# Internship Offer — Master's or 3rd Year Engineering School

## Exploring Pretext Tasks for Building Statistically Constrained Latent Spaces of SAR Time Series

*Self-supervised learning with robust statistical priors*

| | |
|---|---|
| **Level** | ❑ Bachelor+3 ❑ Bachelor+4 ■ Bachelor+5 (Master's 2 or 3rd Year Engineering School) |
| **Duration** | 6 months, starting March 2027 |
| **City, Country** | Strasbourg (Illkirch-Graffenstaden), France |
| **Laboratory** | ICube – Engineering Science, Computer Science and Imaging Laboratory (UMR 7357, Université de Strasbourg / CNRS) |
| **Collaborations** | Exchanges with LISTIC Annecy (M. Gallet, Y. Mhiri) within the GDR IASIS exploratory project APPA-LOOSA. |

## Context

Satellite Image Time Series (SITS) acquired by Synthetic Aperture Radar (SAR), such as Sentinel-1, allow monitoring of the Earth surface day and night, independently of cloud cover, and are now routinely used for environmental monitoring : wet snow and cryosphere mapping [1], or detection of slow-moving landslides from interferograms [2]. Since Sentinel-1 continuously acquires images over the world and come back to its starting point within 6 days, it results in a huge quantity of data that can be studied both temporally and spatially. Such an abundance of data can be easily exploited to train deep neural network and ensure a great generalizability of their performance. However, to ensure such performance, deep networks require labeled data, which in practice is rarely available since expert annotations are costly both in time and ressources.

**Self-supervised learning** addresses this issue by training an encoder to project the data into a generalizable *latent space*. The latter aims to be reused *a posteriori* as input to an additional ligthweight head trained with only few samples to address any downstream task. In the litterature, two main strategies can be applied to perform self-supervised training : constraining the latent space via contrastive learning [3] and/or training the network to address a **label-free pretext task**. While current works in the host laboratories focus on contrastive learning applied to the latent space, this internship offer focuses on the use of pretext tasks. Various pretext tasks have been used in the literature such as spatial, temporal or spatio-temporal masked-autoencoders [4, 5, 6] but also image denoising [7]. In practice, most of representation learning strategies combine several strategies to construct a robust, generalizable and exploitable latent space [8].

These pretext tasks were however designed for optical images or Gaussian-like signals, and are poorly suited to SAR. SAR intensities are corrupted by multiplicative, strongly **non-Gaussian** speckle and are better described by heavy-tailed distributions (Rayleigh, Weibull, compound-Gaussian / elliptical models) [9,10]. the acquisition geometry (layover, shadow, foreshortening in mountainous areas) and the parameters of each acquisition further shape the signal [11]. Standard augmentations (jittering, additive noise, color changes) or Euclidean $\ell_2$ reconstruction losses may thus produce latent spaces that are neither physically meaningful nor robust.

## Internship Topic

The goal of the internship is to **build latent spaces that make sense with respect to the statistical nature of SAR time series**. To this end, the intern will be in charge of designing pretext tasks and losses that embed statistical priors. Two complementary levels will be explored :

1. *Adapted pretext tasks*, where augmentations respect the SAR physics (speckle resampling, temporal resampling [3], geometry-aware masking) and where similarity or reconstruction losses rely on divergences between distributions (Kullback–Leibler, Rényi [13], geometric losses [14], or temporal alignment measures such as soft-DTW [15]) rather than Euclidean distances.
2. *Statistically constrained representations*, where the latent space (or part of it) is tied to the parameters of a robust statistical model of the data (*e.g.* scale/texture parameters of compound-Gaussian distributions, covariance matrices living on a Riemannian manifold [10,12])

The expected outcome is a comparison of these strategies on SAR time series, evaluated on downstream tasks with few labels (linear probing on classification / detection), in terms of performance, robustness and interpretability of the latent space.

