Machine Learning & Deep Learning

in Python

Monday, Sep 18 | 9:00 AM - 5:00 PM CDT

Workshop Goals:

Understanding when to use
Machine Learning vs **Deep Learning**

Being comfortable with how
Scikit-learn & **PyTorch** work

Schedule

Block 1 (09:00 - 10:30 am)

(1) Introduction to machine learning

(2) The scikit-learn API

30 min break (10:30 - 11:00 am)

Schedule

Block 2 (11:00 am - 12:30 pm)

(3) Data preprocessing in Python

(4) Evaluating & tuning machine learning classifiers

Lunch break (12:30 - 1:30 pm)

Schedule

Block 3 (1:30 - 3:00 pm)

(5) Introduction to deep learning

(6) Understanding PyTorch

(7) Training deep neural networks

30 min break (3:00 - 3:30 pm)

Schedule

Block 4 (3:30 - 5:00 pm)

(8) Accelerating PyTorch model training

(9) Finetuning large language models

(10) Conclusion

Schedule

Bonus

(11) Organizing PyTorch code

(12) Advanced features & techniques

100% Open Source

About the Instructor

Sebastian Raschka

(Most people know me from my books)

https://sebastianraschka.com

Lightning AI

Creators of PyTorch Lightning

(2021-present)

I work on LLM research, open source, education

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Asst. Prof. of Statistics

(2018-2023)

I worked on AI & deep learning research (neural nets for ordinal regression problems, etc.)

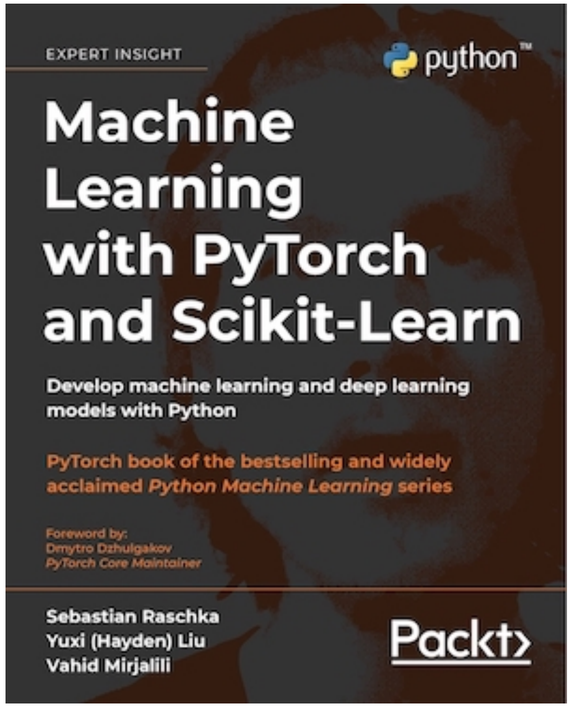

Topics covered in this workshop (but very different presentation)

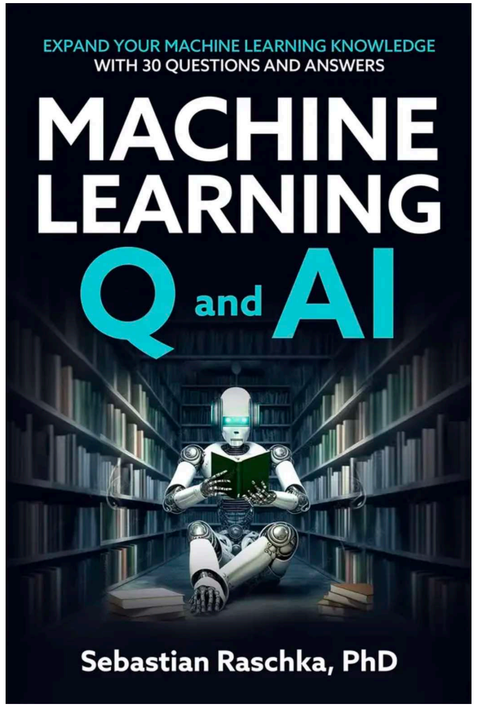

More advanced concepts in ML and AI

https://sebastianraschka.com/books/

Contact

@rasbt

sebastian@lightning.ai

https://sebastianraschka.com

https://lightning.ai

Code & slides

🎮 https://github.com/rasbt/posit2023-python-ml

Let’s Begin!

Information from the organizers

WiFi:
Network: Posit Conf 2023
Password: conf2023

There are gender-neutral bathrooms located are among the Grand Suite Bathrooms.

There are two meditation/prayer rooms: Grand Suite 2A and Grand Suite 2B.
Open Sunday - Tuesday 7:30 a.m. - 7:00 p.m., Wednesday 8:00 a.m. - 6:00 p.m.

The lactation room is located in Grand Suite 1.
Open Sunday - Tuesday 7:30 a.m. - 7:00 p.m., Wednesday 8:00 a.m. - 6:00 p.m.

Participants who do not wish to be photographed have red lanyards; please note everyone's lanyard colors before taking a photo and respect their choices.

The Code of Conduct and COVID policies can be found at https://posit.co/code-of-conduct/. Please review them carefully. You can report Code of Conduct violations in person, by email, or by phone. Please see the policy linked above for contact information.

Schedule

Block 1 (09:00 - 10:30 am)

(1) Introduction to machine learning

(2) Python's scientific computing stack

(3) The scikit-learn API

(4) Data preprocessing in Python

30 min break (10:30 - 11:00 am)

Where do we see machine learning being used in the real world?

