


**LOSCHMIDT
LABORATORIES**



12. Artificial Intelligence in Life Sciences

Outline


- ❑ **Motivation**
- ❑ **Introduction to AI and ML**
- ❑ **Modern challenges in Bioengineering**
- ❑ **Basics of ML**
- ❑ **Recent applications**




Motivation

Motivation

MSLGAKPFGEKKFIEIKGRRMAYIDEGTGDPILFQHGNPTSSYLWRNIMPHC
AGLRLIACDLIGMGDSKLDPSGPERYAYAEHRDYLDALWEALDLGDRVV
LVVHDWGSALGFDWARRHRERVQGIAYMEAIAMPIEWADFPEQDRDLFQ
AFRSQAGEELVLQD



Function




Motivation



Sequence

MSLGAKPFGEKKFIEIKGRRMAYIDEGTG
DPILFQHGNPTSSYLWRNIMPHCAGLGR
LIACDLIGMGDSKLDPSGPERYAYAHR
DYLDALWEALDLGDRVVLVVDWGSAL
GFDWARRHRERVQGIAYMEAIAMPIEW
ADFPEQDRDLFQAFRSQAGEELVLQD

Structure



Evolution



Main approaches:

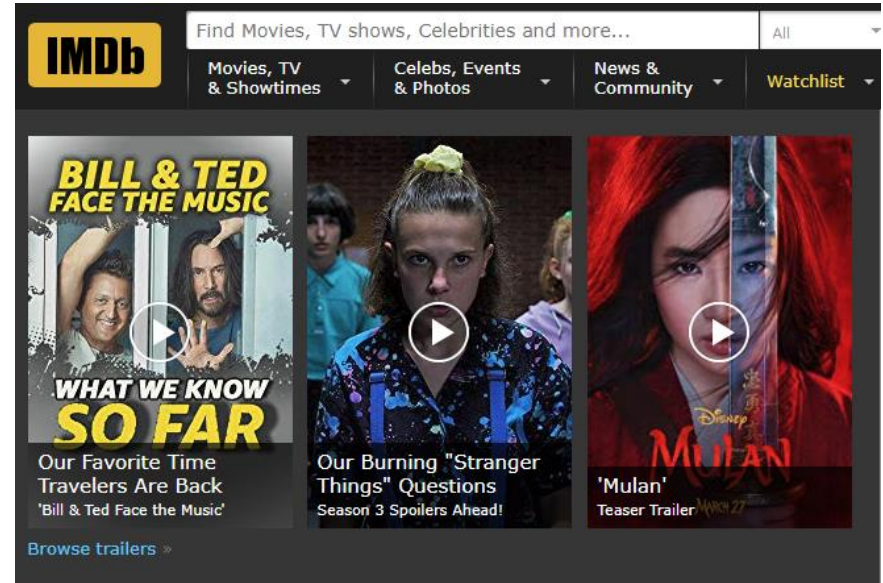
- Experimental
- Rule-based
- Machine learning




Introduction to AI and ML

Introduction to AI and ML

- Recommendation engines
- Gaming
- Image & speech recognition
- Anomaly detection
- Natural language processing
- Data mining
- ...

