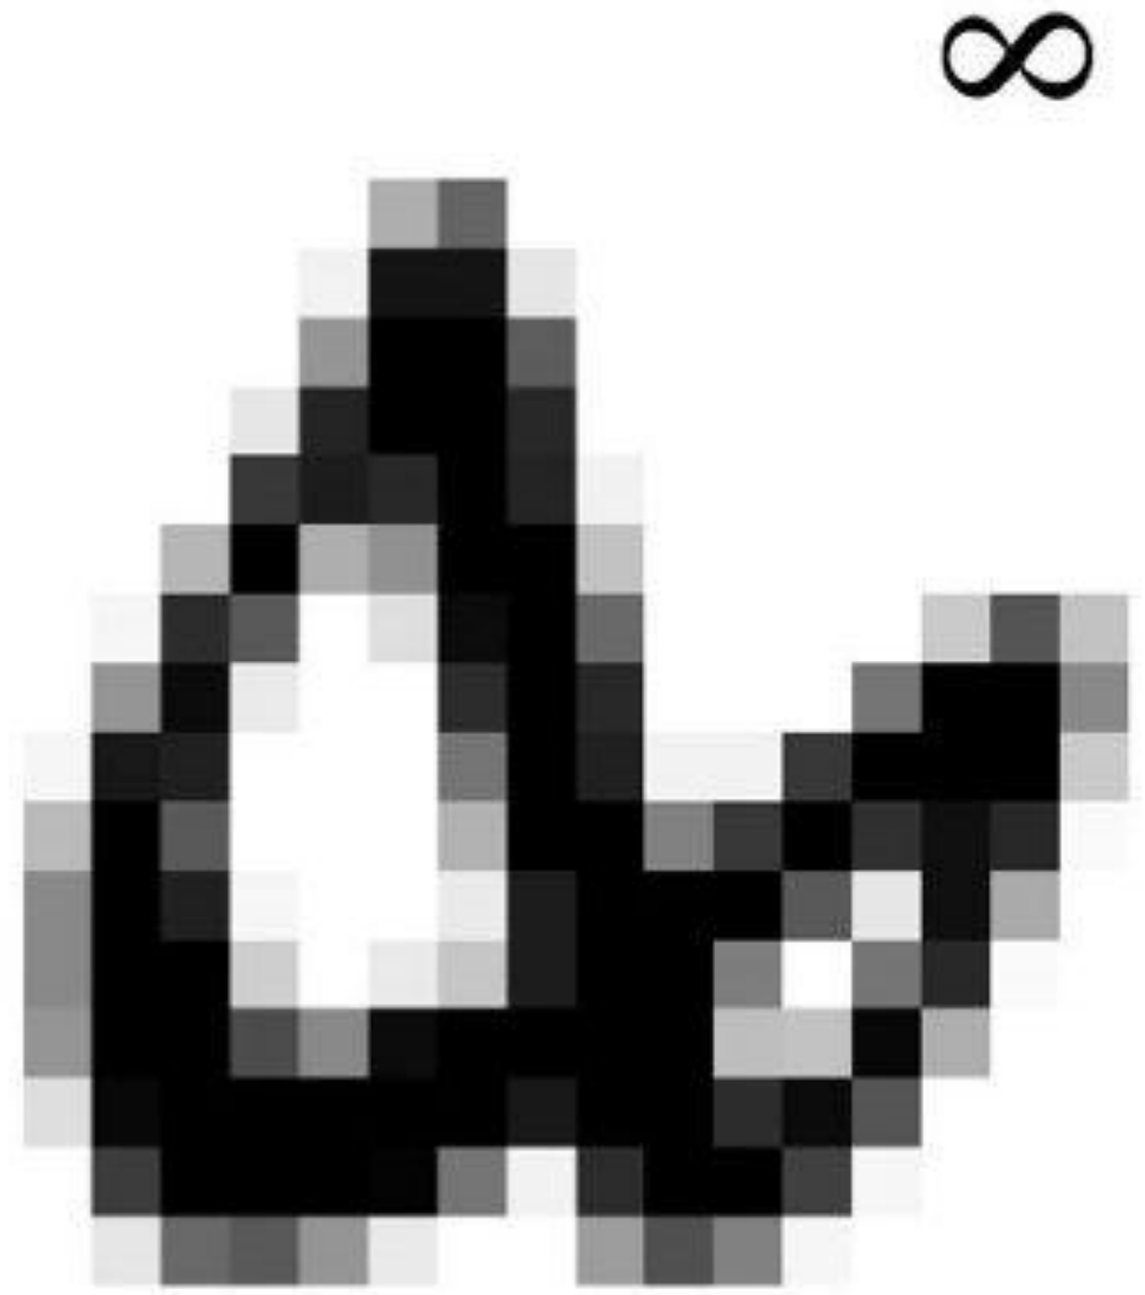

slide by Bernhard Schölkopf

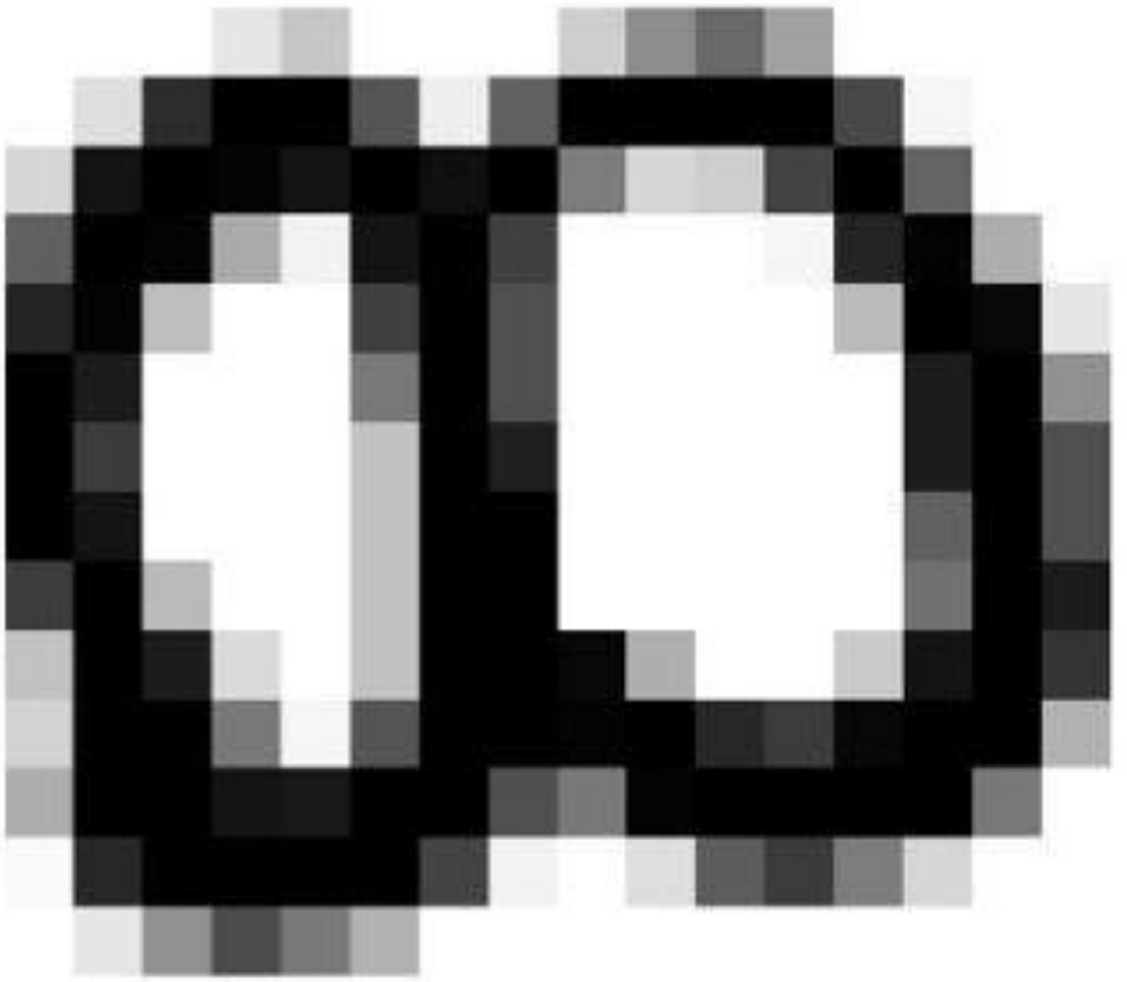

slide by Bernhard Schölkopf

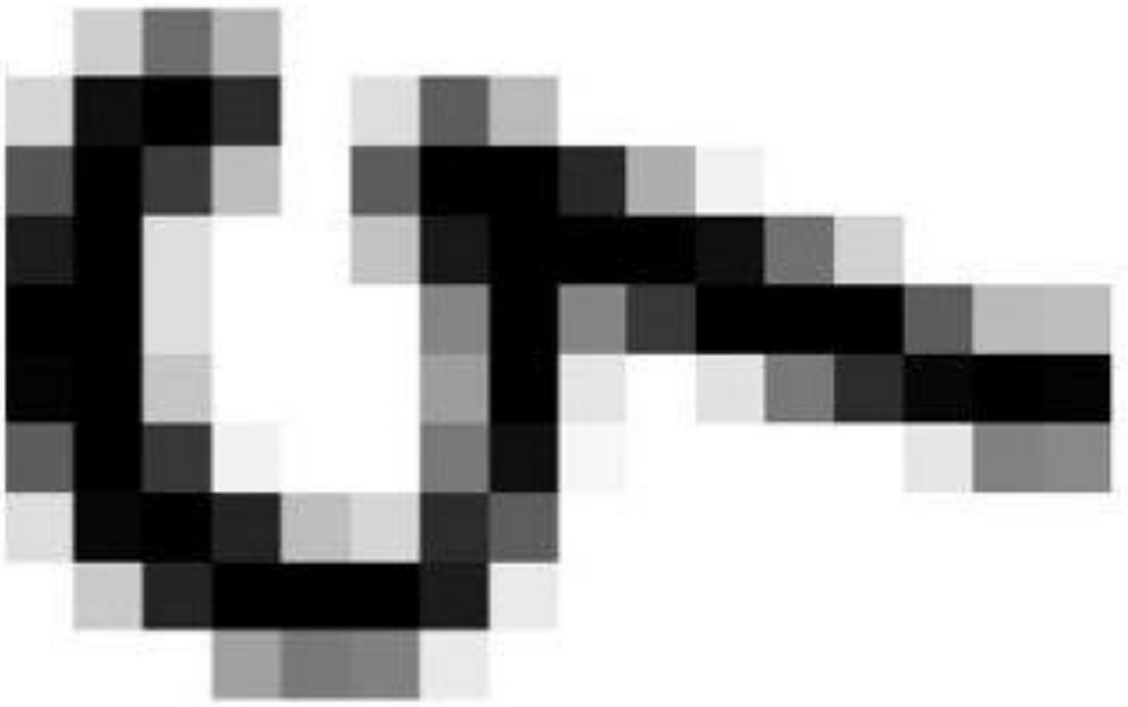

slide by Bernhard Schölkopf

9

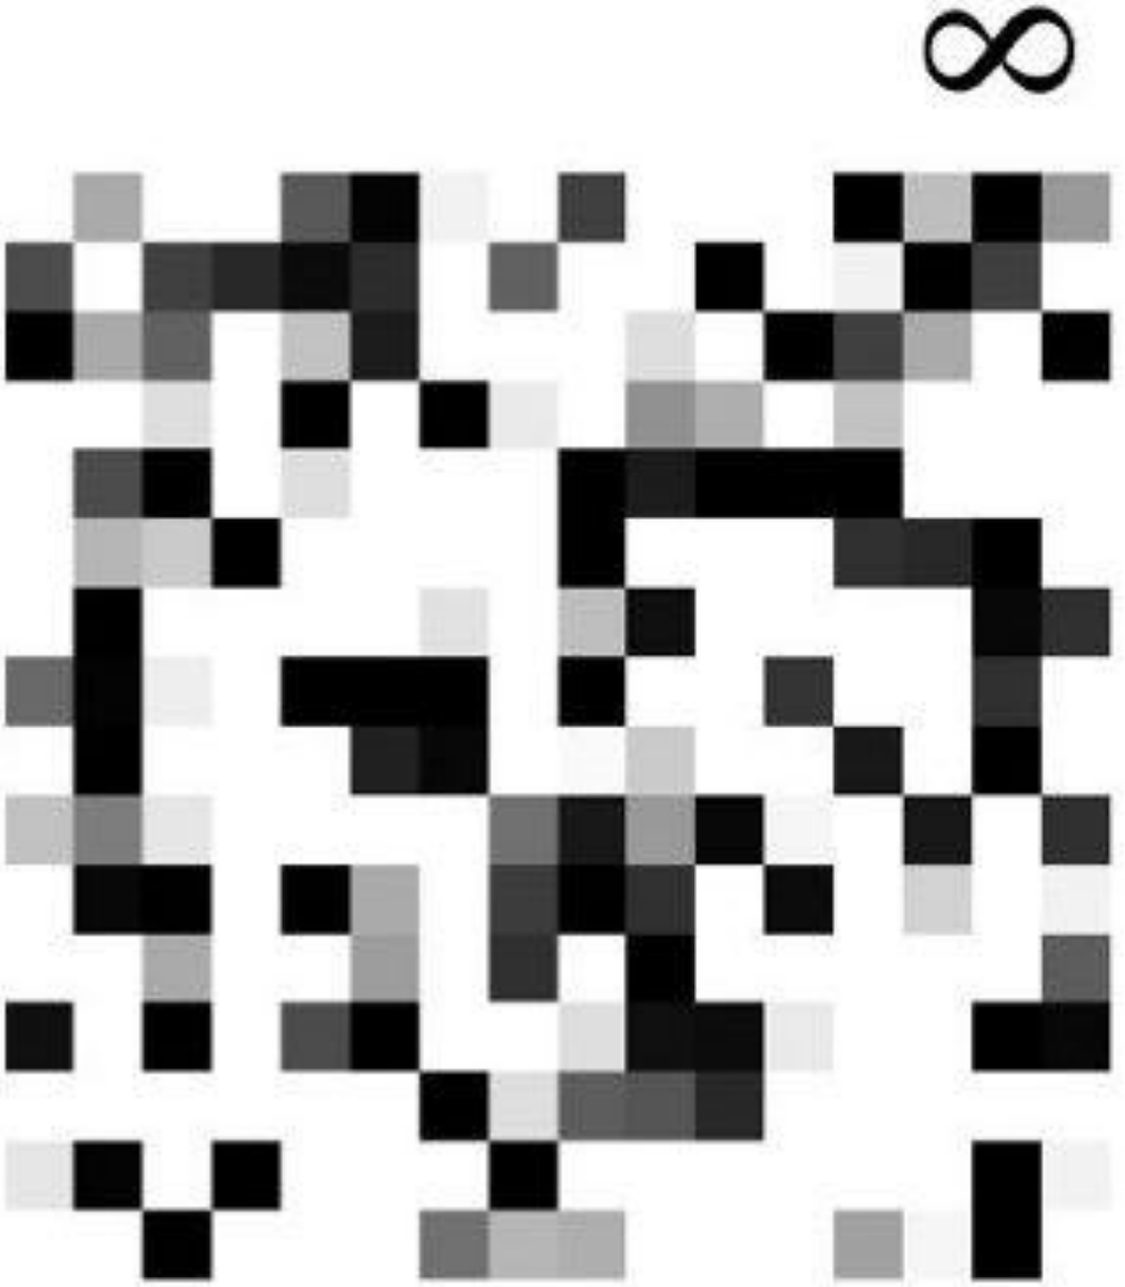

slide by Bernhard Schölkopf

9

