



Supervised Learning Terminology and Concepts

DL4DS Spring 2025

Lecture Outline

- Homeworks and Jupyter Notebooks plan
- Supervised Learning
- More on Projects

Artificial intelligence

Machine learning

Supervised learning

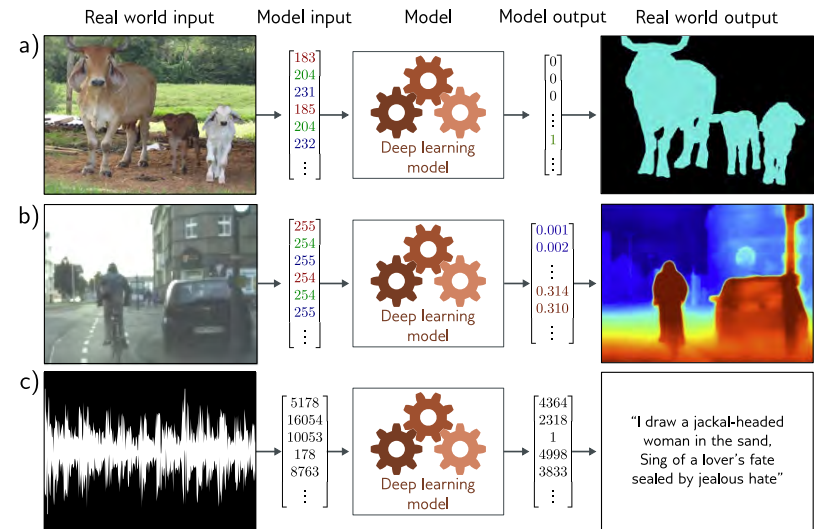
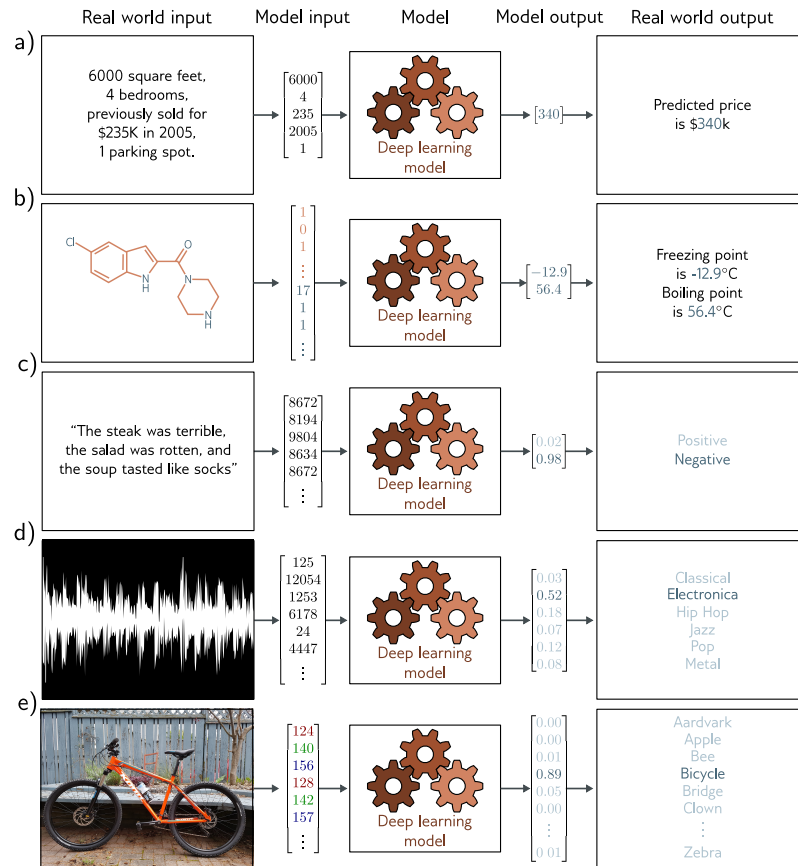
Unsupervised learning

Reinforcement learning




Deep learning

Supervised Learning Classification and Regression Applications



Regression



- Univariate regression problem (one output, real value)

Supervised learning

- Overview
- Notation
 - Model
 - Loss function
 - Training
 - Testing
- 1D Linear regression example
 - Model
 - Loss function
 - Training
 - Testing
- Where are we going?

Supervised learning

- Overview
- Notation
 - Model
 - Loss function
 - Training
 - Testing
- 1D Linear regression example
 - Model
 - Loss function
 - Training
 - Testing
- Where are we going?

