ERCIM fellowship Programme

Final scientific report

Fellow	Ezgi Demircan Türeyen
Host Organisation	Centrum Wiskunde & Informatica (CWI)
Scientific coordinator	Prof. Tristan van Leeuwen

I – SCIENTIFIC ACTIVITY DURING YOUR FELLOWSHIP

During my ERCIM fellowship, my research centered on diffusion models and their application to computed tomography (CT). I specifically focused on out-of-distribution (OOD) detection in CT reconstruction, a topic with considerable potential in tomographic reconstruction systems based on deep generative models. My research followed a comprehensive path, spanning from theoretical and practical understanding of diffusion models to the conceptualization of the project, experimental design, and execution.

At first, I focused on gaining a deep understanding of diffusion models—including score-based generative models, DDPMs, and Score-SDEs—since I had limited prior experience with them. This led to a presentation I gave within our research group, which sparked some engaging discussions afterward. Later, I shifted my attention to exploring the intersection of generative models, inverse problems, and OOD detection. Through an extensive literature review, I realized that while there is significant research showcasing the potential of diffusion models as unsupervised solvers for inverse problems, and as unsupervised tools for OOD detection; studies specifically linking these two areas are lacking. To bridge this gap, we defined a research objective focused on CT reconstruction augmented by an OOD detector, both utilizing the same pretrained diffusion model to capture the prior distribution of the target images. The decision to focus on CT reconstruction, among various inverse problems, was motivated by the group's expertise in this area and the potential to apply the research to real-world scenarios using the experimental facilities.

OOD detection is a versatile task. It can be applied to enhance system reliability by filtering out data beyond its scope, or to intentionally inspect abnormalities (e.g., defects, adverse pathologies) within the data. A CT reconstruction system may require both, and in both cases, due to the limited knowledge of the OOD space, the problem is approached in an unsupervised manner. This aligns with CT reconstruction, where we aim to learn the prior distribution rather than the posterior, as this way, the model is not constrained by specific measurement parameters and can generalize when those parameters change. Therefore, we conceptualized the project so that, given a diffusion model trained on whatever we consider in-distribution, we could explore how to make CT reconstruction a downstream task for an OOD detector. Here, the OOD detector follows a reconstruction-based approach, which utilizes partial diffusions and relies on the model's inability to accurately reconstruct from partially corrupted OOD data. Moreover, we focused solely on sparse-view CT, as it is the most common CT technique for reducing radiation exposure. This setup led to many intriguing questions, which we have distilled into three: (1) How should the reconstruction error be defined for use by the OOD detector in sparse-view tomography? —including considerations such as the input to the diffusion model, the two signals between which the error is computed, and whether to reconstruct from partially diffused input by sampling from the prior or the approximate posterior. (2) Does sampling from approximate posterior consistently aid OOD detection? —given that it is generally better, identifying scenarios where it may mislead OOD detection while aiding CT reconstruction. (3) Can we consider weighting reconstruction errors from both unconditional and conditional samples to enhance the robustness of the OOD score? —where we propose an approach to cope with the cases explored in the research question (2).

Our investigations demonstrated that the OOD detection strategy of introducing varying levels of corruption to the input image and then reconstructing it via a diffusion model can be effectively applied to CT reconstruction as well. The input must be accordingly adjusted, with filtered backprojection (FBP) mappings from sparse-view projections emerging as a practical solution. However, to ensure that OOD images are reconstructed with lower quality than in-distribution images, the missing information must be compensated with an additional source of data fidelity. This sets the stage for conditioning on the measurement, as is done in CT reconstruction. After obtaining the reconstructions, calculating the error in the projection domain by comparing sparse sinograms serves as the most effective surrogate for evaluating reconstruction quality. Overall, such an OOD detector tends to mirror the performance of the CT reconstruction, which we usually evaluate with PSNR/SSIM metrics against ground truth. While this is already useful in assessing the system's confidence, the goal may be to detect measurements arising from any semantic shift from the in-distribution, regardless of how accurately they are CT-reconstructed. In certain contexts where we aim to inspect anomalies, the fact that an image is well reconstructed by the CT system does not necessarily mean that it is anomaly-free. However, when the goal is to detect semantic shifts, more caution is needed when conditioning on the measurements. In this regard, we proposed a weighting mechanism that balances the errors from both conditional and unconditional reconstructions.