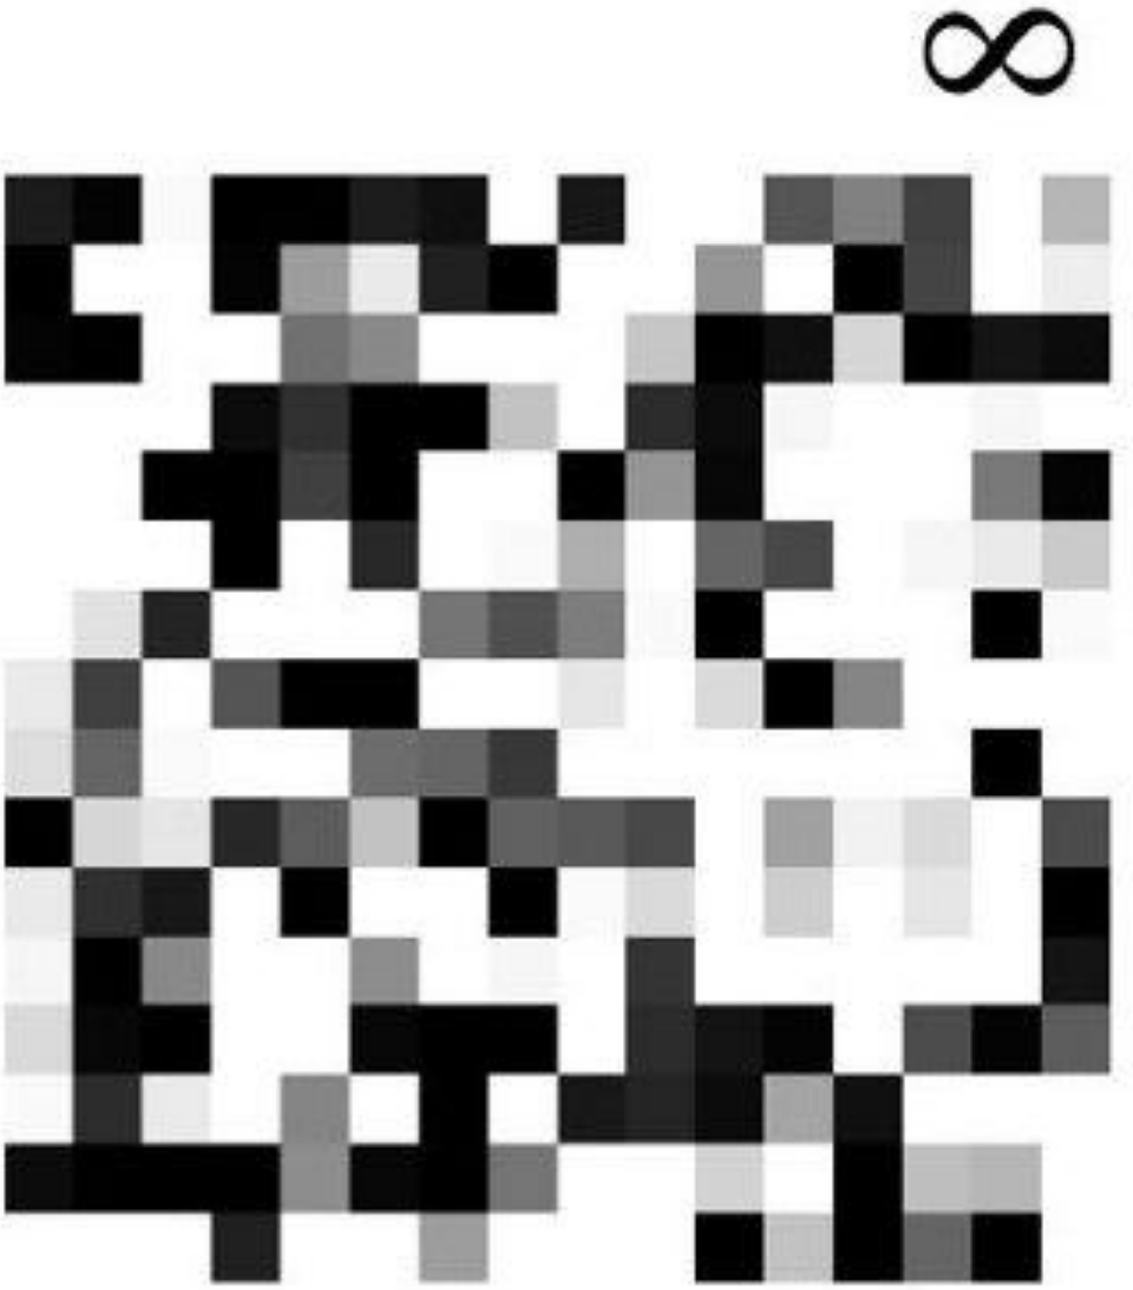

slide by Bernhard Schölkopf

slide by Bernhard Schölkopf

slide by Bernhard Schölkopf

slide by Bernhard Schölkopf

slide by Bernhard Schölkopf

9

?

slide by Bernhard Schölkopf

Empirical Inference

- Example2: perception

"The brain is nothing but a statistical decision organ"