Supervised learning overview

- **Supervised learning model** = mapping from one or more inputs to one or more outputs
- Model is a family of equations → “**inductive bias**”
- Computing the outputs from the inputs → **inference**
- Model also includes **parameters**
- Parameters affect outcome of equation
- **Training** a model = finding parameters that predict outputs “well” from inputs for **training** and **evaluation datasets** of input/output pairs

Supervised learning

- Overview
- **Notation**
 - Model
 - Loss function
 - Training
 - Testing
- 1D Linear regression example
 - Model
 - Loss function
 - Training
 - Testing
- Where are we going?

Also Appendix A of the book.

Notation:

- Input:

x



Variables always Roman letters

- Output:

y

Normal lower case = scalar

Bold lower case = vector

Capital Bold = matrix

- Model:

y = **f**[**x**]



Functions always square brackets

Normal lower case = returns scalar

Bold lower case = returns vector

Capital Bold = returns matrix

Notation example:

- Input:

$$\mathbf{x} = \begin{bmatrix} \text{age} \\ \text{mileage} \end{bmatrix}$$



Vector: Structured
or tabular data

- Output:

$$y = [\text{price}]$$



Scalar output

- Model:

$$y = f[\mathbf{x}]$$



Scalar output function
(with vector input)

Model

- Parameters:

ϕ



Parameters always
Greek letters

- Model :

$$\mathbf{y} = \mathbf{f}[\mathbf{x}, \phi]$$

Data Set and Loss function

- Training dataset of I pairs of input/output examples:

$$\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^I$$

Data Set and Loss function

- Training dataset of I pairs of input/output examples:

$$\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^I$$

- **Loss function** or **cost function** measures how bad model is:

$$L \left[\underbrace{\phi, f[\mathbf{x}, \phi]}_{\text{model}}, \underbrace{\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^I}_{\text{train data}} \right]$$

Dataset and Loss function

- Training dataset of I pairs of input/output examples:

$$\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^I$$

- **Loss function** or **cost function** measures how bad model is:

$$L \left[\underbrace{\phi, f[\mathbf{x}, \phi]}_{\text{model}}, \underbrace{\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^I}_{\text{train data}} \right]$$

or for short:

$$L[\phi]$$

Returns a scalar that is smaller when model maps inputs to outputs better

Training

- Loss function:

$$L[\phi]$$

← Returns a scalar that is smaller when model maps inputs to outputs better

- Find the parameters that minimize the loss:

$$\hat{\phi} = \underset{\phi}{\operatorname{argmin}} [L[\phi]]$$

Testing (and evaluating)

- To test the model, run on a separate **test dataset** of input / output pairs
- See how well it **generalizes** to new data



Supervised learning

- Overview
- Notation
 - Model
 - Loss function
 - Training
 - Testing
- 1D Linear regression example
 - Model
 - Loss function
 - Training
 - Testing
- Where are we going?

Example: 1D Linear regression model

- Model:

$$\begin{aligned}y &= f[x, \phi] \\ &= \phi_0 + \phi_1 x\end{aligned}$$

- Parameters

$$\phi = \begin{bmatrix} \phi_0 \\ \phi_1 \end{bmatrix} \begin{array}{l} \longleftarrow \text{y-offset} \\ \longleftarrow \text{slope} \end{array}$$

Example: 1D Linear regression model


- Model:

$$\begin{aligned}y &= f[x, \phi] \\ &= \phi_0 + \phi_1 x\end{aligned}$$

- Parameters

$$\phi = \begin{bmatrix} \phi_0 \\ \phi_1 \end{bmatrix}$$

← y-offset
← slope



Example: 1D Linear regression model


- Model:

$$\begin{aligned}y &= f[x, \phi] \\ &= \phi_0 + \phi_1 x\end{aligned}$$

- Parameters

$$\phi = \begin{bmatrix} \phi_0 \\ \phi_1 \end{bmatrix}$$

← y-offset
← slope



Example: 1D Linear regression model

- Model:


$$\begin{aligned}y &= f[x, \phi] \\ &= \phi_0 + \phi_1 x\end{aligned}$$

- Parameters


$$\phi = \begin{bmatrix} \phi_0 \\ \phi_1 \end{bmatrix}$$

← y-offset
← slope


[Interactive Figure 2.1](#)



Example: 1D Linear regression training data



Example: 1D Linear regression training data




Loss function:

$$\begin{aligned} L[\phi] &= \sum_{i=1}^I (f[x_i, \phi] - y_i)^2 \\ &= \sum_{i=1}^I (\phi_0 + \phi_1 x_i - y_i)^2 \end{aligned}$$

“Least squares loss function”

Example: 1D Linear regression loss function




Loss function:

$$L[\phi] = \sum_{i=1}^I (f[x_i, \phi] - y_i)^2$$
$$= \sum_{i=1}^I (\phi_0 + \phi_1 x_i - y_i)^2$$

“Least squares loss function”