## Internship Schedule

The recruited student will implement standard pretext tasks (masked autoencoding, denoising, super-resolution) on SAR time series and contribute to the definition of an evaluation protocol on the trained neural network models (linear probing, few-label regime). After this first round of experimentations, he will help design SAR-consistent image augmentations and pretext tasks notably by various losses based on divergences between robust distributions, study their differentiability and numerical stability.

Depending on progress, the results may be published (GRETSI, IGARSS, ICASSP). From a research perspective, the intern could also extend the proposed pretext task by combining them with latent space contrastive strategies [8]. In particular, the intern could develop a method to contrain the construction of the latent space depending on statistical prior dedicated to SAR images. It is worth noting that the internship may be a first step toward a PhD on the topic.

## Required Skills

- **Computing** : Python, PyTorch ; taste for numerical experimentation and GPU computation.
- **Learning** : fundamentals of deep learning
- **Signal & imaging** : signal and image processing, statistics.
- **Appreciated (not required)** : SAR/InSAR imaging.

| | |
|---|---|
| **Stipend** | According to applicable regulations |
| **Contact** | Antoine Bralet : abralet@unistra.fr, Matthieu Gallet (matthieu.gallet@univ-smb.fr), Yassine Mhiri (yassine.mhiri@univ-smb.fr) |

## References


1. M. Gallet, A. Atto, F. Karbou, E. Trouvé, *Wet snow detection from satellite SAR images by machine learning with physical snowpack model labeling*, IEEE JSTARS, 2024.
2. A. Bralet, E. Trouvé, J. Chanussot, A. M. Atto, *ISSLIDE : A new InSAR dataset for slow sliding area detection with machine learning*, IEEE GRSL, 2024. [DOI]
3. A. Saget *et al.*, *Resampling augmentation for time series contrastive learning : application to remote sensing*, 2025. [arXiv:2506.18587]
4. Y. Gandelsman, Y. Sun, X. Chen, A. Efros, *Test-time training with masked autoencoders*, NeurIPS 2022.
5. Y. Yuan, L. Lin, Q. Liu, R. Hang, Z. G. Zhou, *SITS-Former : A pre-trained spatio-spectral-temporal representation model for Sentinel-2 time series classification*, IJAEO 2022
6. Y. Cong, S. Khanna, C. Meng, P. Liu, E. Rozi, Y. He, et al, *Satmae : Pre-training transformers for temporal and multi-spectral satellite imagery*. NeurIPS 2022
7. E. Dalsasso, L. Denis, F. Tupin, *As if by magic : Self-supervised training of deep despeckling networks with MERLIN*, IEEE TGRS, 2021
8. I. Dumeur, S. Valero, J. Inglada, *Paving the way toward foundation models for irregular and unaligned satellite image time series*, 2024. [arXiv:2407.08448]
9. D.-X. Yue, F. Xu, A. C. Frery, Y.-Q. Jin, *Synthetic aperture radar image statistical modeling : Part one – single-pixel statistical models*, IEEE GRSM, 2021. [DOI]
10. E. Ollila, D. E. Tyler, V. Koivunen, H. V. Poor, *Complex elliptically symmetric distributions : survey, new results and applications*, IEEE TSP, 2012.
11. J. Prexl, M. Recla, M. Schmitt, *SARFormer – An acquisition parameter aware vision transformer for synthetic aperture radar data*, CVPR Workshops, 2025.
12. A. Collas, A. Breloy, C. Ren, G. Ginolhac, J.-P. Ovarlez, *Riemannian optimization for non-centered mixture of scaled Gaussian distributions*, IEEE TSP, 2023. [DOI]
13. M. Gallet, A. Mian, A. Atto, *Rényi divergences learning for explainable classification of SAR image pairs*, ICASSP 2024.
14. A. Mensch, M. Blondel, G. Peyré, *Geometric losses for distributional learning*, ICML 2019.
15. M. Cuturi, M. Blondel, *Soft-DTW : a differentiable loss function for time-series*, ICML 2017.