Deep Learning in Practice

Teachers: Guillaume Charpiat (TAU team, INRIA Saclay / LRI)
& Edouard Oyallon (CVN, CentraleSupelec)

Teaching assistants: Victor Berger (TAU team)
& TBA (CVN)

Context: Deep learning methods are now the state of the art in many machine learning tasks, leading to impressive results. Nevertheless, they are still poorly understood, neural networks are still difficult to train, and the results are black-boxes missing explanations. Given the societal impact of machine learning techniques today (used as assistance in medicine, hiring process, bank loans...), it is crucial to make their decisions explainable or to offer guarantees. Besides, real world problems usually do not fit the standard assumptions or frameworks of the most famous academic work (data quantity and quality, expert knowledge availability...). This course aims at providing insights and tools to address these practical aspects, based on mathematical concepts.

Summary: This class will start by emphasizing theoretical misconceptions, by exhibiting cases where neural networks do not reach classical techniques' performance, and also, on the opposite, explicit cases of deep learning success with proofs of depth requirement. We will then study different ways to visualize what a neural network is doing, in order to interpret its decisions, and check that it does neither reproduce undesired biases present in the dataset (such as sensitivity to ethnicity when matching CVs to job offers), nor disclose private information from people in the dataset. The rest of the course will investigate practical issues when training neural networks, in particular data quantity, trials to apply deep learning to real Reinforcement Learning problems, and automatic hyper-parameter tuning.

Organisation and evaluation: The course will comprise lectures as well as practical exercises (three mini-projects, in PyTorch), which will be evaluated. Most sessions will also comprise a lecture by a guest, from the Industry or the Academia, as a showcase of successful applications of deep learning, with practical recommendations.

Schedule: 8 classes of 3 hours, during the second semester (January-March 2019), on Monday mornings (8h30 – 11h45 with a break) at CentraleSupelec (at various places and not every week, check the agenda for details).

Webpage & mailing-list: See `https://www.lri.fr/~gcharpia/deeppractice/`

Prerequisite:

- The introduction to Deep Learning course by Vincent Lepetit (taking place during the 1st semester)
- Notions in information theory, Bayesian statistics, analysis, differential calculus

Outline

1. Introduction: gap between practice and theory – ***January 14th*** *(G. Charpiat & V. Berger)*

   - Going Deep or not?
     - Examples of successes and failures of deep learning vs. classical techniques (random forests)
     - Uselessness of the approximation theorems of (shallow) neural networks [3, 4]
     - Why deep: examples of *depth* vs. *layer size* compromises with explicit bounds [1, 2]
   - Gap between classical Machine Learning and Deep Learning in practice
     - Forgotten Machine Learning basics (Minimum Description Length principle, regularizers, objective function different from evaluation criterion) and incidental palliatives (drop-out)
     - Architectures as priors on function space, initializations as random nonlinear projections, and gradient descent on weights as a partial optimizer
   - Training generative models, by Victor Berger
     - Introduction to GANs and to Variational AutoEncoders (VAE) [6, 7]
     - **Practical session** (mini-project to give before mid-February) ***(bring your laptop!)***

2. Architectures – ***January 21th*** *(G. Charpiat)*

   - Architecture zoo reminder (CNN, auto-encoder, LSTM, adversarial...)
   - Dealing with scale & resolution (fully-convolutional, U-nets, pyramidal approaches...)
   - Dealing with depth (ResNet, Highway networks) and mixing blocks (Inception)
   - Attention mechanisms, R-CNN; "memory"
   - GraphCNN
   - Guest talk by Yaroslav Nikulin & Pierre Fillard (start-up Therapixel): Deep learning for breast cancer detection

3. Interpretability – ***February 11th*** *(G. Charpiat & E. Oyallon)*

   - 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, correlation with explainable features, ablation studies
       * at the network level: low-dimensional representation (t-SNE) of the function learned while training, and information bottleneck
       * by sub-task design: "explainable AI"
     - Adversarial examples & remedies
   - Issues with datasets
     - Biases in datasets
       * Getting invariant to undesirable dataset biases (e.g. gender in CVs / job offers matching)
       * Ensuring errors are uniform over the dataset (vs. better accuracy for certain better-represented categories)

- Differential privacy (database client protection)

- Guest talk by Pierre Stock (Facebook FAIR Paris): Bias detection, fairness & AI ethics

