WS 2025-2026, SMT-EXC, DISAL-MP61 Project Summary

Gas Distribution Modelling using Learning Techniques

Marcel Zelent

Professor : Alcherio Martinoli, Dimitrios Papageorgiou (DTU)
Assistant(s) : Nicolaj Bösel-Schmid

Gas distribution mapping (GDM) aims to reconstruct the dispersion of a gas inside an indoor area from sparse and noisy measurements, typically collected by a robot. Classical methods such as Kernel DM and Gaussian Markov Random Fields rely on statistical approaches, but they struggle to generalize to complex environments with obstacles, disturbed airflow, and limited observations. As a result, recent work has shifted toward learning-based methods, such as Convolutional Neural Networks (CNNs). This thesis investigates whether a newer architecture, the Fourier Neural Operator (FNO), can offer improved performance.

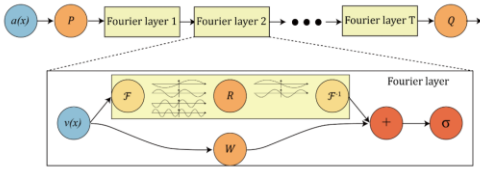

Schematic of the FNO architecture

The evaluation is conducted on two synthetic datasets. The first is generated using the Gaussian plume model over a 2D grid, with varying source position, emission strength, and wind conditions. The second is derived from prior work by Wanting et al., where gas is released in a wind-tunnel-like environment with obstacles. Together, these datasets allow testing under both idealized, controllable conditions and more realistic scenarios where obstacles and flow disturbances increase complexity.

To bridge the gap between simulation and real-world deployment, a realistic metal-oxide sensor model is introduced, incorporating first-order temporal delay and signal-dependent Gaussian noise. In addition, robot data collection is modeled using path-planning strategies, including a simple scanning pattern and a more advanced coverage path planning approach that adapts to obstacles using breadth-first search backtracking.

The CNN is implemented as a standard U-Net with an encoder-decoder structure and skip connections, while the FNO uses spectral convolutions in Fourier space, where low-frequency modes are transformed using learnable weights and high-frequency components are truncated.

Results on the Gaussian plume dataset show that both models perform very well in simple and structured settings. When the source location is fixed, both CNN and FNO accurately reconstruct the gas field even from very sparse inputs (<5% of the original data). However, as the task becomes more challenging, such as with randomized source positions, realistic sensor effects, or reduced observation coverage, the CNN consistently proves to be more robust. While the FNO still achieves good results, its error is generally higher.

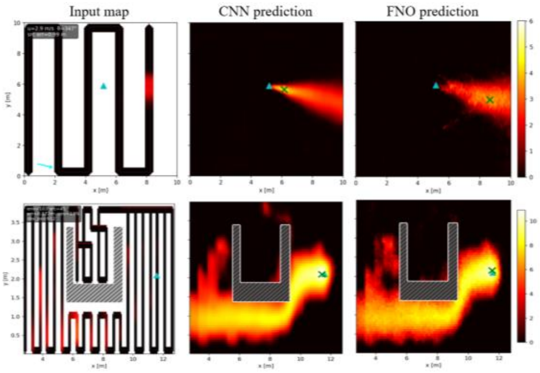

Reconstructions by the models. Top row: GPM dataset with random source placement. Bottom: Wanting dataset with CPP algorithm. Both with MOX sensor model

This trend follows on the more realistic Wanting dataset. Both models produce useful reconstructions, but the CNN typically outperforms the FNO in both reconstruction quality and source localization accuracy. This advantage is especially evident in generalization experiments, where models trained on one obstacle layout are tested on another. Lastly, the obstacle-aware coverage path planning was successfully integrated into the reconstruction pipeline without degrading performance.

Overall, the results suggest that while FNOs are well suited for smooth, globally structured physical fields, they are less robust in scenarios where local irregularities dominate. For gas distribution mapping in realistic robotic settings, this makes the CNN a more reliable choice, despite the theoretical advantages of operator-based learning.

The thesis was a joint project between the Technical University of Denmark and EPFL.