Philips Research Paris - Internship Offers 2021 – AI and/or Software development for medical imaging applications

Duration

5 to 6 months

Preferred start date:

from March 2023 or later

Localization:

Paris (central location)

Gratification:

From 1500€ (Master 1) to 1700€ (Master 2)

How to apply?

Contact Caroline Raynaud caroline.raynaud@philips.com with CV (Cover letter is optional)

Fill also this form https://forms.office.com/r/WHVC1qjtkc

Philips Research Paris

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 Research Paris is an **R&D lab** dedicated to **medical image processing**. It is part of a Philips center for innovation on health technology based in central Paris. The team, with about forty 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.

Internship Offers

PHILIPS

Internship 1: Diffusion models for medical image segmentation

The state-of-the art in generative modeling and image synthesis has greatly improved very recently, reaching new levels of image quality and diversity. GANs, which were on the forefront of this field (1) got outdated by other kinds of models, and particularly Denoising Diffusion Probabilistic Models (DDPMs) (2,3).

Text-to-image diffusion models, such as Stable Diffusion (4) have shown how they manage to learn very abstract concepts from a large amount of unannotated data.

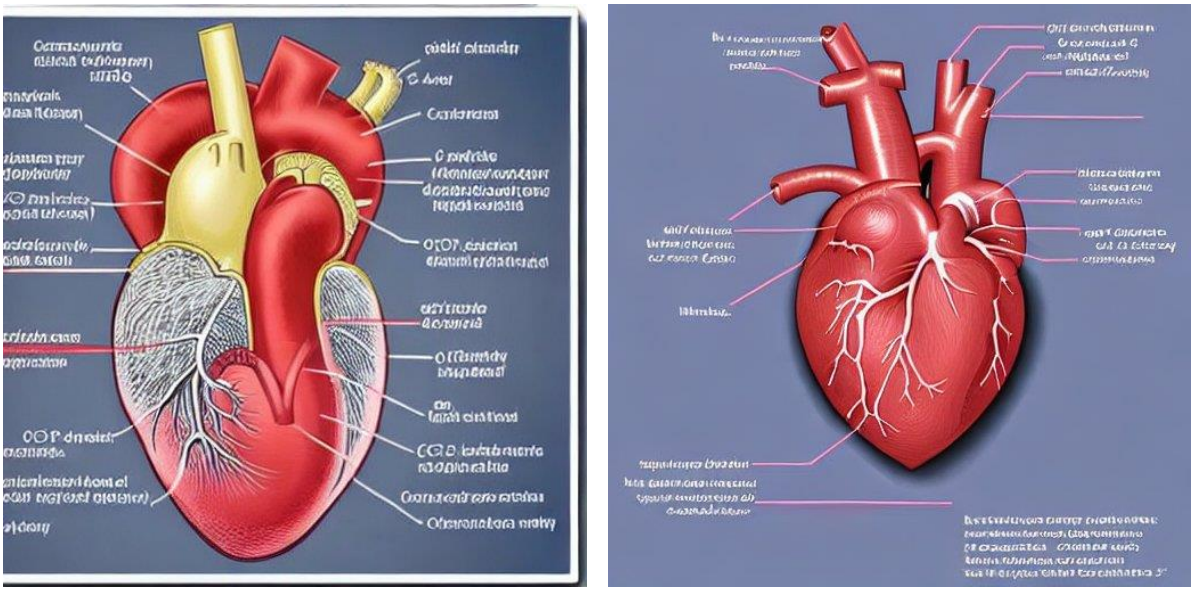

2 images generated by Stable Diffusion using the prompt "Heart Anatomy". The model has been trained on images scrapped from the internet

In medical imaging, segmentation is one of the most studied (and most useful) tasks. Many ideas have been tried to leverage anatomical or topological priors into the training of segmentation networks (5,6), but the very straightforward nnUnet (7) remains state-of-the-art for most segmentation tasks.

The goal of this internship is to study diffusion model as generative models of medical images. In particular, the ability of diffusion models to learn anatomical concepts will be investigated.

As an initial task, diffusion models will be evaluated for the segmentation of vessels in coronary angiography. Particular attention will be drawn on how anatomical concepts can be leveraged to get better segmentation models. Several strategies could be tried. Among them:

- Considering segmentation as an image-to-image task, using DDPMs to generate segmentation masks and conditioning the generation on the input image
- Using DDPMs to generate new training examples, conditioning them on segmentation masks.

In a second phase, the study can be extended to vessel segmentation and/or generation in 3D images.

References

(1) Karras, Tero, et al. "Analyzing and improving the image quality of stylegan." *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020.

