PHILIPS HEALTHTECH INNOVATION PARIS: Internship Proposals

Duration: 6 months
Preferred start date: from March 2024 or later
Localization: Paris 2
Contact: mathieu.bouyrie@philips.com

To apply, please send your résumé to the following address : mathieu.bouyrie@philips.com and indicate which internship(s) you're interested in.

Host entity

Philips is a health technology company focused on improving people's lives through meaningful innovation across the health continuum – from healthy living and prevention to diagnosis, treatment and home care. Applying advanced technologies and deep clinical and consumer insights, Philips partners with customers to deliver integrated solutions that address the Quadruple Aim: improved patient experience, better health outcomes, improved staff experience, and lower cost of care.

Philips Healthtech Innovation Paris (HIP) is based in central Paris (Sentier) and is dedicated to medical image processing. The team, with about thirty researchers and engineers, is focused on delivering the most innovative solutions in the domain and is in close contact with famous universities and clinical sites in France and abroad.

Project 1: Making Multimodality the New Normal

Internship description

Today's medical AI models often make use of a single input modality, such as medical images, clinical notes, or structured data like ICD codes. However, health records are inherently multimodal, containing a mix of provider's notes, billing codes, laboratory data, images, vital signs, and increasingly genomic sequencing, wearables, and more. The multimodality of EHR is only going to grow, having jumped twenty-fold from 2008 to 2015. No modality in isolation provides a complete picture of a person's health state. Analyzing pixel features of medical images frequently requires consulting structured records to interpret findings, so why should AI models be limited to a single modality?

The purpose of this internship is to develop a foundation model that can combine multiple modalities during training. The ability to represent multiple modalities from medical data not only leads to better representations of patient state for use in downstream applications, but also opens more paths for interacting with AI. Clinicians can query databases of medical imaging using natural language descriptions of abnormalities or use descriptions to generate synthetic medical images with counterfactual pathologies.

