

Introduction to Deep Learning (I2DL)

Tutorial 1: Organization

The Team

Lecturer



Prof. Matthias
Niessner

PhDs



Manuel Dahnert



Yujin Chen



Haoxuan Li

Student Tutors



Dan
Halperin



Julius
Horsch



Maximilian
Gollwitzer



Oleg
Magnes



Philipp
Binder



Luca
Fehling-Schuh

Today's Outline

- Organization
 - Lectures, tutorials and schedule
 - Exam
 - Communication

- Exercises
 - General overview
 - Exercise 1



3. Numpy Introduction

Numpy is a high performance framework for operations on multi dimensional arrays. life as a Data Scientist in general.

```
In [12]: # Create a simple (2, 3) numpy array
# [[ 1  2  3]
# [ 4  5  6]]
a = np.array([[1,2,3], [4,5,6]])

# Print the data type of this array
print(a.dtype)

# print the shape of this array
print(a.shape)

int64
(2, 3)
```

Organization


Lecture Material

- Lectures

- This lecture is held virtually, offline
- New lecture recordings and slides every **Tuesday**, pre-recorded
- Available on our webpage <https://niessner.github.io/I2DL/>
- Recommendation: watch in a weekly fashion

- Exercises (Tutorial session + Homework)

- Recordings + slides + homework uploaded **Thursdays 10 AM**
- Homework due following Wednesday 15:59



Tentative Schedule

Introduction



- 0: Organization
- 1: What is Deep Learning?
- 2: Machine Learning Basics

Neural Networks


- 3: Introduction
- 4: Optimization
- 5: Stochastic Gradient Descent
- 6: Training Neural Networks
- 7: Training Neural Networks
- 8: Training Neural Networks

Advanced Architectures

- 9: Introduction to CNNs
- 10: More on CNNs
- 11: Recurrent Neural Networks
- 12: Advanced Deep Learning Topics



Website & TUM-
online update
soon™



Exam FAQ

Us waiting for the exam date

- Date
Who knows, probably beginning / mid of February
- Place
Probably on-site like the old usual exams, no remote exams or alternative available
- Retake?
None this semester, can participate in the exam of next semester's class



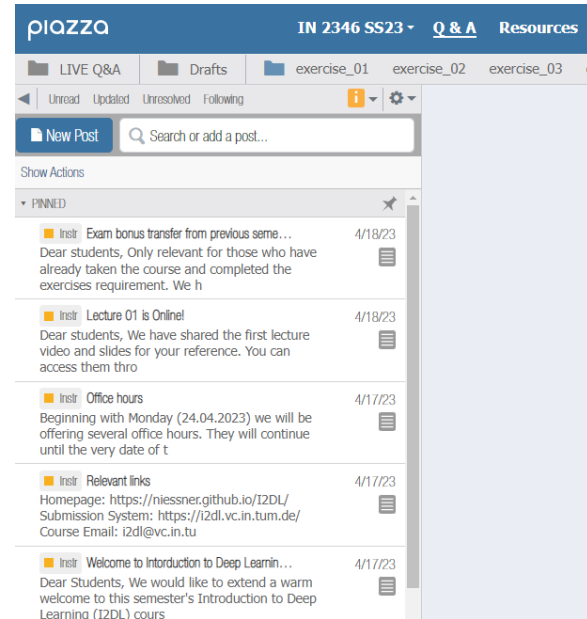
Exam FAQ

- Do we get access to previous exams?
Yes, including a mock exam etc. around week 10
- Do we have to code during the exam?
No, mostly short questions and some multiple choice
- What about my bonus?
Bonus from any previous iteration of I2DL will be transferred to this class and all future exam attempts (regardless of I2DL lecturer)
- For whom?
TUM/LMU students (who will get a certificate)

Availability

- Website ([link](#))
 - Public
 - Will contain all videos, exercise material and solutions

- Piazza ([link](#))
 - Only for LMU/TUM students
 - Links and dates of office hours
 - Exam related information
 - Discussions about content
 - All other communication




Availability - Office Hours

- Daily office hours offered by the tutors
- Via zoom
Link and times - Piazza
- Casual, ask anything, be considerate
- Starting Monday



Availability – contact the staff

- Use piazza for all communication
 - Send private note/question to instructors for individual issues




The screenshot shows the Piazza post creation interface. Under "Post Type", there are three radio buttons: "Question" (selected, with the subtext "if you need an answer"), "Note" (with the subtext "if you *don't* need an answer"), and "Poll/In-Class Response" (with the subtext "if you need a vote"). Under "Post To", there are two radio buttons: "Entire Class" and "Individual Student(s) / Instructor(s)" (selected). Below the "Post To" section, there is a text input field with the placeholder "Enter one or more names..." and a dropdown arrow. To the right of the input field, there is a note: "Select 'Instructors' to include all instructors." The "Individual Student(s) / Instructor(s)" option and the word "Instructors" in the note are highlighted in yellow.

