


TRUSTED
EXTREMELY PRECISE
MAPPING AND PREDICTION
FOR EMERGENCY
MANAGEMENT

D4.3

Report on Phenomenon Prediction and Response Planning







Project Information

Project acronym:	TEMA
Project full title:	Trusted Extremely Precise Mapping and Prediction for Emergency Management
Call identifier:	HORIZON-CL4-2022-DATA-01
Type of action:	HORIZON Research and Innovation Actions
Start date:	1 December 2022
End date:	30 November 2026
Grant agreement no:	101093003

D4.3- Report on Phenomenon Prediction and Response Planning			
Executive Summary:	<p>Deliverable D4.3 Report on Phenomenon Prediction and Response Planning, is the third deliverable of Work Package 4 (WP4) within the TEMA project. This document encapsulates the research achievements of Tasks T4.1 and T4.2 over M19-M36 of the project. See the Section Executive Summary.</p>		
WP:	4		
Author(s):	See table below for a full list of authors		
Editor:	J. Ramiro Martínez de Dios, Abdalraheem Ijeh		
Leading Partner:	USE		
Participating Partners:	All		
Version:	1.0	Status:	final
Deliverable Type:	R Document, report	Dissemination Level:	Public
Official Submission Date:	30 November 2025	Actual Submission Date:	30 November 2025







Disclaimer

This document contains material, which is the copyright of certain TEMA contractors, and may not be reproduced or copied without permission. All TEMA consortium partners have agreed to the full publication of this document if not declared Confidential. The commercial use of any information contained in this document may require a license from the proprietor of that information. The reproduction of this document or of parts of it requires an agreement with the proprietor of that information.


The TEMA consortium consists of the following partners:


No.	Partner Organization Name	Partner Organization Short Name	Country
1	ARISTOTELIO PANEPISTIMIO THES-SALONIKIS	AUTH	GR
2	DEUTSCHES ZENTRUM FUR LUFT UND RAUMFAHRT EV	DLR	DE
3	ENGINEERING - INGEGNERIA INFORMATICA SPA	ENG	IT
4	ATOS IT SOLUTIONS AND SERVICES IBERIA SL	ATOS IT	ES
4.1	ATOS SPAIN SA	ATOS SP	ES
5	UNIVERSIDAD DE SEVILLA	USE	ES
6	TECNOSYLVA SL	TSYL	ES
7	NORTHDOCKS GMBH	ND	DE
9	THE LISBON COUNCIL FOR ECONOMIC COMPETITIVENESS ASBL	LC	BE
10	LATITUDO 40 SRL	LAT40	IT
11	NELEN & SCHUURMANS TECHNOLOGY BV	NS	NL
11.1	NELEN & SCHUURMANS CONSULTANCY BV	NS C	NL
12	FRAUNHOFER GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG EV	FHHI	DE
13	UNIVERSITA DEGLI STUDI DI MESSINA	UNIME	IT
14	KAJAANIN AMMATTIKORKEAKOULU OY	KAMK	FI
16	KENTRO MELETON ASFALEIAS	KEMEA	GR
17	DIMOS MANTOUDIYOU - LIMNIS - AGIAS ANNAS	D.MALIAN	GR
18	REGIONE AUTONOMA DELLA SARDEGNA	RAS	IT





19	BAYERISCHES ROTES KREUZ	BRK	DE
20	KAINUUN HYVINVOINTIALUE	KAHY	FI
21	INTERDISCIPLINARY TRANSFORMATION UNIVERSITY	IT:U	AT







Document Revision History

Version	Description	Contributions
o.1	First draft	J. Ramiro Martinez, A. Ijeh, David Hanny, Bernd Resch, Ehsaneddin Jajilian, Sebastian Schmidt, Dorian Arifi, Afzal Ahmad, Golnesa Karimizindashti, Dmitriy Shutin, Victor Prieto Ruiz, Joost van Dijk, Nicolette Volp, Joep Grispen
o.2	Second draft	J. Ramiro Martinez, A. Ijeh, David Hanny, Bernd Resch, Ehsaneddin Jajilian, Sebastian Schmidt, Dorian Arifi, Afzal Ahmad, Golnesa Karimizindashti, Dmitriy Shutin, Victor Prieto Ruiz, Joost van Dijk, Nicolette Volp, Joep Grispen
1.0	Final version ready for submission.	J. Ramiro Martinez, A. Ijeh

Authors


Name	Partner
J. Ramiro Martinez	USE
Abdalraheem Ijeh	USE
David Hanny	IT:U
Bernd Resch	IT:U
Ehsaneddin Jajilian	IT:U
Sebastian Schmidt	IT:U
Dorian Arifi	IT:U
Afzal Ahmad	IT:U
Joost van Dijk	NSC
Golnesa Karimizindashti	NST
Nicolette Volp	NST
Joep Grispen	NST
Dmitriy Shutin	DLR
Victor Scott Prieto Ruiz	DLR
Miguel Navarrete	TSYL
Margareta Mihalic-Dogan	BRK





Reviewers

Name	Partner
Leila Arras	FHHI
Galip Ümit Yolcu	FHHI
Sebastian Lapuschkin	FHHI









Table of Contents


Table of Contents	7
List of Figures	9
List of Tables	13
List of Terms and Abbreviations	14
Executive Summary	16
1 Introduction	18
1.1 Purpose and scope of the document	18
1.2 Structure of the document	18
2 Summary of the work carried out	19
2.1 Objectives	19
2.2 Summary of the work carried out with respect to the objectives	20
2.2.1 Wildfire smoke dynamics modeling	20
2.2.2 Wildfire behavior modeling	21
2.2.3 Near real-time flood modeling	21
2.2.4 Disaster prediction from user-generated data sources	22
2.2.5 Response planning	22
3 Wildfire smoke dynamics modeling	24
3.1 Introduction	24
3.2 State of the art and TEMA's contribution	24
3.3 Smoke dynamics modeling	27
3.3.1 Signal model for smoke/gas source localization	27
3.3.2 Off-Grid Bayesian smoke source localization	31
3.3.3 Applications and extensions: Cramer-Rao Lower Bound for a Gas Source Localization	38
3.3.4 Applications and extensions: Distributed Gas Source Localization	44
3.4 Experimental campaign in Montiferru	55
3.4.1 Montiferru Data Collection	55
3.4.1.1 18.o6 Wood Pile Burn Setup and Wind Conditions	55
3.4.2 19.o6 Simulated Burn Setup and Ground Truth Collection	57
3.4.3 Smoke Model Evaluation	59
3.4.3.1 Model Definition	59
3.4.3.2 Model Evaluation Methods	60
3.4.3.3 Model Evaluation 18.o6 - Wood Pile Burn	60
3.4.3.4 Model Evaluation 19.o6 - Prescribed Burn	63
3.4.3.5 Discussion	65
3.5 Fulfillment of TEMA KPIs	66
4 Wildfire behavior modeling	68
4.1 State of the art and TEMA's contribution	68
4.2 Summary of advances in the last period	69






4.3	Enhancement of the weather forecast model with real-time data	70
4.4	Calibrated simulations through actual fire fronts locations	72
4.5	Live Fuel Moisture Content (LFMC)	74
4.6	Fulfillment of TEMA KPIs	75
5	Near real-time flood modeling	77
5.1	Introduction	77
5.2	State of the art and TEMA's contribution	77
5.3	Summary of advances in the last period	79
5.4	Hydrodynamic modeling with Clone Cells	79
5.5	Fulfillment of TEMA KPIs	82
6	Disaster prediction from user-generated data sources	85
6.1	Introduction	85
6.2	State of the art and TEMA's contribution	85
6.3	Summary of advances in the last period	86
6.4	Disaster observation with user-generated data	86
6.5	Disaster prediction using user-generated data	88
6.6	Fulfillment of TEMA KPIs	88
7	Response Planning	90
7.1	Introduction	90
7.2	State of the art and TEMA's contribution	90
7.3	Response planning requirements	92
7.3.1	Non-Functional Requirements	92
7.3.2	Functional Requirements	92
7.4	Probability maps as the foundation for response planning	93
7.5	Multi-target search theory and information quantification	94
7.6	From single-look updates to an information footprint	96
7.7	Search-effort unification via an exponential detection law	97
7.8	Practical multi-UAV planning implementation	97
7.9	Integration within the TEMA platform	98
7.10	Fulfillment of TEMA KPIs	99
8	Conclusion	101
	References	103
	Annex A	112







List of Figures


1	2D Green's function for stationary advection-diffusion Partial Differential Equation (PDE) in 2D with source located at $\theta = [-0.7, 0.7]^T$ m, with wind direction set to $v = [0.5, 0.5]^T$ m/s, and diffusion parameter $\kappa = 0.1$ m ² /s.	30
2	Example of a Green's function $G(p, \tilde{\theta})$ for the source at $\tilde{\theta} = [0, 0]^T$. Red marks on the right image indicate random positions of measurements y . The Green's function has a singularity at $\tilde{\theta}$ not shown in the plot.	42
3	Frequency of estimating $\tilde{\alpha} < \infty$ versus number of measurement samples M for different Signal-to-Noise Ratio (SNR) values.	42
4	Conditional Cramer-Rao Lower bound (CRLB) $\text{tr}\{C(\tilde{\theta} \hat{\alpha})\}$ versus number of measurement samples M for different SNR values.	43
5	Conditional CRLB $\text{tr}\{C(\tilde{\theta} \hat{\alpha})\}$ versus SNR for different number of measurements M	43
6	Change in CRLB as a function of a new measurement location $p^* \in \Omega$ with for $M=20$ (left) and $M=60$ (right).	44
7	(a) f_{true} concentration for $L_{\text{true}} = 3$ sources. Dots mark smoke sources. (b) Agent sampling trajectories. (c) Initial source locations ϑ_l . (d) State of the bottom-up initialization for the agent $k=0$. Dots denote source positions.	48
8	Final estimation results at agent $k = 1$. (a) Estimated smoke concentration and source locations; (b) estimated Earth Mover's Distance (EMD), Mean Squared Error (MSE), and the number of sources \hat{L} versus iterations.	49
9	(a) median and 90% confidence for EMD and MSE performance, as well averaged number of detected sources \hat{L} versus M_k for different SNR. (b) Histograms of \hat{L} for different numbers of measurements and SNR values.	50
10	Adaptive Distributed Super-resolution Gas Source Localization (DSR-GSL) algorithm.	52
11	(a, left) f_{true} concentration for $L_{\text{true}} = 3$ sources marked by dots, (a, right) agent sampling trajectories; (b) Estimated EMD, MSE, and the number of sources \hat{L} versus number of measurements per agent.	53
12	(top, middle) Solid lines shows median EMD and MSE performance vs. number of measurements per agent for adaptive DSR-GSL (aDSR-GSL); (bottom) averaged number of detected sources \hat{L} vs. number of measurements. Shaded regions show 90% confidence interval, color indicates the same SNR for all plots. Dashed lines show the non-adaptive DSR-GSL performance.	54
13	Prescribed Burn Setup on 18.o6. This plot shows the physical setup of the prescribed burns on the first day of the Sardinia Trial. The RGB image in the background is a photogrammetric reconstruction of the domain using a DJI drone on site, before the burns occurred. The pale orange dots represent the locations of the static Ground Sensor network DLR-KN implemented. The light pink areas represent the locations of burning wood piles. The background is an OpenStreetMap tile for reference.	55
14	Wood Piles Burned During Sardinia Trial. The left image shows an unburned wood pile, roughly 10×3 m in size. The right shows a burning wood pile, collected from a DJI-M3l drone, and shows the crew of RAS involved in ensuring it burns safely. Note that though the fire is burning visibly in the image, there is little visible smoke.	56







15	Hardware for Ground Truth Collection. The left represents the Ground Sensor units with trace gas sensing capabilities and anemometer. The right unit represents a Holybro X500 drone with a PCB that has trace gas sensors.	56
16	Observed Winds During 18.o6 Prescribed Burn Period. The locations of the wind sensors are represented by red dots and the wind rose graphs fanning the points show both (1) the direction from which the wind came and (2) relative frequencies of these winds and their magnitudes. The wind direction is consistent between the sensors (with slight deviation due to errors in placement).	57
17	Prescribed Burn Setup on 19.o6 (Wood Piles and Grassland Burn). This plot shows the physical setup of the prescribed burns on the second day of the Sardinia Trial. The RGB image in the background is a photogrammetric reconstruction of the domain using a DJI drone on site, after the burns occurred. The pale orange dots represent the locations of the static Ground Sensor network DLR-KN implemented. The dark red dot represents an auxiliary anemometer, which collected additional wind data to compensate for the lack of data from one sensor. The orange areas represent the locations of burning wood piles. Additionally, the blue areas represent an estimate of the burnt regions. The background is an OpenStreetMap tile for reference.	58
18	Observed Winds During 19.o6 Prescribed Burn Period. The locations of the wind sensors are represented by red dots, and the wind rose graphs fanning from the points show both (1) the direction from which the wind came and (2) relative frequencies of these winds and their magnitudes. The third wind station present in the northernmost location is an auxiliary ultrasonic anemometer and was not connected to DLR-KN's GroundSensor Network.	59
19	18.o6 Smoke Model Evaluation. The Smoke Model evaluated at a sample time of drone data collection, with the red dot corresponding to the location of the active wood pile. The wind value corresponds to the prevailing wind calculated over the time of collection. Values lower than 0.001 are not displayed on the plot for clarity to show the problem domain. Note that the northernmost woodpile is not included in this model, due to the fact that it had already burned by the time the drone collected data.	61
20	18.o6 Model v. Drone Trace Gas, Spatial Representation The left image is the model prediction. The right shows trace gas data collected by the human-flown drone. Both models show the (expected, actual) conditions overlaid on the physical location of the drone at the timestamp of the drone path.	62
21	18.o6 Drone and Time Series Evaluation. The simulated smoke field compared to the observed measurements from the drone. Green vertical lines correspond to a plume detection by the drone. Gaps in the data reflect times the drone was not in the air or exceeded a certain altitude, which affects the voltages of the trace gas sensor.	63
22	19.o6 Smoke Model Evaluation. The Smoke Model evaluated at a sample time of drone data collection, with the red dots corresponding to the location of the active wood piles and orange dots the grass burnt. The wind value corresponds to the (average) prevailing wind calculated over the whole time of collection. Values lower than 0.001 are not displayed on the plot for clarity to show the problem domain.	64





23	19.06 Model v. Drone Trace Gas, Spatial Representation The left image is the model prediction. The right shows trace gas data collected by the human-flown drone. Both models show the (expected, actual) conditions overlaid on the physical location of the drone at the timestamp of the drone path.	64
24	19.06 Drone and Time Series Evaluation. The simulated smoke field compared to the observed measurements from the drone. Green vertical lines correspond to a plume detection (peak) by the drone. Gaps in the data reflect times the drone was not in the air or exceeded a certain altitude, which affects the voltages of the trace gas sensor.	65
25	Finite Element Mesh for Domain Representation of the Moniferru region.	66
26	Effect of a single wind measurement (red) on a constant horizontal flow (yellow) using different weights parameter alpha and beta. Main wind from the west in the example.	71
27	Effect of three different wind stations (red) on a constant horizontal flux (yellow). Main wind from the west in the example.	72
28	Example of extraction of control points from a supposed first hours Maps4Fire. . .	73
29	Selection of control points within the standard simulation perimeter.	74
30	On the left, a standard simulation using only the weather forecast and scenario data. On the right, the same simulation is enhancing the weather forecast with actual weather information during the fire and extracting fire fronts' locations as control points from the RAS report.	75
31	Comparison between standard and adjusted simulation using RAS trial inputs. . .	76
32	Illustration of the grid systems used in 3Di; coarser grid (green) as computational grid and finer grid (grey) as subgrid	78
33	Illustration of a computational grid including a quadtree grid refinement	78
34	a) a computational cell without grid cloning, b) a computational cell with grid cloning, c) active domain of each clone cell	80
35	Total number of computational cells using: a) quadtree grid refinement, b) clone cells	81
36	Grid system of the test case (the yellow line represents the cutline and the colored fragments the clone cells)	81
37	The discharge and water level map of the test case	81
38	Overview of computational grid with a flooded domain in Blue and in Green, an obstacle that should block the flow. Left) Original approach, having refinements in combination with a method that blocks the flow on cell edges. Right) New setup with clone cells (marked cells), where active domains can be defined accurately. . .	83
39	Snapshots of flood simulation results in the Mantoudi region. The left image is the simulation results at 2 hours after the start of the flood based on preliminary (basic) data. The second to fourth images are flood simulation results at respectively 2 hours, 4 hours, and 6 hours after the start of the flood, based on enhanced data from the TEMA platform.	84
40	Relation between post count and presence of flooding on H3 grid (level 4). Regions without posts were excluded for legibility.	87
41	Overview of our systematic evaluation of disaster prediction methodologies, including the Foundation Model (FM)-based methods.	89
42	Probability map example for the historical flood case of Ahrtal in Germany. The blue intensity represents the probability of flood presence, with darker blues indicating higher probabilities. This visualization demonstrates the spatial heterogeneity of disaster likelihood across the affected region. (The squares on the grid do not match the actual pixels.)	93