Unlocking our phones

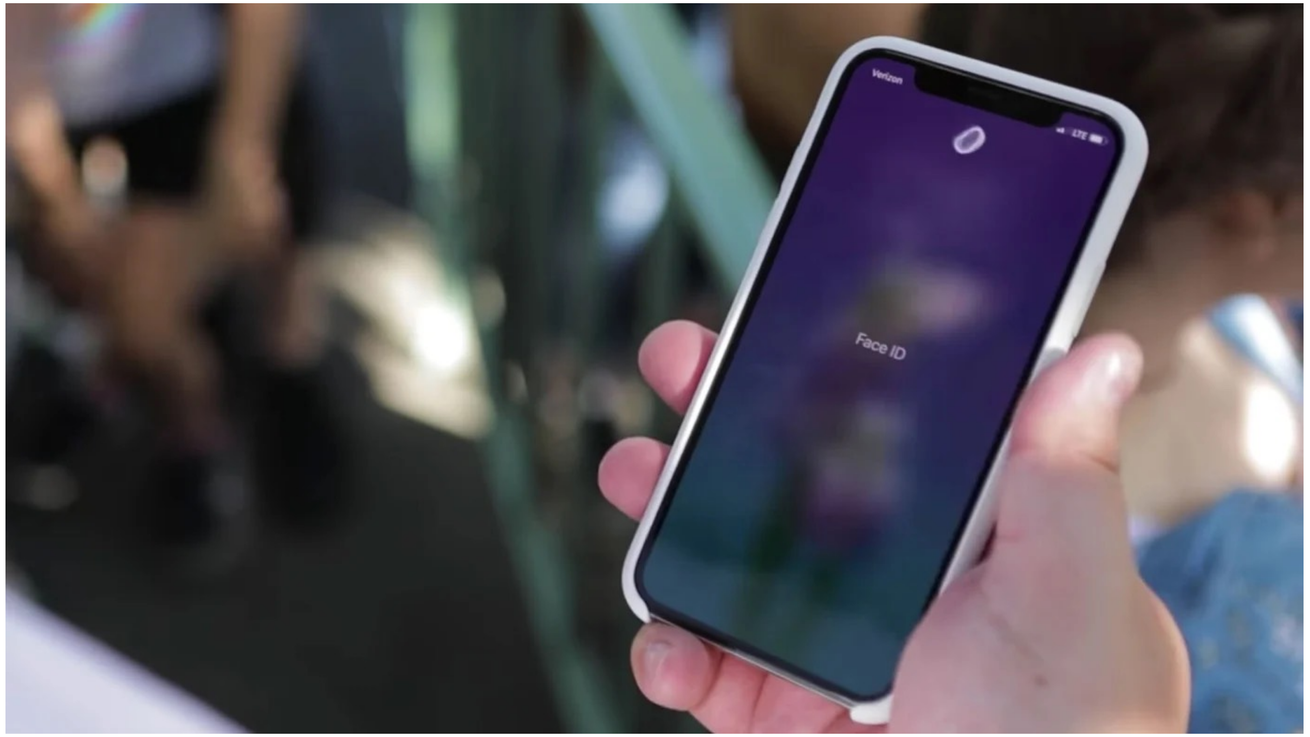

Source: https://techcrunch.com/2017/11/04/a-closer-look-at-the-capabilities-and-risks-of-iphone-x-face-mapping/

Detecting road signs

Source: https://www.researchgate.net/figure/Example-of-autonomous-vehicle-with-speed-sign-recognition-system_fig1_338455384

Recommending movies

Source: https://www.pinterest.com/pin/684758318333854483/

Machine learning is almost
everywhere!

So, how does it work?

