Gas Source Localization using Neural Networks

Jad Benabdelkader

Professor : Alcherio Martinoli

Assistant(s) : Nicolaj Schmid

Gas Source Localization (GSL) is the task of identifying the position of a gas-emitting source based on concentration measurements. It is a key problem in robotics for safety-critical scenarios like leak detection in industrial plants or emergency response in hazardous environments. While classical methods often rely on reactive strategies or physical plume models, they suffer in turbulent, obstacle-rich environments and require strong prior knowledge or assumptions. This project investigates a data-driven alternative using deep learning.

We propose an architecture based on an **Attention U-Net** trained on a large-scale dataset of gas plumes simulated in 200 different indoor environments with varying obstacle geometries. Each input to the model consists of three channels: a Signed Distance Field (SDF) encoding obstacle geometry, a partial gas concentration map simulating noisy sensor data along a robot trajectory, and a visited cell mask indicating which areas have been explored. The model is trained to output a heatmap of likely gas source locations.

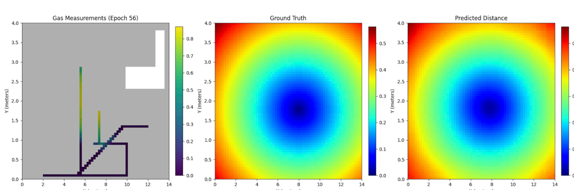

Ground truth and model output on a partial gas map.

Two ground truth representations are explored: (1) a Gaussian heatmap centered on the true source, used with **Kullback-Leibler** divergence loss, and (2) a distance transform map used with **Mean Squared Error** (MSE) loss. We show that both representations enable learning, but exhibit different trade-offs in precision and robustness. Experiments begin with a fully observed gas map to validate the model under ideal conditions. In this setting, the U-Net achieves nearly perfect source localization with sub-decimeter accuracy. We then proceed to more realistic conditions using partial maps, synthetic noise, and varying robot trajectories. A comprehensive **ablation study** using cross-validation confirms that each model component—attention, SDF encoding, and the visited mask—contributes to performance. The full model trained with MSE loss achieves the best trade-off between **accuracy** and **robustness**, outperforming simplified variants.

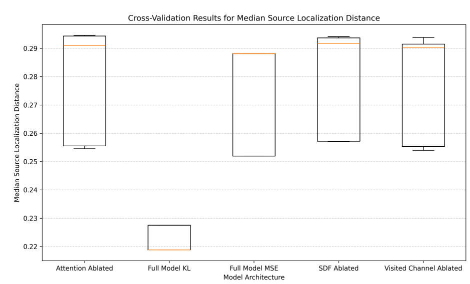

Median localization error across model variants (lower is better).

Finally, we evaluate the model on a **held-out test set** with unseen environments and goal-directed exploration trajectories. Here, the model's performance degrades sharply: mean localization errors rise above 3.5 meters and the Dice score drops to zero. Further experiments show that neither the new geometries nor the different trajectory structures alone are responsible. Instead, the main issue appears to be **sensor modeling discrepancies.**

This work confirms that learning-based GSL is feasible and effective under controlled conditions, but highlights open challenges in generalizing to unseen environments and sensor dynamics. Future directions include integrating physics-informed priors, reinforcement learning for active exploration, and sim-to-real adaptation to bridge the gap between simulated and real-world deployment.