4. Deep Reinforcement Learning – ***February 18th*** *(E. Oyallon)*

- Reminders about RL (The Atari game framework, policy gradient, Deep Q-Networks (DQN))
- Techniques overview:
  - Value-Iteration Networks (VIN)
  - alpha-go, alpha-0
  - Actor-critic: A3C / A2C
- Limitations and applications at stake and state of the art (robotics, self-driving cars) [13]
- Deep Reinforcement Learning by Ludovic Denoyer (FAIR, Paris)

5. Image Retrieval – ***February 25th*** *(NaverLabs)*

- Presentation by Diane Larlus (NaverLabs) on *Visual Search in Large Image Collections*
- **Practical session** by Rafael Sampaio de Rezende (NaverLabs) ***(bring your laptop!)***

6. Scaling Deep learning – ***March 4th*** *(E. Oyallon)*

- Training a model using large batches [10]
- Distributed Deep Learning
- Sacrificing accuracy for inference performances
- Guest talk by Remy Cadène (LIP6)

7. Small data, weak supervision and robustness – ***March 11th*** *(E. Oyallon & G. Charpiat)*

- Information sources
  - Learning from synthetic data [12]
  - Learning from scratch vs. Transfer learning
  - Semi-supervised learning [11]
  - Self-supervised learning (ex: video prediction)
  - Multi-tasking
- Exploiting known invariances, priors or physical properties
  - Example of permutation invariance: "deep sets" [8], applied to people genetics
  - Spatial/Temporal coherence
  - Choosing physically meaningful metrics, e.g. optimal transport (Sinkhorn approximation)[9]
  - Incorporating physical knowledge: data assimilation. Example of learning dynamics equations.
- **Practical session** ***(bring your laptop!)***

8. Auto-DeepLearning – ***March 18th*** *(G. Charpiat & I. Guyon)*
    - 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...
    - Additional real-world difficulties: missing data, unstructured data
    - Presentation of the Auto-ML & Auto-DL challenges by Isabelle Guyon (ChaLearn/LRI-INRIA) ***(bring your laptop!)***

Mini-projects

The exercise sessions will consist of three mini-projects, on the following topics:

A. Training generative models (GAN vs VAE)

B. Image retrieval

C. TBC

A sheet with practical advice for training neural networks will also be given (optimizer choice, overfit, dataset handling, architecture and meta-parameters, weight initialization & normalization, debugging, etc.).

Biographies

- Guillaume Charpiat is an INRIA researcher in the TAU team (INRIA Saclay/LRI/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, formal proofs) and in applications (remote sensing, people genetics, molecular dynamics simulation, brain imagery, weather forecast, skin lesion medical diagnosis, ...).

- Edouard Oyallon is an assistant professor in the lab CVN of CentraleSupélec. He works on the optimization, analysis and interpretation of deep learning algorithms. He has several conference papers in top ranked deep learning conferences, such as ICLR, CVPR or ICCV.

References

[1] *Why does deep and cheap learning work so well?*, Henry W. Lin, Max Tegmark, David Rolnick. `https://arxiv.org/abs/1608.08225`

[2] *Representation Benefits of Deep Feedforward Networks*, Matus Telgarsky. `https://arxiv.org/abs/1509.08101`

[3] *On the structure of continuous functions of several variables*, David A. Sprecher.

[4] *Representation properties of networks: Kolmogorov's theorem is irrelevant*, Federico Girosi and Tomaso Poggio.

[5] *Weapons of Math Destruction*, Cathy O'Neil.

[6] *Practical Variational Inference for Neural Networks*, Alex Graves. https://papers.nips.cc/paper/4329-practical-variational-inference-for-neural-networks

[7] *Auto-Encoding Variational Bayes*, Diederik P. Kingma and Max Welling. https://arxiv.org/abs/1312.6114

[8] *Deep Sets*, Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, Alexander Smola. https://arxiv.org/abs/1703.06114

[9] *Learning Generative Models with Sinkhorn Divergences*, Aude Genevay, Gabriel Peyré, Marco Cuturi. https://arxiv.org/abs/1706.00292

[10] *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

[11] *Temporal Ensembling for Semi-Supervised Learning*, Samuli Laine, Timo Aila. https://arxiv.org/abs/1610.02242

[12] *Learning from Simulated and Unsupervised Images through Adversarial Training*, Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, Russ Webb. https://arxiv.org/abs/1612.07828

[13] https://himanshusahni.github.io/2018/02/23/reinforcement-learning-never-worked.html