Machine learning lets computers learn from data

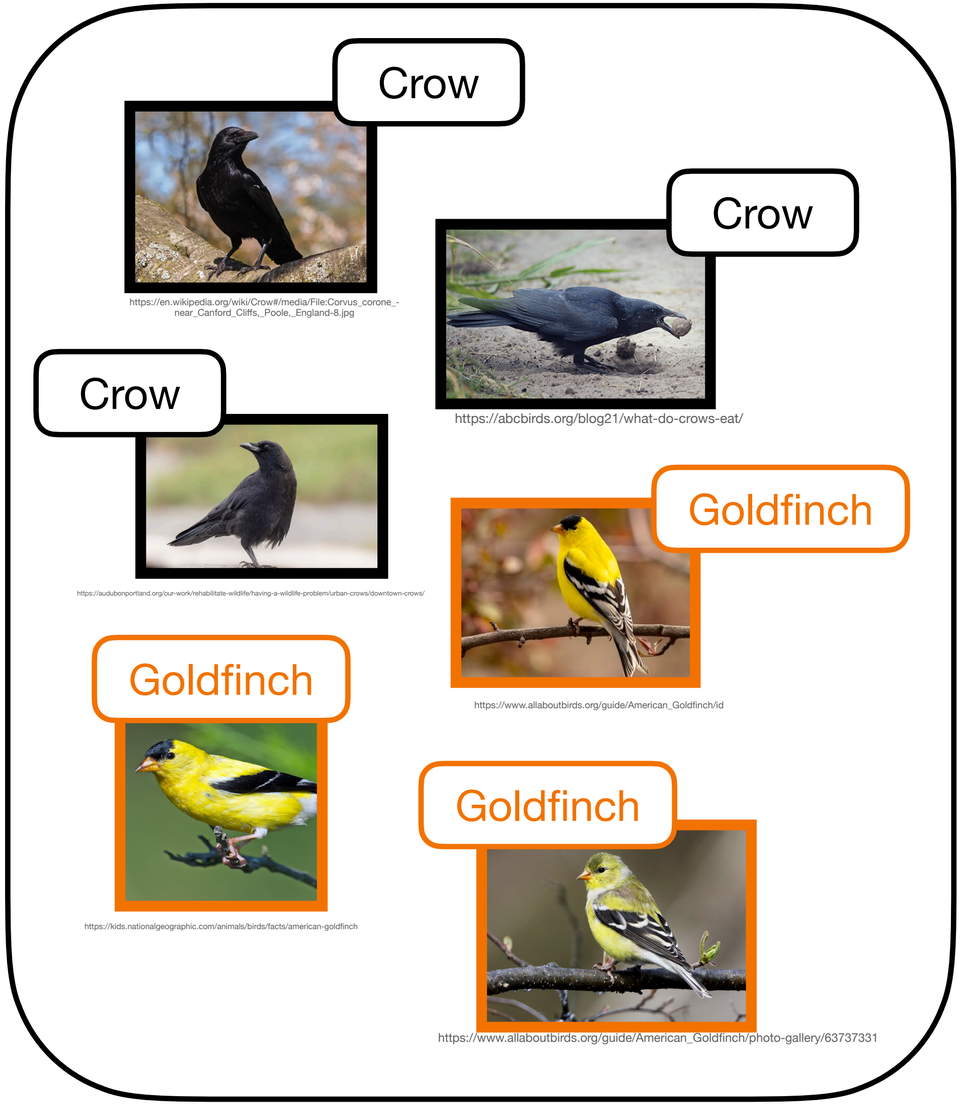

Labeled examples

Machine learning algorithm

Bird classifier

What does this have to do with deep learning?