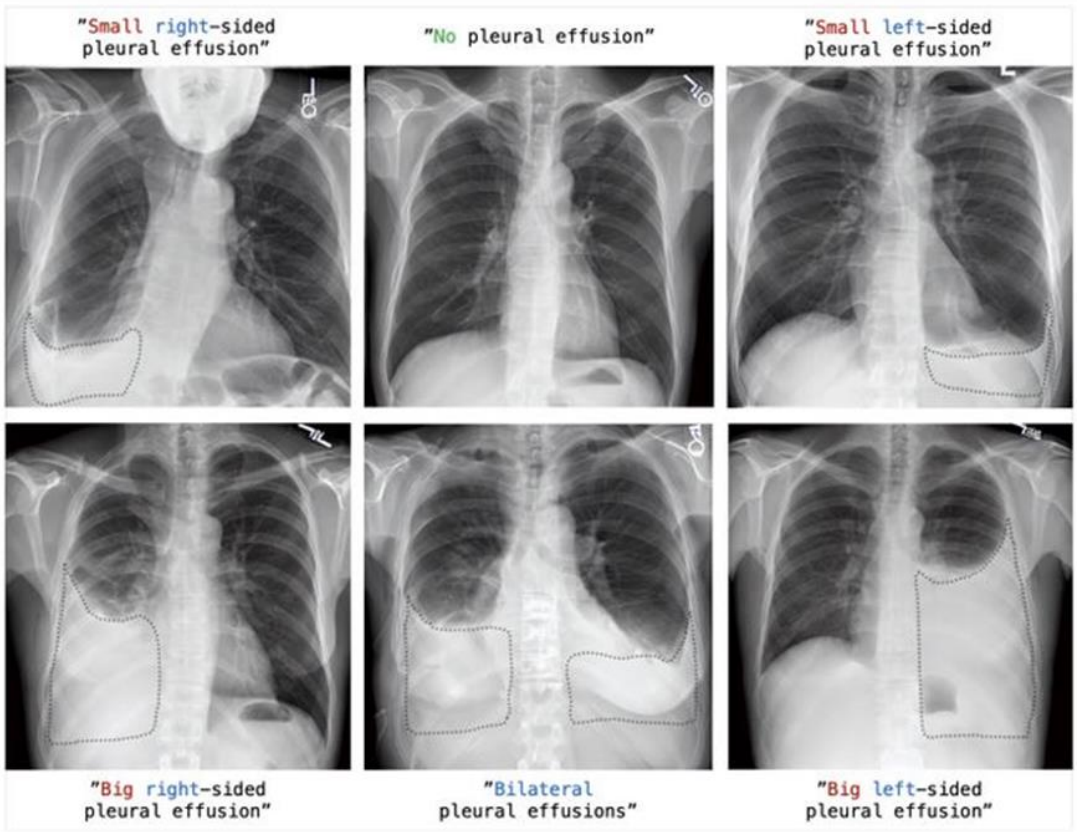

Image Credits:
RoentGen: Vision-Language Foundation Model for Chest X-ray Generation - P. Chambon, C.Bluethgen et al.

References

[1] *Michael Wornow et al. The shaky foundations of large language models and foundation models for electronic health records. npj Digit. Med. 6, 135 (2023)*

[2] *Pierre Chambon et al. ESRGAN: RoentGen: Vision-Language Foundation Model for Chest X-ray Generation. Arxiv (2022)*

Candidate profile

- Third year of engineer school / Master 2 Recherche, with specialty in machine learning, image processing or applied mathematics
- Solid knowledge of statistics, machine learning, deep learning, image processing
- Experience in Python, Tensorflow/Keras/Pytorch
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team

Project 2: Leveraging state-of-the-art motion models for device tracking in Fluoroscopy

Internship description

Image Guided Therapy is a field of medicine where practitioners operate catheters and other devices inside the patient's body, relying on real-time moving x-ray images called *Fluoroscopy.* Tracking devices in those images is of great interest, and presents unique challenges, one of them being the use of the temporal dimension with limited annotated data.

In the video processing community, models such as PIPs [1], and more recently OmniMotion [2] and CoTracker [3] have made groundbreaking progress in tracking points in videos. Those models are trained on synthetic datasets [4] designed to resemble natural videos.

X-ray sequences differ from natural videos in several ways: transparent motion (several organs such as lungs and heart moving independently on top of one another), low image quality, sliding motion of devices, frequent flow of contrast agent... Unsurprisingly, off-the-shelf models cited above work poorly when applied to fluoroscopy.

The goal of this internship is to study how those models can be adapted to work on fluoroscopic sequences. One idea would be to fine-tune them on a synthetic database, specifically designed by the intern to include the challenging features of fluoroscopy (transparent and overlapping motions, sliding catheters...)

References

[1] *Harley, Adam W., Zhaoyuan Fang, and Katerina Fragkiadaki. "Particle video revisited: Tracking through occlusions using point trajectories." European Conference on Computer Vision. Cham: Springer Nature Switzerland, 2022.*

[2] *Wang, Qianqian, et al. "Tracking Everything Everywhere All at Once." arXiv preprint arXiv:2306.05422 (2023).*

[3] *Karaev, Nikita, et al. "Cotracker: It is better to track together." arXiv preprint arXiv:2307.07635 (2023).*

[4] *Doersch, Carl, et al. "Tap-vid: A benchmark for tracking any point in a video." Advances in Neural Information Processing Systems 35 (2022): 13610-13626.*

Candidate profile

- Third year of engineer school / Master 2 Recherche, with specialty in machine learning, image processing or applied mathematics
- Solid knowledge of statistics, machine learning, deep learning, image processing
- Experience in Python
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team