43 Mathematical basis of information-driven planning. 94

44 Complete planning pipeline showing transformation of probability maps through information quantification, search effort optimization, area decomposition, multi-UAV coordination, and trajectory generation. Each stage maintains mathematical optimality while addressing practical constraints. 95


45 Final optimized drone trajectories demonstrate efficient coverage of priority areas while respecting operational constraints. The paths show systematic coverage of high-priority regions with efficient transitions between subareas and safe separation between multiple UAVs. 99






List of Tables


- 1 Average model performance across the two disaster events on the human-engineered base prompt versus the optimised prompts after Automatic Prompt Optimization (APO) with zero-shot learning. All results were computed on a subset of 40 samples per disaster (20 disaster-positive and 20 disaster-negative) to minimise inference time and costs. The best scores are marked in bold, the second-best scores are italic. 89







List of Terms and Abbreviations


Abbreviation	Meaning
aDSR-GSL	adaptive DSR-GSL
AOI	Area Of Interest
API	Application Programming Interface
APO	Automatic Prompt Optimization
ARD	Automatic Relevance Determination
BERT	Bidirectional Encoder Representations from Transformers
CNN	Convolutional Neural Network
CO	Carbon Monoxide
CRLB	Cramer-Rao Lower bound
CTA	Combine-Then-Adapt
DEM	Digital Elevation Model
DoA	Description of the Action
DSR-GSL	Distributed Super-resolution Gas Source Localization
EM	Expectation-Maximization algorithm
EMD	Earth Mover's Distance
FDM	Finite Difference Method
FIRMS	Fire Information for Resource Management System
FM	Foundation Model
FMLM	Fast Marginal Likelihood Maximization
GDACS	Global Disaster Alert and Coordination System
GDELT	Global Database of Events, Language and Tone
GeoJSON	Geographical JSON
GFS	Global Forecast System
GSL	Gas Source Localization
IDW	Inverse Distance Weighting
IR	Infrared
IR-ARD	Incremental Reformulated Automatic Relevance Determination
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
KPIs	Key Performance Indicators
LASSO	Least Absolute Shrinkage and Selection Operator
LFMC	Live Fuel Moisture Content
MAP	Maximum A Posteriori
mCPP	multi-drone Coverage Path Planning
MSE	Mean Squared Error
NASA	National Aeronautics and Space Administration





NCEP	National Centers for Environmental Prediction
ND	Natural Disaster
NDM	Natural Disaster Management
PDE	Partial Differential Equation
pdf	probability density function
R-ARD	Reformulated Automatic Relevance Determination
R&D	Research and Development
RHS	Right-Hand Side
rLASSO	Re-weighted Least Absolute Shrinkage and Selection Operator
ROC-AUC	Receiver Operating Characteristic-Area Under the Curve
SAROPS	Search and Rescue Optimal Planning System
SBL	Sparse Bayesian Learning
SNR	Signal-to-Noise Ratio
SR-GSL	Super-Resolution Gas Source Localization
TSP	Traveling Salesman Problem
TV	Target Values
UAV	Unmanned Aerial Vehicle
UTC	Coordinated Universal Time
WP	Work Package
WRF	Weather Forecast data





Executive Summary

Deliverable D4.3 “Report on Phenomenon Prediction and Response Planning” is the third deliverable of Work Package 4 (WP4) within the TEMA project. It documents the research results of tasks T4.1 “Precise phenomenon prediction” and T4.2 “Response planning and recommendations” carried out between M19-M36. These tasks address the challenges of predicting and responding to natural disasters by developing advanced computational models that leverage heterogeneous data sources, including satellite imagery, ground sensors, drone observations, social media content, and news feeds.

Major technical advancements achieved during this reporting period include:


In Phenomenon Prediction (task T4.1):


- **Wildfire Smoke Dynamics Modeling (DLR-KN):** Developed a novel Sparse Bayesian Learning (SBL) framework for gas/smoke source localization using stationary advection-diffusion equations with Green’s function solutions. The method enables grid-less source localization with arbitrary spatial placement, distributed multi-agent implementation, and theoretical performance characterization through Cramer-Rao bounds. Experimental validation was conducted during the Montiferru campaign in Sardinia (June 2025).
- **Wildfire Behavior Modeling (TYSL):** Enhanced the FireSim simulator with real-time data assimilation capabilities. Implemented integration of field sensor measurements to improve weather forecasting models using inverse distance weighting, and developed calibration methods for fire spread using actual fire front locations extracted from fire probability maps (Maps4Fire from the Information Fusion component, PDM-tech-05). Validation against the historical Montiferru fire demonstrated significant improvements in accuracy.
- **Near Real-time Flood Modeling (NS):** Researched and implemented the Clone Cells method within the 3Di hydrodynamic model to accurately represent disconnected flow regions with minimal computational overhead. This approach allows modeling of subgrid-scale flow separations (e.g., behind barriers or in bifurcations) without requiring fine grid resolution, enabling faster and more accurate flood predictions.
- **Disaster Prediction from User-Generated Data (IT:U, linked to T3.3 and T4.3):** Created a multimodal prediction framework integrating Bluesky social media data, Global Database of Events, Language and Tone (GDELT) news streams, and weather observations. Leveraged foundation models with Automatic Prompt Optimization (APO) for geospatial reasoning, achieving improved performance over statistical hotspot detection methods in both the 2024 Central Europe floods and 2025 Southern California wildfires.

In Response Planning (task T4.2):


- **Optimal Drone Routing (USE):** Developed mathematical frameworks for UAV path planning based on probability maps of disaster phenomena. Formulated the problem using information theory (normalized entropy) to maximize information gain, with approaches including Traveling Salesman Problem (TSP) formulations and search effort density optimization for determining optimal reconnaissance trajectories.


All developed technologies have been successfully integrated into the TEMA platform and have demonstrated fulfillment of relevant TEMA Key Performance Indicators (KPIs), including improvements in prediction accuracy (Objective OA2), responsiveness (OA3), and situational awareness (OC2). Validation through historical case studies (Montiferru fire, Ahrtal and Mantoudi floods)





and the Sardinia experimental campaign (June 2025) confirmed the operational readiness and practical effectiveness of the developed approaches for natural disaster management. These research advances provide significant capabilities for precise phenomenon prediction and optimal response planning in wildfire and flood emergency scenarios, representing substantial progress beyond state-of-the-art approaches in real-time natural disaster management systems.





1. Introduction

1.1. Purpose and scope of the document

Deliverable D4.3 “Report on Phenomenon Prediction and Response Planning” is the third Deliverable of the fourth Work Package (WP4) of the TEMA project. The main purpose of this document is to report the final research results of TEMA Tasks T4.1 “Precise phenomenon prediction” and T4.2 “Response planning and recommendations” between M19-M36. Herein, this deliverable builds on the contents of Deliverable D4.1 “Precise phenomenon prediction”, which was submitted in M18. The TEMA research efforts were focused on the following areas:

- Wildfire smoke dynamics modeling
- Wildfire prediction modeling
- Near real-time flood modeling
- Disaster prediction from user-generated data sources
- Drone response planning


1.2. Structure of the document


The remainder of the document is as follows:

- Section 2 summarizes the work carried out in the involved research areas covered by this deliverable in this reporting period.
- Section 3 summarizes the main research activities performed in wildfire smoke dynamics modeling.
- Section 4 summarizes the main research activities performed in wildfire prediction modeling.
- Section 5 summarizes the main research activities performed in near-real-time flood modeling.
- Section 6 summarizes the main research activities performed in disaster prediction from user-generated data sources.
- Section 7 summarizes the main research activities performed in drone response planning.
- Section 8 summarizes the conclusions of the performed research.

Sections 3-7 describe the work carried out in each of the involved research areas covered by this deliverable in this reporting period, M19-M36. Each section encompasses:

- the state of the art in its field, highlighting the contributions of TEMA,
- the research progress in the TEMA project in the second reporting period M19-M36,
- the status of fulfillment of the TEMA objectives and KPIs in each corresponding field.






2. Summary of the work carried out


2.1. Objectives

TEMA envisions addressing the challenges in extreme data analytics for Natural Disaster Management (NDM) like regional floods, flash floods, and wildfires by leveraging heterogeneous data sources, including edge devices and sensors (e.g., drones, wind sensors, stream flow gauges, etc.), satellite images, geospatial data, meteorological data, and geosocial media. These data sources are heterogeneous, voluminous, frequently updated, complex, multilingual, dispersed, sparse, and extreme in nature. The main objective of TEMA WP₄ "Prediction and decision-making" (M₁₃-M₃₆) is to exploit the outputs of WP₃ analysis methods to enrich and accelerate response to emergencies/disasters, via novel approaches to fast modeling engines for the TEMA use-cases, automated response planning for making optimal sensor placement recommendations, and decision support services for remote sensing.

The specific TEMA objectives linked to Tasks reported in this deliverable are derived from existing challenges in predicting Natural Disaster (ND) phenomena in the TEMA relevant use cases, i.e., regional floods, flash floods, and forest fires/wildfires. They are presented each in one section, along with accompanying Key Performance Indicator (KPI) and Target Values (TV) as defined in Section 1.1.1 of Part B of TEMA Description of the Action (DoA).

Section 2.2 summarizes the Research and Development (R&D) activities conducted under Tasks T_{4.1} and T_{4.2} by TEMA partners. First, it highlights the R&D performed in wildfire smoke dynamics modeling. Second, it outlines the advancements in wildfire behavior modeling. Third, the research performed in near-real-time flood modeling is described. Fourth, it summarizes the advancements in disaster prediction using user-generated data sources. Finally, it describes the research and development in drone response planning.





2.2. Summary of the work carried out with respect to the objectives

2.2.1. Wildfire smoke dynamics modeling


Between M19 and M36, the following R&D work was carried out with respect to the objectives and KPIs of TEMA.


During this project period, significant progress was achieved in advancing the development and evaluation of the smoke propagation model within PDM-tech-03, with a focus on enhancing predictive accuracy, robustness, and adaptability in real-world environmental conditions. Data collection for the calibration and performance assessment of PDM-tech-03 was systematically conducted within the Sardinia pilot trial, leveraging a network of environmental sensors and meteorological monitoring stations. This dataset, characterized by a relatively high temporal and spatial resolution, provided a comprehensive basis for validating model outputs under diverse atmospheric conditions, including variable wind patterns, topographical influences, and fire emission dynamics.

Subsequently, we performed an evaluation through comparative analysis of theoretical methods and models with observed smoke dispersion patterns. The model demonstrated an agreement with field measurements in terms of plume trajectory, concentration gradients, and temporal evolution. Yet we also observed a high dependency of the smoke propagation on the wind direction. The evaluation also highlighted areas for refinement, particularly in capturing these wind fluctuations and higher spatial resolution of smoke measurements, which require further model enhancements.

A key theoretical advancement involved the development of model-based methods for smoke propagation using the method of Greens functions. This approach enabled the formulation of a linear Partial Differential Equation (PDE) framework that captures the spatiotemporal evolution of smoke concentration in response to localized sources. By leveraging Greens functions, the model efficiently computes the influence of emission sources across the domain, facilitating both forward simulation and inverse problem solving. This mathematical foundation proved instrumental in enabling source localization and corresponding uncertainty quantification. In parallel, a rigorous framework for uncertainty quantification based on the Cramer-Rao lower bound was established, integrating the developed PDE model with Sparse Bayesian Learning (SBL) techniques. This hybrid approach allows for the estimation of emission source locations and magnitudes while simultaneously assessing the confidence intervals of these estimates under noisy and incomplete observational data. The SBL framework promotes sparsity in the solution space, effectively identifying dominant sources, thus enhancing the interpretability and reliability of the results.

Furthermore, the methodological framework was extended to accommodate distributed computing architectures and adaptive data collection strategies. The necessity for adaptive sampling was strongly motivated by the insights gained from the Sardinia pilot trial, where dynamic environmental conditions and sensor coverage limitations revealed the importance of intelligent, feedback-driven data acquisition. In fact, we observed that concentration data arrives in discontinuous segments when the sensor and wind were in a right configuration to ensure successful measurement of the concentration. In other words, the sensor makes measurements only intermittently. The proposed adaptive scheme dynamically integrates only “useful” measurements,





thereby optimizing the data processing to realistic sensor outputs. This distributed and adaptive paradigm represents a significant step toward scalable, real-time smoke monitoring systems capable of supporting emergency response and environmental management in complex terrains.

2.2.2. Wildfire behavior modeling

The FireSim component (PDM-tech-01) is responsible for simulating fire behaviour once the alert has been activated in the Fire Mission. Like any simulator, no matter how powerful it may be, if the inputs are not accurate, the outputs will not be either.

FireSim requires inputs relating to fuel models (type of vegetation and structure), terrain, weather information, etc. Knowing that these inputs may not accurately reflect the reality of fuels or terrain, and that weather forecast data are just models, the development work for FireSim focused on:

- Improvement of the weather forecasting model through the integration of real-time weather data from field sensors measured by the PDM-tech-03 component.
- Calibration of simulations based on the actual location of fire fronts in real time, generated from fire probability maps (Maps4Fire) from the Information Fusion component (PDM-tech-05).


In this regard, FireSim calculates a standard simulation upon receiving the alert, using basic information regarding fuels, terrain, ignition point, and weather forecast information. When real-time data becomes available, FireSim reruns another simulation adjusted with the actual data on every update step.


2.2.3. Near real-time flood modeling

The hydrodynamic modeling flood component 3Di (PDM-tech-02) is responsible for simulating flood behaviour once an alert has been activated in a Flood Business Mission. Between M19 and M36, research and development of Clone Cells in 3Di was carried out to meet the requirements of the objectives, stated in the KPIs of TEMA.

The integration of Clone Cells in hydrodynamic modeling is promising because the results are encouraging. Clone Cells can increase the accuracy of a hydrodynamic model much. Accurate hydrodynamic simulations depend very much on the resolution of input data. However, enhanced input data requires more detailed analysis and usually results in higher computational cost. Specifically, in TEMA, detailed information on disconnected flow regions can be obtained during a flood event by drone images or sensor data. When new or updated information is included in the scenario input data of the simulation software, 3Di does not compromise much on calculation cost by adding more detail to the schematization of the real-time event. If, during floods, due to collapsed infrastructure, such as buildings or bridges, areas in the flood region become disconnected, clone cells add detail to the simulation with only a minimal increase in computational cost (see Fig. 35 in Section 5.4).

Besides, 3Di calculates a preliminary flood simulation upon receiving the alert and an AOI, using basic information, such as the low-resolution digital elevation model (DEM) from Copernicus and weather forecast information (ICON EU-forecast). When real-time information becomes available, 3Di reruns the simulation with updated scenario information with enhanced input data.





However, when updated (field) data is unavailable or delayed, 3Di still calculates new simulation output at desired timestep outputs for the end-users, based on the altered weather forecast that also progresses through time.

2.2.4. Disaster prediction from user-generated data sources

To advance precise phenomenon prediction (TEMA Task T4.1) within TEMA, IT:U developed a multimodal disaster prediction framework that integrates heterogeneous, user-generated data sources from Bluesky, GDELT, and weather observations. Building on and linking prior work in Task T3.3 (Geo-social media analysis) and Task T4.3 (Information fusion), the approach leverages large Foundation Model (FM) with in-context learning and Automatic Prompt Optimization (APO) to enable real-time, geospatially explicit disaster prediction. This framework outperforms statistical hotspot and anomaly detection baselines, achieving higher macro F1-scores, improved Receiver Operating Characteristic-Area Under the Curve (ROC-AUC), and increased recall for rare disaster events. IT:U also conducted the first large-scale, geographically explicit assessment of Bluesky for disaster monitoring, analysing 676,337 posts across the 2024 Central Europe floods and the 2025 Southern California wildfires. Results revealed strong temporal alignment with real-world events, meaningful spatial patterns, and consistent emotional responses, confirming Bluesky's potential as a complementary data source for early warning and prediction. These advances significantly contribute to objective OA2 ("Increase accuracy of extreme data analysis algorithms") through improved event detection accuracy and to objectives OA3 and OB1 ("Increase responsiveness/speed") by enabling scalable, cell-based near-real-time disaster phenomenon predictions.


