ANR ASTRAL project: Statistical learning for multi-dimensional SAR imaging

Job offer: 18 months post-doc position

Generative models and representation learning
for synthetic aperture radar imaging

1 Background

Synthetic Aperture Radar (SAR) is an active imaging modality that can operate day and night, even in the presence of cloud cover. Unlike optical imaging, the wavelength range used in radar imaging (from centimeters to meters) does not correspond to visible radiation but to radio waves. Radar images can distinguish smooth surfaces (rivers, lakes, roads) from those that are rough at the scale of the wavelength (grass, low vegetation, trees). Artificial structures (buildings, vehicles, electrical pylons, etc..) send back strong echoes. They appear as very bright structures in the radar images. SAR imaging is therefore a very complementary technique to optical imagery, with many specificities.

The figure below shows two images of the same area, acquired on the left by an airborne radar system and on the right by an optical sensor. Although representing the same scene, these two images are very different. To be able to perform the analysis of radar images, it is therefore necessary to develop specific processing.
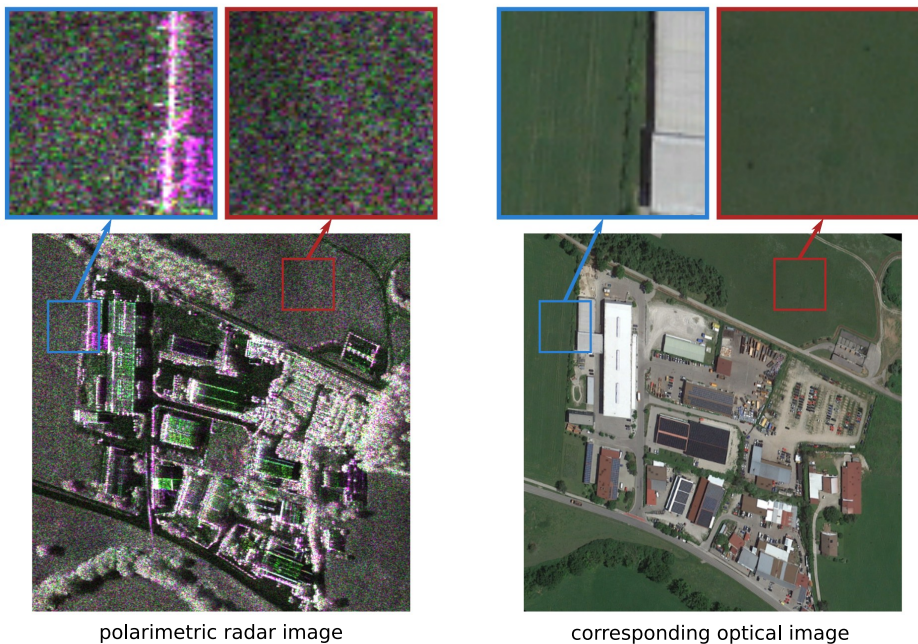

polarimetric radar image
Pauli color representation: (HH-VV,2HV,HH+VV)

corresponding optical image

One of the specific processing for radar imagery is the reduction of speckle (noise that appears on radar images). Spectacular progress has been made recently by our teams by combining deep learning techniques and a self-supervised training strategy [Dalsasso et al., 2021]:
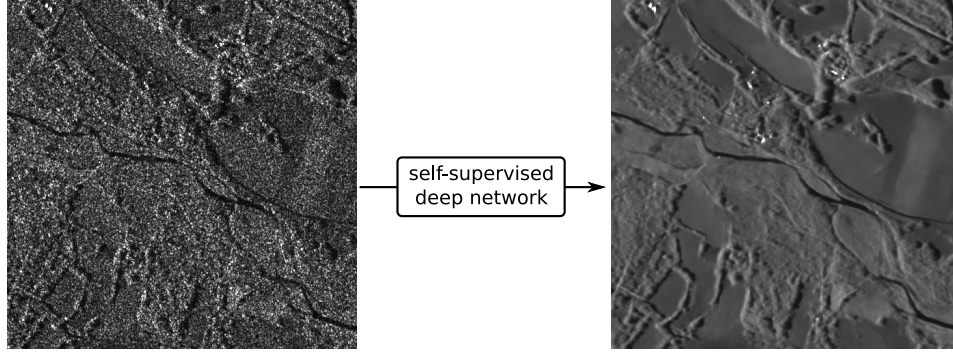
Radar images (here, acquired by the TerraSAR-X satellite) can be strongly enhanced by a deep network.

