



Offer #2025-09152

PhD Position F/M Generative models for unsupervised anomaly detection in spatio-temporal data: Application to medical imaging

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

Level of experience : From 3 to 5 years

About the research centre or Inria department

The Centre Inria de l'Université de Grenoble groups together almost 600 people in 22 research teams and 7 research support departments.

Staff is present on three campuses in Grenoble, in close collaboration with other research and higher education institutions (Université Grenoble Alpes, CNRS, CEA, INRAE, ...), but also with key economic players in the area.

The Centre Inria de l'Université Grenoble Alpe is active in the fields of high-performance computing, verification and embedded systems, modeling of the environment at multiple levels, and data science and artificial intelligence. The center is a top-level scientific institute with an extensive network of international collaborations in Europe and the rest of the world.

Context

Anomaly detection is a challenging task in contexts where abnormalities are not annotated and difficult to detect even for experts. This problem can be addressed through unsupervised anomaly detection (UAD) methods, which identify features that do not match with a reference model of normal profiles. In the context of Parkinson's disease and newly diagnosed patients, the detection task is all the more challenging as abnormalities may be subtle and hardly visible in structural MR brain scans. The goal of this project is to further improve the reliability of the detection by leveraging additional information coming from longitudinal or temporal data.

Assignment

In particular, we would like to investigate new very successful models based on generative diffusion models. Diffusion models have been used for anomaly detection in images on one hand, and on timeseries, on the other hand. In this internship, the goal is to study their use in both contexts, see [Yang et al 2024] for a recent survey on diffusions for spatio-temporal data. Such solutions are often computationally costly. More efficient approaches have been proposed, eg. [Livernoche et al 2024] and the goal is to study their scalability to detect anomalies in images evolving over time and in particular for longitudinal medical image data which present specific challenges, such as very sparse time points, possibly missing data and non-aligned times within the patient population.

Main activities

More specifically, longitudinal data [Hedeker & Gibbons 2006] consist in the repeated observations of patients over time. In practice, we expect to analyse image data at a few different times corresponding to successive visits of patients. Their analysis informs us on the progression of the disease through the evolution of abnormalities, both in size, numbers, or locations. More specifically, when applied to anomaly detection, the expectation is the confirmation of uncertain detections or the discovery of new ones, not visible at early stages.

Modelling longitudinal data presents different types of challenges. First are the methodological challenges related to the design of relevant models to handle all the data and disease's characteristics in order to answer the statistical and medical questions. These modelling difficulties cannot be separated from challenges arising from data with very different modalities and time dependencies, in particular involving different acquisition time-sets and different scales of patient screening, resulting on possibly partially missing data [Couronne et al 2019].

Young et al. data [Young et al 2024] recently performed an exhaustive review of data-driven generative models of how a disease evolves over time. Such models use a generative disease progression model and a set of constraints informed by human insight to infer a data-driven disease time axis and the shape of biomarker trajectories along it.

Within this framework, Sauty et al. [Sauty et al 2022] recently investigated a way to model such longitudinal effects directly in the MR images by training a linear mixed effect model in the latent representation space of a longitudinal variational autoencoder. This design enables to combine the robustness of mixed-effects modelling of clinical biomarkers progression with missing data and, for any timepoint, with that of autoencoders both to learn efficient and compact representation of 3D images and reconstruct the image from the latent variable. This model was shown to successfully model based on 3D T1w MRI normal brains and disease progression in Alzheimer patients. However, it is not clear how to reproduce such results in particular on other images.

In the same line, Puglisi et al. data [Puglisi et al 2024] also recently tackle the issue of progression modelling on medical images by introducing a novel spatio-temporal model that combines a latent diffusion model (LDM) with a ControlNet to generate individualized brain MRIs conditioned on subject-specific data.

Similarly to Sauty et al, this model was shown to successfully model healthy and Alzheimer patients' brains.

Initial directions of research:

Review the state-of-the art in the domain of deep generative progression models, e.g. based on the review by Young et al and Zhang et al 2024 or other recent works.

As a first direction of research, we propose to consider the modalities used in our previous work [Oudoumanessah et al 2023] and investigate the extension of the model and inference technique therein to multiple time data. A first idea would be to use analysis and results at previous times to inform analysis at subsequent times using a Bayesian approach as a way to incorporate information from one time to another.

As a second direction of research, we will focus on accounting for possibly missing time sampling point, considering that the sample size of patients having performed all required analysis at regular time intervals, is often quite small. This task will aim at reporting on the uncertainties associated to the individual prediction in this case. The performances, strengths and weaknesses of various approaches will be compared.

A first generative model to be considered can be the SADM approach of Yoon et al 2023. In particular, we will investigate whether considering a bidirectional version of SADM could be of interest to account both for past and future time stamps when detecting anomalies.