(2) Ho, Jonathan, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models." *Advances in Neural Information Processing Systems* 33 (2020): 6840-6851.

(3) Dhariwal, Prafulla, and Alexander Nichol. "Diffusion models beat gans on image synthesis." *Advances in Neural Information Processing Systems* 34 (2021): 8780-8794.

(4) Rombach, Robin, et al. "High-resolution image synthesis with latent diffusion models." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022.

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

(6) El Jurdi, Rosana, et al. "High-level prior-based loss functions for medical image segmentation: A survey." *Computer Vision and Image Understanding* 210 (2021): 103248.

(7) Isensee, Fabian, et al. "nnu-net: Self-adapting framework for u-net-based medical image segmentation." *arXiv preprint arXiv:1809.10486* (2018).

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

Internship 2: Graph Convolutional Network (GCN) for medical image segmentation

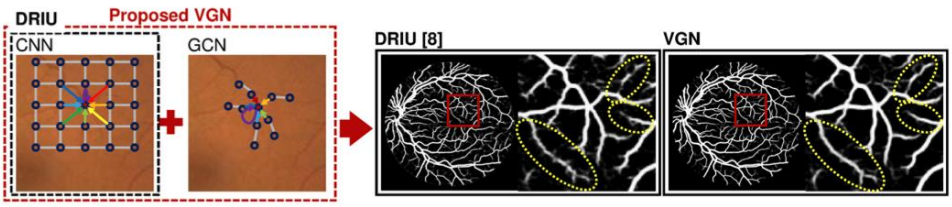

Fig. 1: Motivation of the proposed method. Learning about the strong relationship that exists between neighborhoods is not guaranteed in existing CNN-based vessel segmentation methods. The proposed *vessel graph network* (VGN) utilizes a GCN together with a CNN to address this issue. All figures best 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 Xray imaging of a heart with injected contrast product into the vascular tree. The resulting angiographic videos 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 (1, p. 4.2.6), (2), (3), (4). 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.

Temporal coherence of the segmentation can be also be a part of the research.

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) Saueressig, Camillo, et al. "A joint graph and image convolution network for automatic brain tumor segmentation." International MICCAI Brainlesion Workshop. Springer, Cham, 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, (PyTorch is preferred)
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team

Internship 3: Self-learning methods for 3D Medical Image Segmentation from single-slice annotation

As deep learning methods continue to improve medical image segmentation performance, data annotation is still a big bottleneck due to the labor-intensive and time-consuming burden on medical experts, especially for 3D images. Previous work suggested that a high proportion of pixels in an image (e.g., 80%) is redundant and most pixels can be reconstructed from other pixels with self-learning. In this internship, we propose to explore self-learning methods to deduce a 3D medical image segmentation from only a single annotated slice, which would significantly reduce annotation efforts. The main idea is to learn feature representations that enable semantic correspondence matching among slices, which is used to reconstruct and propagate expert-provided single 2D slices' annotation in testing stage. At the MICCAI conference 2022, promising results were presented in [1], where the proposed self-learning method showed performance close to the fully annotated 3D U-Net.

The goal of this internship is to investigate such methods in the application of liver segmentation in MRI. More precisely, the trainee will implement and study a framework similar to [1], and propose improvements to it. The method to select the reference slice to be annotated will be investigated, and whether the results would be improved with more than one slice in other orientations. Moreover, the effects of the presence of artifacts (e.g.: motion artifacts in T2) will be studied and solutions to handle

them will be proposed. The feasibility to integrate these methods into an improved annotation tool will be investigated. The extension of this method to CT images may also be considered.