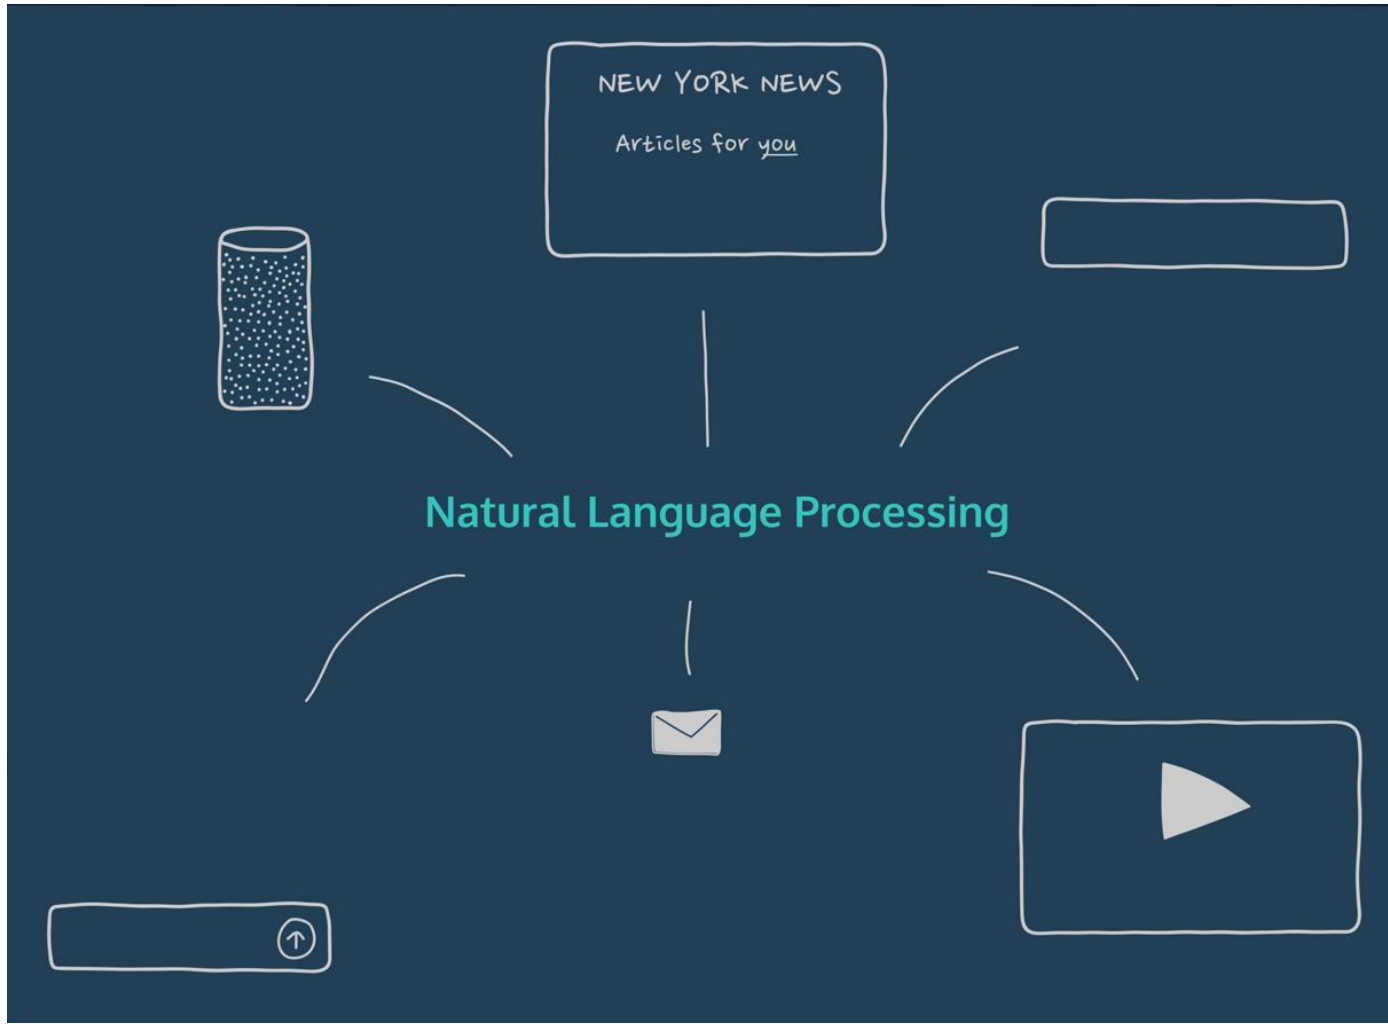


Introduction to AI and ML



| Login | Non- Transactional Activities | | Transaction |
|--|--|--|--|
| <ul style="list-style-type: none"> • Challenges • Device • Cookie • IP Address • Time of day • Network | <ul style="list-style-type: none"> • View balance • View history • Updated address • Update email • Update password | <ul style="list-style-type: none"> • Add new user • Change limits • Set up batch • Set up template • Add payees | <ul style="list-style-type: none"> • ACH • Wire • Bill Pay • Loan Draw |

Introduction to AI and ML



Introduction to AI and ML

DataRobot

PLATFORM ▾

SOLUTIONS ▾

SUCCESS ▾

RESOURCES ▾

PARTNERS ▾

WE'RE HIRING **489** ▾



CONTACT US

Automated Machine Learning: The Competitive Edge You Need

Professional sports is a cutthroat industry, on and off the playing field. The smallest of competitive advantages can