H. Barlow

slide by Bernhard Schölkopf

Color Perception

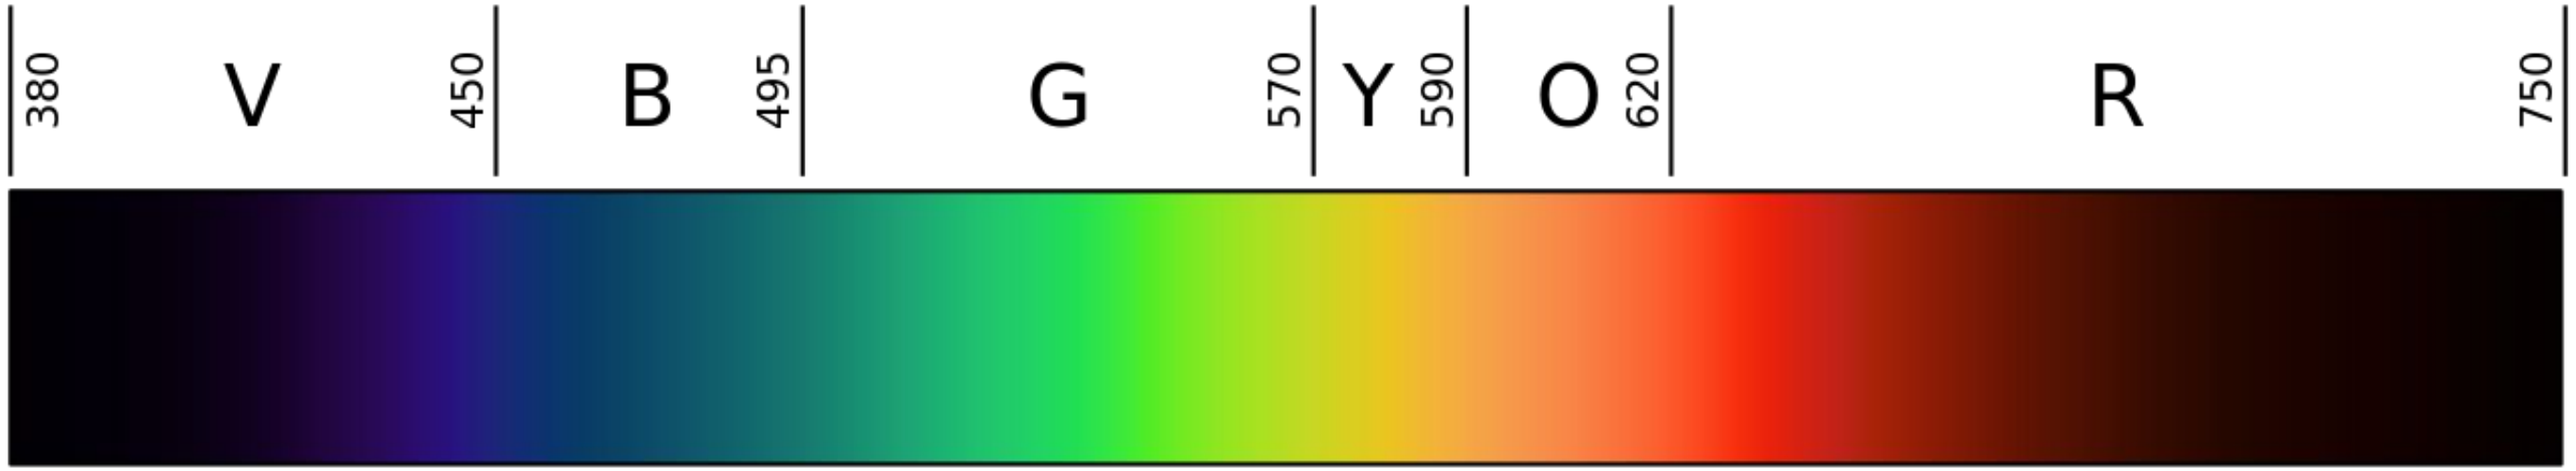

X

slide by Bernhard Schölkopf

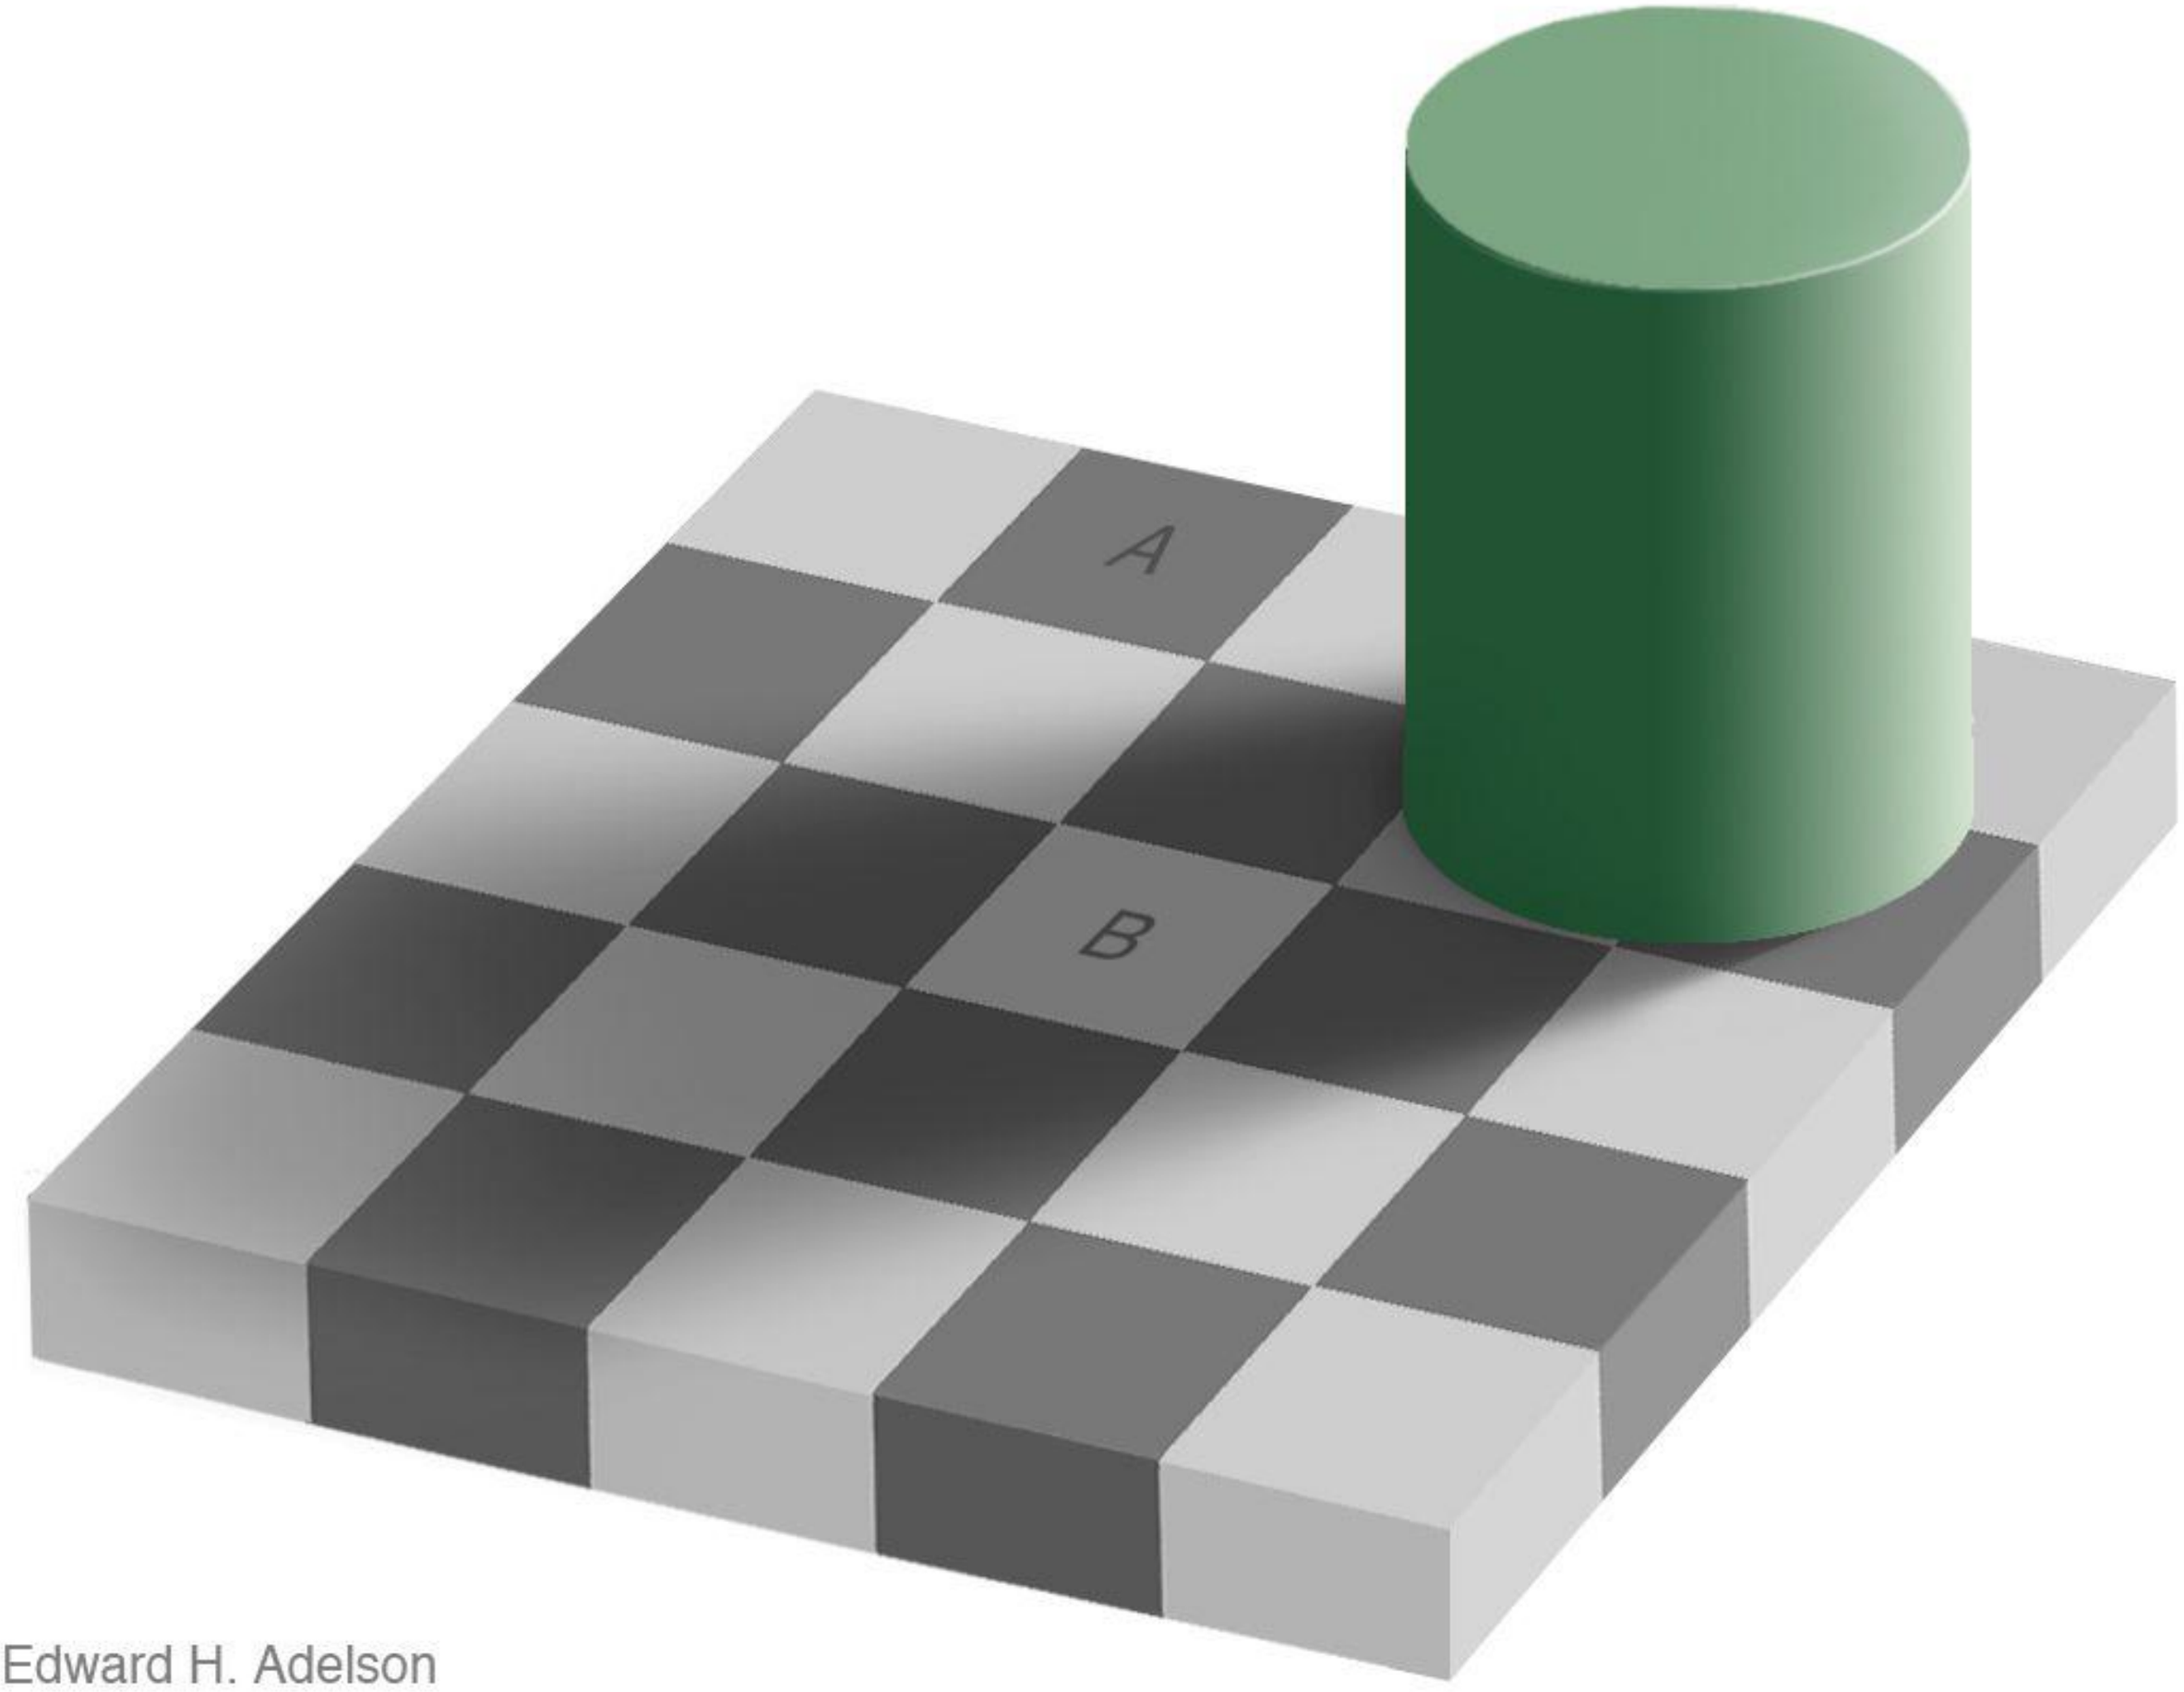

Edward H. Adelson

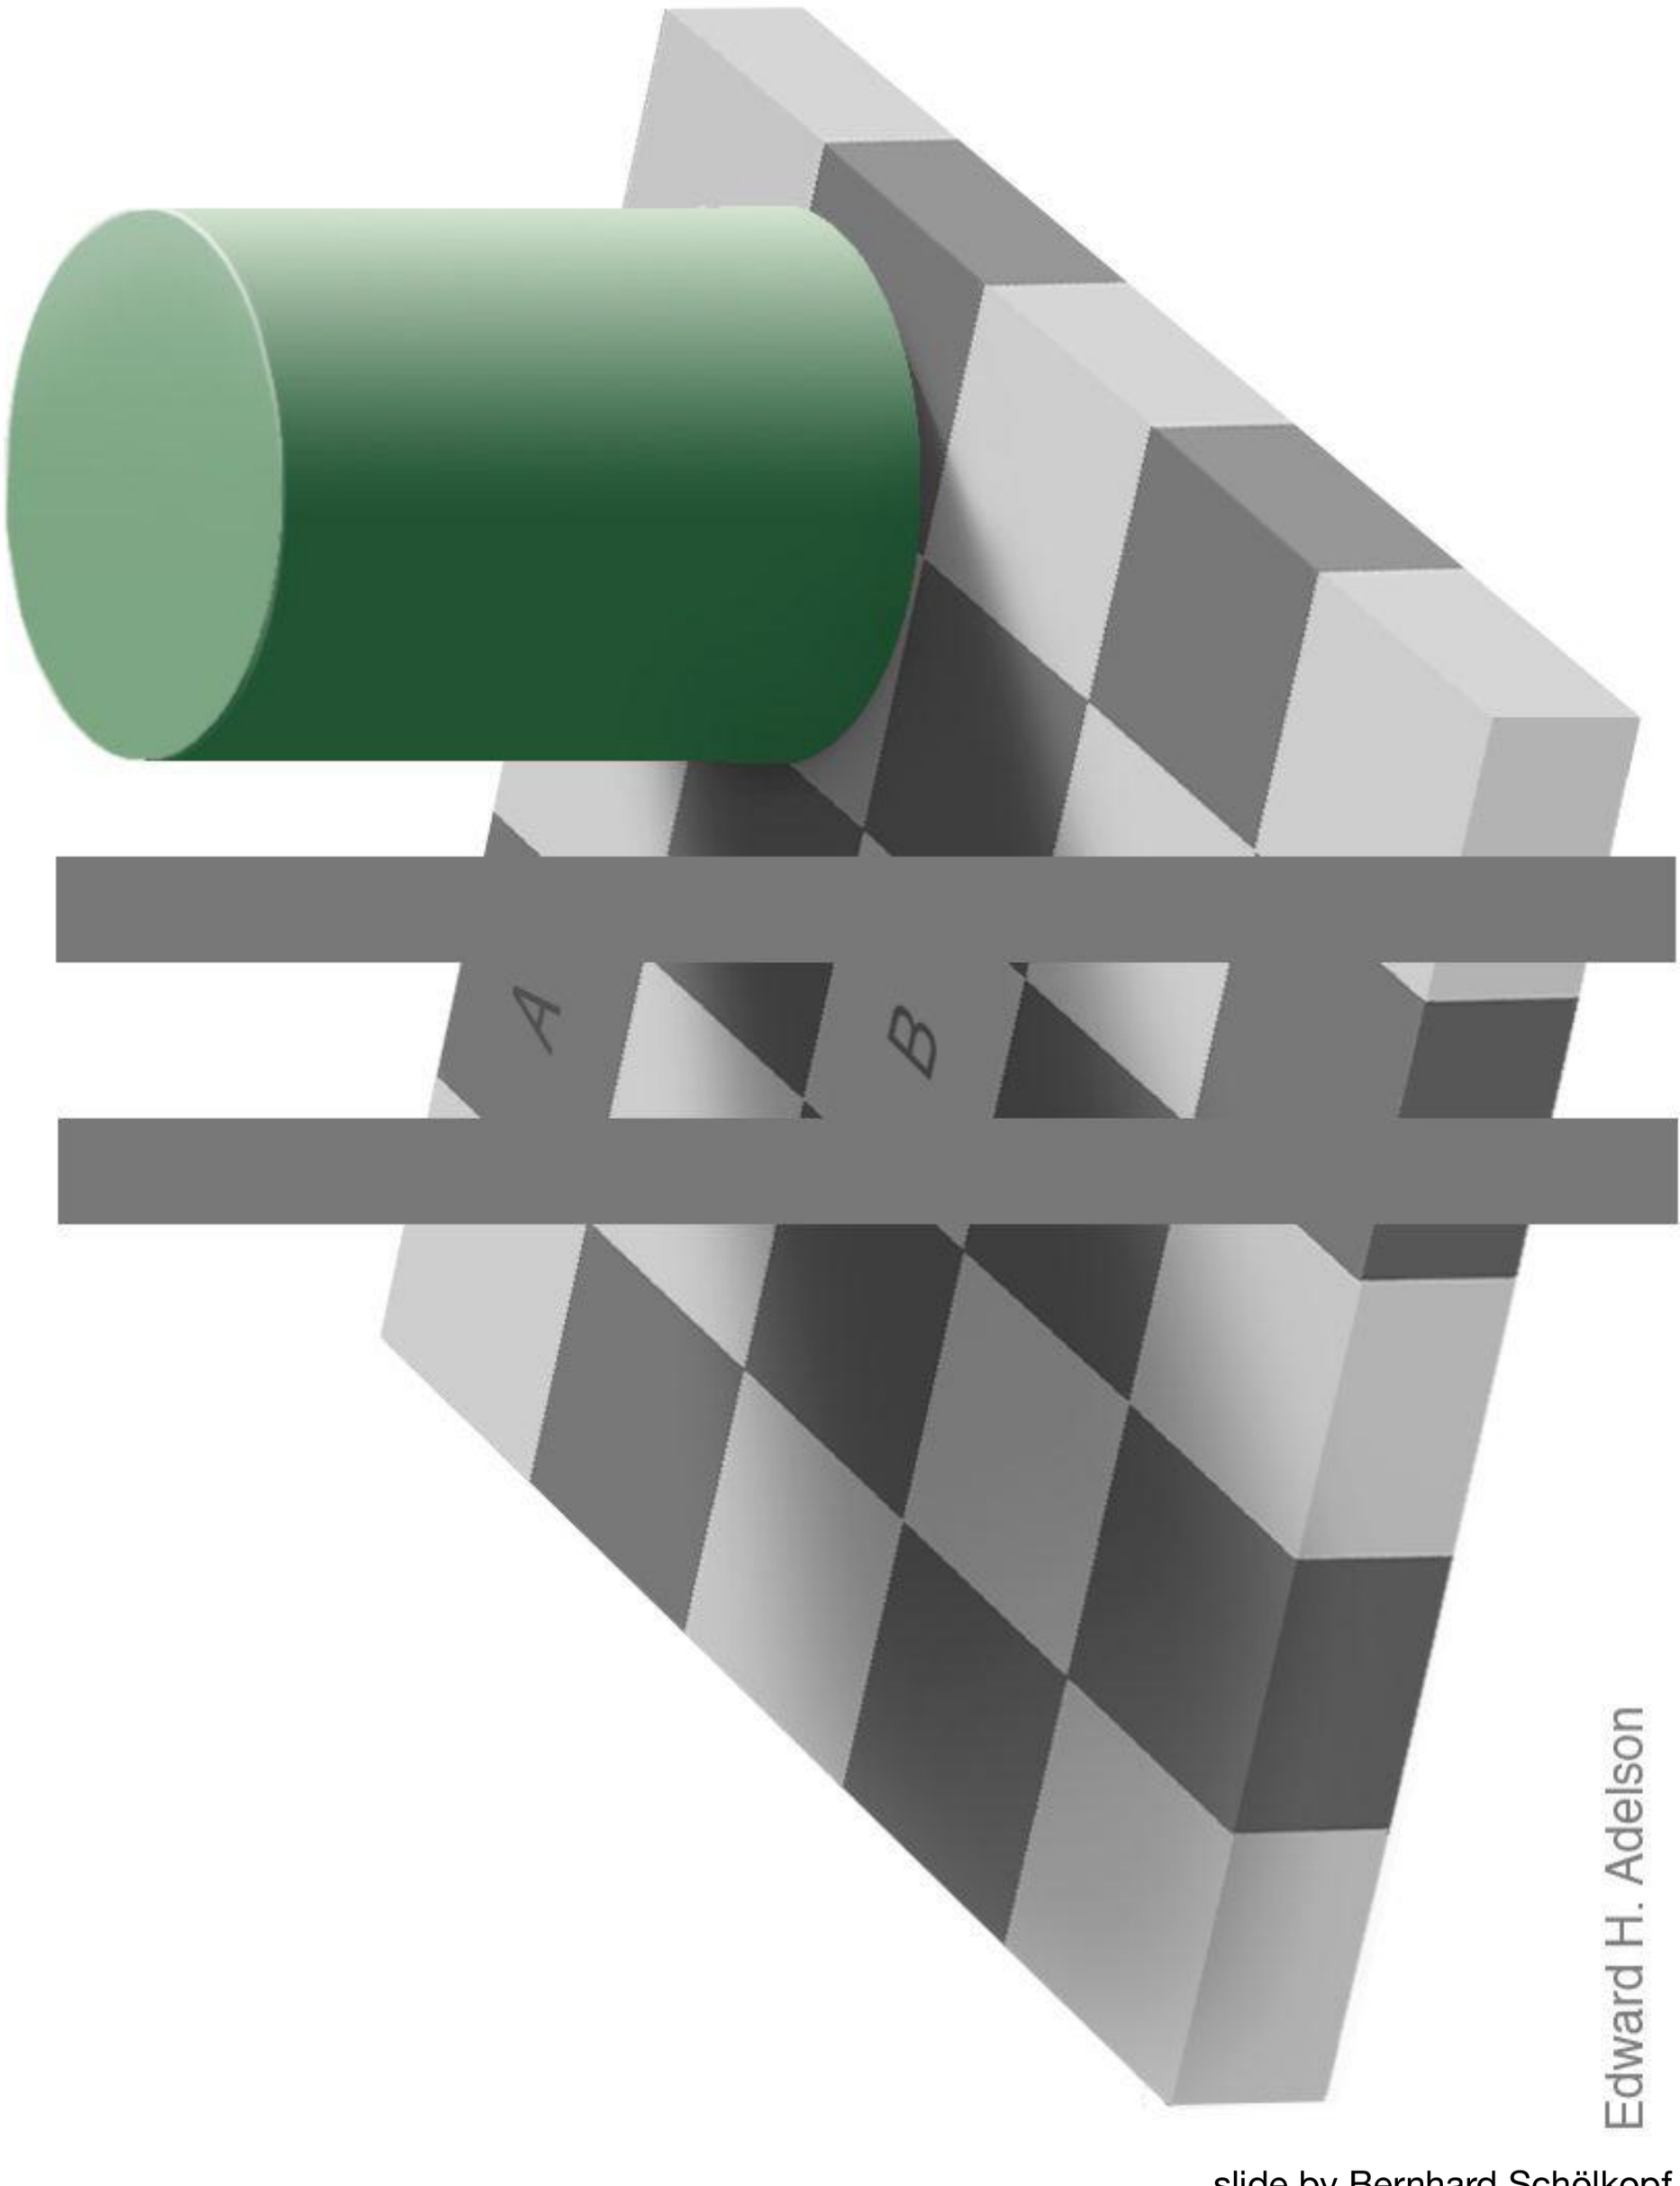

slide by Bernhard Schölkopf

*reflected light = illumination * reflectance*

Hard Inference Problems

- High dimensionality — consider many factors simultaneously to find regularity
- Complex regularities — nonlinear; nonstationary, etc.
- Little prior knowledge — e.g. no mechanistic models for the data
- Need large data sets — processing requires computers and automatic inference methods

slide by Bernhard Schölkopf

What is machine learning?

Example: Netflix Challenge

- Goal: Predict how a viewer will rate a movie
- 10% improvement = 1 million dollars

slide by Yaser Abu-Mostapha

Example: Netflix Challenge

- Goal: Predict how a viewer will rate a movie
- 10% improvement = 1 million dollars
- Essence of Machine Learning:
  - A pattern exists
  - We cannot pin it down mathematically
  - We have data on it

