Course Notes: Deep Learning for Visual Computing

Peter Wonka

August 27, 2022

Contents

1 Syllabus

1.1 Class times and location

- Times: Monday, Thursday 16:00 - 17:30
- Room: Lecture Hall 2 (room 2325) / Building 9
- Start Date: 29.08.2020
- End Date: 08.12.2020

1.2 Who is teaching the class?

- Instructor: Peter Wonka
- Office: building 1, 2nd floor, sea side
- Email: peter.wonka@kaust.edu.sa
- Office hours: please send an email to setup an appointment
- This list is only tentative, there will most likely be fewer TAs.
- Organizing TA: Michael Birsak, can help with setup
- TA: Biao Zhang, Project 1, Image Classification
- TA: Kumail AlHamoud, Project 2, Style Transfer
- TA: Wamiq Reyaz, Project 3, Dense Regression (Segmentation, Depth Estimation)
- TA: Ivan Skorokhodov, Project 4, Generative Modeling
- TA: Ahmed Abdelreheem, TBD

1.3 What is my background?

- Vienna University of Technology (MS, PhD)
  - Grenoble, Rennes
- Georgia Institute of Technology (Post Doc)
- Arizona State University (Assistant / Associate Professor)
- KAUST (Full Professor, Interim Director of the VCC center)

1.4 What are my research topics?

- Deep Learning
- Computer Vision
- Computer Graphics
- Machine Learning, Remote Sensing, Image Processing, Visualization
- Research Keywords: deep learning, 3D computer vision, machine learning for visual computing, generative adversarial networks, generative modeling, diffusion, variational autoencoder, flow models, auto-regressive models, network architectures, 3D reconstruction, depth estimation, NeRF-based reconstruction, laser scanning, urban modeling, urban reconstruction, urban planning, computational design, procedural modeling, layout synthesis, shape modeling, interactive editing, geometry processing, architectural geometry, geo-spatial visualization, surface sampling, surface remeshing, image editing, image analysis, image segmentation, texture synthesis, texture analysis, lighting design.

1.5 Where can I find information about the class?

- Blackboard, Slack, ...
- We try to publish everything online: syllabus, slides, assignments, projects, announcements, ...

1.6 Is there a textbook for the class?

- Unfortunately not
- Deep Learning, Goodfellow et al.
  - Very good description of several basics
  - Outdated with respect to state of the art research
  - Not good for learning how to program

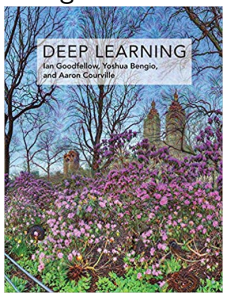

- The following are easy, practical introductions, code examples, ...
  - I do not know these books well
  - I recommend reading implementation heavy literature first (such as these books)
- Deep Learning with PyTorch, Stevens
- Hands-on Neural Networks with PyTorch, Vihar Kurama
- PyTorch Pocket Reference, Joe Papa
- PyTorch Computer Vision Cookbook, Michael Avendi

- Deep Learning with Python, Chollet
  - Easy, practical introduction, many code examples
  - Code is in TensorFlow rather than PyTorch
- Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, Geron
  - General machine learning and deep learning
  - Practical book, TensorFlow rather than PyTorch

1.7 What are the learning goals for the class?

- Learning to implement deep learning algorithms for visual computing applications
- Improving your **Python** and **PyTorch** implementation skills
- Learning the state of the research in selected topics in visual computing

1.8 How is the grade determined for the class?

- Projects (90%)
  - Each student has to work on 5 projects
  - All 5 projects have the same weight
- Assignments and Quizzes (10%)
  - About 10-20 smaller assignments, projects, quizzes, reading assignments
- Participation
  - If you do not attend class or in extreme cases disturb the classroom experience for other students your grade can change downwards
  - This generally does not happen

1.9 What is the late policy?

- If you miss a deadline with justification:
  - If you have a **very important** reason not to be here we can certainly find a solution
  - You need **written documentation(!)** to state your case
  - You need to notify me as soon as possible
- Otherwise:
  - Late projects and assignments get a penalty of one grade step (e.g. A- becomes B+) if you submit within 2 weeks after the deadline.
  - If you submit more than 2 weeks late you get **0 points!**
- Note: not replying to TAs to schedule a demo can also incur a late penalty or result in 0 points awarded for a project.

1.10 Project Grading Procedure?

- After submitting a project you have to send an email to the corresponding TA to setup a grading appointment as soon as possible.
- Last possible appointment for grading a project is the last day of classes!
- During the grading appointment, the grader will ask questions about the project, including the provided code. You should be prepared to explain / reconstruct every line of code.

1.11 Ethics?

- It is unethical to bring to your instructor's attention the possible impact of your course grade on your future plans, including graduation, scholarships, jobs, etc. and argue for grade improvements without merit
- Violations of academic integrity include (but are not limited to) cheating, fabrication, tampering, plagiarism, or facilitating such activities
- **We plan to use automatic plagiarism checks!**
- It is **forbidden** to
  - post project code online
  - share your project code with other students
- You may discuss concepts at a higher level, but you cannot directly give project solutions to other students.

1.12 What is the attendance policy?

- Attendance is required
  - Announcements are made in class
  - Announcements might not be on the web pages, but I will try to post all important information on blackboard
- If you come to class you are expected to participate
- You can occasionally decide to skip class without notifying me if you have important work in other courses

1.13 What are required skills?

- Programming skills
  - Python (or the ability to learn quickly)
  - Some initial experience with deep learning packages such as PyTorch, Keras, Tensorflow (or the ability to learn quickly) is helpful but not required
  - **Ability to write medium size programs**, debug, analyze the results, ...
- Machine Learning and Mathematics
  - It would be good to have background knowledge in Multi-variate Calculus, Linear Algebra, Probability, and Machine Learning
  - Optimization is beneficial for a minor part of the course
  - Example concepts that are used in class: functions, least squares, regression, validation set, loss function, ...