Optimize Player Performance ▾


Predict — and Prevent — Injuries ▾

Project Prospect Improvement ▾

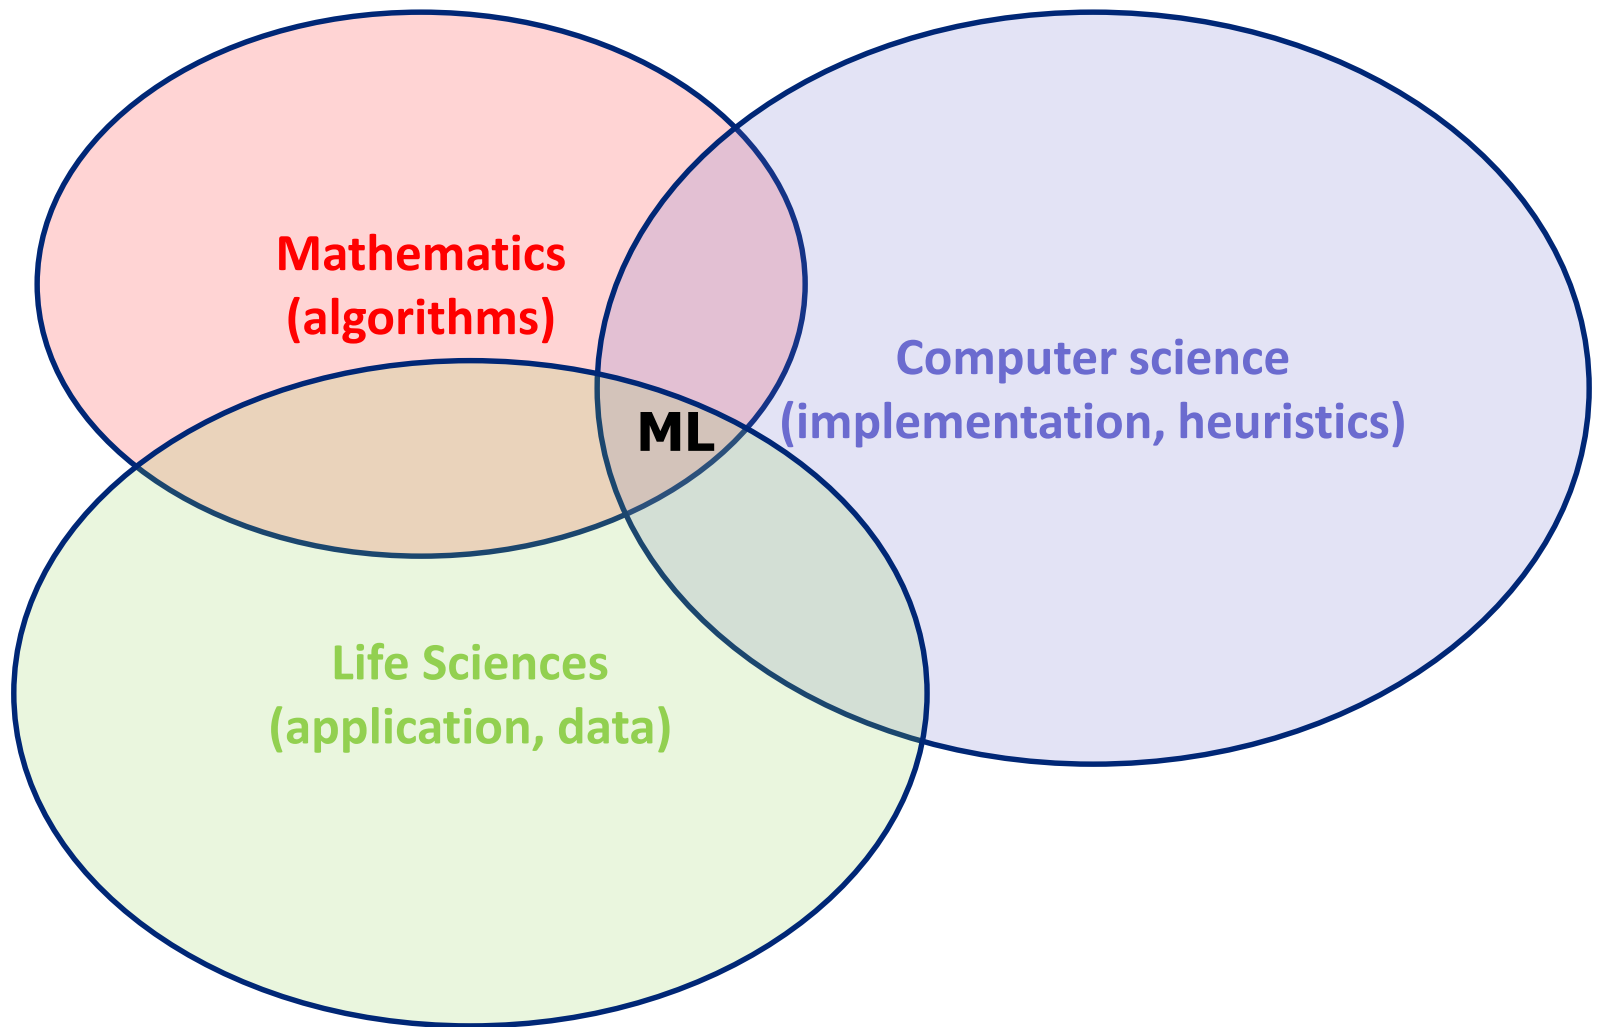
Consolidate Valuable Data ▾

Increase Profit Potential ▾

Improve Operational Efficiency ▾



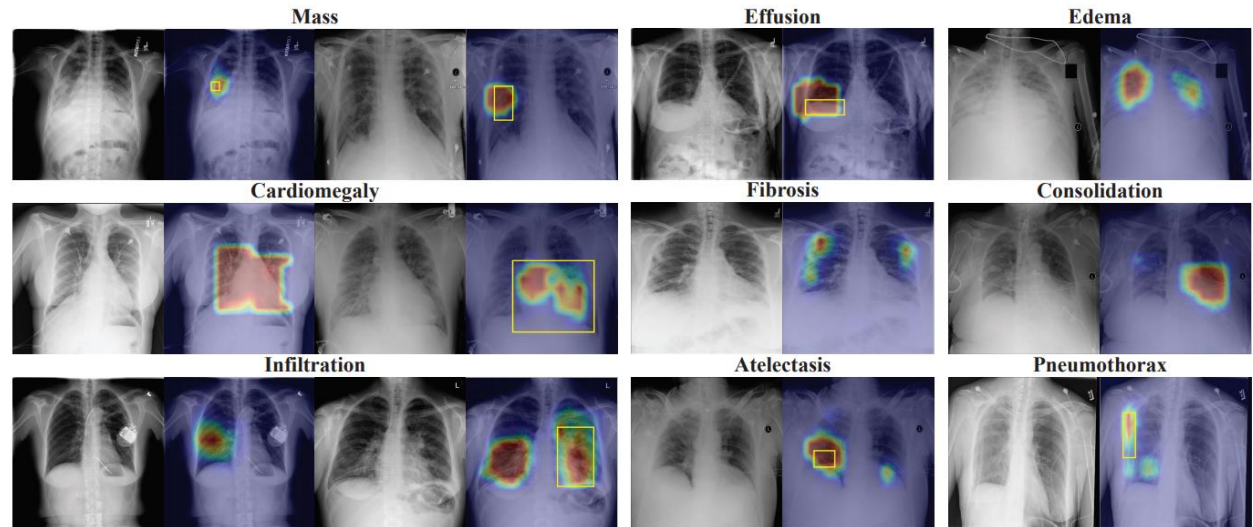


Introduction to AI and ML




A blue-tinted microscopic image of cells, showing their intricate structures and membranes, positioned at the top of the slide.

Modern challenges in Bioengineering

Modern Challenges




Modern Challenges



Chemical Space

92M compounds

~9,600 Drugs




Protein Space


20244 human proteins

~6,200 3D structures

~2700 Drug targets



- **How to annotate:**
 - protein structure
 - protein function
 - protein interactions
 - ...
- **How to increase:**
 - enzymatic activity
 - stability & solubility
 - substrate specificity
 - enantioselectivity



0.2 ns

Modern Challenges

March 2000

Expert panel proposes cohort study of 500,000 adults.

April 2002

Wellcome Trust and U.K. government announce initial funding of £45 million.

March 2012

UK Biobank resource launches.

October 2015

Imaging data available for 5000.