2.2.5. Response planning


The objective of drone response planning is to efficiently assign trajectories to each available drone in order to optimally complete the NDM mission. It is an essential functionality to enable the operation and success of the TEMA technologies. The trajectories are computed and provided to the human operator for validation. Once validated, they are sent to the drones for execution, enabling autonomous operation with a human in the loop. Response planning has been developed in Task 4.2 "Response Planning and Recommendations", which starts in M19 and finishes in M36. The following R&D work was carried out with respect to drone response planning.

Different NDM missions have been considered in the design of the response planning methods, including: surveillance (e.g., for detecting fires as they start), monitoring (e.g., to gather high-quality images of the ND), and search (e.g., to detect and locate people under risk). Accordingly, response planning has been carefully designed to fulfill the end-user requirements.

In particular, the following non-functional and functional requirements apply: EU-RQ-NF-03, EU-RQ-FUNC-03, EU-RQ-FUNC-05, EU-RQ-FUNC-06, EU-RQ-FUNC-07, EU-RQ-FUNC-08, EU-RQ-FUNC-09, and EU-RQ-FUNC-10, as listed in TEMA Deliverable D2.1.

A wide range of existing methods has been developed to coordinate multiple drones in surveillance and monitoring tasks. The drone response planning method developed in TEMA has been focused on the search problem, which can be applied for tasks such as search for people in Natural Disaster environments, search for hot points for wildfires, or search for wildfire re-ignition







points, among many others. Instead of employing a systematic approach that surveys the entire area (e.g., using a zig-zag pattern), the proposed method selects the drones' trajectories to optimize the search and detection probability. The developed method makes use of a statistical formulation and information-theory principles to extend optimal search theory using the normalized entropy to quantify information gain and prioritize search areas.

Drone response planning has been implemented by using efficient and scalable algorithms that strongly help reduce the response time, helping to fulfill KPI OB1. In addition, drone response planning enables fulfilling complex missions that enable gathering high-quality images from a short distance of the ND, enabling the extraction of accurate ND measurements by other components of the TEMA architecture and their integration through the **Information Fusion** component to keep updated the actual status of the ND, which is crucial for the accuracy of NDM predictions. The developed response planning method is flexible and can be applied to flood and fire scenarios (and others), providing a generalizable solution. It is computationally efficient, as it can generate optimal multi-UAV missions within practical time-frames (20 seconds for 100 km^2 areas in a standard PC). It is also scalable with the number of drones considered in the fleet.

The drone response planning method has been developed, coded, and tested in simulation, confirming the required performance. In addition, the drone response planning technology (PDM-tech-04) has been integrated within the TEMA platform architecture, enabling seamless data exchange with the rest of the modules of the TEMA architecture and supporting real-time mission adaptation. As shown in Section 7, the developed drone planning methods and technology satisfy the assigned TEMA KPIs. Although the integration of response planning TEMA technology (PDM-tech-04) is always subject to refinements and updates, the R&D stage of this TEMA component has been completed, and all the assigned objectives and KPIs have been met.





3. Wildfire smoke dynamics modeling

3.1. Introduction

Forest fires have become more frequent and severe due to a number of factors, including climate change and human activity such as land-use change and urban expansion [1, 2]. These combined effects have led to record-breaking wildfire intensities across the globe. To improve environmental monitoring, early warning systems, and real-time disaster response during such events, it is critical to accurately model and localize fire sources, especially during fire early stages [3].


This objective can be achieved through the process of “learning” from data, i.e., using data-driven approaches. For instance, support vector machines [4] for concentration mapping, kernel methods [5], or deep-learning approaches, although primarily in the visual domain [6][7], have been proposed. The key shortcoming of these methods is their dependency on the training data, especially in the case of deep-learning methods. Furthermore, they rely on visual sensors for gas plume registration; they are thus unsuitable when gas is invisible to the camera. Of course, the chemical plume data can also be simulated and used for training [8]. However, to make the simulations realistic, we need to exploit the knowledge about gas propagation – the domain knowledge. Thus, one might wonder why not use this domain knowledge directly.


Indeed, robotic platforms can take advantage of existing knowledge on smoke and gas propagation, such as a model for gas dispersion, and selectively acquire specific facets of this model from available data. This strategic approach mitigates the dependency on extensive training data, fostering more robust and versatile solutions for concentration mapping, in particular in scenarios involving chemical, rather than visual sensors.

Such an approach in practice will thus serve a dual purpose: it will compensate for the lack of data while also adjusting the model to align better with the gathered measurements. This paradigm forms the foundation of the approach that was investigated within TEMA Task 4.1.

3.2. State of the art and TEMA’s contribution

In the broader context, the underlying physics governing the spatio-temporal dynamics of diffused substances can be effectively characterized through Partial Differential Equation (PDE). For instance, the non-homogeneous advection-diffusion equation or scalar transport Eq. [9] serves as a means to simulate the dissemination of materials through the air, originating from multiple emission sources or being absorbed by designated material “sinks.” Given the parameters of these equations, such as boundary and initial conditions as well as the count and positioning of emission sources, the spatio-temporal evolution of material concentration can be numerically computed. This is known as a *forward* modeling problem: it involves solving the corresponding PDE subject to boundary and initial conditions. However, real-world scenarios frequently involve instances where some, if not all, of these parameters are unknown: one might encounter undefined boundary conditions or an unknown count of emission sources. Their inference from







measurements then becomes the objective of the robotic exploratory system, which constitutes an *inverse* problem. Notably, inverse problems are characterized by their ill-posed nature, demanding supplementary constraints or regularization techniques to ensure the stability and existence of a solution. In TEMA Task T4.1, we directed our attention towards incorporating such constraints, specifically sparsity constraints. This allowed us to reliably estimate both the count and locations of dispersed material sources, leveraging the information gathered through measurements, as supported by conducted simulations.

This problem is known in the robotic olfaction community as Gas Source Localization (GSL) [10]. Early methods for gas source localization were quite restrictive in their assumptions. In particular, a fixed number (typically one) of sources was assumed, and the source estimation was a part of the robotic movement strategy. For instance, *chemotaxis*-based methods [11, 12] attempt to move a robot along the gradient of chemical concentration. Yet these approaches are ultimately limited by the structure of the chemical plume – due to turbulence, the gradient can be very erratic, especially far away from the source. This can be corrected by taking the wind into account and following bioinspired algorithms [13, 14], also known as *anemotaxis*. Nonetheless, the smoothness of the concentration gradient remains a major limiting factor for these methods. Standard *anemotaxis* can also be extended to “higher order” approaches, or rather *model-based anemotaxis*, where not only spatial gradients, but also divergence and temporal variations of the concentration values are taken into account [15]. Specifically, the authors estimate divergence and gradients from data, ensuring they adhere to the advection-diffusion PDE, and devise a navigation strategy for multiple robots to follow the plume front. This method permits extension to inverse problems, assuming unknown wind or diffusion parameters [16], and addressing the GSL problem [17, 18] instead of front tracking. It is crucial to emphasize that, while [15, 16, 17] capitalize on domain knowledge to notably enhance standard *anemotaxis* performance, the approach does not directly solve (or invert) the advection-diffusion PDE. Instead, it relies on numerically estimated concentration gradients and divergence using historical observation data or data readings from multiple robots – an intricate task due to noise and the challenges associated with numerical differentiation. Consequently, the smoothness of gradients remains a significant constraint for these strategies. Furthermore, it should be mentioned that the number of gas sources remains fixed.

More promising methods in this respect are probabilistic approaches, which exploit a mathematical structure of the dispersion process. Rather than following gradients of the gas concentrations, the sources are treated as unknown parameters. These are then estimated from the collected measurements. Bayesian approaches towards such parameter estimation are particularly attractive. They can model different uncertainties about source parameters [19], different models [20, 21, 22], unknown environments [23, 24]. Moreover, they enable *infotaxis*-based methods [25] for autonomous navigation of robots to improve GSL and guide the robots to the source using different information-theoretic criteria (see e.g., [25, 26, 27]). Yet in the majority of work on GSL the number of sources remains fixed and known a priori. The main reason for that is an otherwise required numerically challenging integer optimization/inference over the number of sources. First in [28, 29] this assumption was relaxed using a Sparse Bayesian Learning (SBL) approach [30, 31] and a PDE-based dispersion model. There, a SBL approach [30, 31] coupled with a PDE-based dispersion model was employed, representing a departure from the fixed-source count assumption.

The key idea behind the SBL-based approach towards GSL is that a very large number of potential sources is initially assumed. However, the rates of the material release for these sources – the source weights – are *constrained* to be sparse. In other words, it is assumed that only a few source weights are non-zero. This approach was shown to accurately estimate both the count and the





locations of the sources in simulations [29, 32, 28], as well as in real experiments [33, 34] (see also [35] for a more extensive discussion). It thus allows “trading” integer optimization concerning source counts for a non-integer, but sparsity-constrained estimation of source parameters.


One of the features of such PDE-based approach towards GSL is a discretization of the equation required for its numerical solution. Standard approaches, like finite elements or finite difference methods [36], discretize the equation both in space and time; as a result, the unknown functions – gas concentration and source distribution functions – are represented with a set of discrete parameters. As a consequence, the source locations become constrained to the locations of the vertices in the discretization mesh or centers of the discretization cells.¹ This, in turn, implicitly fixes possible locations of the sources as well as their initial number [33]. While in general this can be compensated again by imposing sparsity constraints on the resulting source estimate [29, 35], practically this leads to a high number of unknown parameters that have to be estimated – one parameter for each possible source. Also, the discretization has to be dense enough in order not to miss any potential locations. Alternatively, expensive remeshing has to be used. This inadvertently increases the number of measurements needed to reliably identify the parameters. In particular, in cases where the source signal is very sparse, i.e., when only a few gas or smoke sources are present, such over-parameterization negatively impacts algorithm performance, especially in the early stages, when the number of measurements is low. Furthermore, the discretization limits the source localization accuracy to the size of the mesh elements closest to the actual source location.


The methodology investigated within TEMA alleviated the aforesaid constraint, allowing for arbitrary spatial positioning of sources within the exploration area.

The specific advancement beyond state-of-the-art methods that TEMA contributed to the field can be summarized as follows.

1. We utilized a stationary advection-diffusion equation as a proxy model of the smoke propagation. While this is an approximation, it allows for finding solutions to corresponding PDE in the closed, analytical form both in 2D as well as in 3D. This significantly relaxes computational constraints, making the “forward” modeling of smoke propagation extremely fast.
2. Using a Green’s function-based approach to solve the corresponding PDE, we cast the inference problem in the probabilistic context, and using sparsity-enforcing inference techniques, we proposed an algorithm to identify both the number as well as the locations of the smoke sources.
3. The new algorithm does not constrain the source locations to any discretization grid, thus allowing for an arbitrary placement of the sources.
4. We proposed a modification of the algorithm to perform the inference distributively, over a network of agents, thus eliminating the need to collect large amounts of smoke concentration data. In fact, the developed methods allow processing the data in on-line fashion, as soon as samples become available.
5. Using the method of Green’s function and SBL methods, we also theoretically characterized the achievable estimation performance in terms of Cramer-Rao bounds on the resulting estimate. This development permitted us to compute spatial uncertainty maps that identify locations where data has to be collected to improve the source localization performance of the algorithm.

¹This depends on whether the finite element or finite difference method is used for discretizing the PDE.





3.3. Smoke dynamics modeling

3.3.1. Signal model for smoke/gas source localization

Let us consider a smoke localization, or equivalently, gas source localization GSL problem over some d -dimensional exploration area $\Omega \subset \mathbb{R}^d$. We will assume that the dynamics of the gas or smoke over a time interval $\mathcal{T} = [0, T]$ can be described by an inhomogeneous, linear advection-diffusion PDE:

$$\frac{\partial f(\mathbf{x}, t)}{\partial t} - \kappa \Delta f(\mathbf{x}, t) + \mathbf{v}(\mathbf{x}, t)^\top \nabla f(\mathbf{x}, t) = \sum_{i=1}^{\tilde{L}} w_i \delta_{\theta_i}(\Omega), \quad \mathbf{x} \in \Omega, t \in \mathcal{T} \quad (1)$$

$$\text{s.t. } f(\mathbf{x}, t) = 0, \mathbf{x} \in \partial\Omega, t \in \mathcal{T}, \quad (2)$$

$$f(\mathbf{x}, 0) = 0, \mathbf{x} \in \Omega. \quad (3)$$


Here $f(\mathbf{x}, t)$ represents a spatial gas concentration at position \mathbf{x} and time instance t , $-\kappa \Delta f(\mathbf{x}, t)$ is the corresponding diffusion term, with κ being the diffusion coefficient and Δ is a Laplace operator. The term $\mathbf{v}(\mathbf{x}, t)^\top \nabla f(\mathbf{x}, t)$ models the convective part of the gas propagation process, or the gas transport due to the wind field $\mathbf{v}(\mathbf{x}, t)$. As can be seen, the transport is essentially the gradient $\nabla f(\mathbf{x}, t)$ of the concentration along the wind direction $\mathbf{v}(\mathbf{x}, t)$.


The Right-Hand Side (RHS) of Eq. (1) defines a superposition of \tilde{L} gas sources that “drive” gas propagation. As can be seen from the model, we assume sources to be time-invariant, each located at position $\theta_i \in \Omega$ and having a release rate (amplitude) w_i . This assumption can be relaxed in several ways – either the release rates, source locations, or both can be considered as time-varying. Yet treatment of these cases is left outside the scope of this work. Let us also stress that RHS of the model (1) is a superposition of measures that are arbitrarily placed within Ω . This is important, since it will lead to, on the one hand, non-linear models and will require corresponding estimation schemes. On the other hand, this will allow for grid-less, super-resolving placement of the source within the exploration domain, as we will show later.

Naturally, Eq. (1) is augmented with a boundary condition (2) that defines concentration values at the boundary $\partial\Omega$ of the domain Ω . We will consider an open boundary $\partial\Omega$ in the following, which allows concentration to flow off; This is known as the Dirichlet boundary condition. Alternatively, other boundary conditions can be considered as well. Furthermore, we also specify initial conditions (3) that determine the initial concentration distribution over Ω at time $t=0$.

Surrogate process model

The model expressed in (1) states that the concentration distribution $f(\mathbf{x}, t)$ is governed by \tilde{L} individual sources. We will assume that in general \tilde{L} is rather small, i.e., we have only a few “active” sources. This is a critical assumption, which will eventually allow us to introduce additional regularization terms into the estimator. Our objective is to estimate the number of sources \tilde{L} , as well as the corresponding source parameters w_i and θ_i , $i=1, \dots, \tilde{L}$, by collecting measurements of the process. While source parameters are real-valued, the number of sources is an integer. A joint estimator of these parameters would lead to a numerically challenging problem. On the other hand, SBL can offer an elegant solution to this problem while circumventing mixed optimization over the number of components and component parameters.





To cast the GSL problem into the SBL framework we approximate (1) with a surrogate model

$$\frac{\partial f(\mathbf{x},t)}{\partial t} - \kappa \Delta f(\mathbf{x},t) + \mathbf{v}(\mathbf{x},t)^\top \nabla f(\mathbf{x},t) = \sum_{l=1}^L w_l \delta_{\theta_l}(\Omega), \quad \mathbf{x} \in \Omega, t \in \mathcal{T}, \quad (4)$$

where now $L \gg \tilde{L}$ and some of the weights w_l , $l \in \mathcal{L} \triangleq \{1, \dots, L\}$, are zero. In other words, the source weights are assumed to be sparse. Note that although at first glance the difference between the model (1) and the surrogate (4) seems to be insignificant, it has a major consequence for the estimation approach discussed later: in the former \tilde{L} is unknown, while L is known and fixed in the latter. Such an approach is also sometimes referred to as *max-search* approach [37].

Assume for a moment that RHS of (1), i.e., both w_l and θ_l , $l \in \mathcal{L}$, are fully known. The concentration $f(\mathbf{x},t)$ can then be found by solving (1) subject to the corresponding boundary and initial conditions. The latter is also referred to as a *forward problem*. Corresponding solution is most often found using numerical techniques (see e.g., [9]). These include finite difference or finite element methods and their extensions/modifications. Specifically, the space and time is appropriately discretized, which leads to a set of algebraic equations with respect to a set of parameters that represent a discretized system [9]. While numerical methods are very general and can deal with linear as well as nonlinear equations, diverse types of boundary and initial conditions, they do have a number of shortcomings. In particular, solving inverse problems, such as localizing sources based on collected measurements, or identifying other parameters of the PDE [38], often requires multiple forward solutions. This can potentially lead to schemes that require significant computational resources.

