Deep Learning in Practice

Guillaume Charpiat
Rémy Hosseinkhan, Cyriaque Rousselot & Antoine Szatkownik

TAU team, LISN, Université Paris-Saclay / INRIA Saclay

. . . and guests!

Overview

- Course summary and organization
- Chapters overview

Context

- Deep learning: impressive results in the machine learning literature
- yet difficult to train, and still poorly understood;
  results = black-boxes missing explanations.
- Huge societal impact of ML today (assistance in medicine, hiring process, bank loans...)
  $\Longrightarrow$ explain their decisions, offer guarantees?
- Real world problems: usually do not fit standard academic assumptions (data quantity and quality, expert knowledge availability...).
- This course: aims at providing theoretical insights and tools to address these practical aspects, based on mathematical concepts and practical exercises.

Organisation and evaluation

- Most courses: a lesson + practical exercises (to hand in within 2 weeks, and evaluated)
- Extras: a few guest talks

Schedule

8 classes of 3 hours, at CentraleSupelec (just next to ENS Paris-Saclay), January-March 2023; check the webpage for more up-to-date schedule.

Webpage & subscription:

`https://www.lri.fr/~gcharpia/deeppractice/`

Prerequisite

- The introduction to Deep Learning course by Vincent Lepetit (1st semester), or any similar course
- Notions in differential calculus, Bayesian statistics, analysis, information theory

Links with other Deep Learning courses

- Introduction to Deep Learning (V. Lepetit) : prerequisite

- Fondements Théoriques du deep learning (F. Malgouyres & al)
- Modélisation en neurosciences et ailleurs (J-P Nadal)
- Apprentissage Profond pour la Restauration et la Synthese d'Images (A. Almansa & al)
- Deep learning for medical imaging (O. Colliot & M. Vakalopoulou)
- Object recognition and computer vision (Willow team & al)
- etc. (NLP, graphs...)

- Our course: theoretical understanding and tools to make NN work in practice with a focus on architecture design, explainability, societal impact, real datasets and tasks (e.g. small data, limited computational power vs. scaling up, RL...).
  $\Longrightarrow$ negligible overlap

Outline

Deep learning vs. classical ML and optimization

- **Going Deep or not?**
  - Examples of successes and failures of deep learning vs. classical techniques (random forests)
  - Approximation theorems vs. generalization [3, 4]
  - Why deep: ex. of depth vs. layer size compromises (explicit bounds)
- **Gap between classical Machine Learning and Deep Learning**
  - Forgotten Machine Learning basics (Minimum Description Length principle, regularizers, objective function different from evaluation criterion) and incidental palliatives (drop-out, early stopping, noise)
- **Hyper-parameters and training basics**
  - + list of practical tricks

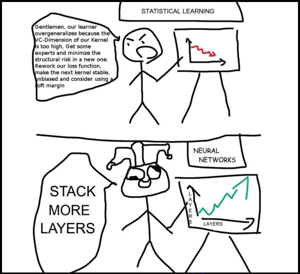

Interpretability

- **At stake:** the example of medical diagnosis, and societal issues with black-box algorithms [5]

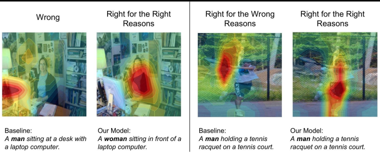

- **Interpretability of neural networks**
  - Analyzing the black-box
    - at the neuron level: filter visualisation, impact analysis
    - at the layer level: layer statistics...
    - at the net level: low-dimensional representation (t-SNE) + IB
    - by sub-task design: "explainable AI"
  - Adversarial examples & remedies
- **Issues with datasets**
  - Biases in datasets : 4 definitions of fairness
    - Getting invariant to undesirable dataset biases (e.g. gender in CVs / job offers matching)
    - Ensuring errors are uniform over the dataset
  - Differential privacy (database client protection)
- Visualization tools: grad-CAM

Architectures

- **Architectures as priors on function space**
  - Change of design paradigm
  - Random initialization
- **Architecture zoo**
  - Reminder (CNN, auto-encoder, LSTM, adversarial...)
  - Dealing with scale & resolution (fully-convolutional, U-nets, pyramidal approaches...)
  - Dealing with depth (ResNet, auxiliary losses) and mixing blocks (Inception)
  - Attention mechanisms
  - GraphCNN

