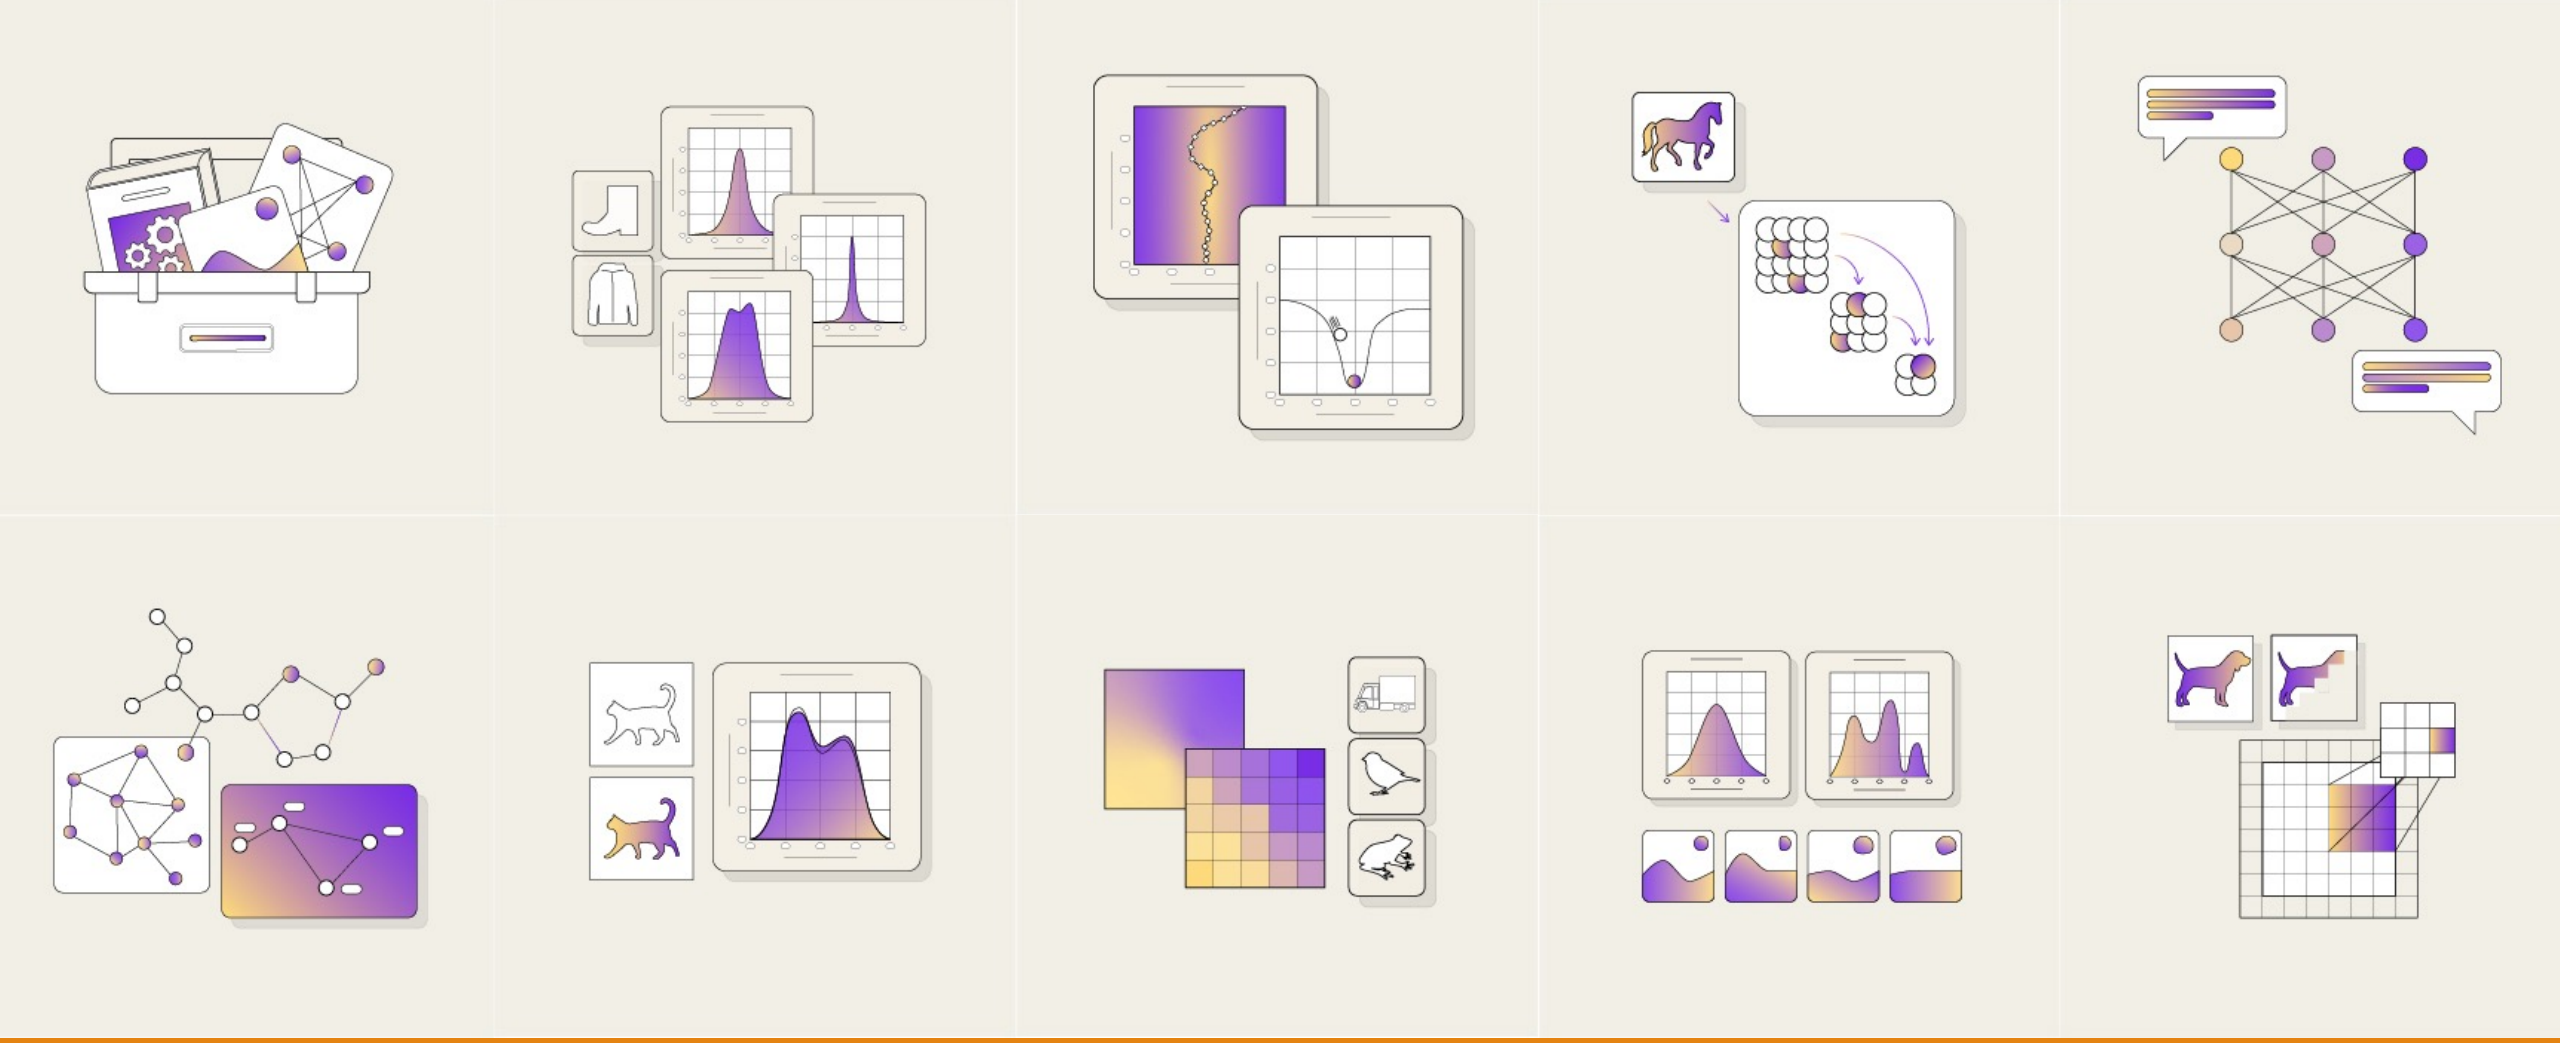

Introduction to the Tutorial Sessions

Phillip Lippe

About Me - Phillip

- PhD student in the QUVA lab, working on:
  - Intersection between Causality and DL
  - Causal Representation Learning
- Alumni MSc AI at UvA
- Co-organizing the Deep Learning 1 since 2020
  - Teaching the implementation side of DL

Schedule

- Werkcollege sessions split into:
  - Tuesday 17.00-18.00 – Tutorial session on one of the notebooks
  - Tuesday 18.00-19.00 – TA hour in one large group; Q&A
  - Fridays – TA hours in smaller groups; Q&A for assignments, lectures, etc.
- TA hours by respective TAs of your group
- Live tutorial sessions by Phillip