May 2015

Genotyping data on 150,000 released.

July 2017


Genotyping data on 500,000 released.

March 2019

Exome data on 50,000 to be released.

2020

All exome data released.




Recruitment of participants

N. DESAI/SCIENCE

Engine of productivity

Published papers based on the UK Biobank's bounty of health and genetics data are piling up fast, in part because the data are freely available.




(GRAPHIC) N. DESAI/SCIENCE; (DATA) UK BIOBANK

UK Biobank Principal Investigator Rory Collins stands amid stored biospecimens from the project's half-million participants. NIGEL HILIER

Huge trove of British biodata is unlocking secrets of depression, sexual orientation, and more

By Jocelyn Kaiser, Ann Gibbons | Jan. 3, 2019, 1:20 PM



1. A large number of relevant features;

- From bio-physico-chemical to textual;
- Heterogeneous data (e.g. clinical, imaging, and genomic data);

2. Complex tasks and large parameter space;

- A single 300-amino-acid-long protein will have $300 \cdot 19 = 5700$ single-point variants!

3. Large datasets available and new data are collected.

Ideal ML setup = complex tasks + relevant features + abundant data




Basics of ML

Basics of ML: features

- Feature vector: $x = (x_1, x_2, \dots, x_n)$
- All features must be converted to numbers:

yes = 1, no = 0



| | | | |
|---|---|---|---|
| 1 | 1 | 1 | 0 |
| 1 | 0 | 1 | 0 |
| 1 | 1 | 0 | 0 |
| 1 | 0 | 0 | 0 |

ATGTA CTGA


One-hot
encoding




| | A | C | G | T |
|---|---|---|---|---|
| 1 | 0 | 0 | 0 | 1 |
| 0 | 0 | 0 | 1 | 0 |
| 0 | 0 | 1 | 0 | 0 |
| 0 | 0 | 0 | 1 | 0 |
| 1 | 0 | 0 | 0 | 0 |
| 0 | 1 | 0 | 0 | 0 |
| 0 | 0 | 0 | 1 | 0 |
| 0 | 0 | 1 | 0 | 0 |
| 1 | 0 | 0 | 0 | 0 |

- **Unsupervised learning**
 - only features are available;
 - goal: cluster the data or reduce their dimensionality;
- **Supervised learning**
 - features and labels are available;
 - goal: learn to predict the label based on the features;
- Reinforcement learning, feature learning, anomaly detection, etc...

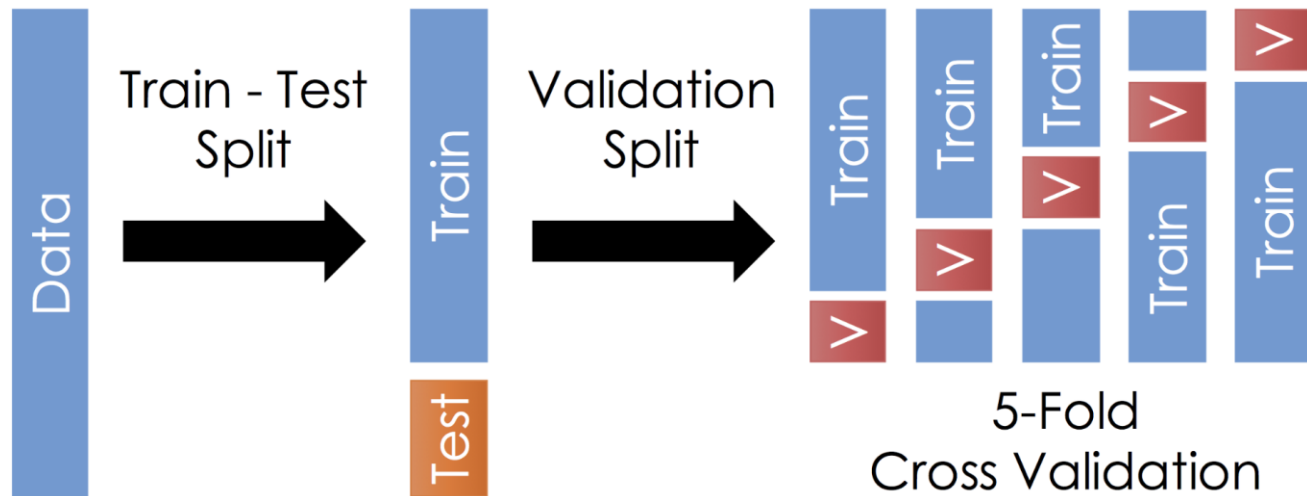
Basics of ML: unsupervised learning





Basics of ML: supervised learning

- We want ML models to be generalizable = good at predicting labels for **previously unseen data**;
- It is essential to split the data into training set and test sets and use the latter for final evaluation only!
- To fine-tune an algorithm, K-fold cross validation is implemented:



Basics of ML: supervised learning



- Let $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ be our data set, where \mathbf{x} are feature values, and \mathbf{y} are labels;
- Any ML predictor is fundamentally a function $f(\mathbf{x})$:

$$f(\text{feature values}) = \text{label}$$

- Usually, a generic group of functions $f(\mathbf{x}, \boldsymbol{\beta})$ is chosen, where $\boldsymbol{\beta}$ is a set of parameters;
- Then we “train” the ML predictor: pick such $\boldsymbol{\beta}^*$ that $f(\mathbf{x}^{(i)}, \boldsymbol{\beta}^*)$ is as close to $\mathbf{y}^{(i)}$ as possible.

Basics of ML: supervised learning

$$f(x) = m \cdot x + b$$





Recent applications

Recent applications: decision tree

| Gene Pair | Interact? | Expression correlation | Shared localization? | Shared function? | Genomic distance |
|-----------|-----------|------------------------|----------------------|------------------|------------------|
| A-B | Yes | 0.77 | Yes | No | 1 kb |
| A-C | Yes | 0.91 | Yes | Yes | 10 kb |
| C-D | No | 0.1 | No | No | 1 Mb |
| ⋮ | | | | | |

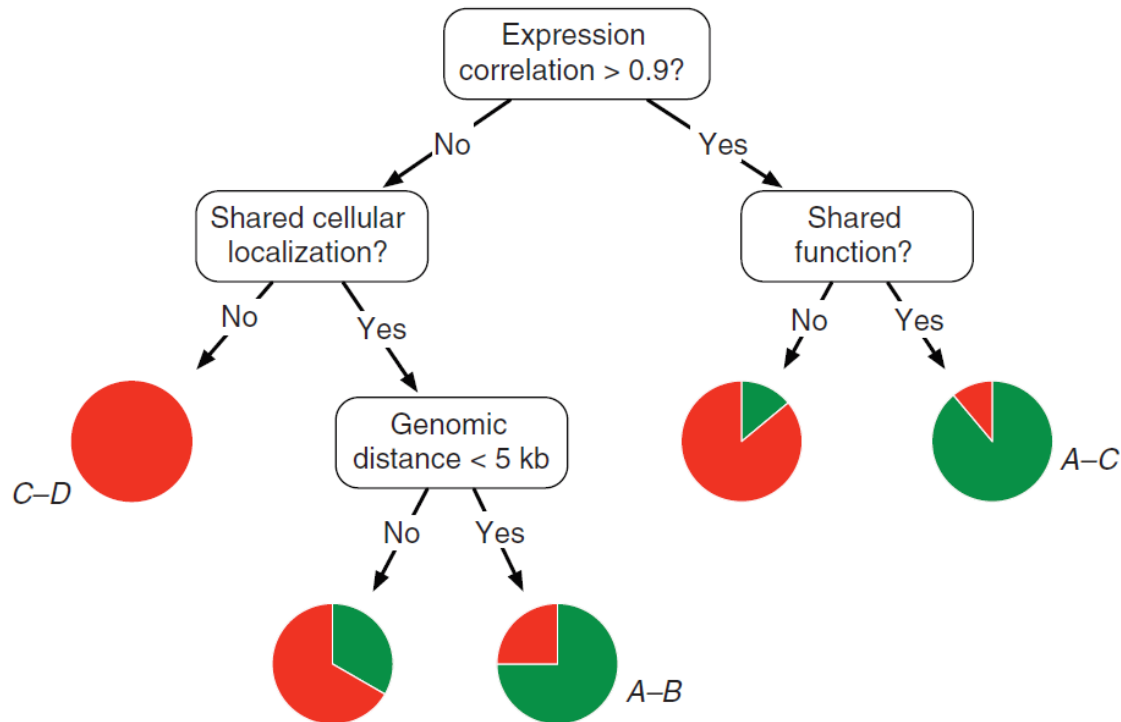



Figure 1 A hypothetical example of how a decision tree might predict protein-protein interactions.

