Machine Learning: Closing

Hao Tang

March 29, 2024

What is a model?

What is a model?

A **model** is an imaginary artifact that describes the inner working of something.

What is a model?

A **model** is an imaginary artifact that describes the inner working of something.

- imaginary

What is a model?

A **model** is an imaginary artifact that describes the inner working of something.

- imaginary
- describes inner working

What is a model?

A **model** is an imaginary artifact that describes the inner working of something.

- imaginary
- describes inner working
- of something

Examples

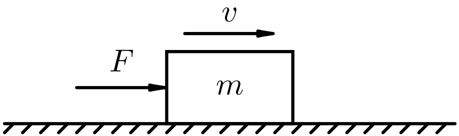

Examples

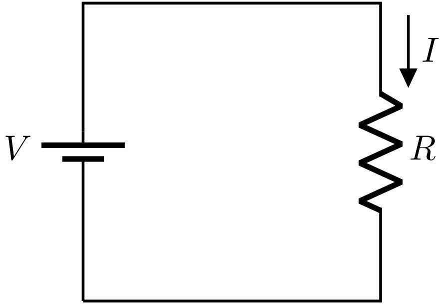

Examples

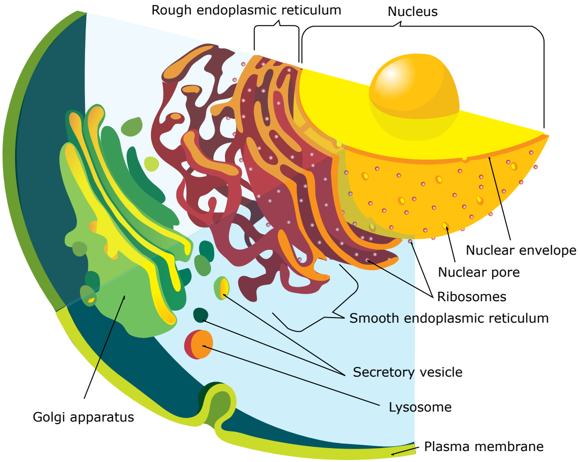

Rough endoplasmic reticulum
Nucleus
Nuclear envelope
Nuclear pore
Ribosomes
Smooth endoplasmic reticulum
Secretory vesicle
Lysosome
Golgi apparatus
Plasma membrane

Examples

The hypothesis class $\mathcal{H}$ is PAC learnable if for any distribution $\mathcal{D}$, $\epsilon > 0$, and $0 < \delta < 1$, there exists N such that for every $n > N$ and for every $h \in \mathcal{H}$,

$$\mathbb{P}_{S \sim \mathcal{D}^n}\Big[|L_{\mathcal{D}}(h) - L_S(h)| > \epsilon\Big] < \delta. \tag{1}$$

What about ...

What about ...

- Perceptrons

What about ...

- Perceptrons
- Linear classifiers

What about ...

- Perceptrons
- Linear classifiers
- Gaussian mixture models

What about ...

- Perceptrons
- Linear classifiers
- Gaussian mixture models
- Transformers

Machine learning

- Machine learning does not eliminate the need to build models.

Machine learning

- Machine learning does not eliminate the need to build models.
- What is machine learning?

Machine learning

- Machine learning does not eliminate the need to build models.
- What is machine learning?
- When to use machine learning?

Machine learning

- Machine learning does not eliminate the need to build models.
- What is machine learning?
- When to use machine learning?
- How to use machine learning?

Types of thinking

- Mathematical
  - What happens if we have the optimal solution?
- Computational
  - What are the exact steps?
- Statistical
  - How many samples do we need?

Connections to other courses

- Foundations of Data Science (FDS)
- Applied Machine Learning (AML)
- Machine Learning and Pattern Recognition (MLPR)
- Probabilistic Modeling and Reasoning (PMR)
- Machine Learning Practical (MLP)
- Machine Learning Theory (MLT)
- Reinforcement learning (RL)

Connections to other courses

- Foundations of natural language processing (FNLP)
- Accelerated natural language processing (ANLP)
- Natural language understanding, generation, and machine translation (NLU+)
- Speech processing (in PPLS)
- Automatic speech recognition (ASR)
- Speech synthesis (in PPLS)
- Image and vision computing
- Advanced robotics

Tasks we haven't talked about

- Information retrieval
- 3D reconstruction
- Text generation
- Protein folding

Evaluation we haven't talked about

- Mean average precision
- Receiver operating characteristic (ROC) curve
- Word error rates (WER)
- Bilingual evaluation understudy (BLEU) score
- Mean opinion score (MOS)

Techniques we haven't talked about

- More optimization
  - There are a lot of problems that cannot be solved with gradient descent.

Techniques we haven't talked about

- More optimization
  - There are a lot of problems that cannot be solved with gradient descent.
- k nearest neighbor
  - A point is like its neighbors.

Techniques we haven't talked about

- More optimization
  - There are a lot of problems that cannot be solved with gradient descent.
- k nearest neighbor
  - A point is like its neighbors.
- Decision trees
  - Decisions are based on conditional statements.

Techniques we haven't talked about

- More optimization
  - There are a lot of problems that cannot be solved with gradient descent.
- k nearest neighbor
  - A point is like its neighbors.
- Decision trees
  - Decisions are based on conditional statements.
- More dimensionality reduction
  - Points in space have structures.

Techniques we haven't talked about

- More optimization
  - There are a lot of problems that cannot be solved with gradient descent.
- k nearest neighbor
  - A point is like its neighbors.
- Decision trees
  - Decisions are based on conditional statements.
- More dimensionality reduction
  - Points in space have structures.
- Bayesian approaches
  - Beliefs are updated and marginalized.

Learning settings we haven't talked about

Learning settings we haven't talked about

- Active learning
  - What is the next sample that is most useful to learn?

Learning settings we haven't talked about

- Active learning
  - What is the next sample that is most useful to learn?

- Online learning
  - What can we learn if we make decisions sequentially and can look back?

Learning settings we haven't talked about

- Active learning
  - What is the next sample that is most useful to learn?
- Online learning
  - What can we learn if we make decisions sequentially and can look back?
- Semi-supervised learning
  - What do we do with unlabeled data?

Learning settings we haven't talked about

- Active learning
  - What is the next sample that is most useful to learn?
- Online learning
  - What can we learn if we make decisions sequentially and can look back?
- Semi-supervised learning
  - What do we do with unlabeled data?
- Federated learning
  - Can learning be distributed?

Aspects

- Privacy
- Fairness
- Interpretability

Differentially private

A learning algorithm A is differentially private if for all data sets S and $S^{(i)}$ that differs in the i-th sample,

$$\mathbb{P}[A(S) \in \mathcal{G}] < e^{\epsilon}\mathbb{P}[A(S^{(i)}) \in \mathcal{G}] + \delta \tag{2}$$

for any subset $\mathcal{G} \subseteq \mathcal{H}$ of hypothesis.

Fairness

- An outcome Y should be independent of groups G. ($Y \perp G$)
- A score R should be independent of the groups G given the outcome Y. ($R \perp G|Y$)
- An outcome Y should be independent of the groups G given the score R. ($Y \perp G|R$)

Artificial Intelligence

Artificial Intelligence

- Building models of intelligence, not solving tasks
- Involving philosophy, psychology, linguistics, etc

What we have learned in this course

- Formalize a ML problem with math
- Read and understand ML theorems
- Turn ML algorithms into programs