AlphaGo vs Lee Sedol

NVIDIA BB8 AI Car

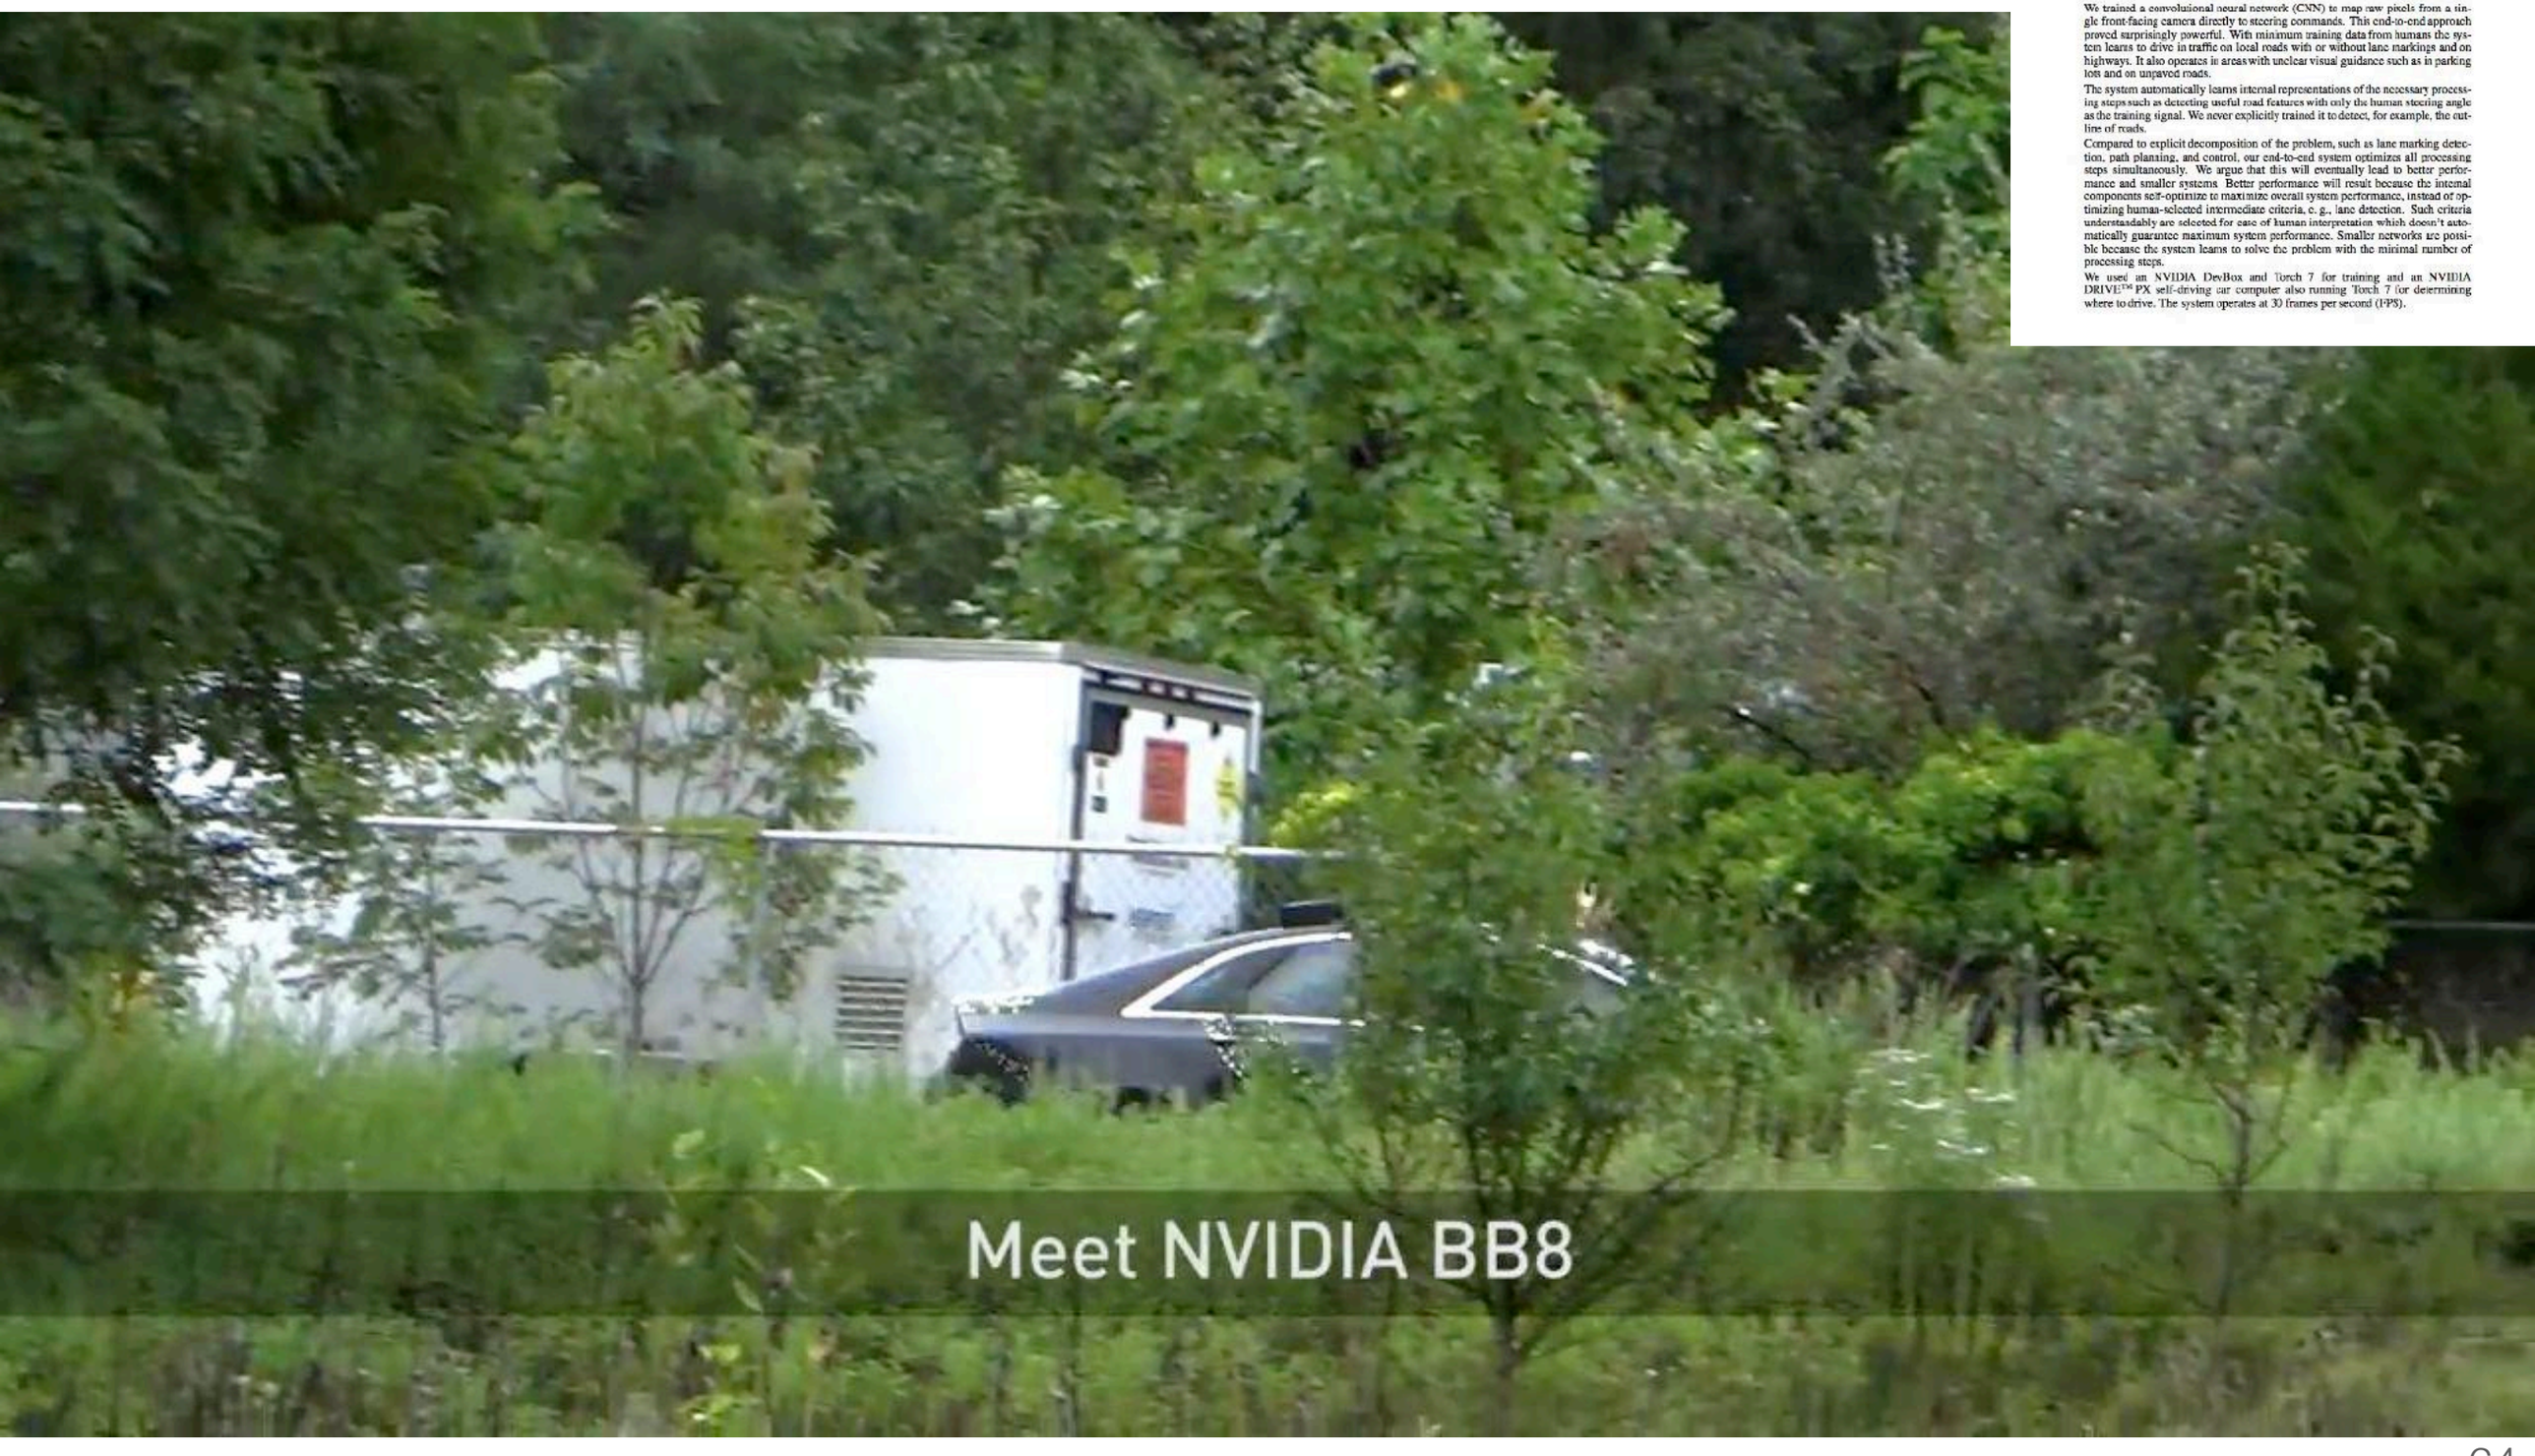

End to End Learning for Self-Driving Cars

Mariusz Bojarski NVIDIA Corporation Holmdel, NJ 07735
Davide Del Testa NVIDIA Corporation Holmdel, NJ 07735
Daniel Dworakowski NVIDIA Corporation Holmdel, NJ 07735
Bernhard Firner NVIDIA Corporation Holmdel, NJ 07735
Beat Flepp NVIDIA Corporation Holmdel, NJ 07735
Prasoon Goyal NVIDIA Corporation Holmdel, NJ 07735
Lawrence D. Jackel NVIDIA Corporation Holmdel, NJ 07735
Mathew Monfort NVIDIA Corporation Holmdel, NJ 07735
Urs Muller NVIDIA Corporation Holmdel, NJ 07735
Jiakai Zhang NVIDIA Corporation Holmdel, NJ 07735
Xin Zhang NVIDIA Corporation Holmdel, NJ 07735
Jake Zhao NVIDIA Corporation Holmdel, NJ 07735
Karol Ziebs NVIDIA Corporation Holmdel, NJ 07735

Abstract

We trained a convolutional neural network (CNN) to map raw pixels from a single front-facing camera directly to steering commands. This end-to-end approach proved surprisingly powerful. With minimum training data from humans the system learns to drive in traffic on local roads with or without lane markings and on highways. It also operates in areas with unclear visual guidance such as in parking lots and on unpaved roads.

The system automatically learns internal representations of the necessary processing steps such as detecting useful road features with only the human steering angle as the training signal. We never explicitly trained it to detect, for example, the outline of roads.

Compared to explicit decomposition of the problem, such as lane marking detection, path planning, and control, our end-to-end system optimizes all processing steps simultaneously. We argue that this will eventually lead to better performance and smaller systems. Better performance will result because the internal components self-optimize to maximize overall system performance, instead of optimizing human-selected intermediate criteria, e. g., lane detection. Such criteria understandably are selected for ease of human interpretation which doesn't automatically guarantee maximum system performance. Smaller networks are possible because the system learns to solve the problem with the minimal number of processing steps.

We used an NVIDIA DevBox and Torch 7 for training and an NVIDIA DRIVE™ PX self-driving car computer also running Torch 7 for determining where to drive. The system operates at 30 frames per second (FPS).

What is Machine Learning?

- [Arthur Samuel, 1959]
  - Field of study that gives computers
  - the ability to learn without being explicitly programmed

- [Kevin Murphy] algorithms that
  - automatically detect patterns in data
  - use the uncovered patterns to predict future data or other outcomes of interest

- [Tom Mitchell] algorithms that
  - improve their performance (P)
  - at some task (T)
  - with experience (E)