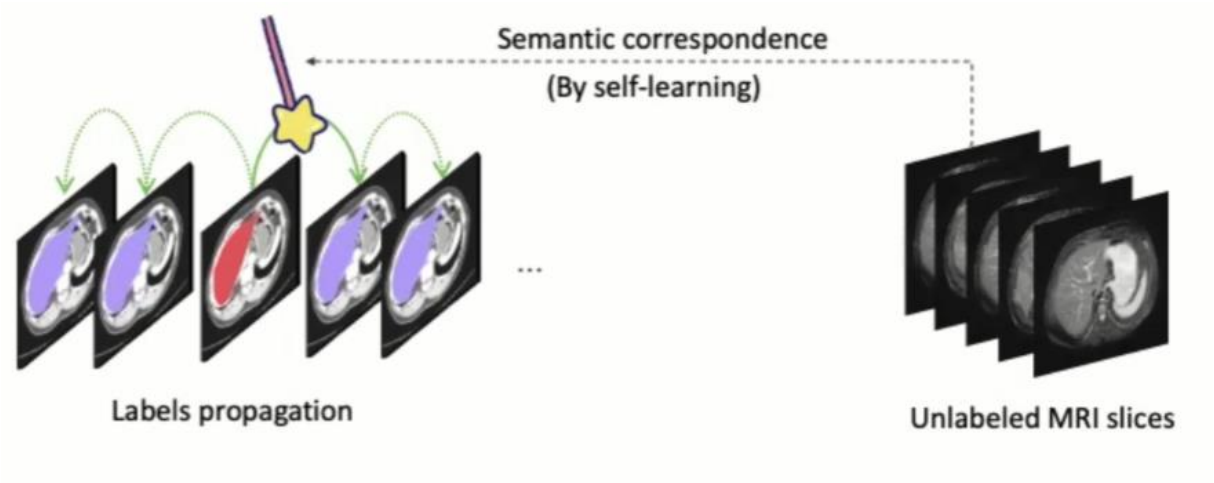

[1] Wu, Y., Zheng, B., Chen, J., Chen, D. Z., & Wu, J. (2022). Self-learning and One-Shot Learning Based Single-Slice Annotation for 3D Medical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 244-254). Springer, Cham.

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

Internship 4: Video processing of fetal ultrasound examinations

Fetal ultrasound examinations are part of routine health care during pregnancy. Those exams include global screening of the fetus, assessment of its position, anatomy and biometric measurements. During the examination, the clinician also makes pauses to explain what is displayed on screen to the parents. The recordings of examinations thus contain meaningful sequences (views which can be used for visual assessment of the anatomy of for measurements) and sequences whose content can be discarded.

The objective of this internship is to conceive and develop video processing methods to automatically select meaningful sequences or frames. It includes:

- Temporal analysis to detect sequences corresponding to pauses in the clinical exam
- Detection of key events (measurements, display of specific views), which can be done with various image processing and/or machine learning techniques, or by using optical character recognition (OCR)
- Detection of sequences which can be discarded (absence of meaningful content)

The internship will start with the selection of the most promising approaches. First implementations will use traditional image and video processing techniques, but extensions to machine learning approaches are envisioned for the most difficult tasks. A large amount of video recordings of complete ultrasound fetal examinations is available to test the methods which will be developed during the internship.

Candidate profile

- Second or Third year of engineer school or master, 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

Internship 5: Cloud-based annotation tools for Deep-Learning

The purpose of the internship is the development of a web-based annotation tool of images and videos acquired in an Education context for training automated (likely Deep Learning-based) annotations solutions. This annotation tool must allow to configure different user groups for the collective annotation of a database of images. The implementation is expected to run on AWS (for both data and app hosting). The backend must allow for multiple users to register on the platform. The annotation jobs need to be flexibly distributed to the pool of users and implement the review of images and annotations uploaded by one user to another user. The development of image and videos annotations should leverage on the LabelStudio platform (or at least its front-end part). For the back end, we favor a Python-based solution (e.g. Django).

The hosting to AWS and the break-down of the application into services will be a key part of the internship and left to the initiative of the candidate.

Candidate profile

- Second or Third year of engineer school or master, with specialty in information systems and software design, computer science
- Solid knowledge of software engineering and web-development (front-end and back-end)
- Experience in Python (Django preferred)

- Experience in Javascript (React preferred)
- Good communication skills and ability to work in a team

Internship 6: Ultrasound image quality, beyond the practical constraints

Image quality is essential for medical image analysis and diagnosis assessment. In ultrasound imaging, the image quality is closely linked to the transmit parameters of the ultrasound waves. However, setting the optimal parameters to maximize the image quality, usually comes with the compromise of the frame rate. Moreover, an ideal transmit configuration may be not compatible with the practical hardware constraints or with the specificities of a given clinical application.

Using an ultrasound software simulator [**1**] allows defining an optimal transmission configuration regardless the practical constraints. Such an optimal transmit sequence could be used as a target for a network that learns how to reproduce high quality images from low quality images, closer to realistic sequences. Denoising diffusion models like SR3 [**2**] could be a good candidate but they may not be sufficient. Exploring how to combine the new network architecture with a Philips AI-based model will be one of the objectives of this internship. The results will be evaluated both qualitatively and quantitatively.

References

[1] Garcia D. **SIMUS: an open-source simulator for ultrasound imaging. Part I: theory & examples.** *Comput Methods Programs Biomed*, 2022;218:106726.

[2] Saharia C. *etal.* **Image super-resolution via iterative refinement**. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 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/signal processing
- Experience in Python, Matlab
- Knowledge of ultrasound acoustics could be a plus
- English speaking, reading and writing is mandatory
- Good communication skills and ability to work in a team