Machine learning started as a subfield of AI

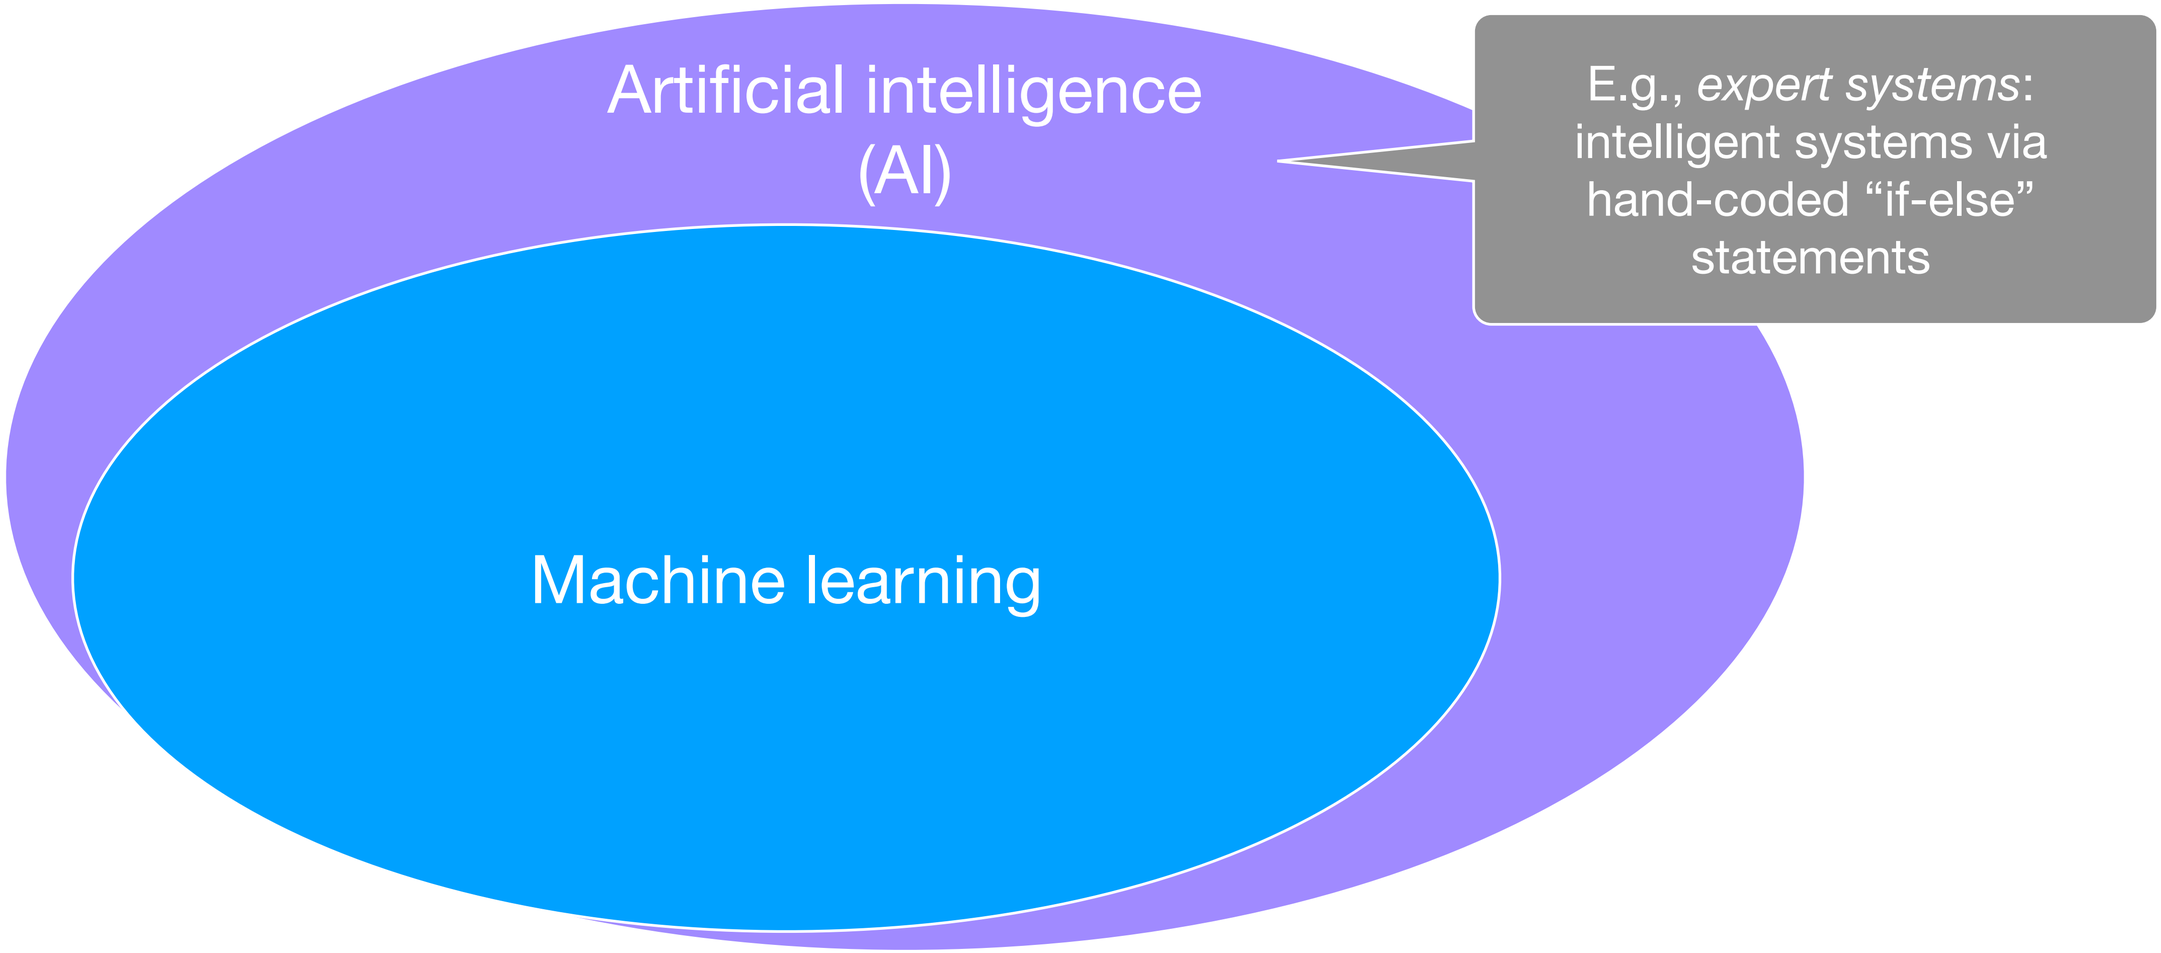

Deep Learning is a Subset of Machine Learning

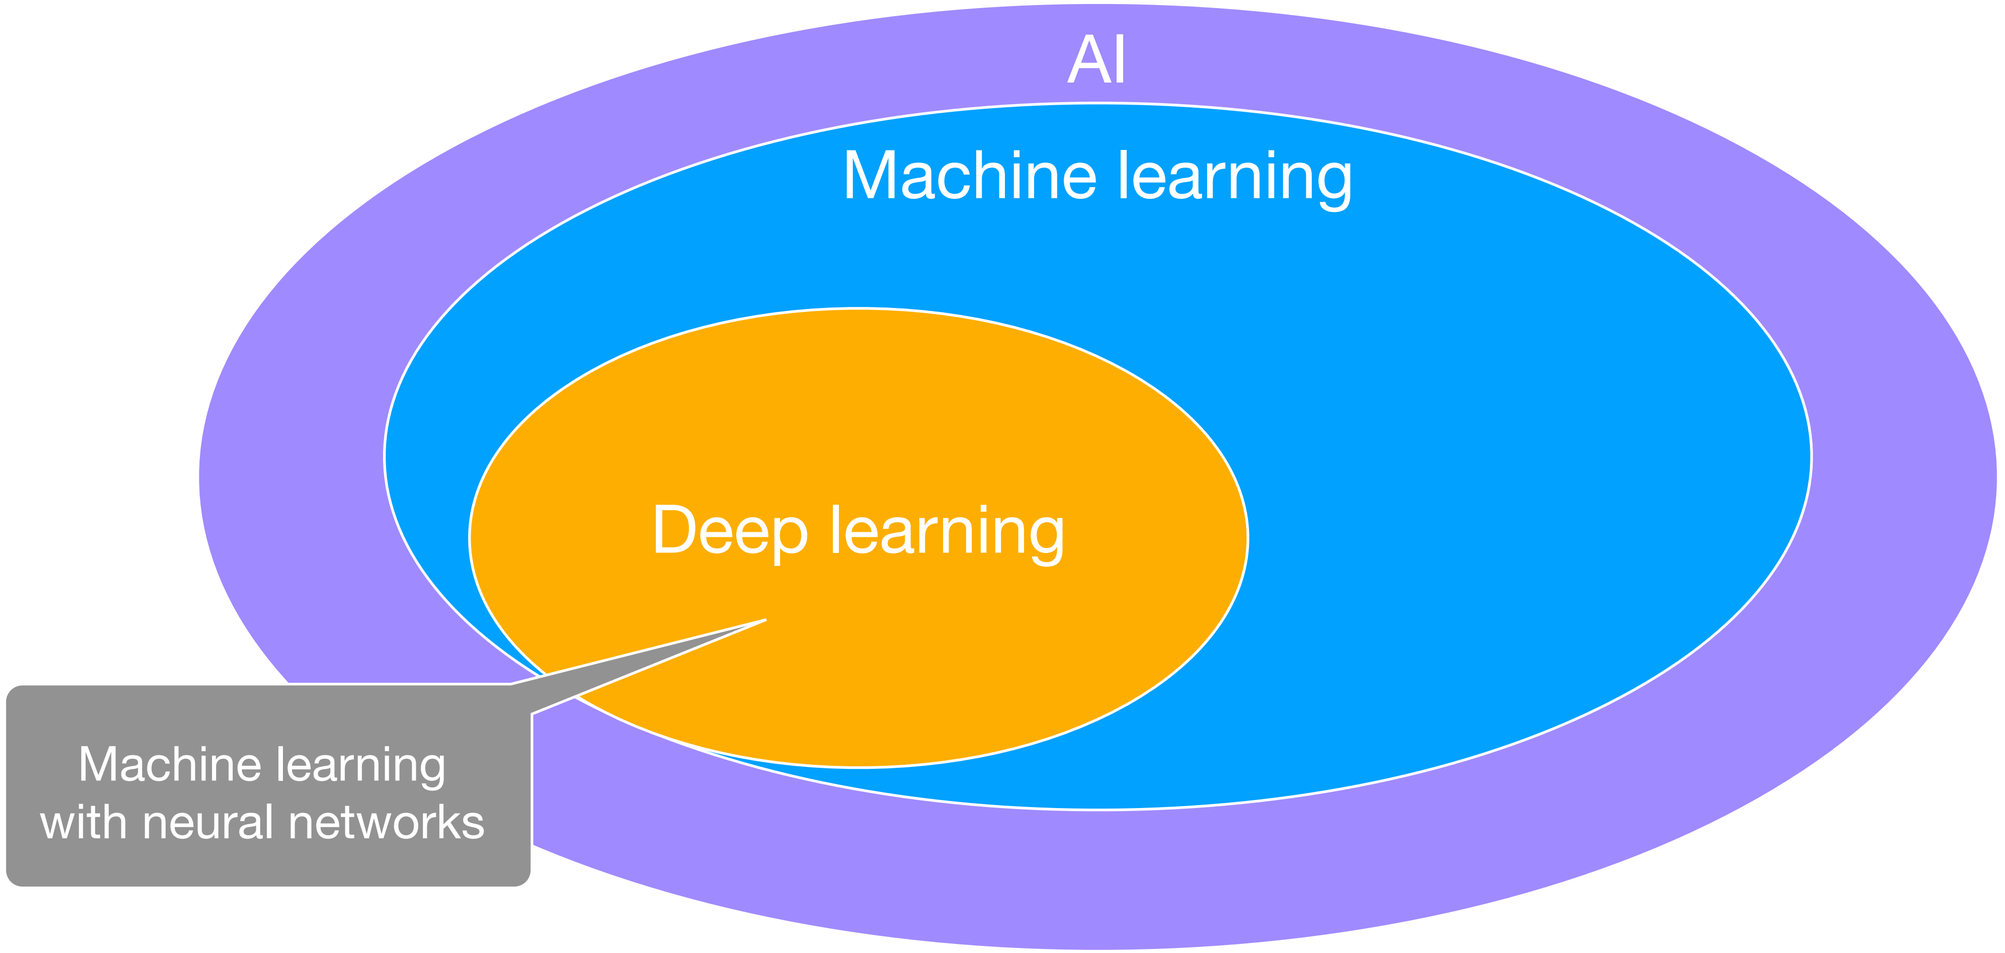

When do we use deep learning vs regular machine learning?

Deep learning

large datasets
in “unstructured” form
(e.g., images and text)

https://en.wikipedia.org/wiki/Crow#/media/File:Corvus_corone_-near_Canford_Cliffs,_Poole,_England-8.jpg

https://abcbirds.org/blog21/what-do-crows-eat/

https://audubonportland.org/our-work/rehabilitate-wildlife/having-a-wildlife-problem/urban-crows/downtown-crows/

https://www.allaboutbirds.org/guide/American_Goldfinch/id

https://kids.nationalgeographic.com/animals/birds/facts/american-goldfinch

https://www.allaboutbirds.org/guide/American_Goldfinch/photo-gallery/63737331

Deep learning

large datasets
in “unstructured” form
(e.g., images and text)

https://en.wikipedia.org/wiki/Crow#/media/File:Corvus_corone_-near_Canford_Cliffs,_Poole,_England-8.jpg

https://abcbirds.org/blog21/what-do-crows-eat/

https://audubonportland.org/our-work/rehabilitate-wildlife/having-a-wildlife-problem/urban-crows/downtown-crows/

https://www.allaboutbirds.org/guide/American_Goldfinch/id

https://kids.nationalgeographic.com/animals/birds/facts/american-goldfinch

https://www.allaboutbirds.org/guide/American_Goldfinch/photo-gallery/63737331

Regular machine learning

Datasets in “structured”