Example: 1D Linear regression loss function




Loss function:

$$L[\phi] = \sum_{i=1}^I (f[x_i, \phi] - y_i)^2$$
$$= \sum_{i=1}^I (\phi_0 + \phi_1 x_i - y_i)^2$$

“Least squares loss function”

Example: 1D Linear regression loss function




Loss function:

$$L[\phi] = \sum_{i=1}^I (f[x_i, \phi] - y_i)^2$$
$$= \sum_{i=1}^I (\phi_0 + \phi_1 x_i - y_i)^2$$

“Least squares loss function”

[Interactive Figure 2.2](#)

Example: 1D Linear regression loss function


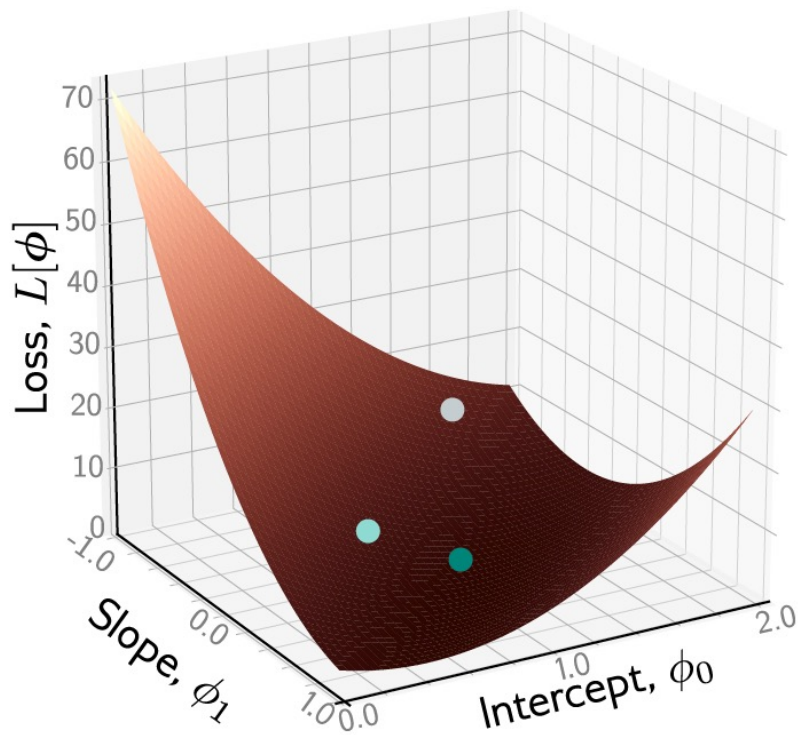


Loss function:



$$\begin{aligned} L[\phi] &= \sum_{i=1}^I (f[x_i, \phi] - y_i)^2 \\ &= \sum_{i=1}^I (\phi_0 + \phi_1 x_i - y_i)^2 \end{aligned}$$

“Least squares loss function”



Example: 1D Linear regression loss function



Example: 1D Linear regression loss function




Example: 1D Linear regression loss function




Example: 1D Linear regression loss function



a)



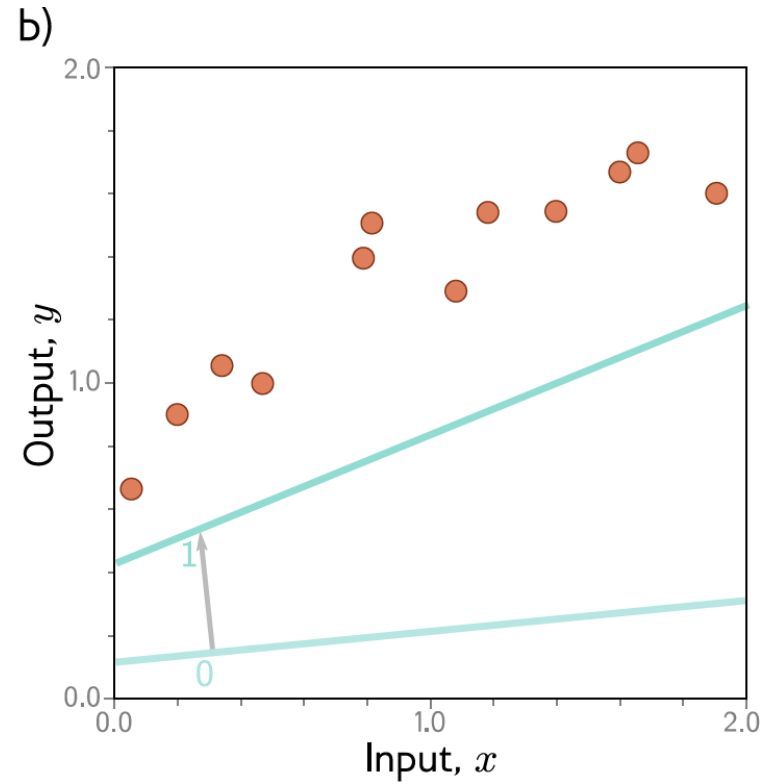

b)