2 Objectives

Generative image models offer a very rich framework in the context of image restoration and interpretation. For example, they can provide access to a more informative characterization of uncertainties than a simple pixel-to-pixel variance [Sankaranarayanan et al., 2022], potentially lead to representations invariant to changes in acquisition (incidence, frequency band), and provide access to multi-modality (optical-radar) fusion strategies.
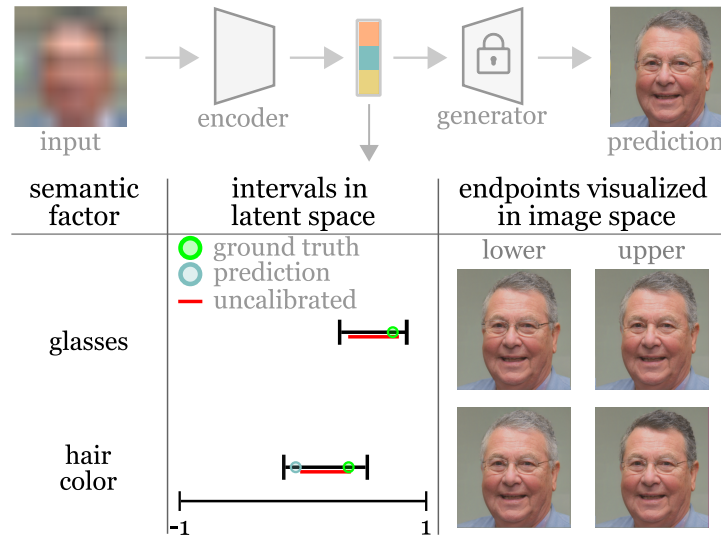
Semantically informative representation of uncertainties
in an image restoration problem (source: [Sankaranarayanan et al., 2022]).

Many generative models have been proposed in optical imaging in the last few years [Bond-Taylor et al., 2021], based for example on GANs [Goodfellow et al., 2014, Karras et al., 2021], invertible networks [Rezende and Mohamed, 2015] or diffusion models [Ho et al., 2020, Dhariwal and Nichol, 2021]. Unsupervised learning of representations, with masking [He et al., 2022] or *contrastive learning* strategies [Chen et al., 2020], is particularly powerful to obtain compact representations of images that can be used in many weakly supervised downstream tasks (e.g. classification, object detection).

The objective of this post-doc position will be to develop models of radar scenes in a self-supervised framework, and then to exploit these models for different tasks (such as restoration, semantic segmentation, change detection, or multi-modality fusion).

3 Work environment

This post-doc position is part of the ANR ASTRAL project involving teams from Télécom Paris, the University of Saint-Etienne, Cnam and ONERA in order to develop statistical learning approaches for radar imagery (https://astral.wp.imt.fr/).

The location of the post-doc is mainly in Saint-Etienne, France, at the Hubert Curien laboratory (`labh-curien.univ-st-etienne.fr/`). The work will be carried out in close collaboration with the teams at Télécom Paris and Cnam, where a stay to carry out part of the work is possible.

The Hubert Curien Laboratory at the University of Saint-Etienne is a laboratory of the CNRS, the French National Research Agency. It gathers 230 staff members, comprising two departments: Computer Science, Security, Image and Optics, Photonics, Surfaces. The laboratory is internationally recognized, particularly for its work in machine learning and inverse problems. The management team for this post-doc project is composed of

- Loïc Denis, Professor (`https://perso.univ-st-etienne.fr/deniloic`)
- Rémy Emonet, Senior Lecturer (`https://home.heeere.com`)
- Amaury Habrard, Professor, Senior Member of the IUF (`https://perso.univ-st-etienne.fr/habrarda`)
- Damien Muselet, Senior Lecturer (`https://perso.univ-st-etienne.fr/muda8804`)

Applications should be sent to Loïc Denis (`loic.denis@univ-st-etienne.fr`) with a CV, a short letter of motivation, a list of publications and the contact information of 3 referees. A nationality of a European Union country is required.

Bibliography

[Bond-Taylor et al., 2021] Bond-Taylor, S., Leach, A., Long, Y., and Willcocks, C. G. (2021). Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models. *IEEE transactions on pattern analysis and machine intelligence.*

[Chen et al., 2020] Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.

[Dalsasso et al., 2021] Dalsasso, E., Denis, L., and Tupin, F. (2021). As if by magic: self-supervised training of deep despeckling networks with merlin. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13.

[Dhariwal and Nichol, 2021] Dhariwal, P. and Nichol, A. (2021). Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794.

[Goodfellow et al., 2014] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.

[He et al., 2022] He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. (2022). Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009.

[Ho et al., 2020] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.

[Karras et al., 2021] Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., and Aila, T. (2021). Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863.

[Rezende and Mohamed, 2015] Rezende, D. and Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR.

[Sankaranarayanan et al., 2022] Sankaranarayanan, S., Angelopoulos, A. N., Bates, S., Romano, Y., and Isola, P. (2022). Semantic uncertainty intervals for disentangled latent spaces. In *Advances in Neural Information Processing Systems*.