Monday	Tuesday	Wednesday	Thursday	Friday
	15-17 Hoorcollege dr. Y.M. Asano SP L1.02			9-11 Hoorcollege dr. Y.M. Asano SP L1.02
	17-19 Werkcollege Phillip Lippe MSc SP L1.02			11-13 Werkcollege E: SP G0.18A F: SP D1.115
				13-15 Werkcollege A: SP D1.114 B: SP D1.115
				15-17 Werkcollege C: SP G0.18A D: SP G0.18B G: SP D1.113

Assignments

- 3 Practicals on different topics aligned with the lectures
  - Published on ans-delft, code submission via Canvas
  - Repository for assignment code: https://github.com/uvadlc/uvadlc_practicals_2022
- Practical 1: MLPs and Backpropagation
- Practical 2: CNNs, Transformers and Graph NNs
- Practical 3: Generative Models
- Compute Resources:
  - Lisa cluster
  - Google Colab

Assignments and Tutorials

- But there is so much more to explore… ⇒ Tutorial notebooks
  - Teaching the concepts from the lectures from an implementation perspective
  - Get familiar with PyTorch and PyTorch Lightning
  - Test your understanding by playing around with the models
- The notebooks are ungraded, but relevant for assignments and exam

Tutorial Notebooks

- Jupyter notebooks where we implement and train our own models
- Ready to be run on Google Colab or locally by yourself
- Accessible via https://uvadlc-notebooks.readthedocs.io/en/latest/
- Integrated in PyTorch Lightnings' documentation and Google's Dev library