form like tables

	Beak length	Wing span	Primary color	...
Bird 1	3 cm	43 cm	black	...
Bird 2	0.5 cm	19 cm	yellow	...
Bird 3	5 cm	48 cm	black	...
Bird 4	4 cm	45 cm	black	...
...	...	...	...	...

Regular machine learning

"Unstructured" data

https://en.wikipedia.org/wiki/Crow#/media/File:Corvus_corone_-near_Canford_Cliffs,_Poole,_England-8.jpg

https://abcbirds.org/blog21/what-do-crows-eat/

https://audubonportland.org/our-work/rehabilitate-wildlife/having-a-wildlife-problem/urban-crows/downtown-crows/

https://www.allaboutbirds.org/guide/American_Goldfinch/id

https://kids.nationalgeographic.com/animals/birds/facts/american-goldfinch

https://www.allaboutbirds.org/guide/American_Goldfinch/photo-gallery/63737331

Manual feature extraction

Datasets in "structured" form like tables

	Beak length	Wing span	Primary color	...
Bird 1	3 cm	43 cm	black	...
Bird 2	0.5 cm	19 cm	yellow	...
Bird 3	5 cm	48 cm	black	...
Bird 4	4 cm	45 cm	black	...
...	...	...	...	...

How can we use machine learning?
common application areas

Common applications areas

1. Making predictions
2. Compressing data
3. Generating new data
4. Learning a series of actions

1 Making predictions

Is this email spam or not?

https://computing.which.co.uk/hc/en-gb/articles/115002560545-Common-email-mistakes-that-lead-to-spam

A label prediction task

1 Making predictions

Where are the cars located?