slide by Dhruv Batra

Comparison

- **Traditional Programming**

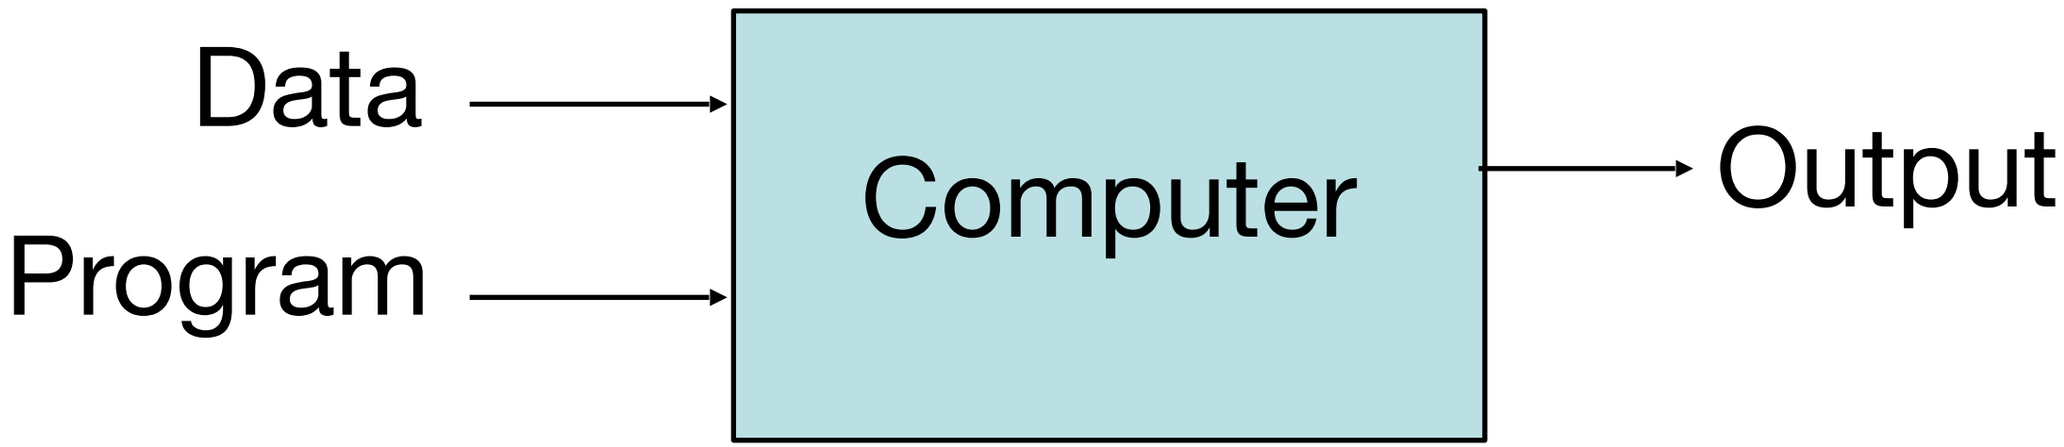

- **Machine Learning**

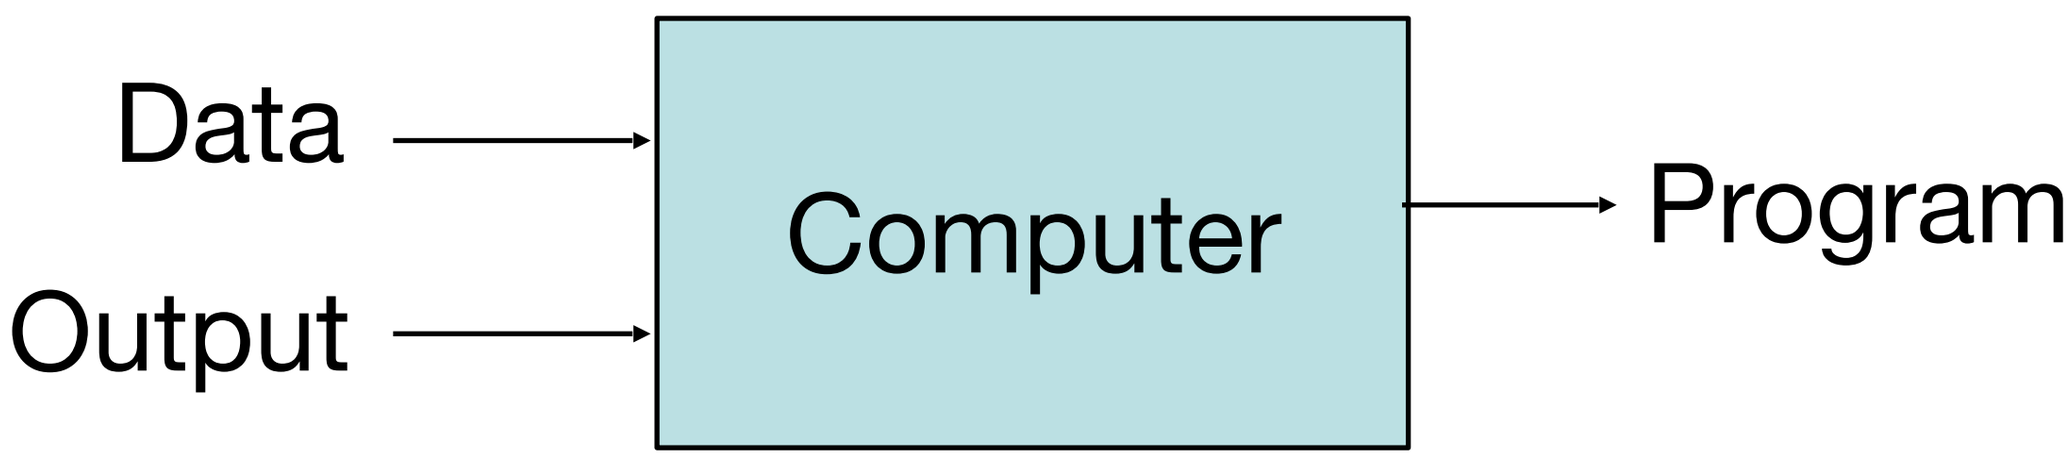

Comparison

- **Traditional Programming**

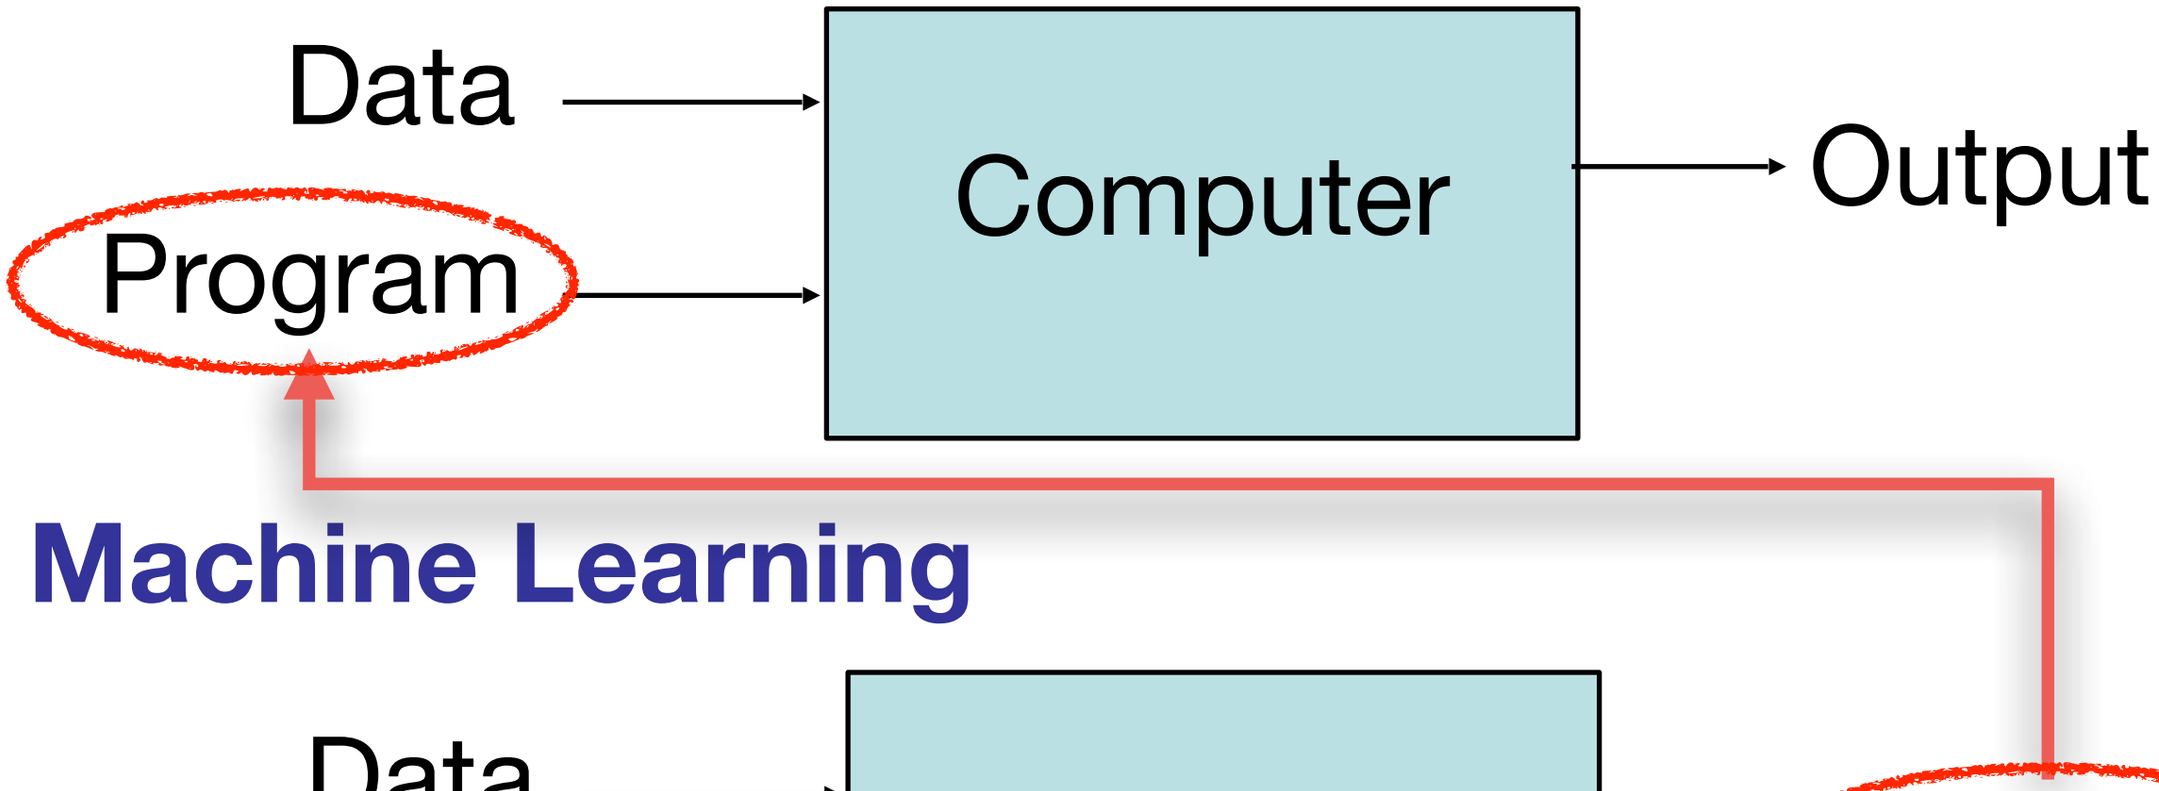

- **Machine Learning**

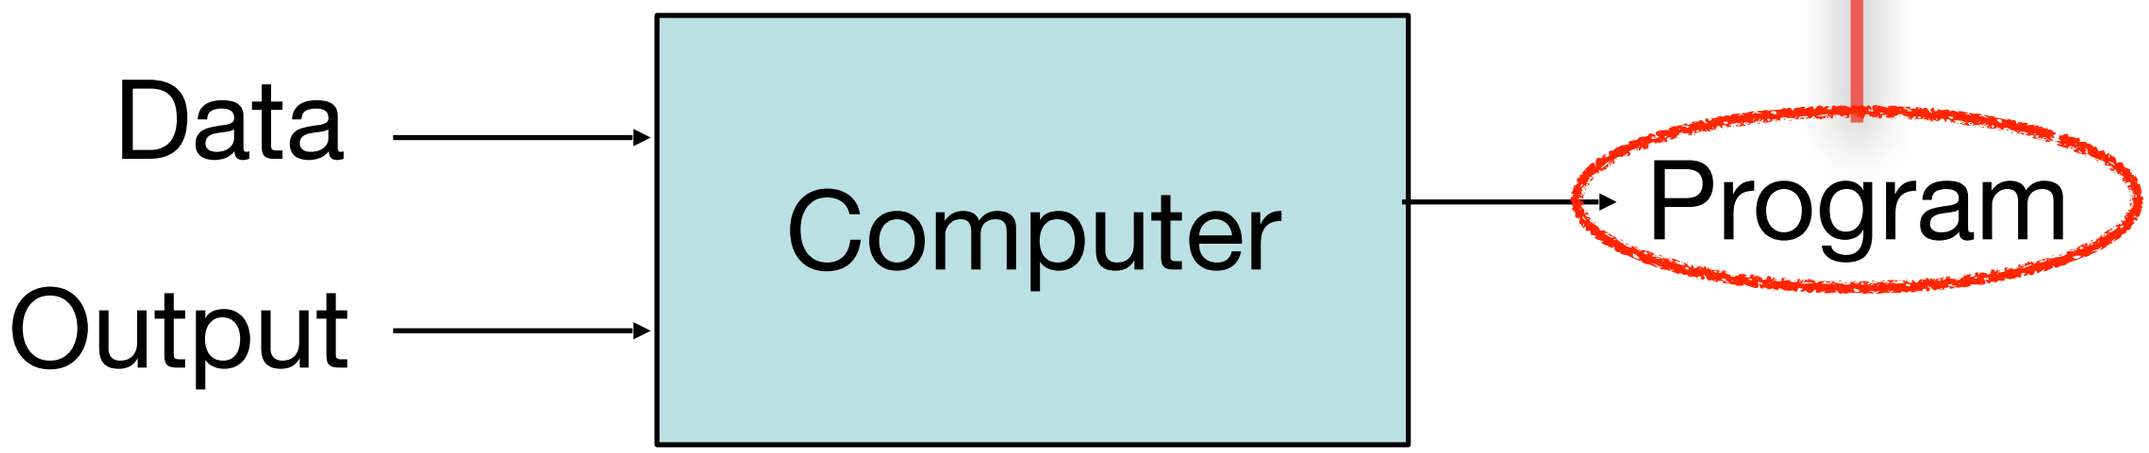

slide by Pedro Domingos, Tom Mitchel, Tom Dietterich

What is Machine Learning?

- If you are a Scientist

- If you are an Engineer / Entrepreneur
  - Get lots of data
  - Machine Learning
  - ???
  - Profit!

slide by Dhruv Batra

Why Study Machine Learning?

Engineering Better Computing Systems

- Develop systems
  - too difficult/expensive to construct manually
  - because they require specific detailed skills/knowledge
  - ***knowledge engineering bottleneck***

- Develop systems
  - that adapt and customize themselves to individual users.
  - Personalized news or mail filter
  - Personalized tutoring

- Discover new knowledge from large databases
  - Medical text mining (e.g. migraines to calcium channel blockers to magnesium)
  - ***data mining***

slide by Dhruv Batra

Why Study Machine Learning?

Cognitive Science

- Computational studies of learning may help us understand learning in humans
  - and other biological organisms.

  - Hebbian neural learning
    - “Neurons that fire together, wire together.”

slide by Dhruv Batra

Why Study Machine Learning?

The Time is Ripe

- Algorithms
  - Many basic effective and efficient algorithms available.
- Data
  - Large amounts of on-line data available.
- Computing
  - Large amounts of computational resources available.

slide by Ray Mooney

Where does ML fit in?