References:

Yiyuan Yang, Ming Jin, Haomin Wen, Chaoli Zhang, Yuxuan Liang[†], Lintao Ma, Yi Wang, Chenghao Liu, Bin Yang, Zenglin Xu, Jiang Bian, Shirui Pan, Qingsong Wen. A Survey on Diffusion Models for Time Series and Spatio-Temporal Data, 2024.

Victor Livernoche, Vineet Jain, Yashar Hezaveh, Siamak Ravanbakhsh. On Diffusion Modeling for Anomaly Detection, ICLR 2024.

L. Young, N. P. Oxtoby, S. Garbarino, N. C. Fox, F. Barkhof, J. M. Schott, and D. C. Alexander, "Data-driven modelling of neurodegenerative disease progression: Thinking outside the blackbox", *Nature Reviews Neuroscience*, vol. 25, no. 2, pp. 111–130, Feb. 2024, issn: 1471-0048. doi:10.1038/s41583-023-00779-6. [Online]. Available: <https://doi.org/10.1038/s41583-023-00779-6>

Raphael Couronne, Marie Vidailhet, Jean-Christophe Corvol, Stephane Lehericy, and Stanley Durlleman. Learning disease progression models with longitudinal data and missing values. In ISBI 2019 - International Symposium on Biomedical Imaging, Venice, Italy, April 2019.

Donald Hedeker and Robert D. Gibbons. Longitudinal data analysis. John Wiley & Sons, Inc, New Jersey, 2006.

Hoehn, M. and Yahr, M. D. Parkinsonism: onset, progression, and mortality, *Neurology* 1998.

Kendall, A and Gal, Y. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? NeurIPS 2017.

Oudoumanessah G, Lartizien C, Dojat M, Forbes F, Frugal unsupervised detection of subtle abnormalities in medical imaging, in: Greenspan H, Madabhushi A, Mousavi P, Salcudean S, James Duncan J, Syeda-Mahmood T, R T (Eds.) Miccai, Springer-Verlag AG Swizerland, Vancouver (Ca), 2023, pp. 411-421.

Kenneth Marek, Sohini Chowdhury, Andrew Siderowf, et al., “The parkinson’s progression markers initiative (ppmi) - establishing a pd biomarker cohort,” Annals of Clinical and Translational Neurology, p. 1460–1477, 2018.

Benoît Sauty, Stanley Durrleman . Progression models for imaging data with Longitudinal Variational Auto Encoders. MICCAI 2022, International Conference on Medical Image Computing and Computer Assisted Intervention, Sep 2022, Singapore, Singapore. hal-03701632

Pinon, G. Oudoumanessah, R. Trombetta, M. Dojat, F. Forbes, and C. Lartizien, Brainsubtle anomaly detection based on auto-encoders latent space analysis : Application to de novoparkinson patients, 2023. arXiv:2302.13593 [eess.IV]. [Online]. Available: <https://arxiv.org/abs/2302.13593>. [11]

Puglisi, D. C. Alexander, and D. Ravì, Enhancing spatiotemporal disease progression models via latent diffusion and prior knowledge, 2024. arXiv:2405.03328 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2405.03328>

Jee Seok Yoon, Chenghao Zhang, Heung-Il Suk, Jia Guo, Xiaoxiao Li, 2023. SADM: Sequence-Aware Diffusion Model for Longitudinal Medical Image Generation. <https://arxiv.org/abs/2212.08228>

Skills

Skills required: computer science, applied mathematics, interest for statistics applied to medical data.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs

- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

General Information

- **Theme/Domain** : Optimization, machine learning and statistical methods Statistics (Big data) (BAP E)
- **Town/city** : Montbonnot
- **Inria Center** : [Centre Inria de l'Université Grenoble Alpes](#)
- **Starting date** : 2025-11-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2025-09-30

Contacts

- **Inria Team** : [STATIFY](#)
- **PhD Supervisor** :
Forbes Florence / florence.forbes@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.

The keys to success

The selected candidate will be supervised by Florence Forbes (Statify) and Agathe Guilloux (Heka). He/she will also benefit from the expertise of Michel Dojat from Grenoble Neurosciences Institute and Carole Lartizien from CREATIS Lyon.

Key words: Statistical and deep learning, Longitudinal analysis, Clustering, Mixed effect model, diffusion models, generative models, Biomarkers.

Theme / Domain / Context: The main topics of this proposal are in statistical and machine learning and big longitudinal data.

Contact: florence.forbes@inria.fr and agathe.guilloux@inria.fr

Main location: INRIA Grenoble, Statify team <https://team.inria.fr/statify/> and Heka team at INRIA Paris

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.