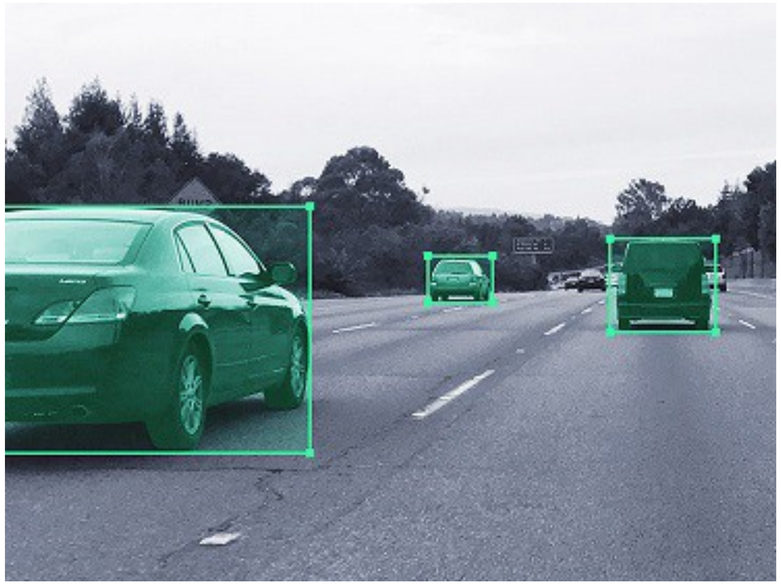

https://medium.com/@harisriaz/building-a-state-of-the-art-autonomous-vehicle-program-in-java-and-python-using-tiny-yolo-3b67301e8bac

Predicting the bounding boxes

1 Making predictions

What is the license plate code?

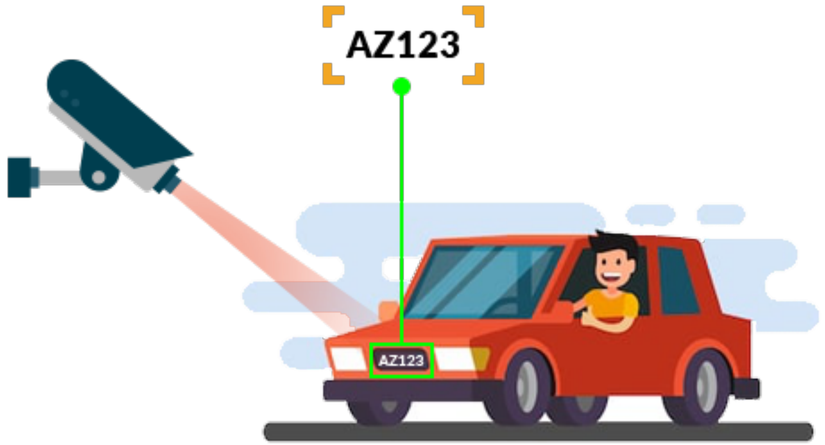

https://www.innominds.com/connected-devices-and-iot/numix-video-analytics-solutions/anpr