- If you cannot use piazza to contact us:
 - i2dl@vc.in.tum.de
 - Please only use it if you cannot access piazza

Piazza

- Be courteous
- Participation in discussions is correlated with success in the exam
- Mark your followup questions as (un-)resolved accordingly if attention is needed by the instructors – resolve otherwise.



LMU Students & TUMOnline Issues

- Check out the google [form](#) we linked on our website.

I2DL WS23/24 - Introduction to Deep Learning (IN2346) - Registration for External Students

Dear students,

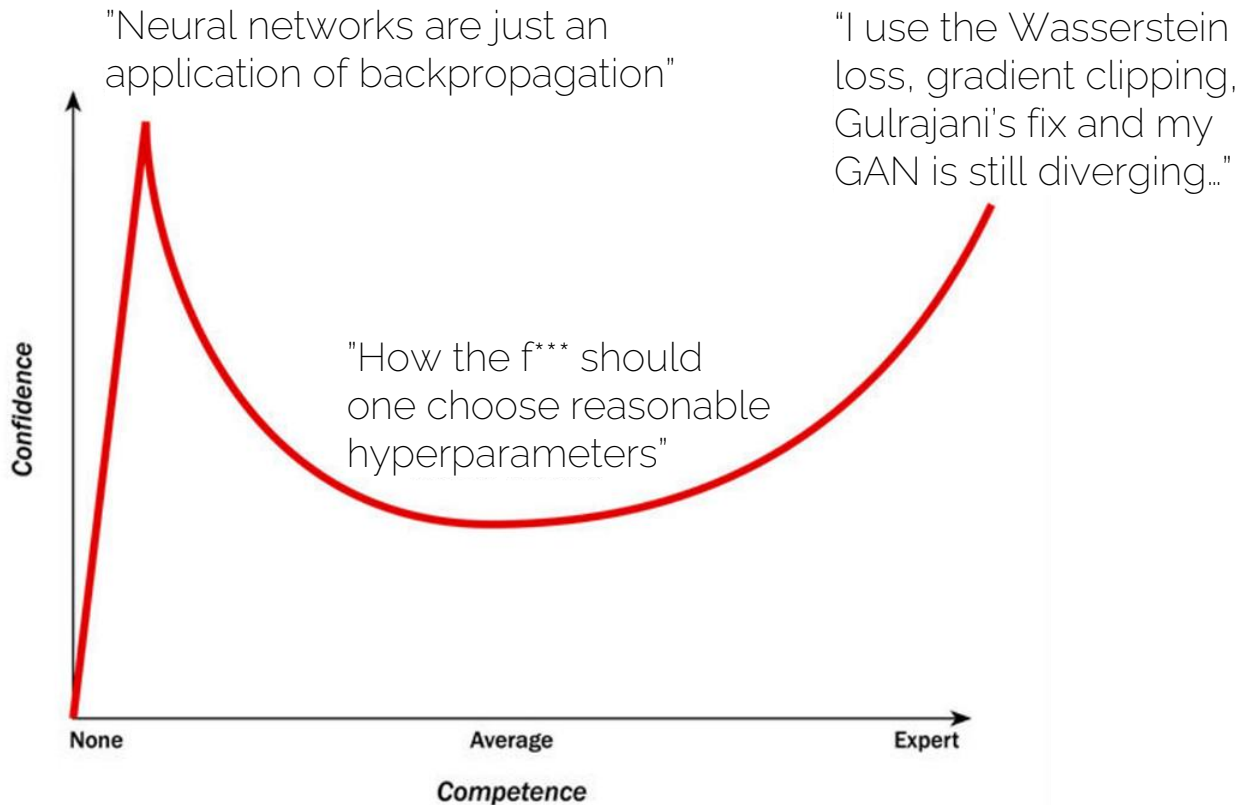
this form is intended for students with TUM-affiliation (e.g., LMU students, TUM students, PhD students from TUM/LMU) who cannot register for courses, etc, who want to participate in our lecture but can not sign up through the TUM course registration system. By filling this form, you will be able to participate in Piazza, as well as the exercises and submit models to our submission portal.

Non TUM/LMU Students

- You can:
 - Participate at lectures as well as exercises
 - Exercises have evaluations in notebooks
- You can't:
 - Have access to our submission system
 - Take part in online discussions on Piazza or attend office hours
 - Participate in the exam

General Exercises & Bonus

From our own experience



Exercises – Tentative Schedule

Lectures

Introduction

0: Organization

1: What is Deep Learning?

