WHAT IS MACHINE LEARNING?

SERTAN ŞENTÜRK DATA SCIENTIST – RESEARCH & DEVELOPMENT LUNCH 'N LEARN, 25 JULY 2018

ARTIFICIAL INTELLIGENCE! MACHINE LEARNING!



Yann LeCun @ylecun

If by "uprising" you mean "being able to fill the dishwasher".

technologyreview.com/s/611424/this- ...



This is how the robot uprising finally begins Combining the latest advances in artificial intelligence with robots could transform manufacturing and warehousing—and take AI to the next level. technologyreview.com

MACHINES ARE LEARNING!!



QUESTIONS TO BE ANSWERED...

- What is machine learning (ML)?
- What is possible with ML?
- How does it (conceptually) work?
- What are the challenges?
- Common misconceptions/hypes
- How can we get into ML?
-
- **Discussion:** What can we do with ML in Kobalt?

WHAT TO EXPECT FROM THIS TALK

Quite high level

No math!

Conversations > Me talking whilst you eat

Emphasis to music/media technologies

(A) DEFINITION

"the field of study that gives computers the ability to **learn** <u>without being explicitly programmed</u>"

Tom Mitchell, 1950s

Kind of early, isn't it?

https://www.coursera.org/lecture/machine-learning/what-is-machine-learning-Ujm7v

LEARNING FROM EXAMPLES







* You don't **code** the movement itself, but make the machine learn to move... somehow [™]

* ACTUAL RESULTS MAY VARY





WHAT IS THE RELATIONSHIP BETWEEN MACHINE LEARNING AI!, MATH, DATA STUFF?*

* roughly



WHERE IS IT USED?

short answer: everywhere ^_^



















WHAT TYPE OF PROBLEMS CAN MACHINE LEARNING SOLVE?

WHAT ML IS GOOD FOR?

- Difficult to formulate
 - What would you like to listen today?
- Many facets affecting the outcome
 - How will the stock market change in 6 months?
- Has a lot of special cases to consider
- Rigorous, repetitive tasks with limited time/resources
- Processing unstructured data
 - audio, video, image, plain text

TYPES OF PROBLEMS

- **Classification:** When you want to answer categorical questions...
- Regression: When you want to predict a quantity...
- **Clustering:** Finding groups with similar characteristics inside the data
- Feature selection/dimension reduction: Obtain (a combination of) most representative information

...

TYPES OF PROBLEMS: CLASSIFICATION

- Image classification
- Genre recognition
- Disease diagnosis
- Text categorization

. .





TYPES OF PROBLEMS: REGRESSION

- House prices
- Temperature
- Artist revenue

. .

Melody tracking





TYPES OF PROBLEMS: CLUSTERING

- Market segmentation
- Community discovery in social networks
- Anomaly detection
 - e.g. server fault prediction, fraud detection





LEARNING APPROACHES

- Supervised Learning
 - You have labeled data, which you feed to the model while training
 - Commonly used in classification, regression etc.
- Unsupervised Learning
 - You don't use labeled data and let machine "discover" some structure
 - Commonly used in clustering and feature learning tasks
- Reinforcement Learning
 - The machine learns some action by rewarding/punishing the decisions
 - Used in autonomous decision making: Games (Chess, Go, DOTA), robotics...
 - Could still be supervised, unsupervised, semi-supervised etc.

SOME PROBLEMS REQUIRE COMBINED SOLUTIONS



Autonomous Driving

Supervised classification: Object recognition (e.g. cars, people) Supervised regression: Location, size etc. of the objects Reinforcement learning: Driving itself

<u>Note</u>: The problem might still be solved by a single model

CONCEPTUAL ML PROCESS



ML STEP-BY-STEP: DATA PROCESSING



Labels are needed for supervised learning.

We are skipping unsupervised learning for simplicity's sake.

PRETTY TASK/DOMAIN SPECIFIC!

ML STEP-BY-STEP: DATA PROCESSING



ML STEP-BY-STEP: MODEL TRAINING



ML STEP-BY-STEP: ITERATIVE DESIGN



HOW CAN WE USE MACHINE LEARNING FOR MUSIC DATA?

- Melody/pitch tracking
- Chord recognition
- Key/Mode detection
- Tuning analysis
- Onset detection
- Beat tracking
- Tempo estimation
- Structural segmentation
- Transcription
- Instrument recognition
- Source-separation







EDM

energetic, danceable, instrumental



distorted







- Music similarity
- Artist similarity
- Genre recognition
- Emotion recognition
- Music auto-tagging
- Music recommendation
- Version identification
- Fingerprinting
- Query by humming
- Audio-score alignment
- Audio-lyrics alignment
- Optical music recognition



haunted, wired, seductive, menacing, mischievous, kind, multi-faceted...

Multimodality



Blackstar - David Bowie

Blackstar is an inky labyrinth of human cruelty and frailty shot through with moments of grace and transcendence, and obsessed with different kinds of transformation. And it's another record on which song carves out its own unique space, with no room for repetitions or redundancies. Even Bowie's voice

never does the same thing twice. It's haunted, wired, seductive, faceted performance from pop's great actor. Read more









Waiting On the Light to Change Ben Rector & Friends



Spread Too Thin

Dirty Heads

🞵 *make me a playlist*

Pittsburgh Symphony Orchestra live from Berlin, Opening concert of Musikfest Berlin 2013

OH, ARE THEY EASY TO DO?

nope ^_^

SOME COMMON PROBLEMS

- Requires A LOT OF computational resources
- Most need A LOT OF (labeled) data
- Needs high-level, in-house, technical expertise
 - Data scientists, data engineers, and other case-specific roles
- Typically requires domain-specific knowledge
- Technical debt
 - Requires rigorous evaluation
 - Outcomes may be difficult to "understand" & "debug"
 - Produces value after setting the appropriate data infrastructure
 - Performance may degrade due to changes in data & requirements...

COMMON MISCONCEPTIONS

RUN LIKE HELL IF YOU HEAR:

- ML is just math!
- ML is science fiction!
- ML is a humanoid robot touching numbers projected to holograms!
- ML is the answer of life, the Universe and everything!
- ML will make humans obsolete!
- ML will destroy the world!
- ML will bring world peace!

HOW CAN I LEARN MACHINE LEARNING?

RESOURCES

- Online courses
 - Coursera (Andrew Ng), fast.ai, Google machine learning resources, Udemy, ...
- Blogs
 - OpenAI, Google DeepMind, Towards DataScience, AnalyticsVidhya, Kdnuggets, Spotify Labs ...
- Software tools & documentation
 - scikit-learn, tensorflow, pytorch, keras, Spark MLlib, ...
- Conferences, Journals, Archives
 - ICML , NIPS, Strata Data Conference, IEEE, ACM, arxiv, ISMIR, ...
- Books
 - Christopher Bishop, Ian Goodfellow, Andrew Ng, David Barber...

CONCLUSION

- Machine learning is not a myth
 - It has been already utilized in every part of our daily lives...
- It is quite powerful
 - after we define <u>clear goals/business problems</u>, make thorough investigation, and build the proper <u>infrastructure</u>
- It's easy to start
 - There are a lot of free and open source training/tools available
- But difficult to master
 - Building & maintaining ML solutions require a fair amount of data, skilled work force, automated processes, and <u>planning</u>...