Recognizing or predicting letters and digits

Supervised learning workflow

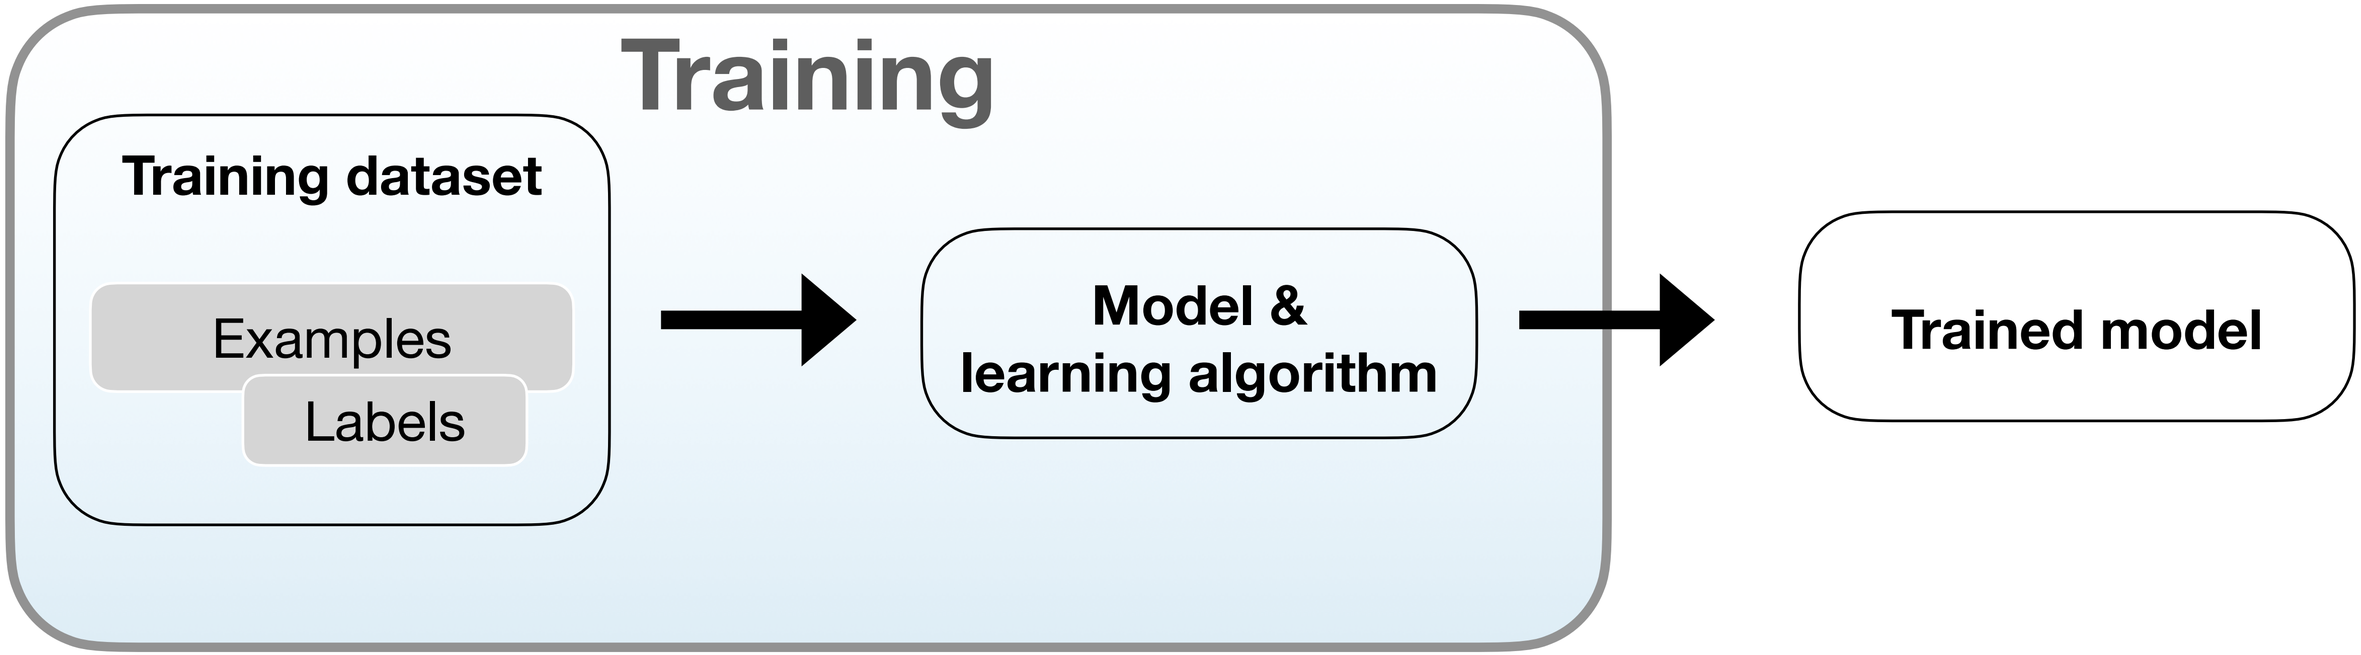

Supervised Learning Workflow