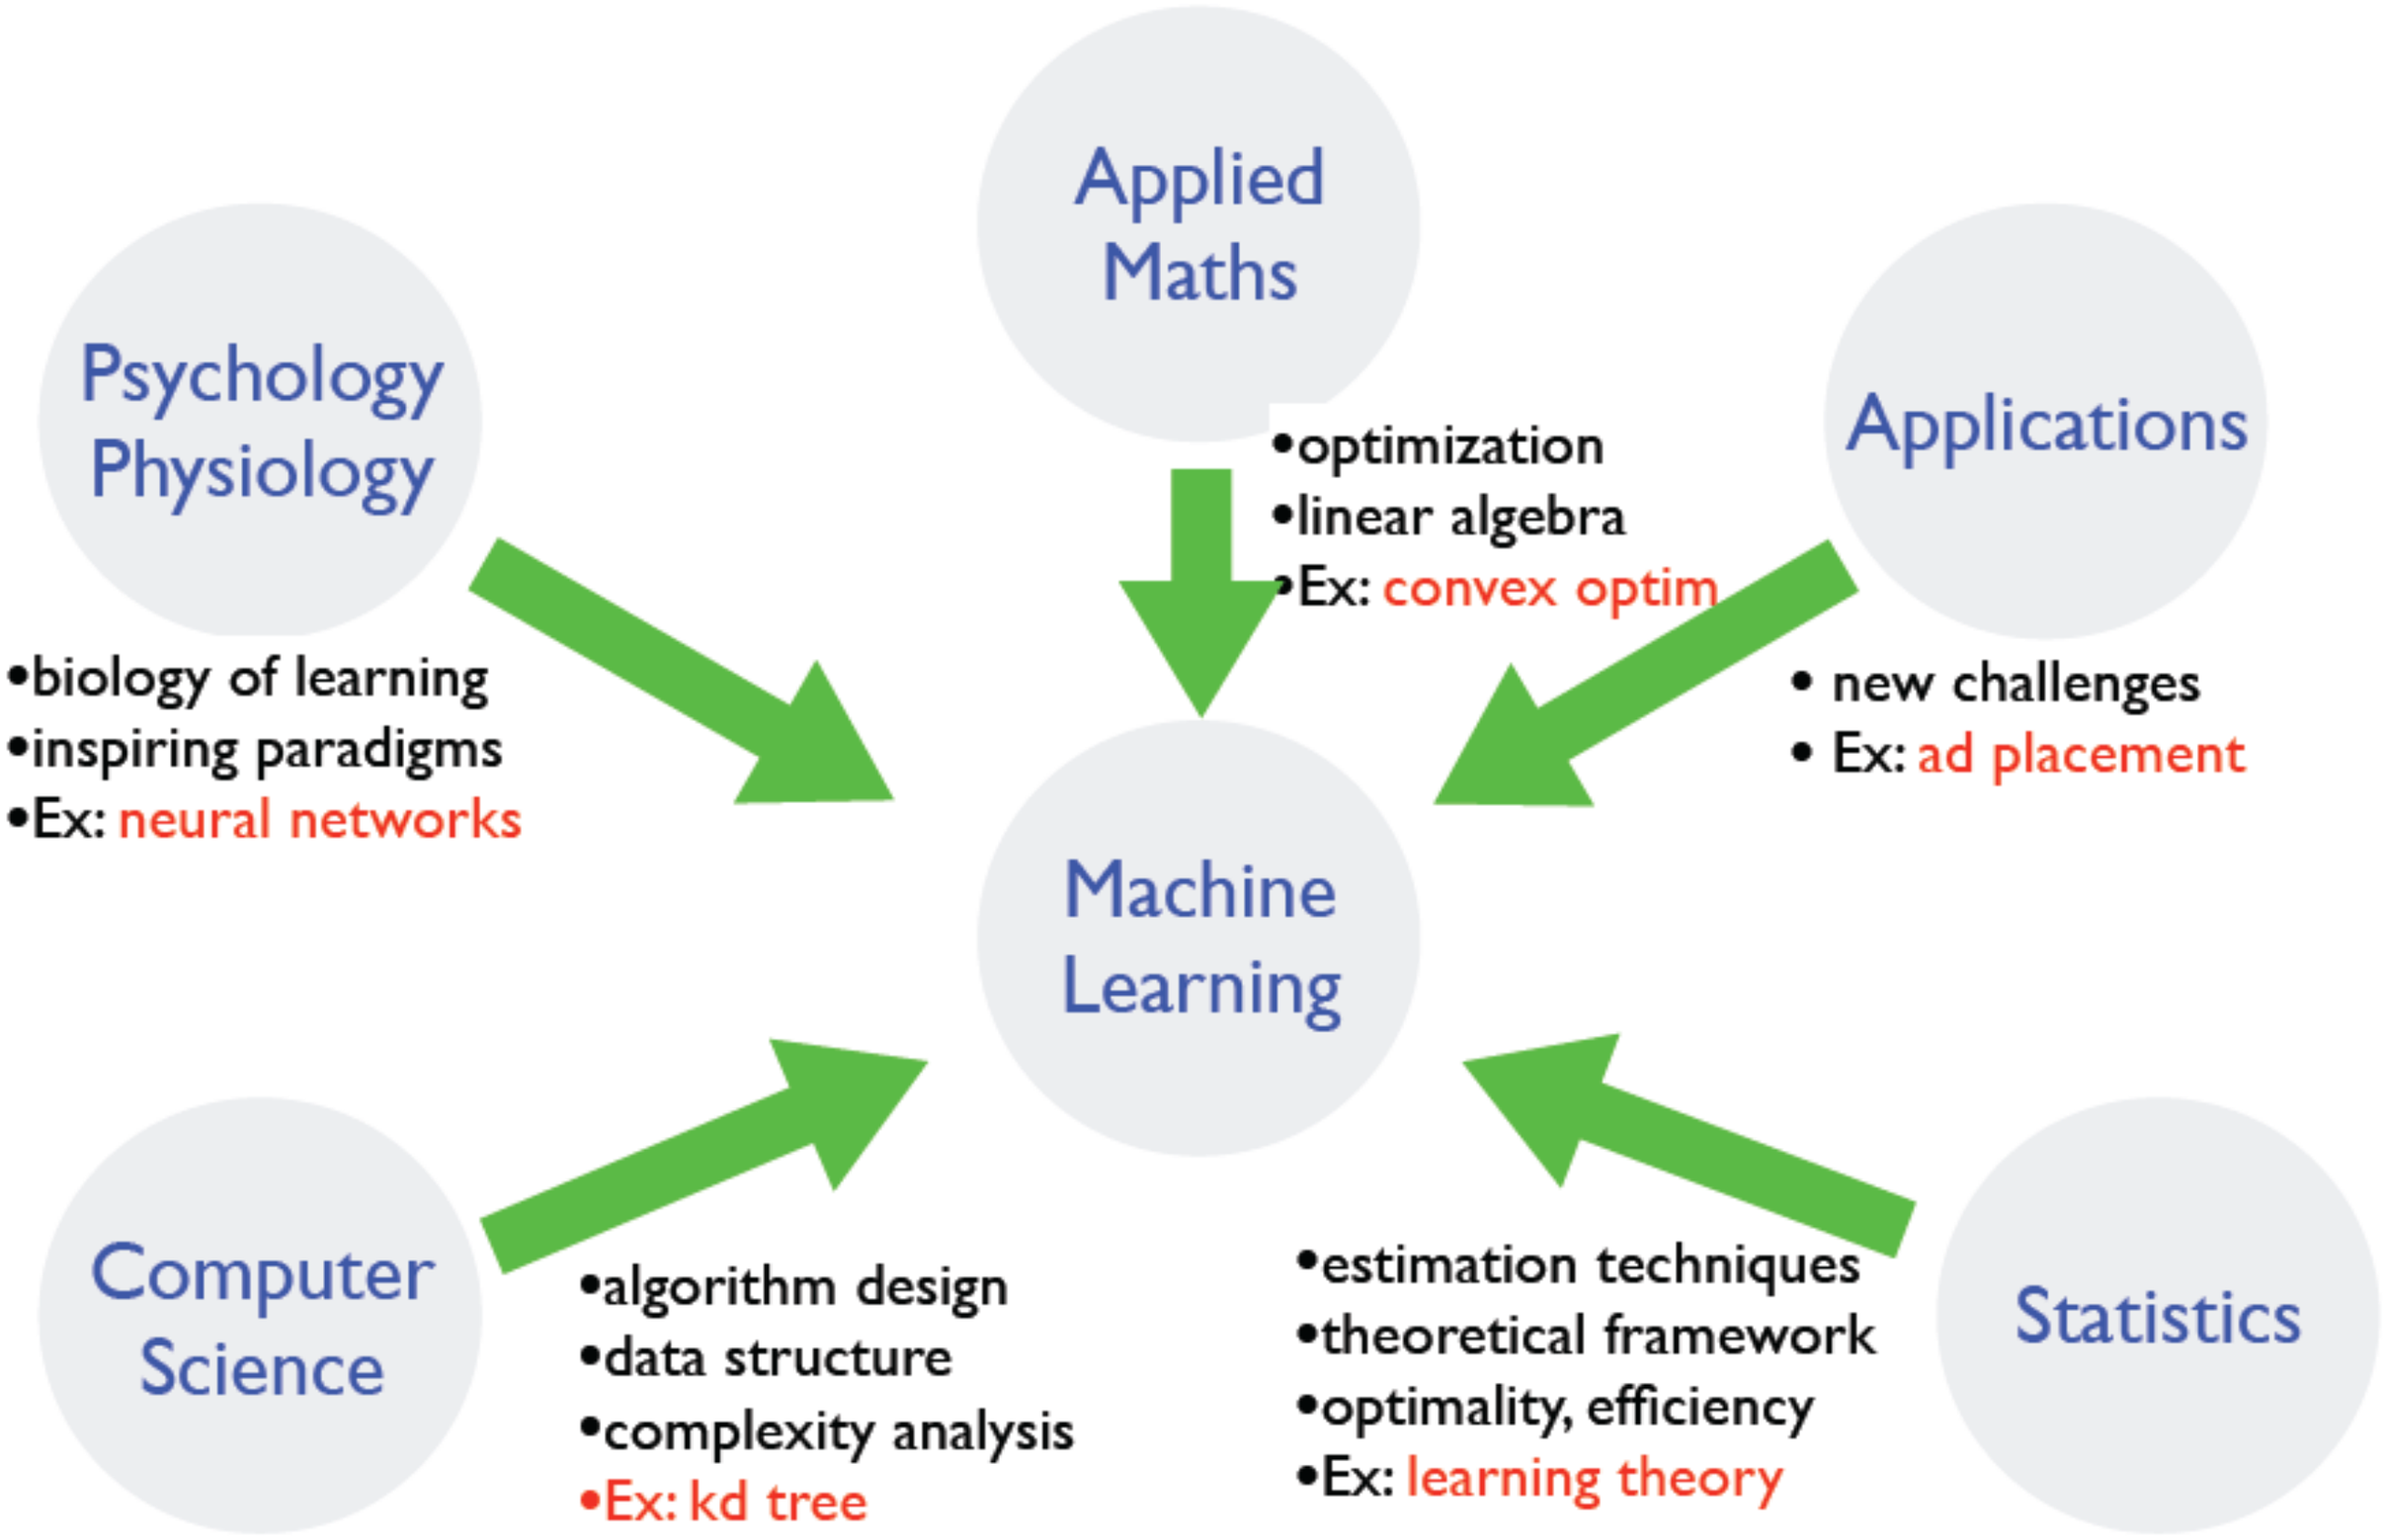

slide by Fei Sha

A Brief History of AI

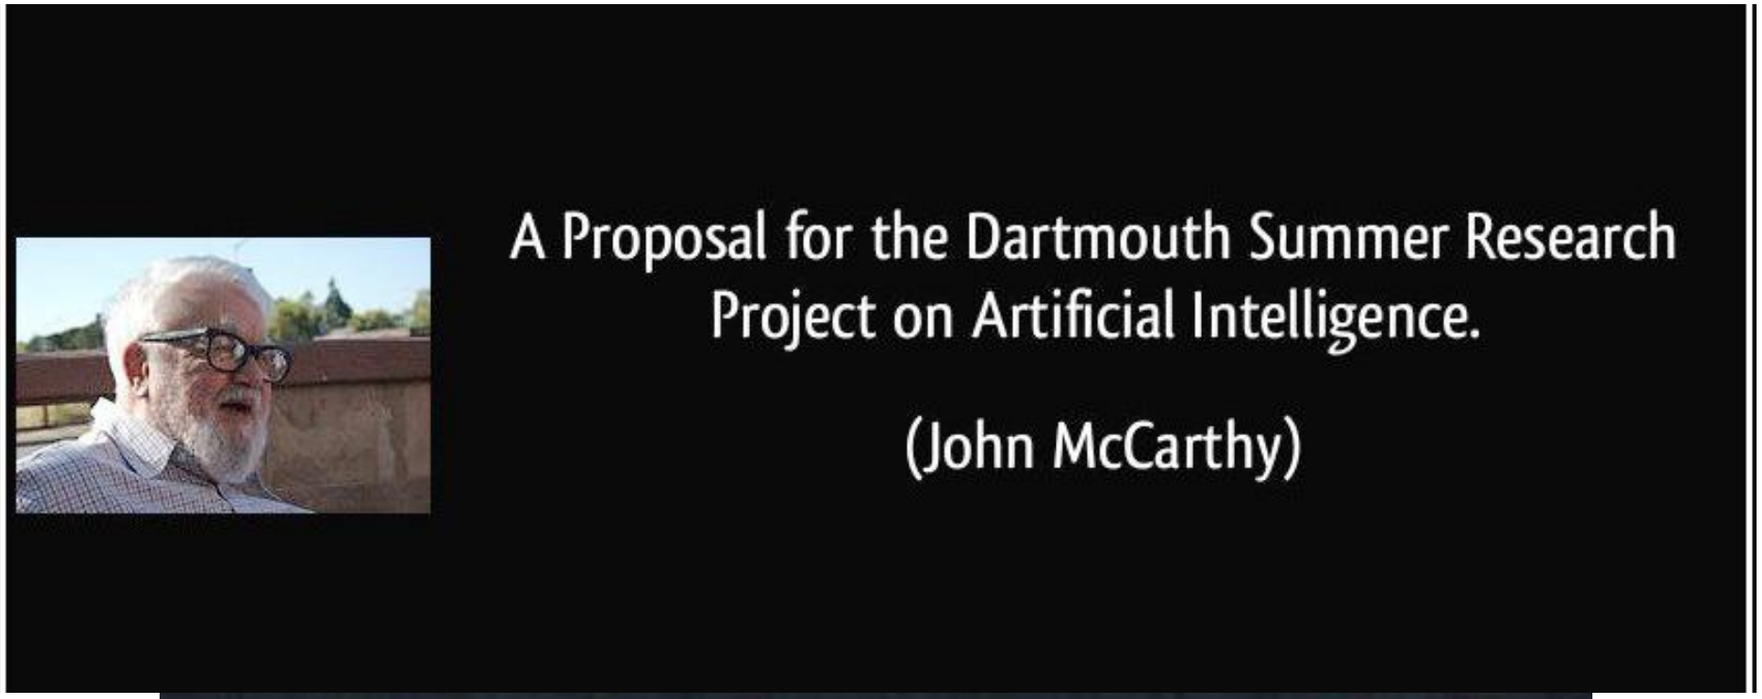

1956

slide by Dhruv Batra

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

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. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

Why is AI hard?