On the other hand, analytical tools can provide an efficient solution to the PDE without complex numerical computations. Moreover, in our case, a special form of RHS in (4) (and thus also in (1)) makes analytical tools particularly appealing. The method we utilize for our purposes is known as Greens function method [39, 40]. In its essence, the Green's function generalizes the concept of a classical impulse response for linear PDE: it is a function $G(\mathbf{x}, \boldsymbol{\theta}, t, \tau)$ that solves the original advection-diffusion equation:


$$\frac{\partial G(\mathbf{x}, \boldsymbol{\theta}, t, \tau)}{\partial t} - \kappa \Delta G(\mathbf{x}, \boldsymbol{\theta}, t, \tau) + \mathbf{v}(\mathbf{x}, t)^\top \nabla G(\mathbf{x}, \boldsymbol{\theta}, t, \tau) = \delta_{\boldsymbol{\theta}}(\Omega) \delta_{\tau}(\mathcal{T}), \quad (5)$$


$$\mathbf{x}, \boldsymbol{\theta} \in \Omega, t, \tau \in \mathcal{T}$$

subjected to the same boundary and initial conditions as the original equation. Note that RHS of this equation is a Dirac source, located at θ_l in space, and active at time instance τ . In other words, $G(\mathbf{x}, \boldsymbol{\theta}, t, \tau)$ is a "response" of the PDE to a single source at the location $\boldsymbol{\theta} \in \Omega$ and moment of time $\tau \in \mathcal{T}$. Provided the latter exists (see e.g., [40] for more details), the solution $f(\mathbf{x}, t)$ of (4) can then be found as

$$f(\mathbf{x}, t) = \int_{\mathcal{T}} \int_{\Omega} G(\mathbf{x}, \boldsymbol{\theta}, t, \tau) \underbrace{\sum_{l=1}^L w_l \delta_{\theta_l}(\Omega)}_{\text{RHS of (4)}} d\boldsymbol{\theta} d\tau = \sum_{l=1}^L w_l \int_{\mathcal{T}} G(\mathbf{x}, \boldsymbol{\theta}_l, t, \tau) d\tau. \quad (6)$$

It can be seen that the concentration $f(\mathbf{x}, t)$ for an arbitrary number of sources is computed as a linear combination of the Green's functions centered at source locations θ_l , $l \in \mathcal{L}$. While the method in general provides an analytic solution to a PDE, finding the function $G(\mathbf{x}, \boldsymbol{\theta}, t, \tau)$ for arbitrary boundary and initial conditions can be quite challenging. In fact, closed-form solutions are known for a rather limited set of cases and additional assumptions.





One possible avenue is to approximate Green's functions with, e.g., a neural network [41] – currently an active research area. We will, however, leave the aspects related to such approximations outside the scope of this work and focus on illustrating the use of SBL on examples when Green's function is available analytically.

Stationary advection-diffusion with constant wind

In order to bring (4) into a form that permits an analytical solution, we need to make a few assumptions. First, we assume a stationary regime, in which wind $\mathbf{v}(\mathbf{x}, t)$ and concentration $f(\mathbf{x}, t)$ are time-invariant, i.e., independent of t . Furthermore, we will also assume the wind vector to be constant over Ω , i.e., $\mathbf{v}(\mathbf{x}) \equiv \mathbf{v}$. Finally, we consider an unbounded domain Ω . With these assumptions, it follows that the surrogate model becomes

$$-\kappa \Delta f(\mathbf{x}) + \mathbf{v}^T \nabla f(\mathbf{x}) = \sum_{l=1}^L w_l \delta_{\theta_l}(\Omega), \quad \mathbf{x} \in \Omega, \quad (7)$$

With a Dirichlet boundary constraint over the unbounded domain now implies that $f(\mathbf{x}) \rightarrow 0$ as $\mathbf{x} \rightarrow \infty$. Under these assumptions, the Green's function should solve

$$\begin{aligned} -\kappa \Delta G(\mathbf{x}, \boldsymbol{\theta}) + \mathbf{v}^T \nabla G(\mathbf{x}, \boldsymbol{\theta}) &= \delta_{\boldsymbol{\theta}}(\Omega), \quad \mathbf{x}, \boldsymbol{\theta} \in \Omega \\ \text{s.t. } G(\mathbf{x}, \boldsymbol{\theta}) &= 0, \mathbf{x} \in \partial\Omega \end{aligned} \quad (8)$$

The corresponding solution can be obtained analytically. For the 2D case, it is given by:

$$G(\mathbf{x}, \boldsymbol{\theta}) = \frac{1}{2\pi\kappa} e^{\frac{\mathbf{v}^T(\mathbf{x}-\boldsymbol{\theta})}{2\kappa}} K_0\left(\frac{\|\mathbf{x}-\boldsymbol{\theta}\| \|\mathbf{v}\|}{2\kappa}\right), \quad (9)$$

where $\|\cdot\|$ is an ℓ_2 norm and $K_0(\cdot)$ is a zero-order modified Bessel function of the second kind. We visualize the corresponding solution in Fig. 1. Note that the solution is discontinuous at the source, signifying that RHS of (7) is a superposition of Dirac measures, rather than functions².


For the 3D case, the Green's function can also be computed in closed form:

$$G(\mathbf{x}, \boldsymbol{\theta}) = \frac{1}{(4\pi\kappa)} \frac{1}{\|\mathbf{x}-\boldsymbol{\theta}\|} \exp\left(\frac{\mathbf{v}^T(\mathbf{x}-\boldsymbol{\theta})}{2\kappa} - \frac{\|\mathbf{v}\|}{2\kappa} \|\mathbf{x}-\boldsymbol{\theta}\|\right). \quad (10)$$

Similarly to the 2D case, it is discontinuous at the location of the source, yet this time it decays proportionally to the distance from the source, rather than to its logarithm. Again, we stress that the obtained results assume an unbounded domain Ω and a constant wind vector. While in 2D this might seem to be a reasonable assumption, in 3D such an assumption is difficult to justify: the wind in vertical direction changes with the altitude, thus making a constant wind assumption less meaningful. Furthermore, the ground naturally bounds the domain Ω in the vertical direction.

Let us re-iterate again that the analytical computation of Green's function is not necessarily needed. The corresponding solution can also be approximated with suitable techniques, also for more general areas Ω , as well as non-constant wind conditions. For our purposes, we will, however, use the 2D case as an example to illustrate the use of SBL techniques.

²Note that in practical implementation this discontinuity needs to be accounted for. To do this we substitute $\|\mathbf{x}-\boldsymbol{\theta}\|=0$ with some small number $\|\mathbf{x}-\boldsymbol{\theta}\|=10^{-4}$ to stabilize the evaluation.



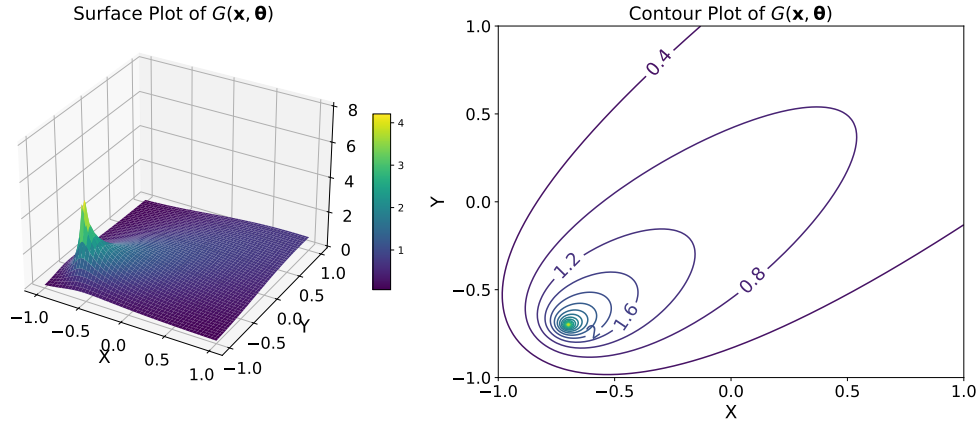



Figure 1. 2D Green's function for stationary advection-diffusion PDE in 2D with source located at $\theta = [-0.7, 0.7]^T$ m, with wind direction set to $v = [0.5, 0.5]^T$ m/s, and diffusion parameter $\kappa = 0.1$ m²/s.

Discretization of the domain

While we do allow sources to be placed arbitrarily in the domain, their estimation requires a numerical procedure. This in turn requires an appropriate discretization of the model (7). To this end, we partition Ω into N smaller subdomains or cells, which form a (not necessarily regular) grid. For each cell with center coordinates $x_i, i \in \mathcal{N} \triangleq \{0, \dots, N-1\}$, we then assume a constant concentration value $f(x_i) = \text{const}$. The corresponding concentrations are then aggregated into a vector $f \triangleq [f(x_0), \dots, f(x_{N-1})]^T$. Using this discrete representation, we can now define a discretization of Green's function as follows:

$$G(\Theta) \triangleq \begin{bmatrix} G(x_0, \theta_1) & \dots & G(x_0, \theta_L) \\ \vdots & \ddots & \vdots \\ G(x_{N-1}, \theta_1) & \dots & G(x_{N-1}, \theta_L) \end{bmatrix}$$

where the columns correspond to different sources of the surrogate model and the rows are the evaluations of the Green's function at different discretization elements. Furthermore, to keep the notation compact, we define $\Theta \triangleq [\theta_1, \dots, \theta_L]$ to aggregate all potential sources. Note that $G(\Theta)$ can be computed either using the analytic form of the Green's function or using its numeric approximation. This allows us to rewrite (6) in a matrix vector form as


$$f = G(\Theta)w, \quad (11)$$


with $w \triangleq [w_1, \dots, w_L]^T$. As such, the solution is represented as a linear combination of discretized Green's functions centered at (arbitrary) source locations.

Let us point out that the method we will propose can easily account for other discretization approaches, such as finite elements to represent $f(x)$.

Measurement model

The source parameters are not observed directly and have to be estimated from the measurements. Assume now that M noisy samples of the concentration $f(x)$ have been obtained at





locations x_m , $m = 1, \dots, M$. Without loss of generality, we will assume that $N \gg M$ and that x_m are a subset of discretization cells. Note that we can always select such a discretization of Ω that will include locations where measurements are taken. Now collect the measurements in a vector $\mathbf{y} \in \mathbb{R}^M$ as

$$\mathbf{y} = \Phi \mathbf{f} + \boldsymbol{\xi}, \quad (12)$$

where $\Phi \in \mathbb{R}^{M \times N}$ is a selection (or measurement) matrix, which contains zeros and a single 1 at a position that “picks” an element from \mathbf{f} that corresponds to a cell where measurement is taken. The additive term $\boldsymbol{\xi}$ is a random, zero-mean Gaussian perturbation with precision matrix $\lambda_\xi \mathbf{I}$ and some $\lambda_\xi > 0$.

Our goal now is to use \mathbf{y} to estimate a sparse vector \mathbf{w} , locations Θ , and recover \mathbf{f} from (11). To this end, we pursue a Bayesian approach towards parameter estimation as detailed in the following.

3.3.2. Off-Grid Bayesian smoke source localization

As the next step, we need to cast the constructed model into the probabilistic framework. This will define the required inference structure and show how SBL is introduced into the estimator.

To cast the estimation problem into the probabilistic framework, consider the following posterior probability density function (pdf) of the variables of interest:

$$p(\mathbf{f}, \mathbf{w}, \Theta | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \mathbf{w}, \Theta) p(\mathbf{w}) p(\Theta), \quad (13)$$


where we explicitly assume that source locations Θ and the corresponding weights \mathbf{w} are independent. Let us now discuss all the factors in (13) in more detail.


Based on (12) we immediately see that the first factor – the concentration likelihood – is given by $p(\mathbf{y} | \mathbf{f}) = \mathcal{N}(\mathbf{y} | \Phi \mathbf{f}, \lambda_\xi^{-1} \mathbf{I})$. The pdf $p(\mathbf{f} | \mathbf{w}, \Theta)$ represents the relationship between sources and concentration \mathbf{f} based on (11). This deterministic relationship can be represented with a Dirac distribution $p(\mathbf{f} | \mathbf{w}, \Theta) = \delta_{\mathbf{G}(\Theta) \mathbf{w}}(\mathbb{R}^N)$. The other two terms $p(\Theta)$ and $p(\mathbf{w})$ are the prior distributions of the source parameters. Concerning the former, we will assume a uniform prior over Ω , i.e., $p(\Theta) \propto \text{const}$. Let us stress here that we explicitly treat Θ as a random variable. If Θ is fixed to some predefined grid, we would end up with a “linearized” representation. Such an approach was explored in [33, 35] within the context of numerical solutions of advection-diffusion PDE. In the analytical approach we follow here, parameters Θ are random unknown variables that are inferred along with other unknown variables.

Now, consider source release rates \mathbf{w} . These will enter the inference model linearly but are constrained to be sparse, following the surrogate modeling approach. It is here where SBL comes into play. In case of $p(\mathbf{w})$ we employ a modeling approach used in SBL by imposing the corresponding L -dimensional hierarchical prior: $p(\mathbf{w}, \boldsymbol{\alpha}) = p(\mathbf{w} | \boldsymbol{\alpha}) p(\boldsymbol{\alpha}) = \prod_{l=1}^L p(w_l | \alpha_l) p(\alpha_l)$, where $p(w_l | \alpha_l) = \mathcal{N}(w_l | 0, \alpha_l^{-1})$, $l \in \mathcal{L}$. The hyperpriors $p(\alpha_l)$, $l \in \mathcal{L}$, we select to be uniform. As we will show, this will lead to Fast Marginal Likelihood Maximization (FMLM) type of inference procedures. Now, the resulting parameter posterior pdf becomes

$$p(\mathbf{f}, \mathbf{w}, \boldsymbol{\alpha}, \Theta | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \mathbf{w}, \Theta) p(\mathbf{w}, \boldsymbol{\alpha}) p(\Theta). \quad (14)$$

Note that in contrast to linear models, we now have a nonlinear dependency of the posterior (14) on the location parameters Θ .





The posterior (14) captures the joint probability distribution of all unknown parameters. However, not all of them are actually “equally” important for the inference. Let us explain what exactly is meant here. First, we point out that in (14) the discretized concentration \mathbf{f} is treated as a random variable. However, we know that due to the form of $p(\mathbf{f}|\mathbf{w}, \Theta)$ it is actually deterministically dependent on \mathbf{w} and Θ . Second, the release rates \mathbf{w} enter the problem linearly; as such, these parameters can be estimated in closed form given estimates of Θ , α , and α , as we will show later. Thus, it is the parameters α and Θ that ultimately determine all other unknowns. It therefore makes sense to consider a marginalized posterior

$$p(\alpha, \Theta | \mathbf{y}) = \iint p(\mathbf{f}, \mathbf{w}, \alpha, \Theta | \mathbf{y}) d\mathbf{f} d\mathbf{w} \quad (15)$$

$$\propto p(\mathbf{y} | \alpha, \Theta) p(\Theta) p(\alpha)$$

instead of a full posterior (14). Also, for the nonlinear case, it can be computed in closed form by noting that

$$p(\mathbf{y} | \alpha, \Theta) = \iint p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \mathbf{w}, \Theta) p(\mathbf{w}, \alpha) d\mathbf{f} d\mathbf{w} \quad (16)$$

$$\propto \frac{1}{|\Sigma_\alpha(\Theta)|^{\frac{1}{2}}} e^{-\frac{1}{2} \mathbf{y}^T \Sigma_\alpha(\Theta)^{-1} \mathbf{y}} \quad (17)$$

where the covariance matrix $\Sigma_\alpha(\Theta)$ is given by

$$\Sigma_\alpha(\Theta) = \lambda_\xi^{-1} \mathbf{I} + \Phi \mathbf{G}(\Theta) \mathbf{A}^{-1} \mathbf{G}(\Theta)^T \Phi^T,$$

with $\mathbf{A} = \text{diag}(\alpha)$. Despite the fact that (17) is available in closed form, maximizing it directly is computationally challenging, especially due to the nonlinear dependency on the location parameters. We therefore turn to the variational Bayesian inference to find an approximate solution to this optimization problem. The details of this approach are outlined in the following section.