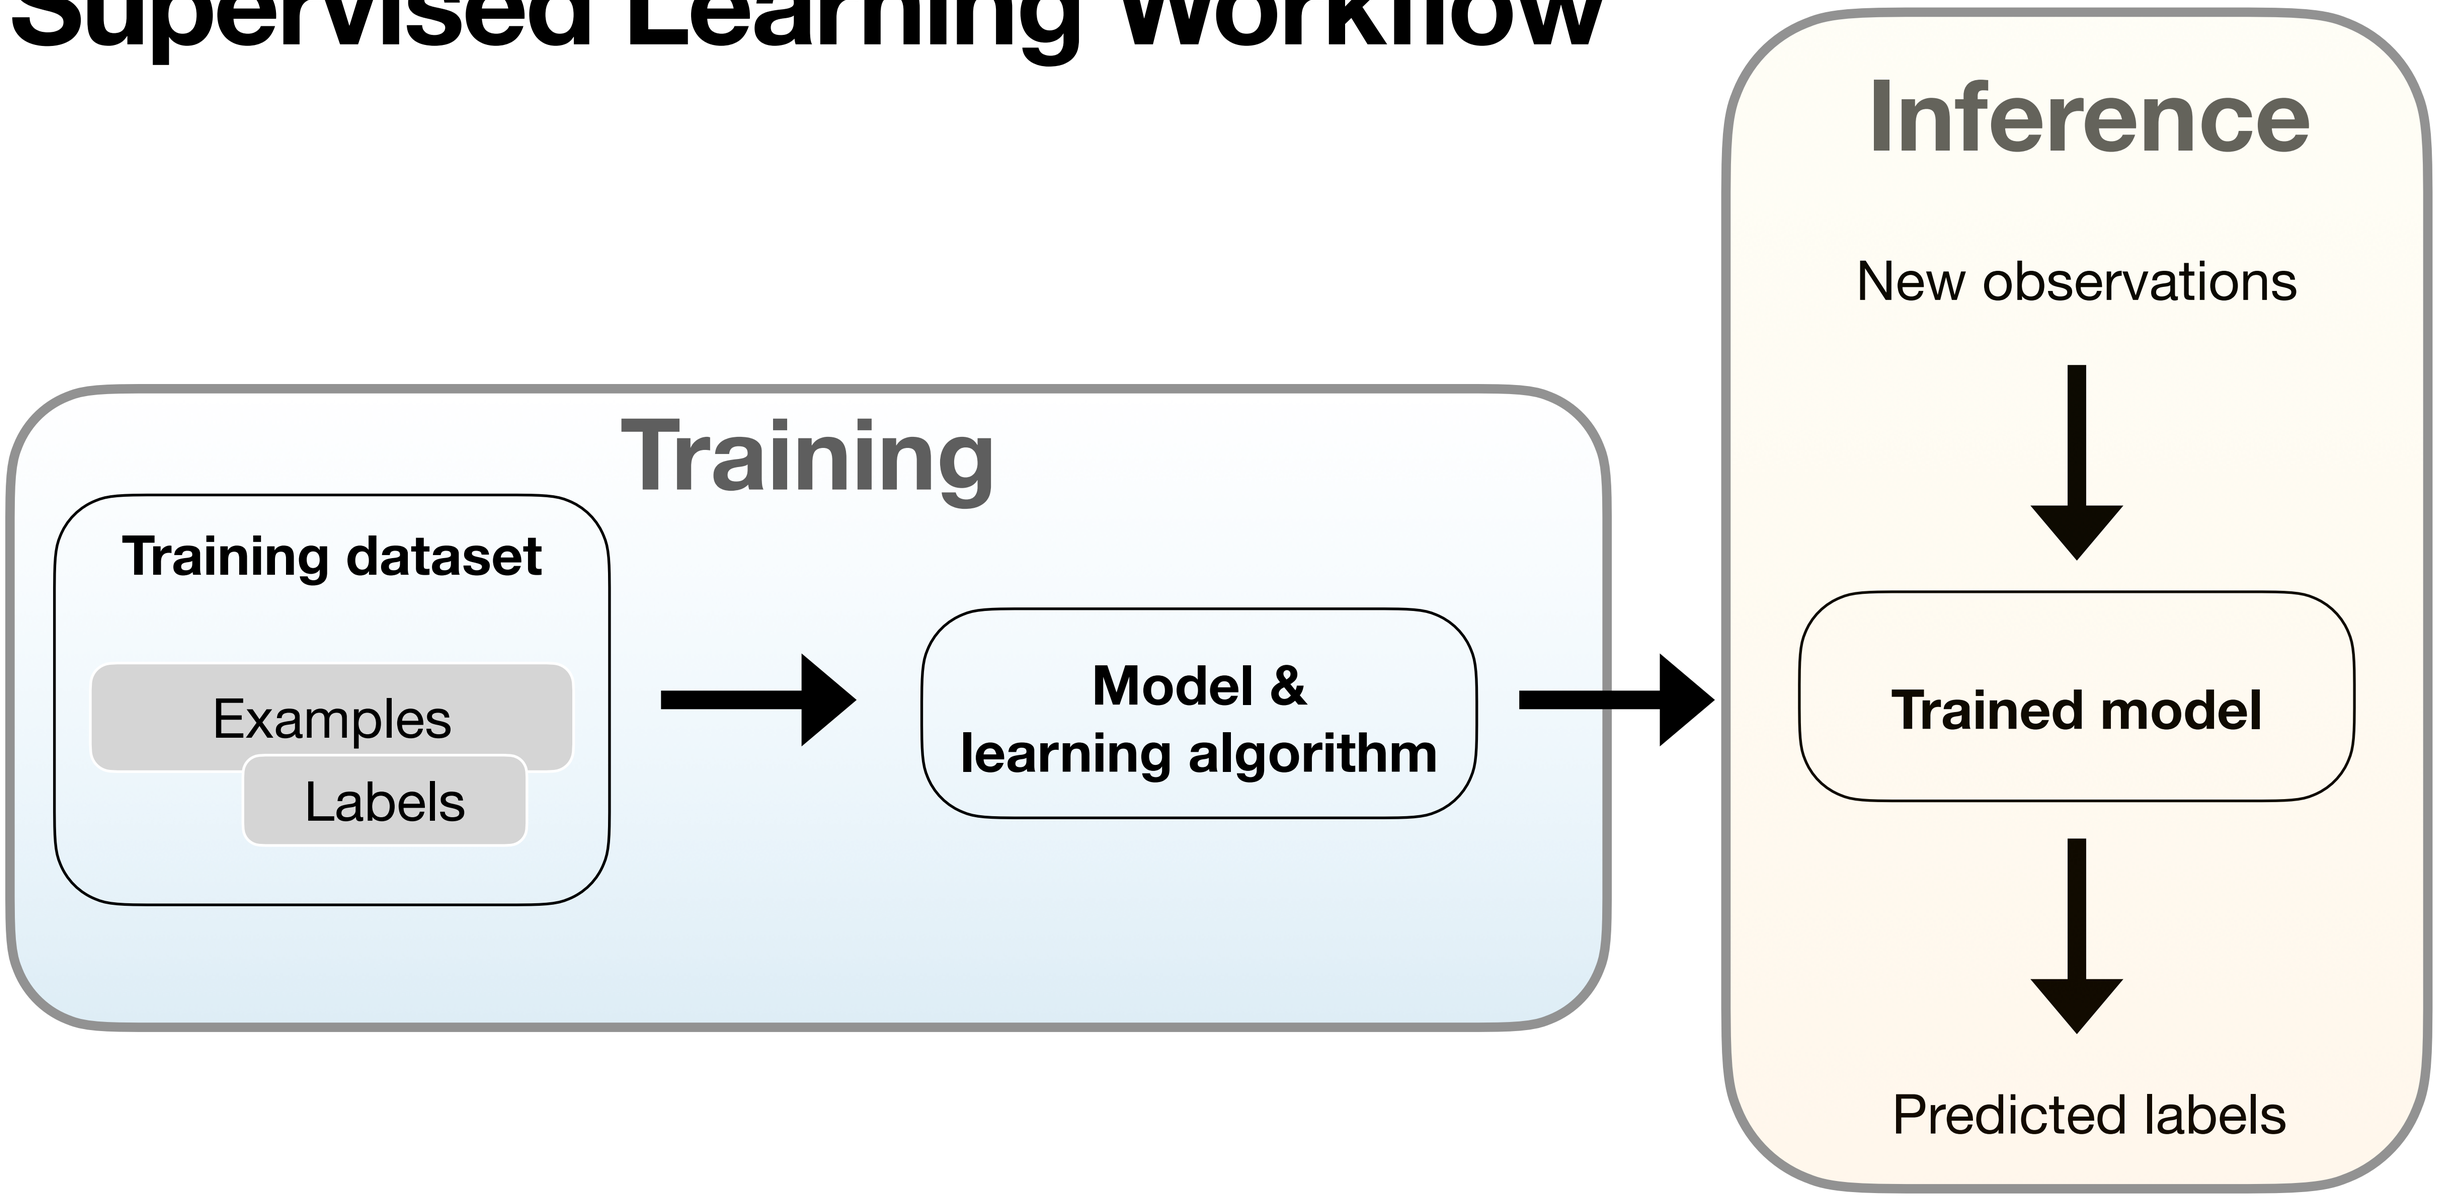

Exercise:

Download and set up environment
(if not done already)

🎮 https://github.com/rasbt/posit2023-python-ml