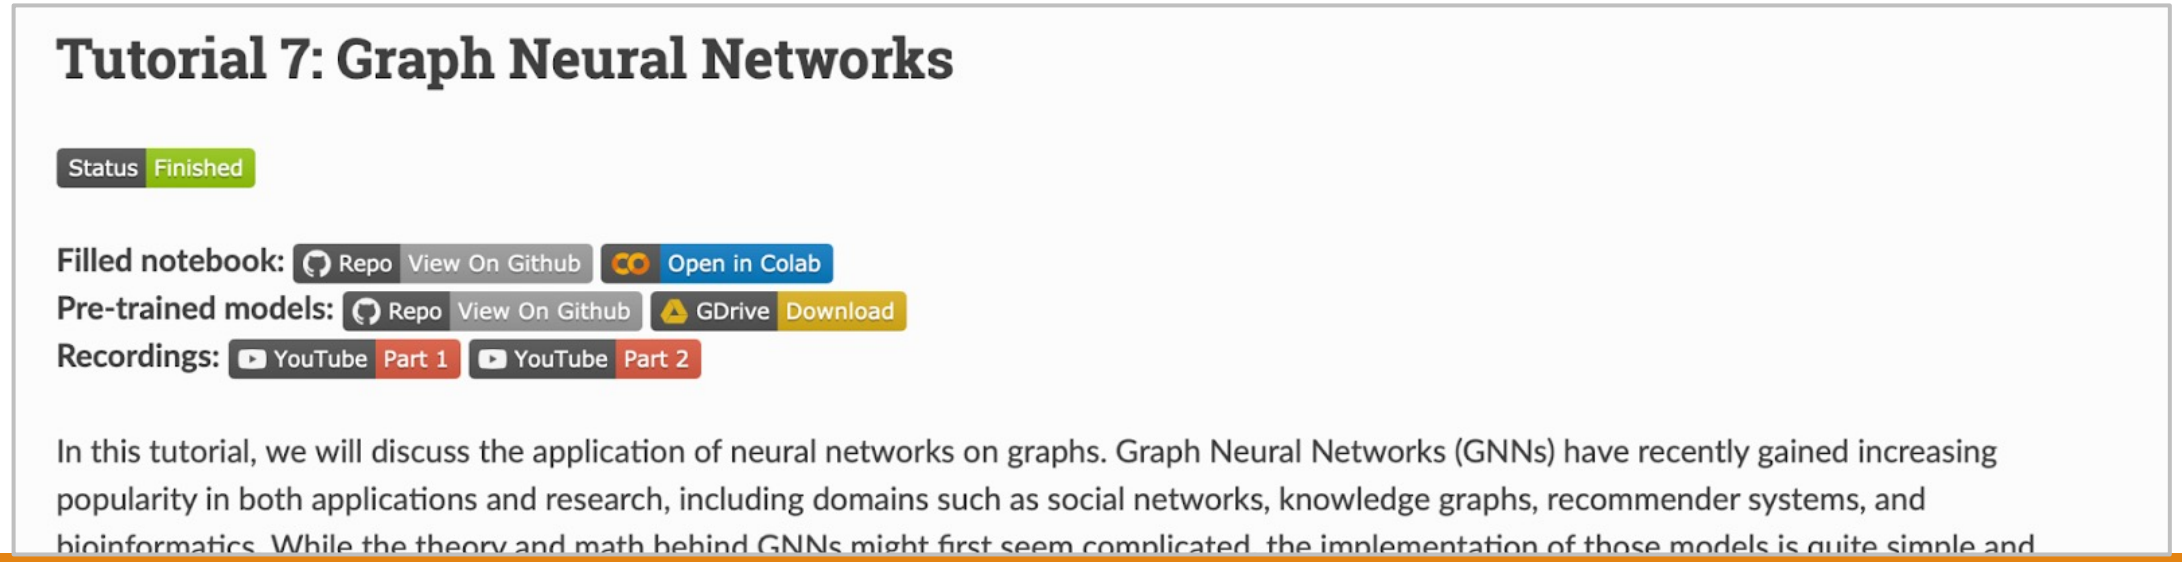

Tutorial Notebooks

- How to interact/learn from the notebooks:
  - Reading
  - Running code yourself
  - Video recordings
  - Live tutorial sessions on Tuesdays