Project 3: GCN- for medical image segmentation

Internship description

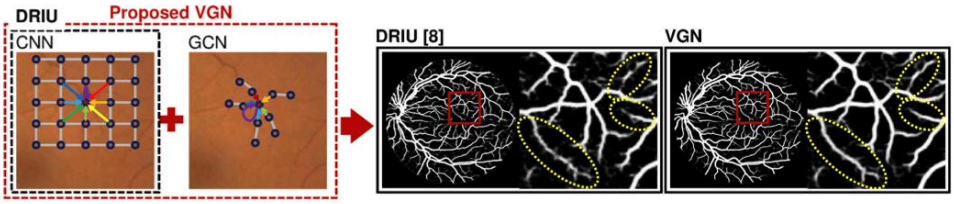

Fig. 1: Motivation of the proposed method. Learning about the strong relationship that exists between neighborhoods is not guaranteed in existing CNN-based vessel segmentation methods. The proposed *vessel graph network* (VGN) utilizes a GCN together with a CNN to address this issue. All figures best viewed in color.

Image source : https://arxiv.org/pdf/1806.02279.pdf

Examination of a coronary arterial tree before treatment can be done in multiple ways, but the fastest is the coronary angiography. This approach consists of X-ray imaging of a heart with injected contrast product into the vascular tree. The resulting angiographic video sequence with contrast vessels is used for localization and quantification of the vessel lesion for subsequent report. Automatization of this process would facilitate the work of the cardiologist a lot.

Vessel segmentation (both binary or semantic) plays significant role in the task of localization of the lesion and estimation of the diameter of the lumen. The vascular structure provides prior anatomical information for segmentation, this fact brought up several studies on this topic, which tried to leverage this prior knowledge. Graph-like structure of the vessel tree implies connectivity information; however, it was rarely used in application to cardiovascular diseases and angiography.

The goal of this internship is to study the usage of anatomical prior knowledge and graph-like structure of the vascular tree in the training of neural networks on coronary angiographies for segmentation task.

References

[1] *Liu, Lu, et al. "Anatomy-aided deep learning for medical image segmentation: a review." Physics in Medicine & Biology 66.11 (2021): 11TR01.*

[2] *Mishra, Suraj, et al. "VTG-Net: A CNN Based Vessel Topology Graph Network for Retinal Artery/Vein Classification." Frontiers in Medicine 8 (2021).*

[3] *Shin, Seung Yeon, et al. "Deep vessel segmentation by learning graphical connectivity." Medical image analysis 58 (2019): 101556.*

[4] *Zhao, Gangming, et al. "Graph Convolution Based Cross-Network Multiscale Feature Fusion for Deep Vessel Segmentation." IEEE Transactions on Medical Imaging 42.1 (2022): 183-195.*

Candidate profile

- Third year of engineer school / Master 2 Recherche, with specialty in machine learning, image processing or applied mathematics
- Solid knowledge of statistics, machine learning, deep learning, image processing
- Experience in DL in CV, Linux, git, Docker, PyTorch, data analysis stack (pandas, mpl, etc.)
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team

Project 4: Radiance-field based 3D reconstruction from sparse X-ray data

Internship description

X-Ray radiography is used extensively in Image guided therapy to acquire accurate information about a patient’s organs prior to intervention. In its usual mode of operation, this requires a densely sampled rotational scan which leads to high radiation exposure for the patient.

The objective of the internship is to explore novel computer vision and deep learning techniques, such as *Neural Radiance Fields (NERFs)* and *Gaussian Splatting*, with the end-goal of speeding up 3D reconstruction of organ structures, providing fast rotational scans from fewer scans that traditional ones.

The envisioned objectives of the internship will be

- to review the current literature about 3D scene rendering/reconstruction using sparse input data
- to implement state-of-the-art methods, such as gaussian splatting / NERFs, on X-ray data
- to compare the quality of the produced 3D reconstructions to traditional methods, and
- to evaluate the potential impact of the method in a clinical context.

Candidate profile

- Third year of engineer school / Master 2 Recherche, with specialty in machine learning, image processing or applied mathematics
- Solid knowledge of statistics, machine learning, deep learning, image processing
- Experience in Python
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team