2: Machine Learning Basics

Neural Networks

3: Introduction

4: Optimization

5: Stochastic Gradient Descent

6: Training Neural Networks

7: Training Neural Networks

8: Training Neural Networks

Advanced Architectures

9: Introduction to CNNs

10: More on CNNs

11: Sequence Models

12: Advanced Deep Learning Topics

Exercises

Exercise 01: Organization

Exercise 02: Math Recap

Exercise 03: Datasets

Exercise 04: Linear Regression

Exercise 05: Neural Networks

Exercise 06: Hyperparameter Tuning

Exercise 07: Introduction to Pytorch

Exercise 08: Autoencoder


Exercise 09: Convolutional Networks

Exercise 10: Semantic Segmentation

Exercise 11: Sequence Models

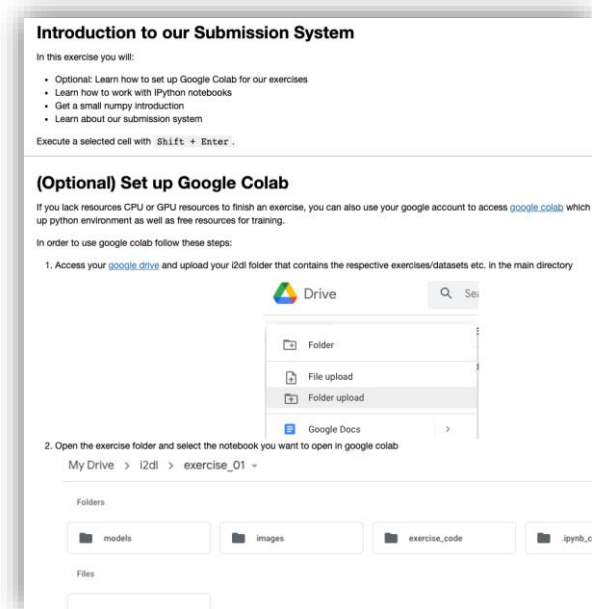
Exercises - Goal

- Goal
Gather enough experience to start your own individual (research) project
- Focus
 - Reimplementations
 - Introduction to common libraries
 - Applications



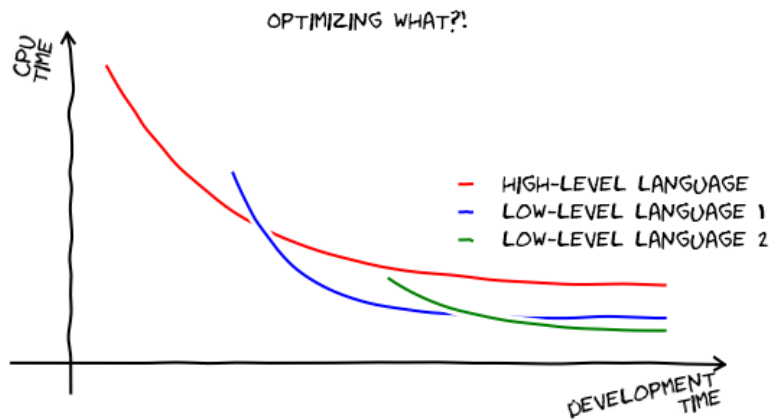
Tutorials Sessions & Exercises Format

- Presentation: video w/ slides
Posted to webpage and Piazza
- Video length
Tutorial sessions are usually much shorter than the lecture
- Exam
All tutorial sessions and optional exercises (incl. coding) are exam-relevant
- Main work
Interactive coding notebooks (don't underestimate them – start early!)



Why Python?

- Very easy to write development code thanks to an intuitive syntax
- Most used language used in deep learning research



Exercises – Tech Stack

- Python
 - Jupyter notebooks
 - Numpy
- Deep Learning library
 - Pytorch
- Hardware requirements
 - Minimum: CPU
 - Preferred: Nvidia GPU
 - Alternative: Google Colab



Extra Spice: Leaderboard (Ex6, 9-11)

- Compete against your fellow students on our datasets for a variety of tasks
- Some of the best solutions will be featured in exercise discussions

Leaderboard

The leaderboard shows for each exercise the highest scoring submission from each user. Only valid submissions are displayed.