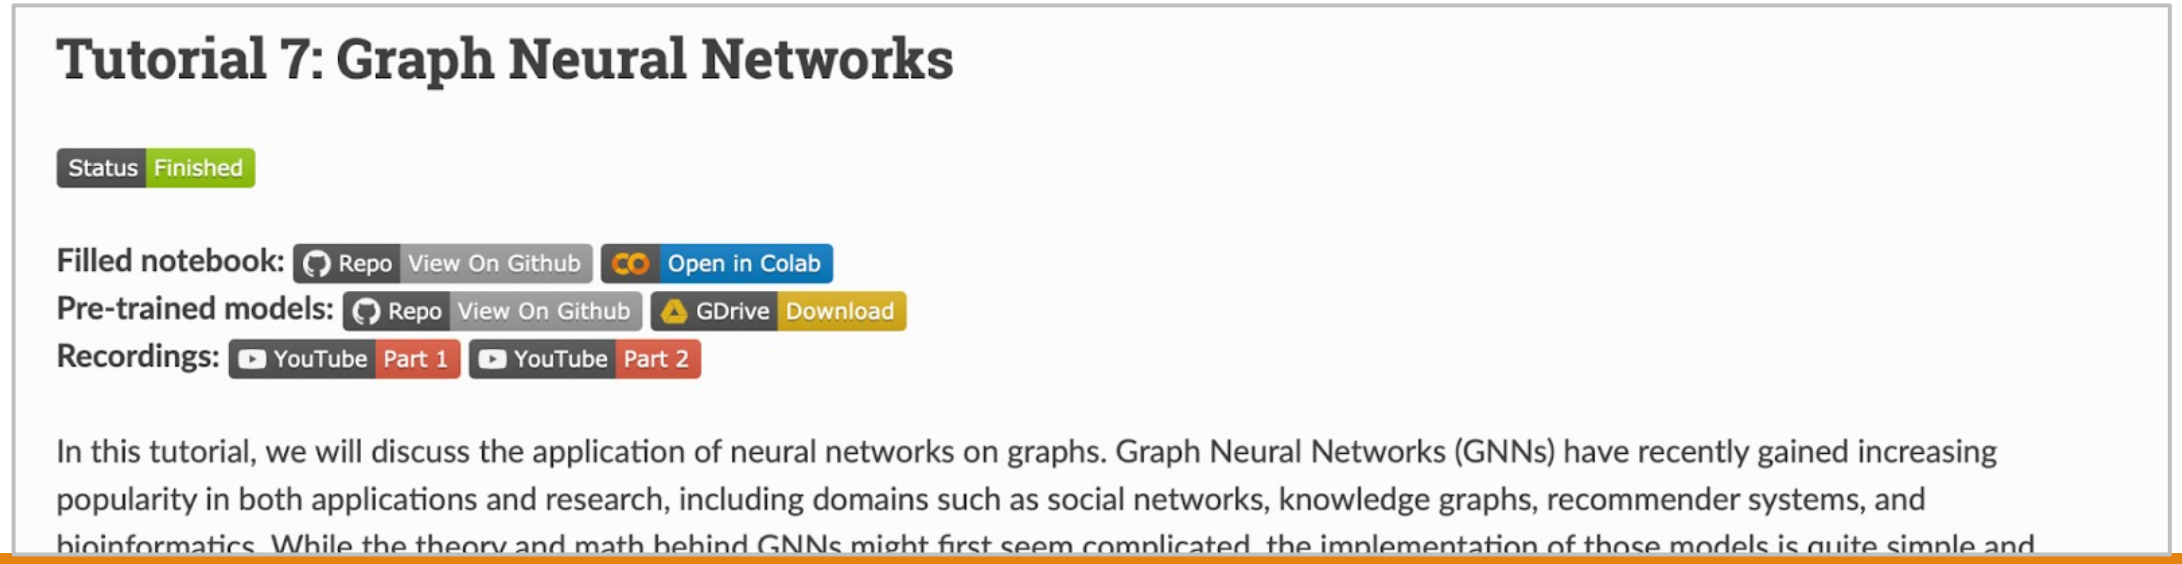

Why PyTorch?

- There is a great variety of other DL frameworks out there

TensorFlow
- Production-level code
- Used in companies
- Large community

PyTorch

PyTorch
- Most popular in research
- Simple, easy to debug
- Large community

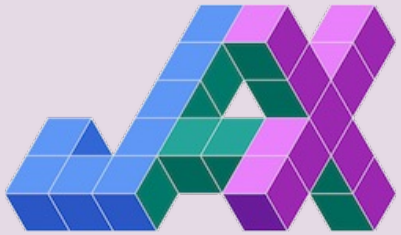

JAX
- Function-oriented
- JIT-compiled, very fast
- Currently on the rise

In this course:
- Main framework: PyTorch
- All notebooks provided in both PyTorch and JAX
- If you already know PyTorch, give JAX a try!