Recent applications: neural networks



Recent applications: ANNs for drug design

- QSAR modeling for ligand binding activity;
- Structure generation;
- Synthetic pathway generation;






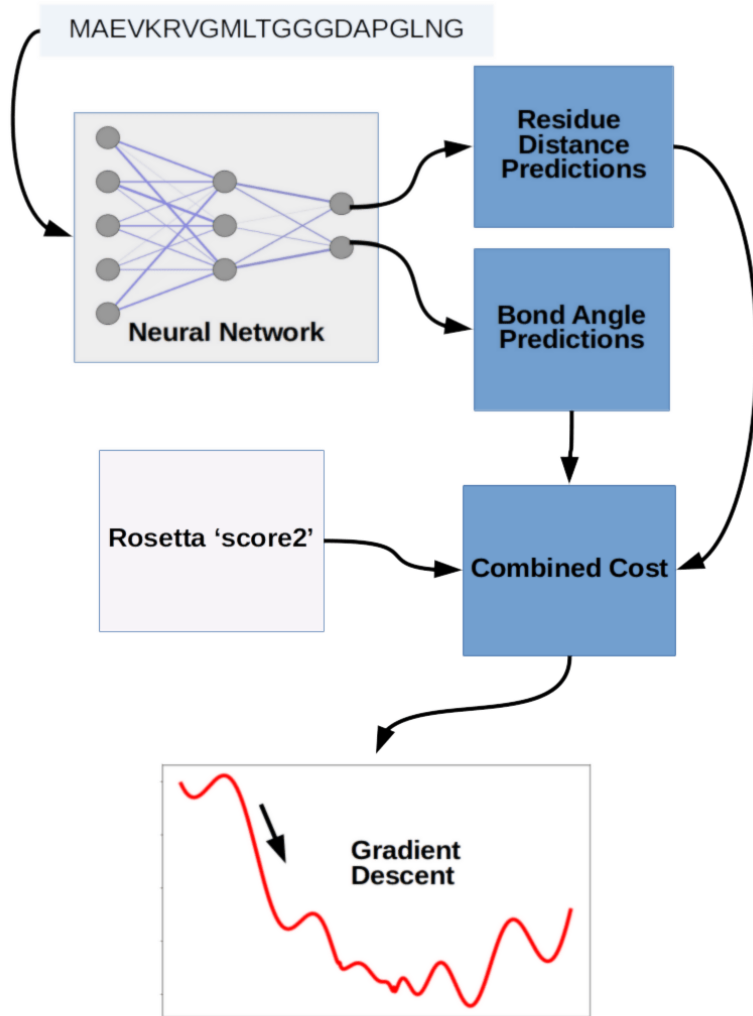



Figure 2: Illustration of the Chemception framework. After a SMILES to structure conversion, the 2D images are mapped onto an 80 x 80 image that serves as the input image data for training a deep neural network to predict toxicity, activity, and solvation properties.