Exercise 1	Exercise 3	Exercise 4	Exercise 5	Exercise 6	Exercise 7	Exercise 8	Exercise 9	Exercise 10	Exercise 11
#	User	Score							
1	a0005	68.77							
2	a0010	64.41							

11 Exercises - Content

Exercise 01: Organization
Exercise 02: Math Recap

Intro

Exercise 03: Dataset and Dataloader
Exercise 04: Solver and Linear
Regression
Exercise 05: Neural Networks
Exercise 06: Hyperparameter Tuning

Numpy
(Reinvent the wheel)

Exercise 07: Introduction to Pytorch
Exercise 08: Autoencoder

Pytorch/Tensorboard

Exercise 09: Convolutional Neural
Networks
Exercise 10: Semantic Segmentation
Exercise 11: Sequence Models

Applications
(Hands-off)

Submissions & Bonus

- Submissions
 - 9 (out of 11) exercises labelled as *submissions*
 - submissions have a fixed, mostly one week due date until they have to be solved and successfully uploaded.
- Bonus
 - Passing exam grade get's a -0.3 bonus
 - Need to have passed **8/9** submissions

11 Exercises – 9 Submissions

Exercise 01: Organization
Exercise 02: Math Recap

Intro

Exercise 03: Dataset and Dataloader
Exercise 04: Solver and Linear
Regression
Exercise 05: Neural Networks
Exercise 06: Hyperparameter Tuning

Numpy
(Reinvent the wheel)

Exercise 07: Introduction to Pytorch
Exercise 08: Autoencoder

Pytorch/Tensorboard

Exercise 09: Convolutional Neural
Networks
Exercise 10: Semantic Segmentation
Exercise 11: Recurrent Neural Networks

Applications
(Hands-off)

How to pass this class

- The Efficient
Binges the lecture videos at x2 speed
Reads through all lecture notes, exercise slides and screenshots from their friend's jupyter notebook
- The 'I just want to do applications'
Visits every lecture, reads through exercise material
Works through the pytorch exercise and applications during the semester break
- The Alpha
Works through weekly lecture and exercise content, gets passing submissions in, asks questions on piazza and tries to answer other student's posts

Exercise 1

Exercise 1 Goal

- Get your python environment up and running
- Familiarize yourself with our notebooks
- Familiarize yourself with our submission system
- It's short and technical

Overview

- Start date/availability
With the release of this tutorial
Find it on the Website/Piazza
- Submission deadline
Always planned for Wednesday 15:59, updated deadline always in the submission system.
- Content
 - First experience with our environment and submission system
 - Optional: python refresher
- Counts for the bonus (i.e it's one of the 9 submissions for the bonus)

Instructions

- Download the zipped folder
- Extract zip locally and follow the instructions in the README.md file to:
 - Install anaconda to create a new python environment
 - Set up your repository using the requirements.txt file
 - Execute `jupyter notebook` in the i2dl folder. We recommend using VSCode if you want an IDE.

New python users: <http://nbviewer.jupyter.org/github/jrjohansson/scientific-python-lectures/blob/master/Lecture-1-Introduction-to-Python-Programming.ipynb>

Hitchhiker's Guide: Notebooks

1. Run cells from top to bottom
2. Be careful when changing notebook cells
3. Edit external code in browser or in your favorite IDE
 - Don't code outside our boxes in `exercise_code` files
4. Checking other code
 - Generally optional
 - Look out for green boxes

Task: Check Code

Please read `make_dataset(directory, class_tasks)`. Additionally, it would be wise decision to get projects. As it is not beginner friendly, we removed i


Google Colab Instructions

(Optional) Set up Google Colab

If you lack resources CPU or GPU resources to finish an exercise or have issues setting up a local python environment, you can also use your google account to access [google colab](#) which provides an already set up python environment as well as free resources for training.

In order to use google colab follow these steps:

1. Access your [google drive](#) and upload your i2dl folder that contains the respective exercises/datasets etc. in the main directory




2. Open the exercise folder and select the notebook you want to open in google colab

My Drive > i2dl > exercise_01 ▾

Folders



Files



You really don't need colab at all for the first five exercises

Exercise Submission

- Register on submission webpage
<https://i2dl.vc.in.tum.de/>
 - Sign up with valid matriculation number
 - Get id and password via email (it will display the email address)
- Submit models:
 - Upload created zip file
 - Note: You will submit your whole code folder as well

FAQ

- I don't want to code in notebooks. Can I use my favourite IDE?
 - Yes
- Cool, so I can just change the whole code structure?
 - No
 - You can write any helper functions, but keep the skeleton classes intact (i.e., don't rename important functions or variables)
 - You will upload all files and those will be archived on our end

FAQ

- How do I know that I passed?
 - Once you submit a score that surpasses the threshold, you will receive an email that contains a message which tells you that you passed this submission
- Help, I got this message a second time!?
 - You will receive this message every time you submit an exercise that exceeds the score
 - Your best performing model will count, not the latest submitted one

FAQ

- Where can I get assistance?
 - Piazza. Ask there. Maybe someone has asked the same question before.
 - Office hours
 - Other students

Upcoming Lecture

- Next lecture: Lecture 2: Machine Learning basics
- Next Tutorial: Math Recap (+ Math exercise sheet)

See you next week