SBL for smoke source localization

Consider the posterior $p(\alpha, \Theta | \mathbf{y})$ from (15). On the most general level, we now attempt to utilize variational Bayesian methods [42, 43] to approximate it. To this end We construct the bound on the log-evidence $p(\mathbf{y})$ as follows [43, 42]


$$\log p(\mathbf{y}) \geq \mathbb{E}_{q(\Theta, \alpha)} \log \frac{p(\mathbf{y}, \alpha, \Theta)}{q(\Theta, \alpha)} \triangleq -\mathcal{F}(q(\Theta, \alpha) || p(\mathbf{y}, \alpha, \Theta)), \quad (18)$$


where the term $\mathcal{F}(q(\Theta, \alpha) || p(\mathbf{y}, \alpha, \Theta))$ is known as the variational free energy [44, 43, 42].

Variational Bayesian inference methods are then realized as follows: by suitably choosing $q(\Theta)$ such as to allow a tractable $\mathcal{F}(q(\Theta, \alpha) || p(\mathbf{y}, \alpha, \Theta))$, the lower bound on $\log p(\mathbf{y})$ is maximized, i.e., the free energy $\mathcal{F}(q(\Theta, \alpha) || p(\mathbf{y}, \alpha, \Theta))$ is minimized. The latter procedure is equivalent to the minimization of the Kullback-Leibler divergence $D_{\text{KL}}(q(\Theta, \alpha) || p(\alpha, \Theta | \mathbf{y}))$ between an approximant $q(\Theta, \alpha)$ and the actual posterior pdf $p(\alpha, \Theta | \mathbf{y})$. Obviously, the minimum of the free energy is achieved when $q(\Theta, \alpha) = p(\alpha, \Theta | \mathbf{y})$, which turns the bound into an equality.

When this solution is intractable, there are two possible (and mutually non-exclusive) approaches to solve the inference problem:

1. take a restricted family of distributions $q(\Theta, \alpha)$ and seek a member of this family that minimizes the free energy, and





2. resorts to optimization strategies that facilitate the optimization of the bound.

A typical strategy to realize the first approach consists in assuming a factorization of $q(\Theta, \alpha)$ that neglects some of the statistical dependencies between the variables, as well as constraining the pdf of some of the factors to some predefined classes of pdf, e.g., of exponential families. For our case we explicitly assume that a variational proxy pdf $q(\Theta, \alpha)$ factors as

$$q(\Theta, \alpha) = q(\Theta)q(\alpha).$$

Moreover, we will also constrain both $q(\Theta)$ and $q(\alpha)$ to a class of Dirac measures on the corresponding domains, i.e.,

$$q(\Theta) = \prod_{l=1}^L q(\theta_l) = \prod_{l=1}^L \delta_{\hat{\theta}_l}(\Omega), \text{ and } q(\alpha) = \prod_{l=1}^L q(\alpha_l) = \prod_{l=1}^L \delta_{\hat{\alpha}_l}(\mathbb{R}_0^+). \quad (19)$$

This is a simplifying assumption. It will essentially lead to an approximation of Maximum A Posteriori (MAP) estimates of the corresponding parameters. Now the minimization of the variational free energy $\mathcal{F}(q(\Theta, \alpha) \| p(\mathbf{y}, \alpha, \Theta)) = \mathcal{F}(q(\Theta)q(\alpha) \| p(\mathbf{y}, \alpha, \Theta))$ is realized iteratively, by alternating between two steps:

Step 1: Minimizing $\mathcal{F}(q(\Theta)q(\alpha) \| p(\mathbf{y}, \alpha, \Theta))$ with respect to $q(\alpha)$ with $q(\Theta)$ fixed, and then

Step 2: Minimizing $\mathcal{F}(q(\Theta)q(\alpha) \| p(\mathbf{y}, \alpha, \Theta))$ with respect to source locations $q(\Theta)$ with the source support $q(\alpha)$ kept fixed at the new, updated value.

Steps 1 and 2 are then iteratively repeated until sequences of $q(\Theta)$ and $q(\alpha)$ converge to some limits. Under assumptions (19), these steps can be formulated more explicitly.

Consider now the estimation of $q(\alpha)$ in Step 1. Using (19) the minimum of $\mathcal{F}(q(\Theta)q(\alpha) \| p(\mathbf{y}, \alpha, \Theta))$ with respect to $q(\alpha)$ amounts to finding³ a minimizer $q(\alpha)'$ of the Kullback-Leibler divergence $D_{\text{KL}}(q(\alpha) \| \tilde{p}(\alpha))$

$$q(\alpha)' = \underset{q(\alpha)}{\operatorname{argmin}} D_{\text{KL}}(q(\alpha) \| \tilde{p}(\alpha)),$$

where $\tilde{p}(\alpha)$ is given as

$$\tilde{p}(\alpha) \propto \exp(\mathbb{E}_{q(\Theta)} \log p(\mathbf{y}, \alpha, \Theta)) \propto p(\mathbf{y} | \alpha, \hat{\Theta}) p(\alpha). \quad (20)$$


The latter is essentially the original posterior (15), yet with Θ set to the support $\hat{\Theta}$ of $q(\Theta)$. Moreover, $\tilde{p}(\alpha)$ can thus be recognized as an approximation of the posterior $p(\alpha | \mathbf{y}, \hat{\Theta})$. Since $q(\alpha)$ is assumed to be a Dirac measure, the minimum of $D_{\text{KL}}(q(\alpha) \| \tilde{p}(\alpha))$ will be achieved when its support $\hat{\alpha}$ coincides with the maximum of $\tilde{p}(\alpha)$. In other words, the support $\hat{\alpha}'$ of $q(\alpha)'$ can be found as


$$\hat{\alpha}' = \underset{\alpha}{\operatorname{argmax}} p(\mathbf{y} | \alpha, \hat{\Theta}) p(\alpha). \quad (21)$$

Similarly, we proceed for the estimation of locations $q(\Theta)$. To this end we now let $q(\alpha) \leftarrow q(\alpha)'$ and minimize $\mathcal{F}(q(\Theta)q(\alpha) \| p(\mathbf{y}, \alpha, \Theta))$ with respect to $q(\Theta)$. The latter is found as a solution to the minimization

$$q(\Theta)' = \underset{q(\Theta)}{\operatorname{argmin}} D_{\text{KL}}(q(\Theta) \| \tilde{p}(\Theta)), \quad (22)$$

³We will use a subscript symbol (\cdot)' to indicate a new "updated" value of a function or parameters to distinguish it from the corresponding function of values prior to optimization of the free energy.





where $\tilde{p}(\Theta)$ is given as

$$\tilde{p}(\Theta) \propto \exp(\mathbb{E}_{q(\alpha)} \log p(\mathbf{y}, \alpha, \Theta)) \propto p(\mathbf{y} | \hat{\alpha}, \Theta) p(\Theta). \quad (23)$$

Note that here $\hat{\alpha}$ is now kept fixed, and $\tilde{p}(\Theta)$ approximates the posterior distribution $p(\Theta | \mathbf{y}, \hat{\alpha})$. Due to $q(\Theta)$ being set to a Dirac measure, the support of this measure should be chosen so as to minimize the Kullback-Leibler divergence $D_{\text{KL}}(\tilde{p}(\Theta) \| q(\Theta))$ by solving the following optimization problem:

$$\hat{\Theta}' = \underset{\Theta}{\operatorname{argmax}} p(\mathbf{y} | \hat{\alpha}, \Theta) p(\Theta), \quad (24)$$

which leads to a new estimate of $q(\Theta)'$. Then, we set $q(\Theta) \leftarrow q(\Theta)'$ and repeat the estimation step for $q(\alpha)$ until some convergence criterion is met.

This procedure establishes a general approach towards nonlinear parameter estimation within SBL. However, neither the support estimation step (24) nor the source localization (24) is an easily tractable optimization. In fact, they form what we will call an *outer* update iteration loop. Solutions to both (21) and (24) will be iterative in nature and will form *internal* update iterations. We will discuss in detail in the following section.

Estimation of sparsity parameters and source support pattern

Let us begin with the solution to (21) to estimate the support of $q(\alpha)$. Also, we now explicitly make use of the assumption that $p(\alpha) \propto \text{const}$. Thus, we consider the maximization of the marginal likelihood $p(\mathbf{y} | \alpha, \hat{\Theta})$ with respect to α . With $q(\Theta)$ being fixed, it is easy to see that the former becomes (see also (17))

$$p(\mathbf{y} | \alpha, \hat{\Theta}) \propto \frac{1}{|\Sigma_{\alpha}(\hat{\Theta})|^{\frac{1}{2}}} e^{-\frac{1}{2} \mathbf{y}^T \Sigma_{\alpha}(\hat{\Theta})^{-1} \mathbf{y}} \quad (25)$$

with

$$\Sigma_{\alpha}(\hat{\Theta}) = \lambda_{\xi}^{-1} \mathbf{I} + \Phi \mathbf{G}(\hat{\Theta}) \mathbf{A}^{-1} \mathbf{G}(\hat{\Theta})^T \Phi^T.$$

By keeping $q(\Theta)$ fixed, we effectively “linearize” the estimation problem, reducing the estimation of support $q(\alpha)$ to a classical SBL inference problem. Indeed, the form of the likelihood (25) now coincides exactly with that employed in a classical SBL for linear models. As such, we can reuse whatever SBL algorithms we are happy with to maximize it. Here we will use a FMLM method discussed in [45], making use of the assumption that $p(\alpha) \propto \text{const}$. Let us sketch this method here for completeness, now explicitly accounting for the locations $\hat{\Theta}$.

Define a negative log-likelihood function as follows


$$\ell(\alpha) \triangleq -\log p(\mathbf{y} | \alpha, \hat{\Theta}) = \frac{1}{2} \log |\Sigma_{\alpha}(\hat{\Theta})| + \frac{1}{2} \mathbf{y}^T \Sigma_{\alpha}(\hat{\Theta})^{-1} \mathbf{y}. \quad (26)$$


Then, we partition α into two sets: a single component α_l and all other entries $\alpha_{\bar{l}}$. This allows us to partition $\ell(\alpha)$ in (26) as

$$\ell(\alpha, \alpha_l) = \text{const}(\alpha_{\bar{l}}) + \frac{1}{2} \left[\log(\alpha_l) - \log(\alpha_l + \varsigma_l^{-1}) + \frac{\mu_l^2 \varsigma_l^{-2}}{\alpha_l + \varsigma_l^{-1}} \right], \quad (27)$$

where now

$$\varsigma_l(\hat{\Theta}) \equiv \varsigma_l \triangleq (\mathbf{c}_l^T \Psi_l^{-1} \mathbf{c}_l)^{-1}, \quad \mu_l(\hat{\Theta}) \equiv \mu_l \triangleq \varsigma_l \mathbf{c}_l^T \Psi_l^{-1} \mathbf{y}, \quad (28)$$





with $c_l(\hat{\Theta}) \equiv c_l \triangleq \Phi[G(\hat{\Theta})]_l$ and

$$\Psi_l \triangleq \lambda_\xi^{-1} \mathbf{I} + \sum_{m \in \mathcal{L}; m \neq l} \alpha_m^{-1} c_m c_m^\top.$$

The vector c_l represents the response of the l th Green's function, projected onto the measurement space via the measurement matrix Φ . The optimum of (27) with respect to α_l is given by:

$$\hat{\alpha}_l = \begin{cases} (\mu_l^2 - \varsigma_l)^{-1} & \frac{\mu_l^2}{\varsigma_l} > 1 \\ \infty & \text{otherwise.} \end{cases} \quad (29)$$

Now to find the gas source support, we evaluate (29) with l ranging in \mathcal{L} , in a round-robin fashion, until some stopping criterion is met. These iterative update cycles partition all components in the surrogate model (7) into *active* components with finite sparsity parameters, and *inactive* ones, for which the sparsity parameter is infinity. These updates form one of the *internal* update loops – an iterative estimation of $\hat{\alpha}$, or rather of $q(\alpha)$ with $q(\Theta)$ fixed.

Once convergence of this inner loop is established, we can readily use an estimate $\hat{\alpha}$ to infer a sparse weight vector w from the conditional marginal posterior pdf

$$p(w|\hat{\alpha}, \hat{\Theta}, \mathbf{y}) \propto p(\mathbf{y}|w, \hat{\Theta})p(w|\hat{\alpha}).$$

The corresponding pdf is Gaussian, with its mean \hat{w} and covariance matrix $\hat{\Sigma}_w$ given by

$$\begin{aligned} \hat{\Sigma}_w &= \left(\lambda_\xi G(\hat{\Theta})^\top \Phi^\top \Phi G(\hat{\Theta}) + \hat{\mathbf{A}}^{-1} \right)^{-1}, \\ \hat{w} &= \lambda_\xi \hat{\Sigma}_w G(\hat{\Theta})^\top \Phi^\top \mathbf{y}, \end{aligned} \quad (30)$$

where $\hat{\mathbf{A}} = \text{diag}(\hat{\alpha})$. Naturally, the latter is computed only for active components, since for inactive ones the weight estimates become numerically zero.

Estimation of source location parameters

Let us now consider the estimation of $q(\Theta)$ with (24). This step is a bit more elaborate due to the nonlinear dependency of $p(\mathbf{y}|\hat{\alpha}, \Theta)$ on Θ . As such, this step will require a nonlinear optimization, since we do not discretize Θ . We realize the latter using a gradient-descent technique [36]. Again, we will make use of an assumption that Θ is flat.⁴ Thus, solving (24) amounts to finding support $\hat{\Theta}$ of $q(\Theta)$ that maximizes the marginal likelihood $p(\mathbf{y}|\hat{\alpha}, \Theta)$.

Maximizing $p(\mathbf{y}|\hat{\alpha}, \Theta)$ with respect to Θ directly is rather complicated. Instead, we propose to utilize the Expectation-Maximization algorithm (EM) algorithm to perform this optimization via complete data. Let us again consider (16). For a fixed $q(\alpha)$ we can state that


$$p(\mathbf{y}|\hat{\alpha}, \Theta) = \iint p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|w, \Theta)p(w, \hat{\alpha})d\mathbf{f}dw \quad (31)$$


$$\equiv \int p(\mathbf{y}|\mathbf{f} = G(\Theta)w)p(w, \hat{\alpha})dw, \quad (32)$$

where we explicitly account for the fact that \mathbf{f} deterministically depends on Θ and w via PDE $\mathbf{f} = G(\Theta)w$.⁵ It can be observed that w (as well as \mathbf{f}) can thus be both treated as complete data for Θ [46]. We can now formulate the following iterative EM procedure to maximize $p(\mathbf{y}|\hat{\alpha}, \Theta)$ with respect to Θ .

⁴Alternatively, $p(\Theta)$ can be incorporated into the estimation quite trivially.

⁵This is needed to account for the fact that $p(\mathbf{f}|w, \Theta)$ is a measure, rather than a function.





E-step of the algorithm

Define the following function

$$Q(\Theta|\hat{\Theta}) \triangleq \mathbb{E}_{p(\mathbf{w}|\hat{\alpha}, \hat{\Theta}, \mathbf{y})} -\log\{p(\mathbf{y}|\mathbf{f} = \mathbf{G}(\Theta)\mathbf{w})p(\mathbf{w}|\hat{\alpha})\}, \quad (33)$$

which is known as the Q -function of the EM algorithm. Notice that in contrast to a usual convention, we define it as an expectation of the negative logarithm of complete likelihood with respect to the distribution of the complete data given a current estimate of locations $\hat{\Theta}$ and, of course, fixed sparsity parameters $\hat{\alpha}$. This is done to formulate a minimization problem instead of a more usual maximization. The distribution of complete data $p(\mathbf{w}|\hat{\alpha}, \hat{\Theta}, \mathbf{y})$ in (33) can be computed quite easily.