“I saw her duck”

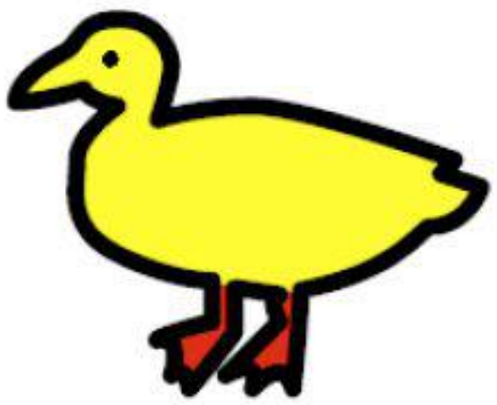

slide by Liang Huang

Image Credit: Liang Huang

“I saw her duck”

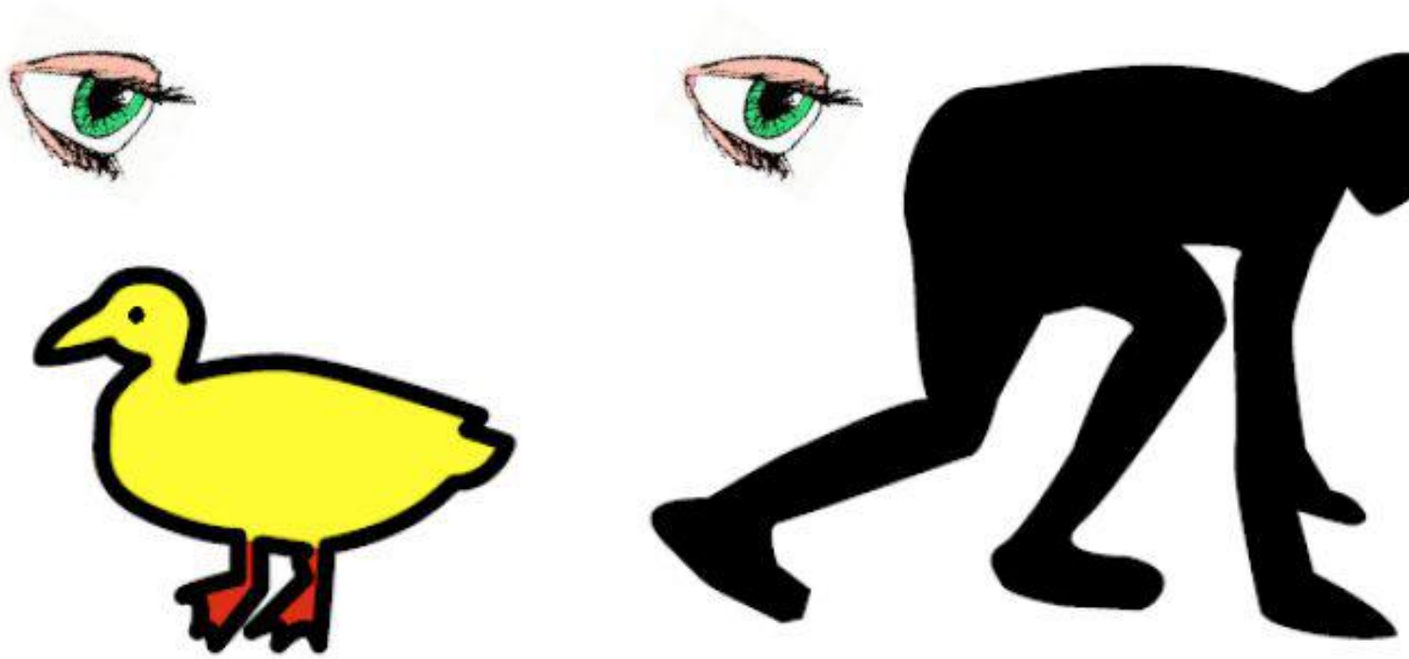

slide by Liang Huang

Image Credit: Liang Huang

“I saw her duck”

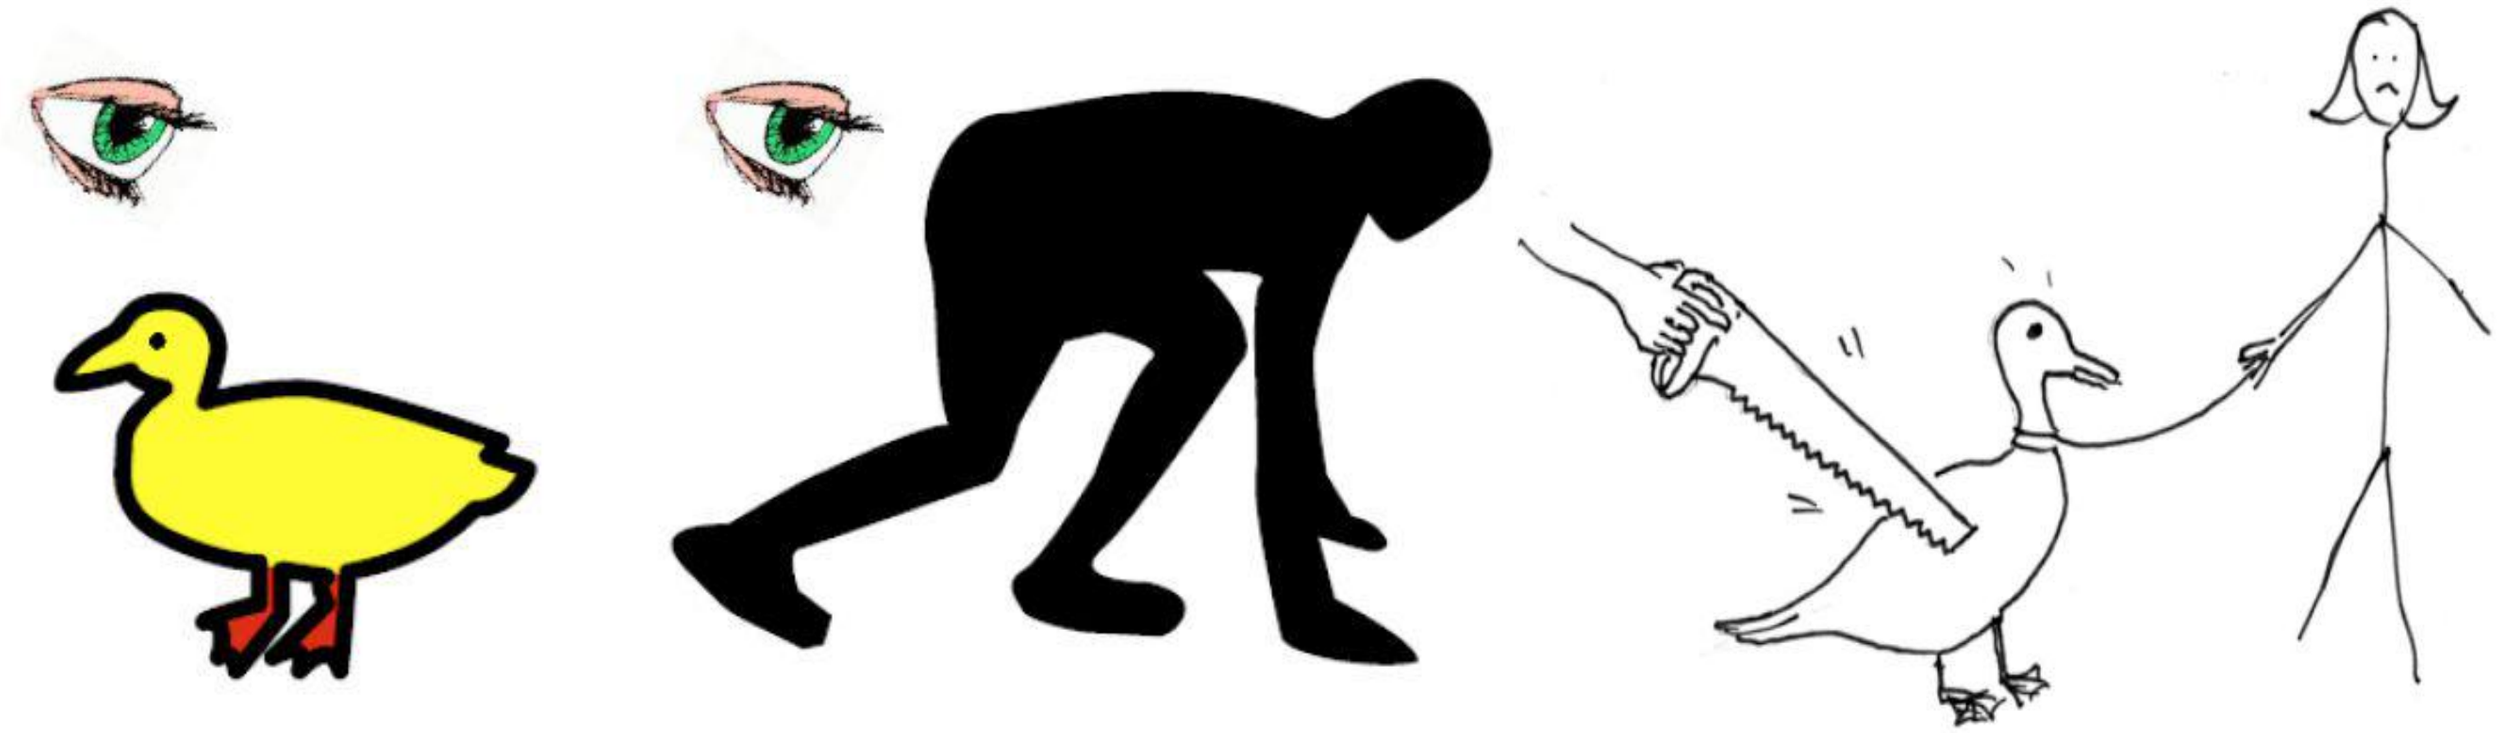

Image Credit: Liang Huang

slide by Liang Huang

Why are things working today?

- More compute power
- More data
- Better algorithms/ models

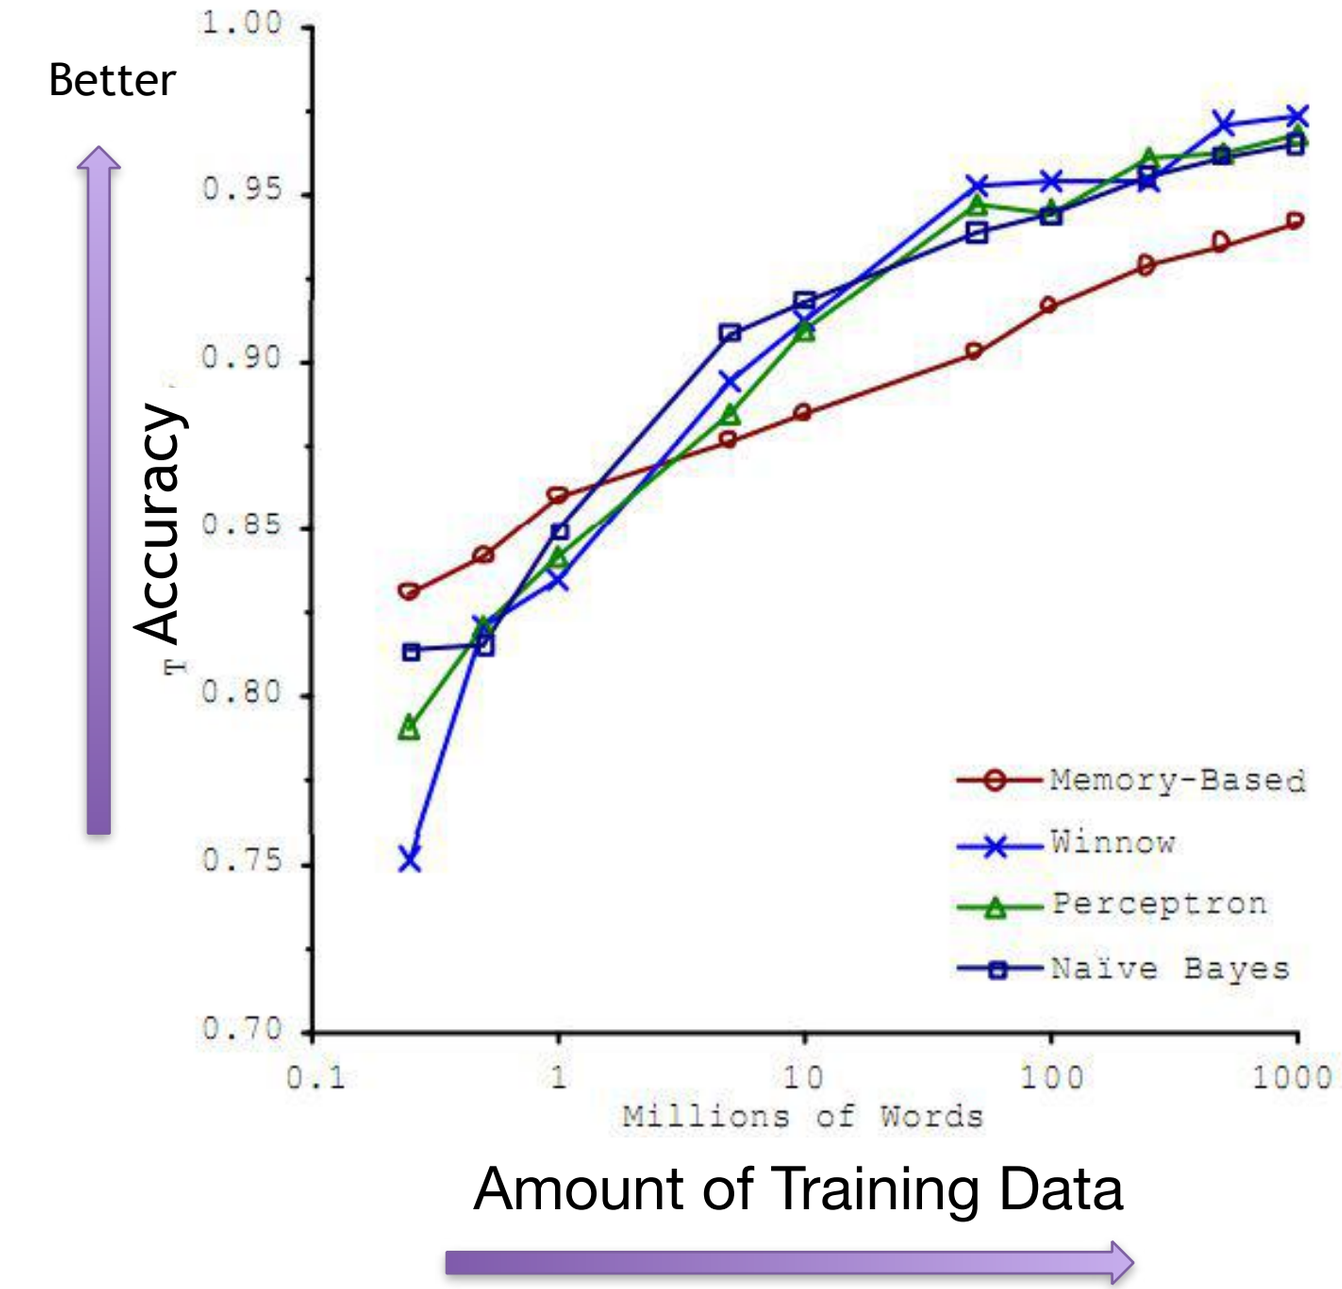

Figure Credit: Banko & Brill, 2011

slide by Dhruv Batra