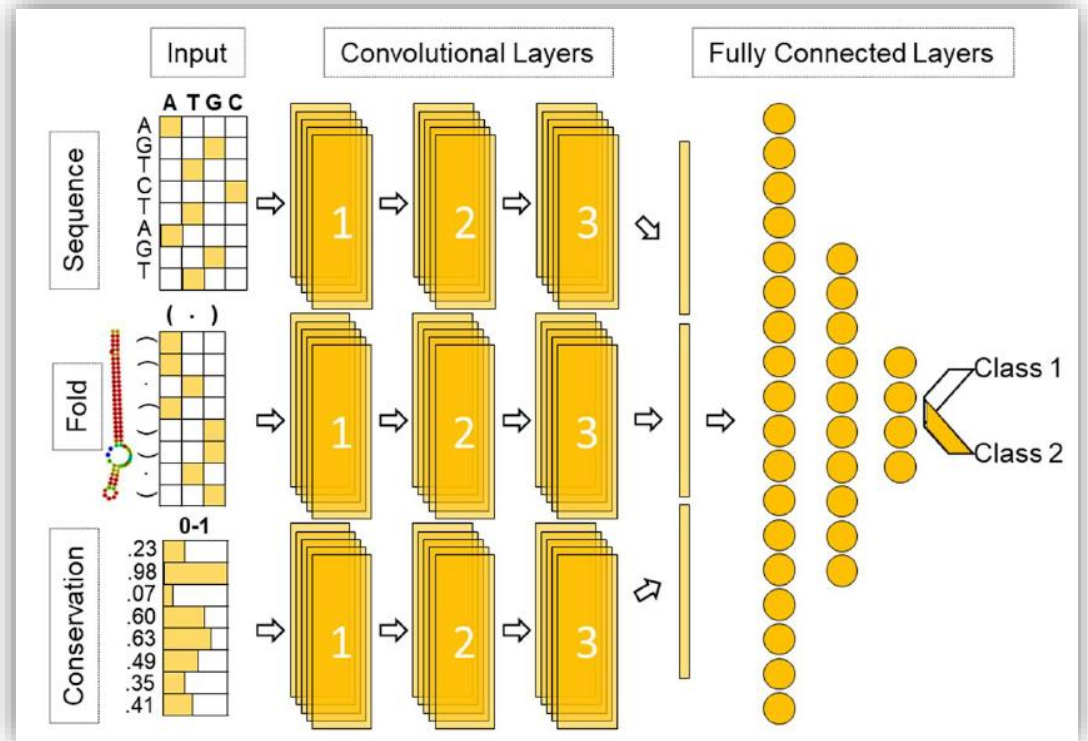
Recent applications: AlphaFold for protein folding



Recent applications: MuStARD for genome annotation

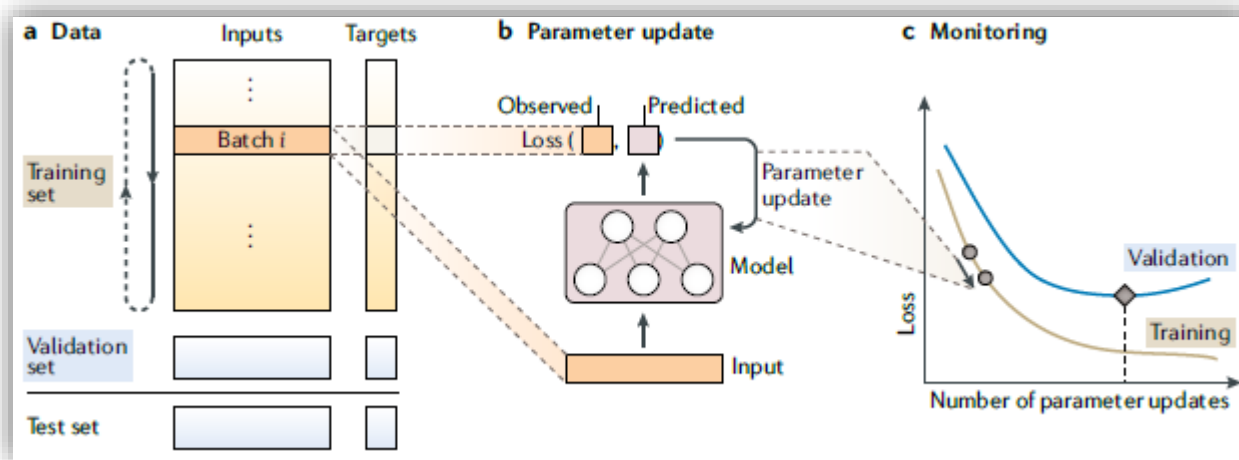


© 2012 Pearson Education, Inc.



Reading

- ❑ Eraslan G et al. 2019 "Deep learning: new computational modelling techniques for genomics." Nature Reviews Genetics. (especially pages 2-6)



REVIEWS

Deep learning: new computational modelling techniques for genomics

Gökçen Eraslan^{1,2,5}, Žiga Avsec^{3,5}, Julien Gagneur^{3*} and Fabian J. Theis^{1,2,4*}

Abstract | As a data-driven science, genomics largely utilizes machine learning to capture dependencies in data and derive novel biological hypotheses. However, the ability to extract new insights from the exponentially increasing volume of genomics data requires more expressive machine learning models. By effectively leveraging large data sets, deep learning has transformed fields such as computer vision and natural language processing. Now, it is becoming the method of choice for many genomics modelling tasks, including predicting the impact of genetic variation on gene regulatory mechanisms such as DNA accessibility and splicing.