First, we note that

$$p(\mathbf{w}, \hat{\alpha}, \hat{\Theta}|\mathbf{y}) = \int p(\mathbf{f}, \mathbf{w}, \hat{\alpha}, \hat{\Theta}|\mathbf{y}) d\mathbf{f} = \frac{p(\mathbf{y}|\mathbf{w}, \hat{\Theta})p(\mathbf{w}|\hat{\alpha})p(\hat{\Theta})}{p(\mathbf{y})}, \quad (34)$$

where

$$p(\mathbf{y}|\mathbf{w}, \hat{\Theta}) \propto \exp\left(-\frac{\lambda_\xi}{2} \|\mathbf{y} - \Phi \mathbf{G}(\hat{\Theta})\mathbf{w}\|^2\right).$$

It is now easy to see that

$$p(\mathbf{w}|\hat{\alpha}, \hat{\Theta}, \mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{w}, \hat{\Theta})p(\mathbf{w}|\hat{\alpha})}{p(\mathbf{y}|\hat{\alpha}, \hat{\Theta})} \quad (35)$$

where both $p(\mathbf{y}|\mathbf{w}, \hat{\Theta})$ and $p(\mathbf{w}|\hat{\alpha})$ are Gaussian pdf. It then follows that $p(\mathbf{w}|\hat{\alpha}, \hat{\Theta}, \mathbf{y})$ is also Gaussian. Its mean $\hat{\mathbf{w}}$ and covariance $\hat{\Sigma}_w$ can be trivially found to equal to

$$\hat{\Sigma}_w = \left(\lambda_\xi \mathbf{G}(\hat{\Theta})^\top \Phi^\top \Phi \mathbf{G}(\hat{\Theta}) + \hat{\mathbf{A}}^{-1}\right)^{-1} \quad (36)$$

$$\hat{\mathbf{w}} = \hat{\Sigma}_w \lambda_\xi \mathbf{G}(\hat{\Theta})^\top \Phi^\top \mathbf{y}. \quad (37)$$

Taking now the expectation in (33) and ignoring all the terms that are independent of Θ we obtain

$$Q(\Theta|\hat{\Theta}) = -\log\{p(\mathbf{y}|\mathbf{f} = \mathbf{G}(\Theta)\hat{\mathbf{w}})\}, \quad (38)$$


which is now optimized during the M-step of the algorithm.


M-step of the optimization

At the M-step we now minimize the Q -function (38) with respect to Θ to find a new support $\hat{\Theta}'$ of $q(\Theta)$ by solving the following optimization problem:

$$\hat{\Theta}' = \underset{\Theta}{\operatorname{argmin}} Q(\Theta|\hat{\Theta}). \quad (39)$$

This is clearly a nonlinear optimization problem. We find its solution via gradient descent, which naturally requires computing the gradient of (38). To simplify this computation, we reformulate





(39) as a constrained optimization, keeping the dependency on \mathbf{f} explicit. To this end, we define the following cost function

$$J(\Theta, \mathbf{f}) \triangleq -\log p(\mathbf{y}|\mathbf{f}) = \frac{\lambda_\xi}{2} \|\mathbf{y} - \Phi \mathbf{f}\|^2, \quad (40)$$

$$\text{s.t. } \zeta(\Theta, \mathbf{f}) \triangleq \mathbf{f} - \mathbf{G}(\Theta) \hat{\mathbf{w}} = \mathbf{0}, \quad \Theta \in \Omega. \quad (41)$$

Obviously, \mathbf{f} is an implicit function of Θ . Then, (39) can be equivalently represented as the following constrained optimization problem:

$$\hat{\Theta}' = \underset{\Theta}{\text{argmin}} J(\Theta, \mathbf{f}) \quad (42)$$

$$\text{s.t. } \zeta(\Theta, \mathbf{f}) = \mathbf{0}, \quad \Theta \in \Omega.$$

Now, instead of solving (42) exactly, we resort to reducing $J(\Theta, \mathbf{f})$ within a single gradient update step:

$$\hat{\Theta}' = \hat{\Theta} - \rho \nabla_{\Theta} J(\hat{\Theta}, \hat{\mathbf{f}}), \quad (43)$$

where $\nabla_{\Theta} J(\hat{\Theta}, \hat{\mathbf{f}})$ is a gradient of (42) with respect to Θ evaluated at current estimates $\hat{\Theta}$ and $\hat{\mathbf{f}}$. The parameter ρ is an appropriately chosen step size. It is the computation of the gradient in (43) that is the actual reason we keep the variable \mathbf{f} explicit in the optimization problem. To be more specific, a constrained form of optimization (42) will allow us to use an adjoint state method [9, 47] to compute $\nabla_{\Theta} J(\hat{\Theta}, \hat{\mathbf{f}})$. As such, the whole estimation approach will remain valid even if we do not have an explicit solver for \mathbf{f} , as we do in the case of Green's function method, but instead are forced to resort to numerical PDE solution approaches. This is how the adjoint method is formulated.

Let $\boldsymbol{\eta}$ be an adjoint variable defined as a solution to

$$\left(\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \mathbf{f}} \right)^{\top} \boldsymbol{\eta} = - \left(\frac{\partial J(\Theta, \mathbf{f})}{\partial \mathbf{f}} \right)^{\top}. \quad (44)$$

It is rather straightforward to show that

$$\boldsymbol{\eta} = -\lambda_\xi \Phi^{\top} (\Phi \mathbf{f} - \mathbf{y}). \quad (45)$$

Obviously, $\boldsymbol{\eta}$ is an implicit function of source weights $\hat{\mathbf{w}}$ and locations Θ , since $\mathbf{f} = \mathbf{G}(\Theta) \hat{\mathbf{w}}$ by our definition. Thus, at the current estimate of $\hat{\Theta}$ and $\hat{\mathbf{w}}$, the concentration estimation is given by $\hat{\mathbf{f}} = \mathbf{G}(\hat{\Theta}) \hat{\mathbf{w}}$ and the adjoint variable can be computed as

$$\hat{\boldsymbol{\eta}} = -\lambda_\xi \Phi^{\top} (\Phi \hat{\mathbf{f}} - \mathbf{y}) = -\lambda_\xi \Phi^{\top} (\Phi \mathbf{G}(\hat{\Theta}) \hat{\mathbf{w}} - \mathbf{y}), \quad (46)$$


which is an N -dimensional vector, where N is the number of discrete elements we use to tessellate Ω .⁶ Now, we can evaluate the gradient $\nabla_{\Theta} J(\hat{\Theta}, \hat{\mathbf{f}})$.


First, let us recall that we defined Θ as a collection of L source locations θ_l , $l \in \mathcal{L}$. Instead of computing the gradient for all components at the same time, we consider an update for a single component $l \in \mathcal{L}$. In this case (43) becomes⁷

$$\hat{\theta}_l' = \hat{\theta}_l - \rho \nabla_{\theta_l} J(\hat{\Theta}, \hat{\mathbf{f}})^{\top}, \quad \forall l \in \mathcal{L}, \quad (47)$$

⁶We note that the solution \mathbf{f} to a general PDE (1) can also be found numerically and inserted directly into (46). In this case, N would be a number of finite elements of discretization cells, depending on the approach used to solve the PDE.

⁷Here, a gradient is considered as a row vector, against a common notation for vectors as columns adopted in this work.





where

$$\nabla_{\theta_l} J(\hat{\Theta}, \hat{\mathbf{f}}) = (\hat{\boldsymbol{\eta}})^\top \frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_l} \Big|_{\Theta = \hat{\Theta}, \mathbf{f} = G(\hat{\Theta})\mathbf{w}} \quad (48)$$

Note that the derivative $\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_l}$ has to be computed for each element of the location θ_l . Indeed, for a 2D domain Ω , the source location has two components: $\theta_l = [\theta_{l,x}, \theta_{l,y}]^\top$, one for each dimension. Similarly, for a 3D domain there will be three components $\theta_l = [\theta_{l,x}, \theta_{l,y}, \theta_{l,z}]^\top$. Here, the subscripts (l,x) , (l,y) , and (l,z) refer to the corresponding components of the source location. The derivative $\frac{\partial \zeta(\Theta, \mathbf{f}, \mathbf{w})}{\partial \theta_l}$ can then be computed as

$$\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_l} = \left[\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,x}}, \frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,y}} \right] \quad (49)$$

for the 2D case, and

$$\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_l} = \left[\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,x}}, \frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,y}}, \frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,z}} \right] \quad (50)$$

for the 3D case, respectively. Taking the structure of $G(\Theta)$ into account, it is easy to see that in general

$$\frac{\partial \zeta(\Theta, \mathbf{f})}{\partial \theta_{l,\#}} = - \left[\hat{w}_l \frac{\partial G(\mathbf{x}_0, \theta_l)}{\partial \theta_{l,\#}}, \dots, \hat{w}_l^{[i]} \frac{\partial G(\mathbf{x}_{N-1}, \theta_l)}{\partial \theta_{l,\#}} \right]^\top$$

where the symbol $\# \in \{x, y, z\}$ is used to refer to the corresponding coordinate. In other words, it requires a corresponding derivative of the Green's function evaluated over the discretization cells of the domain Ω , weighted by the current weight of the source \hat{w}_l – the mean of the complete data given by. The latter weights are computed from (37) during the E-step. Naturally, the update (47) is evaluated for all active components $l \in \mathcal{L}$.

The E- and M-steps are then repeated multiple times, leading to a sequence of estimates


$$\underbrace{\hat{\mathbf{w}}}_{\text{E-step (37)}} \rightarrow \underbrace{\hat{\Theta}'}_{\text{M-step (43)}} \rightarrow \underbrace{\hat{\mathbf{w}}'}_{\text{E-step (37)}} \rightarrow \underbrace{\hat{\Theta}''}_{\text{M-step (43)}} \rightarrow \underbrace{\hat{\mathbf{w}}''}_{\text{E-step (37)}} \rightarrow \dots$$


This sequence of updates can be shown to lead to a non-increasing sequence of free energies $\mathcal{F}(q(\Theta)q(\boldsymbol{\alpha})||p(\mathbf{y}, \boldsymbol{\alpha}, \Theta))$, thus making the optimization optimal.

3.3.3. Applications and extensions: Cramer-Rao Lower Bound for a Gas Source Localization

Cramer-Rao Lower Bound for a Gas Source Localization

While the algorithm presented above focuses primarily on detecting and estimating smoke source parameters, little attention has been given to quantifying the theoretical limits of such estimators. The main challenge here is the high complexity of the model, which makes detailed theoretical





analysis often intractable. However, investigations performed within the TEMA project showed that the Green's function approach for solving the PDE, when combined with SBL, can be used to compute Cramer-Rao Lower bound (CRLB) for estimating source location, thus quantifying achievable source localization precision.

The summary presented here makes use of a Poisson's equation in 2D – a simplified advection-diffusion model under the assumption of zero wind and Ω constrained to a unit circle – with a single source modeled as a scaled Dirac measure with unknown support. By inspecting the source estimation within the SBL framework, we were able to derive the corresponding CRLB for the source location. The latter is shown to depend on the Green's function of the PDE and its derivatives. Finally, we show how the computed CRLB can be utilized for identifying new informative sampling locations. This provides a basis for autonomous drone exploration towards areas of high interest for source localization. In other words, we show how the computed bound can be used in information seeking [48] – a sampling strategy that will increase the information about the source position.

In the following, we thus will assume that the set of locations Θ includes only a single smoke source at location $\theta = [\theta_x, \theta_y]$ for the 2D case. Similarly, there is only one sparsity parameter α and weight w associated with this single source.

From (17) we can immediately recognize that conditioned on the parameters α and θ the observation \mathbf{y} follows a Gaussian distribution with zero mean and covariance matrix $\Sigma_\alpha(\theta)$. For this case, the general form of the elements of a 3×3 information matrix $\mathbf{I}(\alpha, \theta)$ for the parameters α and θ is given as [49]:

$$I_{\eta, \epsilon}(\alpha, \theta) = \text{tr} \left\{ \Sigma_\alpha(\theta)^{-1} \frac{\partial \Sigma_\alpha(\theta)}{\partial \eta} \Sigma_\alpha(\theta)^{-1} \frac{\partial \Sigma_\alpha(\theta)}{\partial \epsilon} \right\}, \quad (51)$$

where $\eta, \epsilon \in \{\alpha, \theta_x, \theta_y\}$ represent parameters with respect to which derivatives in (51) are computed. Let us now compute the corresponding expressions for evaluating the information matrix. First we note that for a single source the inverse $\Sigma_\alpha(\theta)^{-1}$ can be efficiently computed using ShermanMorrison formula [50]:

$$\Sigma_\alpha(\theta)^{-1} = \lambda_\xi \mathbf{I} - \frac{\lambda_\xi^2 \Phi \mathbf{g}(\theta) \mathbf{g}(\theta)^\top \Phi^\top}{\alpha + \lambda_\xi \mathbf{g}(\theta)^\top \Phi^\top \Phi \mathbf{g}(\theta)} \quad (52)$$


Furthermore, taking the partial derivatives of $\Sigma_\alpha(\theta)$ with respect to the parameters of interest we obtain


$$\frac{\partial \Sigma_\alpha(\theta)}{\partial \alpha} = - \frac{\Phi \mathbf{g}(\theta) \mathbf{g}(\theta)^\top \Phi^\top}{\alpha^2} \quad (53)$$

$$\frac{\partial \Sigma_\alpha(\theta)}{\partial \theta_x} = 2\alpha^{-1} \Phi \frac{\partial \mathbf{g}(\theta)}{\partial \theta_x} \mathbf{g}(\theta)^\top \Phi^\top \quad (54)$$

$$\frac{\partial \Sigma_\alpha(\theta)}{\partial \theta_y} = 2\alpha^{-1} \Phi \frac{\partial \mathbf{g}(\theta)}{\partial \theta_y} \mathbf{g}(\theta)^\top \Phi^\top \quad (55)$$

It is important to stress that the elements of the information matrix naturally depend on the derivatives of the Green's function $\mathbf{g}(\theta)$. As we mentioned, the Green's function is not always available in a closed form. However, physics-inspired neural networks can be used to learn the Green's function as well as its derivatives from the data [51, 52]. In our case – for zero wind and Ω constrained to a unit circle – we can compute it analytically. Specifically, the Green's function of





the advection-diffusion equation with zero wind (also known as the Poisson equation) over the unit circle can be computed as (see [40, 53])

$$G(\mathbf{p}, \boldsymbol{\theta}) = \frac{1}{4\pi} \log \left(\frac{1 - 2\mathbf{p}^\top \boldsymbol{\theta} + \|\mathbf{p}\|^2 \|\boldsymbol{\theta}\|^2}{\|\mathbf{p} - \boldsymbol{\theta}\|^2} \right) \quad (56)$$

To this end, consider a discrete cell in Ω with coordinates (x, y) , and define $\ell(\mathbf{p}, \boldsymbol{\theta}) \triangleq (1 - 2\mathbf{p}^\top \boldsymbol{\theta} + \|\mathbf{p}\|^2 \|\boldsymbol{\theta}\|^2)$. Then, the derivative of the Green's function at the location $\mathbf{p} = [x, y]^\top$, i.e., a single element in the vector $\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \theta_x}$ or $\frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \theta_y}$, can be computed as

$$\left. \frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \theta_x} \right|_{\mathbf{p}} = \frac{\|\mathbf{p}\|^2 - 1}{2\pi \ell(\mathbf{p}, \boldsymbol{\theta}) \|\mathbf{p} - \boldsymbol{\theta}\|^2} \times \left[\theta_x (1 + \|\mathbf{p}\|^2 - 2y\theta_y) + x(\theta_y^2 - \theta_x^2 - 1) \right] \quad (57)$$

$$\left. \frac{\partial \mathbf{g}(\boldsymbol{\theta})}{\partial \theta_y} \right|_{\mathbf{p}} = \frac{\|\mathbf{p}\|^2 - 1}{2\pi \ell(\mathbf{p}, \boldsymbol{\theta}) \|\mathbf{p} - \boldsymbol{\theta}\|^2} \times \left[\theta_y (1 + \|\mathbf{p}\|^2 - 2x\theta_x) + y(\theta_x^2 - \theta_y^2 - 1) \right] \quad (58)$$

Now, (57) and (58) are inserted into (54) and (55), respectively. Using the latter two expressions together with (53) we can now evaluate (52) and compute the CRLB $\mathbf{C}(\alpha, \boldsymbol{\theta})$ as a lower bound on the variance of an unbiased estimator of $(\alpha, \boldsymbol{\theta})$:

$$\mathbf{C}(\alpha, \boldsymbol{\theta}) \triangleq \begin{pmatrix} I_{\alpha, \alpha}(\alpha, \boldsymbol{\theta}) & I_{\alpha, \theta_x}(\alpha, \boldsymbol{\theta}) & I_{\alpha, \theta_y}(\alpha, \boldsymbol{\theta}) \\ I_{\theta_x, \alpha}(\alpha, \boldsymbol{\theta}) & I_{\theta_x, \theta_x}(\alpha, \boldsymbol{\theta}) & I_{\theta_x, \theta_y}(\alpha, \boldsymbol{\theta}) \\ I_{\theta_y, \alpha}(\alpha, \boldsymbol{\theta}) & I_{\theta_y, \theta_x}(\alpha, \boldsymbol{\theta}) & I_{\theta_y, \theta_y}(\alpha, \boldsymbol{\theta}) \end{pmatrix}^{-1} = \begin{pmatrix} C_{\alpha, \alpha}(\alpha, \boldsymbol{\theta}) & \mathbf{c}_{\alpha, \boldsymbol{\theta}}(\alpha, \boldsymbol{\theta})^\top \\ \mathbf{c}_{\boldsymbol{\theta}, \alpha}(\alpha, \boldsymbol{\theta}) & \mathbf{C}_{\boldsymbol{\theta}, \boldsymbol{\theta}}(\alpha, \boldsymbol{\theta}) \end{pmatrix}. \quad (59)$$

Note that in (59) we partitioned the covariance matrix $\mathbf{C}(\alpha, \boldsymbol{\theta})$ into four blocks corresponding to parameter α and location parameter $\boldsymbol{\theta}$. This is done to account for the value of the sparsity parameter on the location estimate, as will be discussed in the following.

