



Offer #2020-02553

PhD Position F/M Metric learning for instance- and category-level visual representations

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

Level of experience : Recently graduated

About the research centre or Inria department

The Inria Rennes - Bretagne Atlantique Centre is one of Inria's eight centres and has more than thirty research teams. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PME's, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

The PhD will be supervised by Yannis Avrithis, Ewa Kijak and Laurent Amsaleg. The position has a duration of three years and is part of a national research grant in collaboration with a number of academic partners. The overall goal of the project is to study visual and text representations with the purpose of one disambiguating the other and both being used for multimodal question answering over large-scale knowledge bases. Work will be carried out within Inria team LinkMedia. The team specializes in multimedia content processing for analytics, gathering specialists from different fields: natural language processing, image processing and computer vision, data mining, databases.

Assignment

The goal of this PhD is to revisit the connection between classification and metric learning in visual representation learning and to extend the study of metric learning in supervision and localization settings that have mostly been studied in terms of classification.

There are many tasks where supervised metric learning appears to have a similar objective with supervised classification, but classes at inference are different from classes at learning. These include e.g. fine-grained classification [SXJ16], face recognition [SKP15], person re-identification [AGM18], local descriptor learning [HLJ15] and instance retrieval [RIT18]. Few-shot learning [LAP19] also includes two training stages with different classes and is treated as either metric learning or classification. A better understanding of the properties of the two approaches will allow a smoother progress towards more challenging problems like long-tail [WRH17] and open-set recognition [LMZ19].

Ideally, metric learning should be explored in all supervision settings where classification has been explored, e.g. semi-supervised [ITA19], few-shot [LAP19] and incremental learning [RKS17], on seen or unseen categories. This would allow e.g. self-learning to rank [CHX19] in the unsupervised setting [ITA18] or training a student to rank like a teacher in distillation [HVD15]. It is also natural to extend the study of metric learning to localization tasks including spatial attention [SPC16], object detection [RHG15] and instance segmentation [ZZY18]: Different supervision settings have not been explored as much as in classification.

It is the objective of this PhD to investigate such ideas in the broad context of searching knowledge bases consisting of visual and text data, using queries that consist of images and text as well. Detected objects or generic visual categories can help enrich the representation of a knowledge base or disambiguate text queries; conversely, cues originating in text queries can guide detection by means of attention [SPC16] and priming [RBT18].

References:

[AGM18] Jon Almazan, Bojana Gajic, Naila Murray, and Diane Larlus. Re-ID done right: towards good practices for person re-identification. In arXiv preprint arXiv:1801.05339, 2018.

[CHX19] Fatih Cakir, Kun He, Xide Xia, Brian Kulis, and Stan Sclaroff. Deep Metric Learning to Rank. In CVPR, 2019.

[HLJ15] Xufeng Han, Thomas Leung, Yangqing Jia, Rahul Sukthankar, and Alexander C Berg. MatchNet: Unifying Feature and Metric Learning for Patch-Based Matching. In CVPR, 2015.

[HVD15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In arXiv preprint arXiv:1503.02531, 2015.

[ITA18] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Mining on Manifolds: Metric Learning without Labels. In CVPR, 2018.

[ITA19] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Label Propagation for Deep Semi-Supervised Learning. In CVPR, 2019.

[LAP19] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. Dense Classification and Implanting for Few-shot Learning. In CVPR, 2019.

[LMZ19] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-Scale Long-Tailed Recognition in an Open World. In CVPR, 2019.

[RIT18] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting Oxford and Paris: Large-Scale Image Retrieval Benchmarking. In CVPR, 2018.

[RKS17] Sylvester-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In CVPR, 2017.

[RHG15] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In NIPS, 2015.

[RBT18] A. Rosenfeld, M. Biparva, and J. K. Tsotsos. Priming neural networks. In CVPRW, 2018.

[SKP15] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In CVPR, 2015.

[SXJ16] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In CVPR, 2016.

[SPC16] C. Sun, M. Paluri, R. Collobert, R. Nevatia, and L. Bourdev. Pronet: Learning to propose object-specific boxes for cascaded neural networks. In CVPR, 2016.

[WRH17] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In NIPS, 2017.

[ZZY18] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly Supervised Instance Segmentation Using Class Peak Response. In CVPR, 2018.

Main activities

Not applicable.

Skills

The candidate should ideally have a degree in Computer Science, Applied Mathematics or Electrical Engineering; solid mathematical background and programming skills; fluency in English language; preferably, prior experience in computer vision and deep learning, optionally as well in natural language processing.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs

Remuneration

Monthly gross salary amounting to 1982 euros for the first and second years and 2085 euros for the third year

General Information

- **Theme/Domain** : Vision, perception and multimedia interpretation
Scientific computing (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2020-10-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2020-06-30

Contacts

- **Inria Team** : [LINKMEDIA](#)
- **PhD Supervisor** :

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Please submit online: your resume, transcripts, motivation letter and letters of recommendation

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.