Our study has resulted in a preprint, which will soon be made openly accessible on arXiv. We plan to submit it to a journal, and we believe it will attract interest from various fields, including inverse problems, machine learning, signal processing, and tomography. The findings have already been presented in a seminar during my research exchange program and in a talk at the MAC-MIGS workshop, where the work was well-received by participants from various backgrounds. During the fellowship period, I was in strong collaboration with Dr. Felix Lucka, along with my scientific supervisor Prof. Tristan van Leeuwen. They offered me an opportunity to collaborate further by applying the knowledge I've gained on deep generative models to a new research area in dynamic CT imaging, which I gladly accepted. This presents a chance for me to develop a new specialization, explore potential future collaborations, shape a research-focused career path, and continue learning from their expertise.

II – PUBLICATION(S) DURING YOUR FELLOWSHIP

Ezgi Demircan-Tureyen, Felix Lucka, and Tristan van Leeuwen, *Exploring Out-of-distribution Detection for Sparse-view Computed Tomography with Diffusion Models* (to be submitted).

Abstract: Recent works demonstrate the effectiveness of employing diffusion models as unsupervised solvers for inverse imaging problems. Sparse-view computed tomography (CT) has greatly benefited from these advancements by offering better generalizability to unknown measurement processes, as these models are independent of measurement parameters. However, these benefits come at the cost of hallucinations, especially when confronted with out-of-distribution (OOD) data. Therefore, there is a driving need to study OOD detection for CT reconstruction in both clinical and industrial applications, in order to ensure reliability and enable inspection. Fortunately, once a diffusion model is trained to capture the distribution of interest, it can also function as an OOD detector that assesses OOD-ness based on reconstruction error. Nonetheless, since in sparse-view tomography the input is incomplete, the concept of reconstruction error needs to be redefined. In this paper, we conceptually explore the feasibility of

posing sparse-view CT as a downstream task within a reconstruction-based OOD detection scheme. We delve into the intricacies involved in augmenting CT reconstruction machinery with an OOD detector, with both employing the same diffusion model. Our proof-of-concept experiments on MNIST dataset showcase various failure and success scenarios to clearly understand the limitations and potential. Moreover, we introduce a novel approach to measuring reconstruction error, which improves the robustness of the OOD detector and suggests a promising direction for research in sparse-view CT.

III – ATTENDED SEMINARS, WORKHOPS, CONFERENCES

- Speaker, "Exploring Out-of-distribution Detection for Sparse-view Computed Tomography with Diffusion Model", Dutch Inverse Problems Meeting, Delft, The Netherlands (to be presented in November 2024) [https://www.aanmelder.nl/dutch-inverse-problems-meeting/new-page]
- Speaker, "Exploring Out-of-distribution Detection for Sparse-view Computed Tomography with Diffusion Model", MAC-MIGS Workshop on Scientific Computation, Statistics, and PDEs, Amsterdam, The Netherlands (September 2024) [https://www.cwi.nl/nl/events/mac-migs-amsterdamutrecht-workshop-on-scientific-computation-statistics-and-pdes/]
- Participant, Lorentz Center workshop: Integrating Acquisition and AI in Tomography, Leiden, The Netherlands (November 2023) [https://www.lorentzcenter.nl/integrating-acquisition-and-ai-in-tomography.html]

IV – RESEARCH EXCHANGE PROGRAMME (REP)

I had the opportunity to visit Dr. Thomas Moreau and the Models and Inference for Neuroimaging Data (MIND) team at INRIA, France, from July 15th to July 19th, 2024. During my visit, I presented a seminar titled "Exploring Diffusion-based Out-of-distribution Detection for Diffusion-based Sparse-view Computed Tomography" to the team, which was well-received. The constructive feedback I received played a significant role in refining and shaping my research further.

Moreover, the time I spent there allowed me to take a step back from my usual work and shift my focus towards the practical, implementational aspects of my research. I explored a PyTorch-based deep learning library* for inverse imaging developed by INRIA researchers, which gave me hands-on experience with tools relevant to my work. This exploration helped broaden my technical understanding and inspired new ideas for potential future projects and collaborations.

* https://deepinv.github.io/deepinv/