Conditioning the bound using SBL


As expected, CRLB matrix $\mathbf{C}(\alpha, \boldsymbol{\theta})$ depends on the parameters α and $\boldsymbol{\theta}$ due to the nonlinear dependency of the likelihood (17) on these parameters. As such, CRLB is a local bound, characterizing the estimator performance around some "working" location of interest within the parameter space $(\alpha, \boldsymbol{\theta})$. Furthermore, due to the sparsity-inducing property of SBL, we know that α can become infinite for some test locations $\boldsymbol{\theta}$ [54]. Let us, for the moment, fix $\boldsymbol{\theta}$ and consider the value of $\hat{\alpha}$ that maximizes the likelihood (17). Specializing results from [54] to our case we find


$$\hat{\alpha} = \begin{cases} (\mu^2 - \varsigma)^{-1} & \frac{\mu^2}{\varsigma} > 1 \\ \infty & \text{otherwise,} \end{cases} \quad (60)$$

where

$$\varsigma \triangleq \left(\lambda_\xi \mathbf{g}(\boldsymbol{\theta})^\top \boldsymbol{\Phi}^\top \boldsymbol{\Phi} \mathbf{g}(\boldsymbol{\theta}) \right)^{-1}, \mu \triangleq \varsigma \lambda_\xi \mathbf{g}(\boldsymbol{\theta})^\top \boldsymbol{\Phi}^\top \mathbf{y}. \quad (61)$$

As we see, for a single source, μ and ς are the mean and variance of the source weight obtained by finding the maximum likelihood fit of Green's function centered at $\boldsymbol{\theta}$ to the collected measurements \mathbf{y} . Indeed, the product $\boldsymbol{\Phi} \mathbf{g}(\boldsymbol{\theta})$ samples the corresponding Green's function at the locations where measurements were taken. Parameter μ then measures the projection of the measurements \mathbf{y} onto these samples. The ratio μ^2/ς used in (60) to test the presence of a source can thus





be interpreted as an Signal-to-Noise Ratio (SNR) estimate. In other words, according to (60) the parameter $\hat{\alpha}$ is finite when the corresponding source SNR is above 1, or equivalently, above 0dB. Moreover, given a hypothetical source location θ we can (i) tell if measured data \mathbf{y} supports source at this location, i.e., $\hat{\alpha} < \infty$, and (ii) condition the information about source position θ on the corresponding value $\hat{\alpha}$ at the location of interest. Specifically, we can compute the conditional CRLB as a Schur's complement of $C_{\alpha,\alpha}$ in $C(\alpha,\theta)$:

$$C(\theta|\hat{\alpha}) = C_{\theta,\theta}(\hat{\alpha},\theta) - \frac{c_{\theta,\alpha}(\hat{\alpha},\theta)c_{\theta,\alpha}(\hat{\alpha},\theta)^T}{C_{\alpha,\alpha}(\hat{\alpha},\theta)} \quad (62)$$

Thus, given measurements \mathbf{y} , we can now test whether these support the existence of a source at location θ with finite $\hat{\alpha}$ and then compute the corresponding uncertainty bound $C(\theta|\hat{\alpha})$ for this location. Note that for $\hat{\alpha} \rightarrow \infty$, we have $\mathbf{I}(\alpha,\theta) \rightarrow 0$, making the CRLB diverge toward infinity.

Objective for CRLB-based information seeking

Using the conditional CRLB, we can now assess the decrease in CRLB when planning measurements at new locations. This will allow avoiding non-informative locations and thus build the foundation for the information-seeking type of data collection [48], as explained in the following.

Let us assume now that given measurements \mathbf{y} a source has been detected at location $\hat{\theta}$ with the corresponding sparsity parameter $\hat{\alpha}$. Can we find new sampling locations that will decrease the bound $C(\hat{\theta}|\hat{\alpha})$, i.e., make the gas source localization more precise? To address this question, consider now a new test location $\mathbf{p}^* \in \Omega$ for a potential measurement. This location generates a selection vector Φ^* : vector ϕ^* contains 1 at the entry that corresponds to the spatial location \mathbf{p}^* and zeros at all other entries. Now, since CRLB does not depend on the actual measurement z^* at \mathbf{p}^* , we can compute a prediction $C^*(\hat{\theta}|\hat{\alpha},\mathbf{p}^*)$ at $\hat{\theta}$ by augmenting the selection matrix Φ as

$$\Phi^* = \begin{pmatrix} \Phi \\ (\phi^*)^T \end{pmatrix} \quad (63)$$


We recompute (52)–(55) by exchanging Φ with Φ^* , which allows to compute a prediction $C^*(\alpha,\theta|\mathbf{p}^*)$ from (59) and, finally, compute $C^*(\hat{\theta}|\hat{\alpha},\mathbf{p}^*)$ from (62) as a predicted bound for the source location $\hat{\theta}$. The latter can be used as an input into information-seeking algorithms for autonomous navigation. In particular, by selecting a new measurement location as a minimizer of a predicted CRLB, e.g.,

$$\hat{\mathbf{p}}^* = \underset{\mathbf{p}^* \in \Omega}{\operatorname{argmin}} \operatorname{tr} \left\{ C^*(\hat{\theta}|\hat{\alpha},\mathbf{p}^*) \right\}. \quad (64)$$

Criterion (64), also known as the A-optimality, thus minimizes the average variance of (or, equivalently, maximizes the averaged information about) the source location at $\hat{\theta}$.

Simulation results

Let us now analyze the computed CRLB numerically. To this end, we discretize Ω , generating a uniform grid with 50 elements in the interval $[-1,1]$ for both x - and y - directions. This results in $N = 1876$ spatial cells C_i , $i \in \mathcal{N}$, within a unit circle; each cell has a spatial dimension 0.04×0.04 units. This defines the discretization resolution of the GSL problem. For simulations, we will place a gas source at location $\theta = [0,0]^T$. Without loss of generality, we will also select $\kappa = 1$ and $w = 1$ in (4). The Green's function for this scenario is shown in Fig. 2, where we also show an example







Figure 2. Example of a Green's function $G(p, \tilde{\theta})$ for the source at $\tilde{\theta} = [0, 0]^T$. Red marks on the right image indicate random positions of measurements y . The Green's function has a singularity at $\tilde{\theta}$ not shown in the plot.

of a sampling pattern with $M = 20$ samples. To study the impact of measurement noise on CRLB we also define SNR as $\text{SNR} = 10 \log_{10}(w^2 \lambda_{\xi})$, where λ_{ξ} is the precision of the additive noise.

We begin investigations by studying the dependency of CRLB on the number of measurements M . To this end, we first consider the probability of detecting a source at the location $\tilde{\theta}$ with SBL, i.e., the frequency of cases when $\hat{\alpha}$ computed from (6o) is finite. The corresponding results, averaged over 1000 Monte Carlo runs, are shown in Fig. 3. Note that for very low SNR most of the

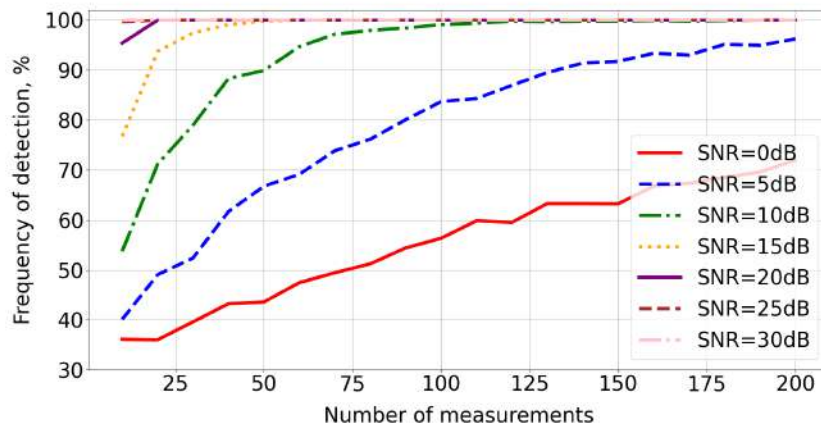




Figure 3. Frequency of estimating $\tilde{\alpha} < \infty$ versus number of measurement samples M for different SNR values.

time we will not be able to detect a source at $\tilde{\theta}$ with a low number of measurements: although the detection frequency increases as more samples are collected, at low SNR more samples are needed to make any source detection possible. Furthermore, although $M = 200$ is only about 10% of all possible sampling locations, collecting these samples in practice with robotic platforms will require significant planning resources.

Assuming a source can be detected, let us study how $\text{tr}\{C(\tilde{\theta}|\hat{\alpha})\}$ changes as the number of samples grows. We restrict ourselves to SNR in the range above 14dB. The corresponding results,





averaged over 1000 Monte Carlo runs, are shown in Fig. 4. We also show on the plot the horizon-




Figure 4. Conditional CRLB $\text{tr}\{C(\tilde{\theta}|\hat{\alpha})\}$ versus number of measurement samples M for different SNR values.

tal line that corresponds to the discrete cell resolution of 0.04 and cut the graph at the size of the domain Ω . As expected, with increasing SNR and M , the CRLB decreases. We also see that for SNR below ≈ 20 dB the achievable bound can be larger than the whole domain Ω , making GSL useless in this regime as the expected gas localization uncertainty exceeds the whole modeling domain. Furthermore, in a high SNR regime with roughly $M = 40$, we can already achieve sub-cell localization accuracy. The plot also shows that increasing SNR by 4dB allows for reducing the number of measurement samples needed to achieve sub-cell localization precision by half. In Fig. 5 we plot the same results but now as a function of SNR. In this plot, we see that localization






Figure 5. Conditional CRLB $\text{tr}\{C(\tilde{\theta}|\hat{\alpha})\}$ versus SNR for different number of measurements M .





precision increases linearly with an increase in SNR. Also, with only $M = 20$ samples, at least 20dB SNR is needed to reduce localization standard deviation below the radius of Ω .

Let us now evaluate the prediction $C^*(\tilde{\theta}|\hat{\alpha}, p^*)$ as we foresee it for information-seeking purposes to improve source estimation precision. To this end we will evaluate the change $\Delta\text{CRLB}(p^*)$ in the bound as

$$\Delta\text{CRLB}(p^*) = \sqrt{\text{tr}\{C(\tilde{\theta}|\hat{\alpha})\}} - \sqrt{\text{tr}\{C^*(\tilde{\theta}|\hat{\alpha}, p^*)\}}$$

as a function of a potential measurement location p^* and fix SNR at 14dB. The results are sum-






Figure 6. Change in CRLB as a function of a new measurement location $p^* \in \Omega$ with for $M = 20$ (left) and $M = 60$ (right).

marized in Fig. 6. As we can see, the most “relevant” areas are around the location of the source: measurements in these areas will decrease the CRLB by more than the size of the discretization cell. Areas that are further away from the source do not contribute that much. For instance, measurements close to the boundary will not impact the CRLB at all. We also note an asymmetry of the resulting uncertainty precision gain on the left plot in Fig. 6. Exactly at the location of the source, no prediction can be made, yet away from θ , the areas most relevant to exploration protrude into the domains with fewer available samples. These indicate more informative areas, and as such can be utilized to guide further exploration of Ω . Next, we generate additional 40 samples, increasing the total number of measurements to $M = 60$ while keeping SNR fixed at the same value. It can be observed (see Fig. 6 right) that the uncertainty pattern becomes more centered around the source location. This is due to measurements being now almost uniformly scattered over Ω . Moreover, we also see that the decrease of CRLB becomes less prominent: additional samples will no longer increase the estimation precision by much.

3.3.4. Applications and extensions: Distributed Gas Source Localization

An iterative distributed algorithm to maximize the lower bound (18) was proposed in [55, 54] for the 1D case employing the Poisson equation. Its extension to a 2D case is quite straightforward, yet with a few minor exceptions that we will outline in the following.





Multi-agent smoke sampling network

First, the measurement model needs to be modified to account for the fact that smoke concentration measurements are now collected by multiple drones. Consider a network of K robotic agents that take measurements of smoke concentrations and communicate wirelessly with each other. We model such a network with a strongly connected, weighted graph, with weights along the graph edges set according to the Laplacian rule [56]. Assume now that each agent collects M_k , $k \in \mathcal{K} \triangleq \{1, \dots, K\}$, noisy samples of the concentration $f(\mathbf{x})$ at locations $\mathbf{x}_{m,k}$, $m = 1, \dots, M_k$. Furthermore, w.l.o.g. we assume that $N \gg M_k$ and that $\mathbf{x}_{m,k}$, $\forall m, k$, are a subset of discretization cells. The measurements of the agent k we collect in a vector $\mathbf{y}_k \in \mathbb{R}_k^M$ such that

$$\mathbf{y}_k = \Phi_k \mathbf{f} + \boldsymbol{\xi}_k, \quad k \in \mathcal{K}, \quad (65)$$

where $\Phi_k \in \mathbb{R}^{M_k \times N}$ is a 0-1 sensing matrix that “selects” measured elements of \mathbf{f} . The perturbation $\boldsymbol{\xi}_k$ is assumed to be a homoscedastic zero-mean Gaussian random variable, with precision matrix $\lambda_\xi \mathbf{I}$ for some $\lambda_\xi > 0$.

Our goal now is to use $\mathbf{y} \triangleq [\mathbf{y}_1, \dots, \mathbf{y}_K]^\top$ to cooperatively estimate a sparse vector \mathbf{w} , locations Θ , and thus recover $\mathbf{f} = \mathbf{G}(\Theta)\mathbf{w}$. To this end, we pursue a distributed Bayesian approach towards parameter estimation proposed in [55], yet extend it to a 2D exploration domain.

Optimization strategy in a distributed setting

The optimization strategy discussed in Section 3.3.2 consists of two key steps: (i) estimating the support parameters α from (21), and (ii) estimating source locations Θ along with the weights \mathbf{w} as a solution to the optimization problem (40). It has been shown [55] that both of these optimizations can be solved distributively based on two quantities \mathbf{D} and \mathbf{d} defined as

$$\mathbf{D} \triangleq \mathbf{G}(\hat{\Theta})^\top \bar{\Phi} \mathbf{G}(\hat{\Theta}), \quad \text{and} \quad \mathbf{d} \triangleq \mathbf{G}(\hat{\Theta})^\top \bar{\mathbf{y}}, \quad (66)$$

where

$$\bar{\Phi} \triangleq \lambda_\xi \sum_{k \in \mathcal{K}} \Phi_k^\top \Phi_k, \quad \text{and} \quad \bar{\mathbf{y}} \triangleq \lambda_\xi \sum_{k \in \mathcal{K}} \Phi_k^\top \mathbf{y}_k \quad (67)$$

are two variables that “aggregate” all agents’ measurements. These quantities have to be computed cooperatively, using e.g., an averaged consensus algorithm [57]. Moreover, they do not depend on the parameters of interest in (26) and (40) and thus can be computed only once before the optimization.


The maximizer of (26) with respect to α_l is achieved at (see also [45, 58, 55])


$$\hat{\alpha}_l = \begin{cases} \frac{s_l^2}{q_l^2 - s_l}, & \text{if } q_l^2 > s_l, \\ 0, & \text{otherwise.} \end{cases} \quad (68)$$

where parameters s_l and q_l are computed as

$$s_l = S_l(1 - \alpha_l^{-1} S_l)^{-1}, \quad q_l = Q_l(1 - \alpha_l^{-1} S_l)^{-1},$$

with $S_l = [\mathbf{D} - \mathbf{D} \boldsymbol{\Sigma}_w \mathbf{D}]_{ll}$, $Q_l = [\mathbf{d} - \mathbf{D} \boldsymbol{\Sigma}_w \mathbf{d}]_l$ and $\boldsymbol{\Sigma}_w = (\mathbf{D} + \text{diag}\{\hat{\alpha}\})^{-1}$. We then evaluate (68) for all L components in a *round-robin* fashion, removing sources for which $\hat{\alpha}_l$ is zero. Upon convergence we then compute $\hat{\mathbf{w}} = \boldsymbol{\Sigma}_w \mathbf{d}$.





