MOSIG – MSIAM – 2017-2018 Information Access and Retrieval – GBX9MO23
Georges Quénot – Philippe Mulhem – Jean-Pierre Chevallet
22 January 2018 – 14h00-16h00 (2:00pm-4:00pm) – 2 hours

Course materials, the three papers related to the examinations, personal notes, and calculators (without network capabilities) are allowed.

The examination consists in questions related to three scientific papers and/or to the contents of the course. Paper [2] was given only as a context for paper [1]. Please use separate examination sheets for questions related to paper [1] and questions related to paper [3].

[1] Fanghong Jian, Jimmy Xiangji Huang, Jiashu Zhao, Tingting He, and Po Hu. 2016. A Simple Enhancement for Ad-hoc Information Retrieval via Topic Modelling. ACM SIGIR '16. ACM, New York, NY, USA, 733-736.
[2] Jiaul H. Paik. 2013. A novel TF-IDF weighting scheme for effective ranking. In Proceedings of the 36th international ACM SIGIR '13. ACM, New York, NY, USA, 343-352.
[3] Huang, Gao and Liu, Zhuang and van der Maaten, Laurens and Weinberger, Kilian Q., Densely connected convolutional networks, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.

Notice: On average, you may spend 5 minutes per question, so write only few sentences per question.

Part 1. Questions related to paper [1]:

Q1.1: Explain (by giving examples, if needed) why the proposal of paper [1], that integrates topics from LDA, may be interesting for information retrieval. Focus on the generality of the approach.
Q1.2: Explain the role of the lambda in formula (1) of paper [1]. What is happening when lambda=0? When lambda = 1?
Q1.3: Explain a potential problem with the linear combination of query terms in formula (3) during the fusion of the classical tf and the LDA-based sub-scores?
Q1.4: Paper [1] presents, in section 4, results for several test collections. Explain why such multiple experiments are useful when evaluating information retrieval proposals.
Q1.5: Explain shortly the role of the different evaluation measures MAP, P@10 and P@20, and their objective for information retrieval evaluations.
Q1.6: Explain why the authors of paper [1] write about the Figure 1: "It is clear that the hybrid models perform better that either traditional models or topic models on all datasets".
Q1.7: What conclusions can we draw about the impact of the number of extracted topics from LDA, according to the Figure 2?
Q1.8: Comment the specific behavior on the collection WT2G in Figure 2.
Q1.9: Open question: What would you suggest as an extension for the work of [1]?

Part 2. Questions related to paper [3]:

Q2.1: What is the main innovation in the architecture of deep networks that made ultra-deep (> 100 layers) networks trainable and efficient?

Q2.2: In experiments on the base version of DenseNets, what are the hyper-parameters related to the architecture and what are the hyper-parameters related to their training?

Q2.3: How many dense blocks were used for the DenseNet implementations for each of the four data sets?

Q2.4: What is the role of 1x1 convolutions inside the dense blocks and in the transition layer?

Q2.5: Regarding experiments on CIFAR and SVHN, why is dropout used in some cases and not into others?

Q2.6: According to the authors, what are the benefit brought by "skip connections" in neural networks?

Q2.7: What is the main goal of dropout?

Q2.8: Why is the performance of C10+/C100+ experiments significantly better than the performance of C10/C100 experiments?

Q2.9: According to results displayed in table 2 and to the publication dates of the cited papers, build a table indicating year by year the best performance on the CIFAR-100 data-set along with the name of the corresponding method.