Tutorial schedule

Week 1
Introduction to PyTorch

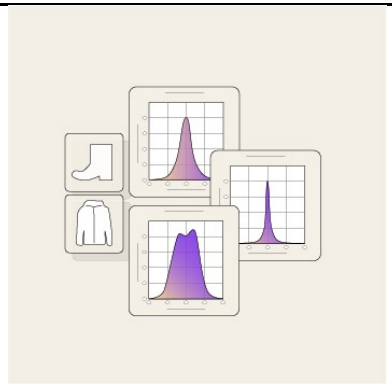

Week 2
Activation Functions

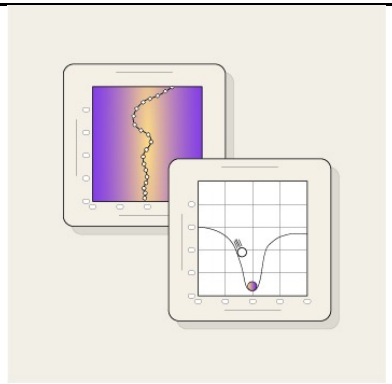

Week 3
Optimization & Initialization

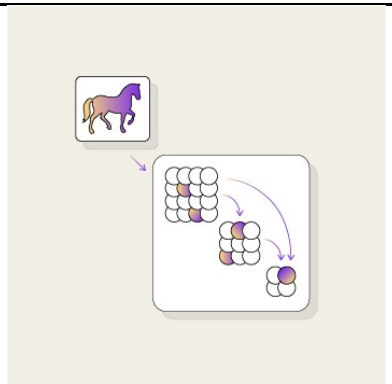

Week 4
Inception, ResNet & DenseNet

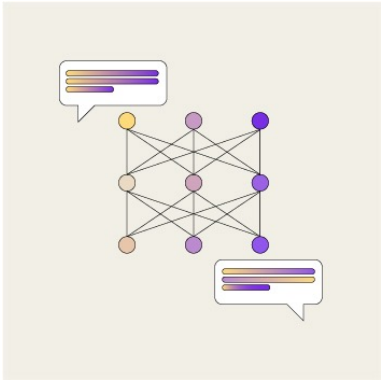

Week 5
Transformers & Attention

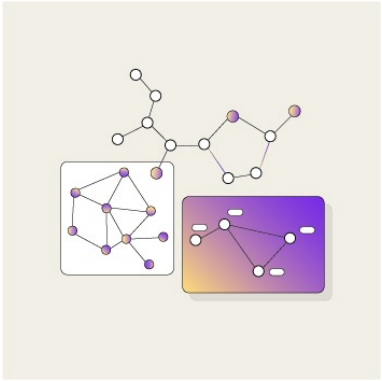

Week 6
Graph Neural Networks

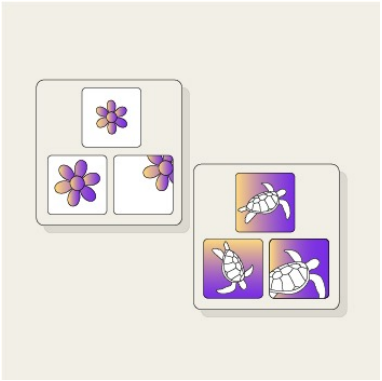

Week 7
Self-Supervised (and Causal) Representation Learning

Figure credit:
PyTorch Lightning team

Tutorial-Lecture Alignment

Week	Lecture - Tuesday	Tutorial - Tuesday	Lecture - Friday
Week 1	Introduction to DL	Introduction to PyTorch	Modular Learning
Week 2	DL Optimizations 1	Activation Functions	DL Optimizations 2
Week 3	Convolutional NNs	Optimization & Initialization	Modern ConvNets
Week 4	Transformers	Inception, ResNet and DenseNet	Graph Neural Networks
Week 5	Generative Modelling	Transformers and Attention	Deep Variational Inference
Week 6	Neural Modelling of 3D	Graph Neural Networks	Deep Learning for Physics
Week 7	Self-Supervised Learning 1	Self-Supervised (and Causal) Representation Learning	Self-Supervised Learning 2

Plan for today

- First look at notebooks
- Tutorial 2: Introduction to PyTorch
- Afterwards: Introduction to Assignment 1

- Friday: Working with the Lisa cluster