



Offer #2022-04862

Post-Doctoral Research Visit F/M Postdoc Position F/M [2 years]: 3D Transformers for Autonomous Driving

Contract type : Fixed-term contract

Renewable contract : Yes

Level of qualifications required : PhD or equivalent

Fonction : Post-Doctoral Research Visit

Level of experience : Recently graduated

Context

The job position is in central Paris, France. As it is part of a joint team project, the candidate will be colocated on [Inria](#) and [Valeo.ai](#) sites.

At Inria the candidate will work in RITS team that has approx. 20 people working on computer vision, decision, planning, modeling for autonomous driving. The [computer vision group](#) publishes in top-tier scientific venues of computer vision and machine learning.

Valeo.ai is the [artificial intelligence research center](#) for automotive applications of Valeo Group. The center started in 2017 to conduct ambitious research projects, regarding assisted and autonomous driving. The Valeo.ai group has about 25 people, all publishing in top-tier scientific venues in artificial intelligence.

Work environments are nice and lively with people from worldwide origins. Social skills is appreciated as collaborations with other researchers/PhDs is expected.

Assignment

3D Transformers for Autonomous Driving

Main activities

In the same way that learning with transformers has somehow brought a revolution in natural language processing [Vas17], and then in image processing [Car20, Dos21], recently emerging transformers [Guo21, Zha21] for point clouds are now also starting to significantly boost the performance of 3D processing. Yet, the general concepts of transformers and the reasons for their success are not fully understood, and there are still many open ways to adapt them efficiently to 3D.

The emergence of transformers opens a wide field of research subjects that matches well with the research area “Vision and 3D Perception for Scene Understanding”. The goal of this postdoc is to attack some of them.

1. A first objective is to **investigate the specificity of transformers for 3D**. It includes the question of positional encoding, which is widely different in 3D from the 1D (text) and 2D (image) perspectives. It also includes the square complexity issues with cross- and self-attention, taking into account the sparsity of point clouds. More generally, it concerns as well the time and space complexity of point cloud processing, in the perspective of real-time, embedded software.

This study of 3D transformers will be made through downstream tasks that include object detection, semantic or panoptic segmentation, and denser depth predictions, which are key tasks for autonomous driving. We will study in particular the impact of the specific sparsity and data patterns induced by vehicle sensors. We will also consider a stream of point clouds, as available from a lidar, taking time into account in a 4D perspective.

2. Besides, we will investigate the use of **transformers backbones regarding self-supervision** [Sim21], which is a key approach to bring supervised learning to a new level by saving the cost of annotating large datasets. We will study new pretext tasks in 3D that transformers more specifically leverage, as well as contrastive learning techniques that are tightly linked to the attention mechanism of transformers.

A PhD student is about to start on self-supervision for 3D at Valeo and another works on 2D supervision for 3D at Inria, though none of them focusing on transformers. The postdoc will be able

to work jointly with either PhD student on self-supervision issues.

3. Moreover, we will study the **transferability of learned transformer models in the perspective of domain adaptation** [Vu19a, Vu19b]. In particular, we will investigate the disentangling of latent space representations, working towards domain-invariant features by enforcing orthogonality of the domain features while enabling the discovery of exclusive task or domain features, through their realization via multi-head attention.

Another PhD student in Valeo is about to start working on 3D domain adaptation, although with a different perspective (using optimal transport). There will nonetheless be a number of collaboration opportunities between this PhD student and the postdoc regarding adapting transformer-based features.

4. A last research direction concerns **multi-modality**, when lidar point clouds are acquired together with camera images, to leverage the similarity and complementarity of sensor information [Jar20]. One technical subject concerns a possible interplay between the forms of transformer attention used in 2D and the kinds of attention that are and will be developed in 3D. Another more general question is the joint self-supervision from the interaction of 2D and 3D, or from cross-task representations. Last, we will study the intertwined relation of geometry and semantics through the semantic scene completion task [Rol20, Rol21].

[Car20] End-to-End Object Detection with Transformers. Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, Sergey Zagoruyko. European Conference on Computer Vision (ECCV), 2020.

[Dos21] An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, Neil Houlsby. International Conference on Learning Representations (ICLR), 2021.

[Guo21] PCT: Point cloud transformer. Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R. Martin, Shi-Min Hu. Computational Visual Media, 2021.

[Jar20] xMUDA: Cross-Modal Unsupervised Domain Adaptation for 3D Semantic Segmentation. Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Émilie Wirbel, Patrick Pérez. Conference on Computer Vision and Pattern Recognition (CVPR), 2020.

[Rol20] LMSCNet: Lightweight Multiscale 3D Semantic Completion. Luis Roldao, Raoul de Charette, Anne Verroust-Blondet. International Conference on 3D Vision (3DV), 2020

[Rol21] 3D Semantic Scene Completion: a Survey. Luis Roldao, Raoul de Charette, Anne Verroust-Blondet. International Journal of Computer Vision (IJCV), 2021.

[Sim21] Localizing Objects with Self-Supervised Transformers and no Labels. Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, Jean Ponce. British Machine Vision Conference (BMVC), 2021.

[Vas17] Attention Is All You Need. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin. Conference on Neural Information Processing Systems (NeurIPS), 2017.

[Vu19a] ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic Segmentation. Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, Patrick Pérez. Conference on Computer Vision and Pattern Recognition (CVPR), 2019.

[Vu19b] DADA: Depth-aware Domain Adaptation in Semantic Segmentation. Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, Patrick Pérez. International Conference on Computer Vision (ICCV), 2019.

[Zha21] Point Transformer. Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, Vladlen Koltun. International Conference on Computer Vision (ICCV), 2021.

Skills

Applicants should have defended or be finishing their PhD and have a strong publications record. They should have a solid background in computer vision (including 3D processing) and machine learning, particularly in deep learning, with strong PyTorch coding skills.

Applicants should apply [on this platform](#):

- a cover letter explaining their interest and adequacy for the postdoc topic,
- their CV/resume,
- possibly, references or recommendation letters.

Apply as soon as possible. Applications are reviewed on a rolling basis. Starting date October 1st 2022 or before (firm)

Benefits package

- Subsidised catering service
- Partially-reimbursed public transport
- Flexible working hours
- Sports facilities

General Information

- **Theme/Domain** : Robotics and Smart environments
Scientific computing (BAP E)
- **Town/city** : Paris
- **Inria Center** : [Centre Inria de Paris](#)
- **Starting date** : 2022-09-01
- **Duration of contract** : 2 years
- **Deadline to apply** : 2022-08-15

Contacts

- **Inria Team** : [RITS](#)
- **Recruiter** :
De Charette Raoul / raoul.de-charette@inria.fr

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

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.