Example: 1D Linear regression training





Example: 1D Linear regression training




Example: 1D Linear regression training



Example: 1D Linear regression training



Example: 1D Linear regression training




This technique is known as **gradient descent**

[Interactive Figure 2.3](#)

Possible objections


- But you can fit the line model in closed form!
 - Yes – but we won't be able to do this for more complex models
- But we could exhaustively try every slope and intercept combo!
 - Yes – but we won't be able to do this when there are a million parameters



Here's a visualization of the loss surface for the 56-layer neural network [VGG-56](#) (from [Visualizing the Loss Landscape of Neural Networks](#) -- <https://losslandscape.com/explorer>)

Example: 1D Linear regression testing

- Test with different set of paired input/output data (**Test Set**)
 - Measure performance
 - Degree to which *Loss* is same as training = **generalization**
- Might not generalize well because
 - Model too simple: **underfitting**
 - Model too complex
 - fits to statistical peculiarities of data
 - this is known as **overfitting**



Piazza Poll

- <https://piazza.com/class/m5v834h9pcatx/post/12>

Supervised learning

- Overview
- Notation
 - Model
 - Loss function
 - Training
 - Testing
- 1D Linear regression example
 - Model
 - Loss function
 - Training
 - Testing
- Where are we going?

Where are we going? Next lectures...

- Shallow neural networks (a more flexible model)
- Deep neural networks (even more flexible with fewer parameters)
- Loss functions (where did least squares come from?)
- How to train neural networks (gradient descent and variants)
- How to measure performance of neural networks (generalization)

Course Project

- Work in teams of 2-3
- Can be application, algorithmic, theoretical or combination thereof
- Project proposal due ~Feb. 16
- Deliverables:
 - Code in GitHub repo
 - Report/paper
 - 3-4 minute video
- More info later, but feel free to brainstorm with me now

Spring 2024 Project Mini Conference



<https://dl4ds.github.io/sp2024/miniconf.html>

Look at Kaggle, Conferences, Workshops, Datasets....

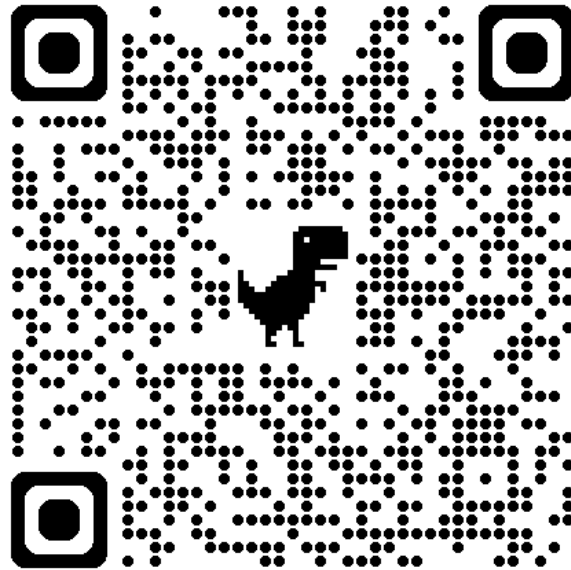
- Application workshops at major conferences can be good sources of ideas. Often times they are associated with new and interesting datasets. Some potential conferences include:
 - [NeurIPS](#),
 - [CVPR](#),
 - [ICML](#),
 - [ICMLA](#),
 - [SPIE](#)
- [Kaggle](#) and other competition websites can be a source of ideas.
- You might find some interesting datasets at [Papers with Code](#)
- Lot of applications are posted on X/Twitter, Reddit, LinkedIn, etc.

Project Grading (45% of course grade)

% of Project Grade	Category	Criteria
20%	Project Report	Conference style paper with complete sections (per template), well written, no typos or formatting issues.
20%	Project Repo/Software	Repo is well documented. Code is reproducible. Top level readme giving project overview, roadmap to directories/files, summary of results.
20%	Final Presentation and Video	Video/presentation is clear and concise, adheres to time limits. Introduces the problem/project, approach, dataset, conclusions, etc.
30%	Individual contribution	Is there clear evidence of project contributions such as commit history or co-authored commits, document revisions. Leave bread crumbs!!
10%	Individual contribution to collaboration and teamwork	Is there indication, e.g. from peer surveys, of collaboration and constructive teamwork?

Project proposal and mid-point check-in will count as homework.

Feedback?



[Link](#)