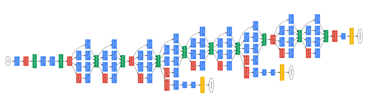

- Problem modeling: molecular dataset using graph-NN

Small data, weak supervision and robustness

- Small data
  - Data augmentation / synthetic data
  - Multi-tasking
  - Transfer learning
- Few labeled examples: forms of weak supervision
  - Semi-supervision
  - Weak supervision
  - Self-supervision
  - Active learning
- Noisy data
  - Denoising auto-encoder
  - Classification with noisy labels
  - Regression with noisy labels
- Exploiting known invariances or priors
  - Permutation invariance: “deep sets” [8], applied to people genetics
  - Choosing physically meaningful metrics, e.g. optimal transport (Sinkhorn approximation)[9]
- Transfer learning

Incorporating physical knowledge / Learning physics

- Course by Michele Alessandro Bucci and Lionel Mathelin (Safran/LISN)
  - Data assimilation
  - Learning a PDE (equation not known)
  - Incorporating invariances/symmetries of the problem
  - Knowing an equation that the solution has to satisfy: solving PDEs!
  - Deep for physic dynamics : learning and controlling the dynamics

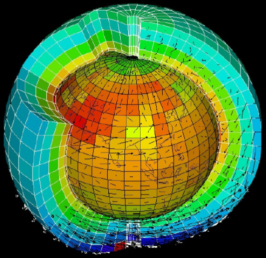

- Learning a dynamical system

Generative models + Modeling tasks and losses

- **Generative models**
  - GAN, VAE (Variational Auto-Encoder), Normalizing Flows, and Diffusion Models
- **Modeling tasks and losses**
  - KL, optimal transport, MMD...
- GAN vs. VAE vs. NF vs. DM

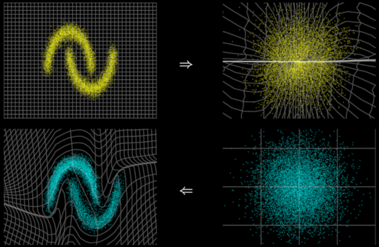

Guarantees? Generalization and formal proofs + Auto-DL

- **Guarantees?**
  - Generalization: double gradient descent and Neural Tangent Kernel
  - formal proofs of (very small) neural networks
- **Auto-DeepLearning** by Isabelle Guyon's group
  - Overview of recent approaches for automatic hyper-parameter tuning (architecture, learning rate, etc.): classical blackbox optimisation, Reinforcement Learning approaches, constrained computational time budget, self-adaptive architectures...
  - Presentation of the Auto-ML & Auto-DL challenges

To attend the course

- go see the website and **subscribe to the course**
  https://www.lri.fr/~gcharpia/deeppractice/
- install PyTorch, Jupyter and matplotlib
- See you... on 22nd (DSBA) / 25th (MVA) of January

Biographies

- Guillaume Charpiat is an INRIA researcher in the TAU team (INRIA Saclay/LISN/Paris-Sud). He has worked mainly in computer vision, optimization and machine learning, and now focuses on deep learning. He conducts studies on neural networks both in theory (self-adaptive architectures, frugal learning, formal proofs) and in applications (population genetics, PDEs and dynamical systems, and, formerly, remote sensing, molecular dynamics simulation, brain imagery, weather forecast...).
- Rémy Hosseinkhan, Cyriaque Rousselot and Antoine Szatkownik are PhD students in the TAU team or in the Dataflot team (LISN), working on deep learning for dynamical systems, on causality, and on deep generative models for population genetics, respectively.

Bibliography

- *Why does deep and cheap learning work so well?*, Henry W. Lin, Max Tegmark, David Rolnick. https://arxiv.org/abs/1608.08225
- *Representation Benefits of Deep Feedforward Networks*, Matus Telgarsky. https://arxiv.org/abs/1509.08101
- *On the structure of continuous functions of several variables*, David A. Sprecher.
- *Representation properties of networks: Kolmogorov's theorem is irrelevant*, Federico Girosi and Tomaso Poggio.
- *Weapons of Math Destruction*, Cathy O'Neil.
- *Practical Variational Inference for Neural Networks*, Alex Graves. https://papers.nips.cc/paper/4329-practical-variational-inference-for-neural-networks
- *Auto-Encoding Variational Bayes*, Diederik P. Kingma and Max Welling. https://arxiv.org/abs/1312.6114
- *Deep Sets*, Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, Alexander Smola. https://arxiv.org/abs/1703.06114
- *Learning Generative Models with Sinkhorn Divergences*, Aude Genevay, Gabriel Peyré, Marco Cuturi. https://arxiv.org/abs/1706.00292
- *Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour*, Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, Kaiming He. https://arxiv.org/abs/1706.02677
- *Temporal Ensembling for Semi-Supervised Learning*, Samuli Laine, Timo Aila. https://arxiv.org/abs/1610.02242