The optimization (40), on the other hand, has to be solved numerically, using the gradient descent method [36]. In a distributed version of the algorithm, both the gradient computation and location parameters update are realized over the network, with a so-called diffusion strategy [56]. In particular, in [55] the Combine-Then-Adapt (CTA) rule is used [56], which amounts to computing the following update at the iteration j of the gradient descent:

$$\Psi_{k,c} = \sum_{l \in \mathcal{N}(k)} a_{lk} \hat{\Theta}_{l,c}^{[j]}, \quad (\text{combination}) \quad (69)$$

$$\hat{\Theta}_{k,c}^{[j+1]} = \Psi_{k,c} - \mu \nabla_{\Theta_c} J(\Theta)|_{\Theta=\Psi_k} \quad (\text{adaptation}). \quad (70)$$

Here subindex $c \in \{x, y\}$ indicates the x - or y -coordinate of the source location, a_{lk} , $k, l \in \mathcal{K}$, are graph edge weights, $\nabla_{\Theta_c} J(\Theta)$ is a gradient of (40) evaluated along the coordinate c at the combined estimate Ψ_k for both x - and y -coordinates, and μ is an appropriately chosen step size. Note that in [55] the expression for the gradient was specified for a 1D case, while here a 2D case is considered. Thus, except for the subscript c , which indicates the differentiation direction, the expressions have the same form. The CTA update iterations (69) and (70) and then repeated until some suitable convergence criterion is satisfied.

Bottom-up cooperative initialization

In the original method [54, 55] a simple initialization strategy was proposed, which consisted of selecting initial locations $\hat{\Theta}^{[0]}$ by suitably partition Ω into $\tilde{L} \geq N$ grid points ϑ_l , $l = 1, \dots, \tilde{L}$, followed by the computation of $\bar{\Phi}$ and \bar{y} from (67) with averaged consensus and initialization of the support $\hat{\gamma}$ using the strategy proposed in [45]. This fast, non-cooperative initialization requires no agent interaction. However, in the 2D case, it tends to fail: by treating all sources independently, an initial model is constructed that poorly approximates the true solution. To address this, we propose a *cooperative* strategy inspired by orthogonal matching pursuit [59] – a greedy signal recovery method. In contrast to [54, 55], the new strategy builds the model from bottom up, incrementally adding sources to best “match” the observed concentration values. It leads to better initial models, despite requiring a cooperative stage.

Assume that $\bar{\Phi}$ and \bar{y} are available and let $g(\theta) \triangleq [G(x_1, \theta), \dots, G(x_N, \theta)]^T$. Furthermore, at the iteration i of the initialization, let


$$r^{[i]} = y - \Phi f^{[i]} = y - \Phi G(\hat{\Theta}^{[i]}) \hat{w}^{[i]}, \quad (71)$$


define the residual obtained by canceling the initialized model at iteration i . Similarly to (67) it can be shown that

$$\bar{r}^{[i]} \triangleq \lambda_\xi \Phi^T r^{[i]} = \bar{y} - \bar{\Phi} f^{[i]}. \quad (72)$$

Thus, given $\bar{\Phi}$ and \bar{y} , the current residual \bar{r} can be evaluated by each agent $k \in \mathcal{K}$, which during initialization then proceeds as in Algorithm 1.

An agent begins with an empty model by setting all source weights to 0 and sets the residual $\bar{r}^{[0]}$ to \bar{y} . Then, in line 5, we identify the source location that best fits the residual data – i.e., the observed concentration with “canceled” responses from already detected sources. The quantities in line 6 are scalar counterparts of those in (66), computed for a single Green’s function centered at the best-matching location ϑ_{\max} . Line 7 mirrors the test in (68). Here, the bottom-up strategy pays off: if this condition fails, SBL refrains from including a source at ϑ_{\max} and the initialization stops. Otherwise, a new source is “detected” and it is included in the model. Subsequently, each





Algorithm 1 Bottom-Up Initialization for agent k

```

1: Set:  $w_k \leftarrow 0$ ,  $\tilde{r}^{[0]} \leftarrow \bar{y}$ , number of sources  $l \leftarrow 0$ 
2: Initialize search space  $\vartheta_l$ , for  $l = 1, \dots, \tilde{L}$ 
3: while Continue initialization do
4:   # Finding the best fitting source
5:   Solve:  $\vartheta_{\max} = \operatorname{argmax}_{\vartheta} |g(\vartheta)^\top \tilde{r}^{[l]}|^2$ 
6:    $D_{\max} \leftarrow g(\vartheta_{\max})^\top \Phi g(\vartheta_{\max})$ ,  $d_{\max} \leftarrow g(\vartheta_{\max})^\top \tilde{r}$ 
7:   if  $\frac{d_{\max}^2}{D_{\max}} > 1$  then
8:     # Source detected at location  $\vartheta_{\max}$ 
9:      $l \leftarrow l + 1$ 
10:     $\theta_l \leftarrow \vartheta_{\max}$ 
11:     $\gamma_l \leftarrow \left( \frac{d_{\max}}{D_{\max}} \right)^2 - \frac{1}{D_{\max}}$ 
12:     $\hat{f}^{[l]} \leftarrow [g(\theta_1), \dots, g(\theta_l)] (D + \hat{\Gamma}^{-1})^{-1} d$ 
13:    Update residual:  $\tilde{r}^{[l]} \leftarrow \bar{y} - \Phi \hat{f}^{[l]}$ 
14:  else
15:    # No source detected. Stop initialization.
16:    break
17:  end if
18: end while

```


agent evaluates the combined response $f^{[l]}$ for all l detected sources (line 12), and updates the residual $\tilde{r}^{[l]}$ (line 13) by canceling the corresponding contribution.

Simulation results for distributed inference approach

Let us now demonstrate the algorithm's performance with some simulations. We set $\Omega = [-1, 1] \times [-1, 1]$ and $\kappa = 0.03$ and consider a constant wind vector $v = v(x)$ over Ω . As a performance measure, we use Mean Squared Error (MSE) criterion $\frac{1}{N} \|\hat{f} - f_{\text{true}}\|^2$ between the estimated and true concentration value, as well as Earth Mover's Distance (EMD) [60] – a discrete equivalent of a Wasserstein distance between probability distributions [61]. The EMD provides an indication of how well the source support is estimated, while the MSE criterion reveals how accurately the concentration values are reconstructed.

We first illustrate a single run of the algorithm. To this end, we partition Ω into $N = 2500$ discrete cells and draw $L_{\text{true}} = 3$ sources randomly distributed around the origin, with the release rates all set to 1. The components of the wind v are uniformly drawn from the interval $[-1, 1]$. We set SNR, defined as $10 \log_{10} \frac{\lambda_{\xi}}{N} \|f_{\text{true}}\|^2$ to 25dB. The corresponding ground truth concentration f_{true} is shown in Fig. 7a. For the estimation, we use a network of $K = 3$ fully connected agents, each collecting $M_k = 100$ samples along a Hilbert curve trajectory [62] within Ω . Such a path avoids collisions between the agents and effectively partitions the whole area into K non-overlapping regions (see also Fig. 7b).

To initialize the algorithm we set $\tilde{L} = 250$ and generate initial source locations ϑ_l , $l = 1, \dots, \tilde{L}$, using the Vogel's algorithm [63]. This uniformly covers the area around the origin where true sources



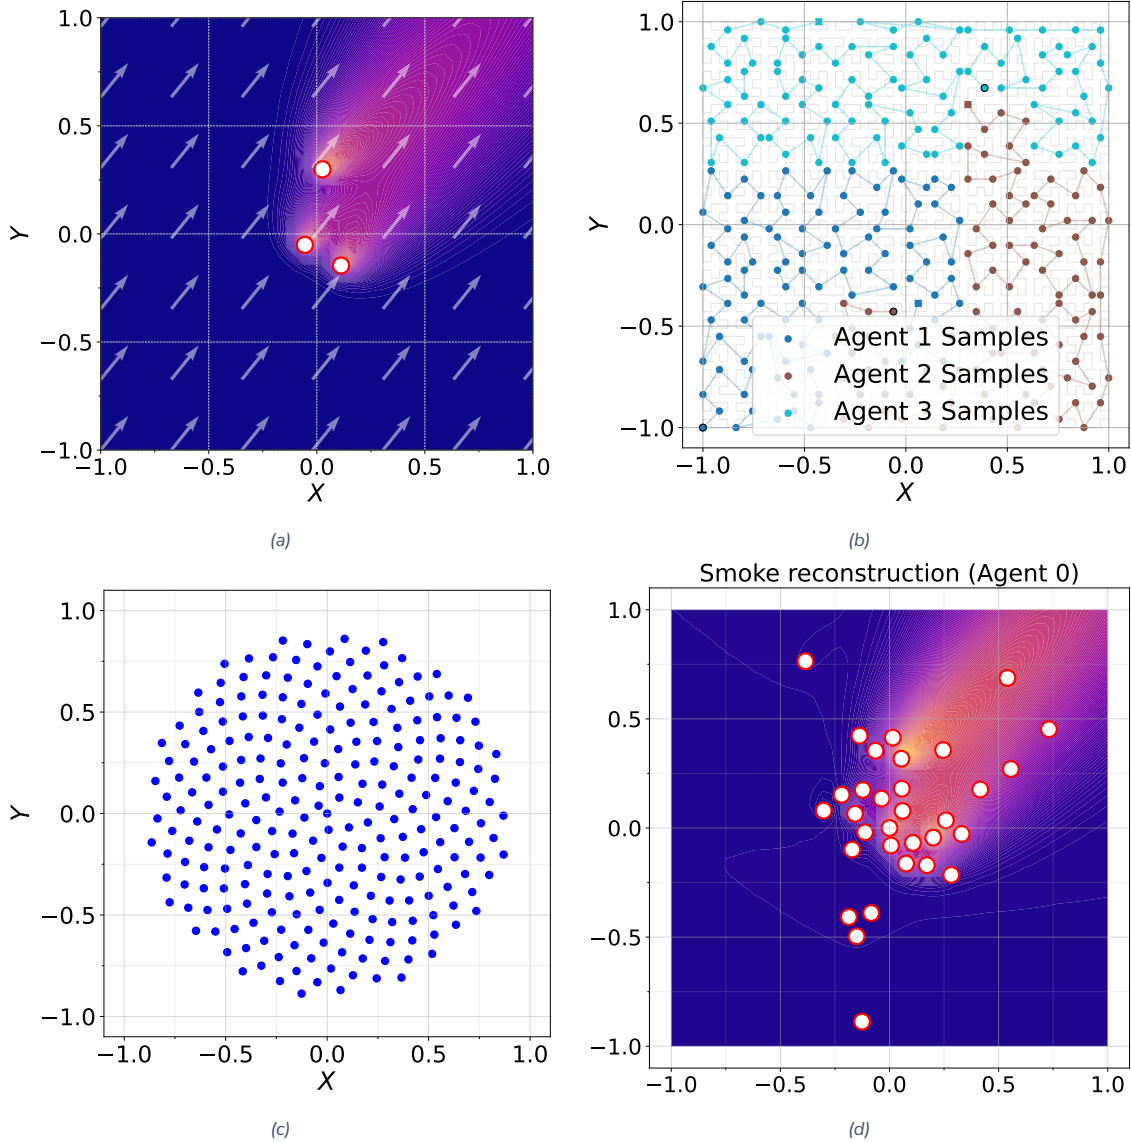




Figure 7. (a) f_{true} concentration for $L_{\text{true}} = 3$ sources. Dots mark smoke sources. (b) Agent sampling trajectories. (c) Initial source locations ϑ_i . (d) State of the bottom-up initialization for the agent $k = 0$. Dots denote source positions.

are expected. The resulting initial configuration is shown in Fig. 7. This is the starting point for the bottom-up initialization. Applying Alg. 1 we find $\hat{L} = 26$ initial sources; those are shown in Fig. 7d. along with the corresponding estimated concentration. We then begin with the main algorithm iterations, which are interrupted when both MSE and the relative change of EMD between two consecutive iterations drop below 10^{-2} . The final estimation results are shown in Fig. 8. As we can see, in the high SNR regime with a sufficient number of measurements M_k correct locations and number of sources can be detected. We should emphasize, however, that the obtained detection results are achieved with a modification of the SBL pruning condition (68), akin to the approach adopted in [64]. The standard SBL pruning criterion (68) alone proves insufficient to



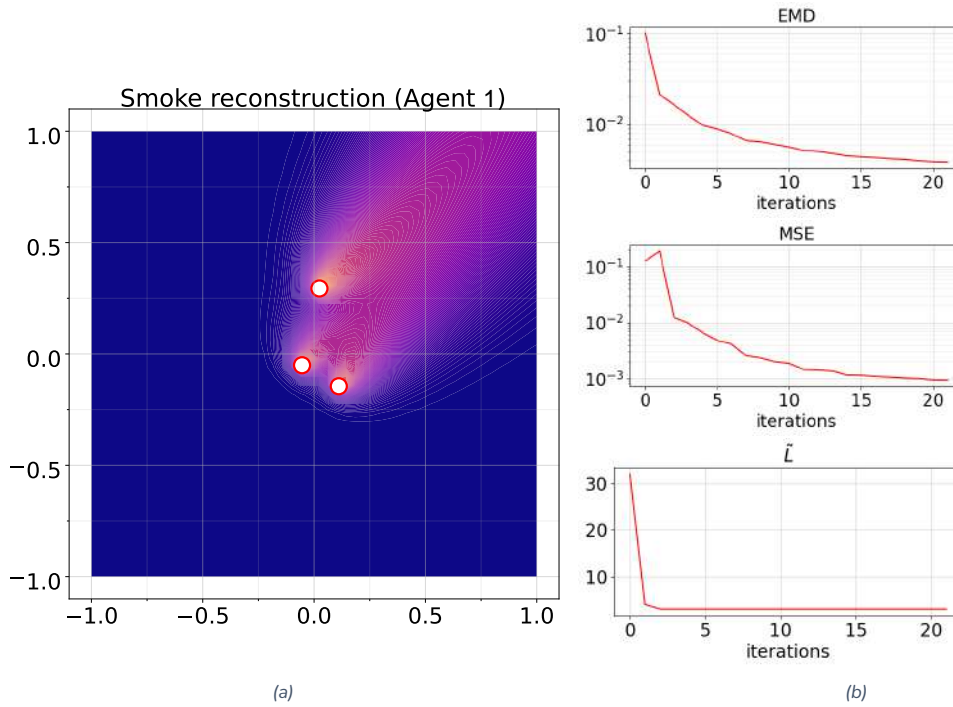




Figure 8. Final estimation results at agent $k=1$. (a) Estimated smoke concentration and source locations; (b) estimated EMD, MSE, and the number of sources \hat{L} versus iterations.

eliminate estimation artifacts. These arise due to the optimization problem (24) interlaced with the support estimation (21). This seems to be inherent to SBL applied to nonlinear models, and has been previously observed, e.g., in [65, 66]. Due to space constraints, we, however, omit the corresponding detailed discussion of this effect here.

Let us now look into the dependency of the method on SNR and the number of measurements M_k agents take. To this end, we vary SNR between 0dB and 30dB and compute EMD, MSE as well as the final number of estimated components \hat{L} over multiple Monte Carlo runs. We use the setting as before, but now reduce the discretization to $N=900$. The corresponding results, averaged over 100 independent runs, are shown in Fig. 9. In Fig. 9a we plot EMD and MSE performance versus the number of measurements per agent M_k for different SNR, with a thick line indicating the median and shaded regions showing 90% confidence. As we see, at low SNR SBL rejects all the sources with the adjusted condition (68). In other words, the algorithm prefers to keep reliably detected sources. Although the artifacts can be added to the model, their release rates are very small, making their contribution insignificant. As SNR increases, the support can be estimated more accurately. At the same time, low SNR cannot be adequately compensated with more measurement samples M_k . Let us also point out that the number of estimated sources \hat{L} , shown in Fig. 9b is indeed random: artifact sources or missed detections are observed. However, as the number of samples grows, we see the mode of the distribution peaks at the true number of sources $L_{\text{true}}=3$. This also happens more rapidly with growing SNR, as more sources can be detected.



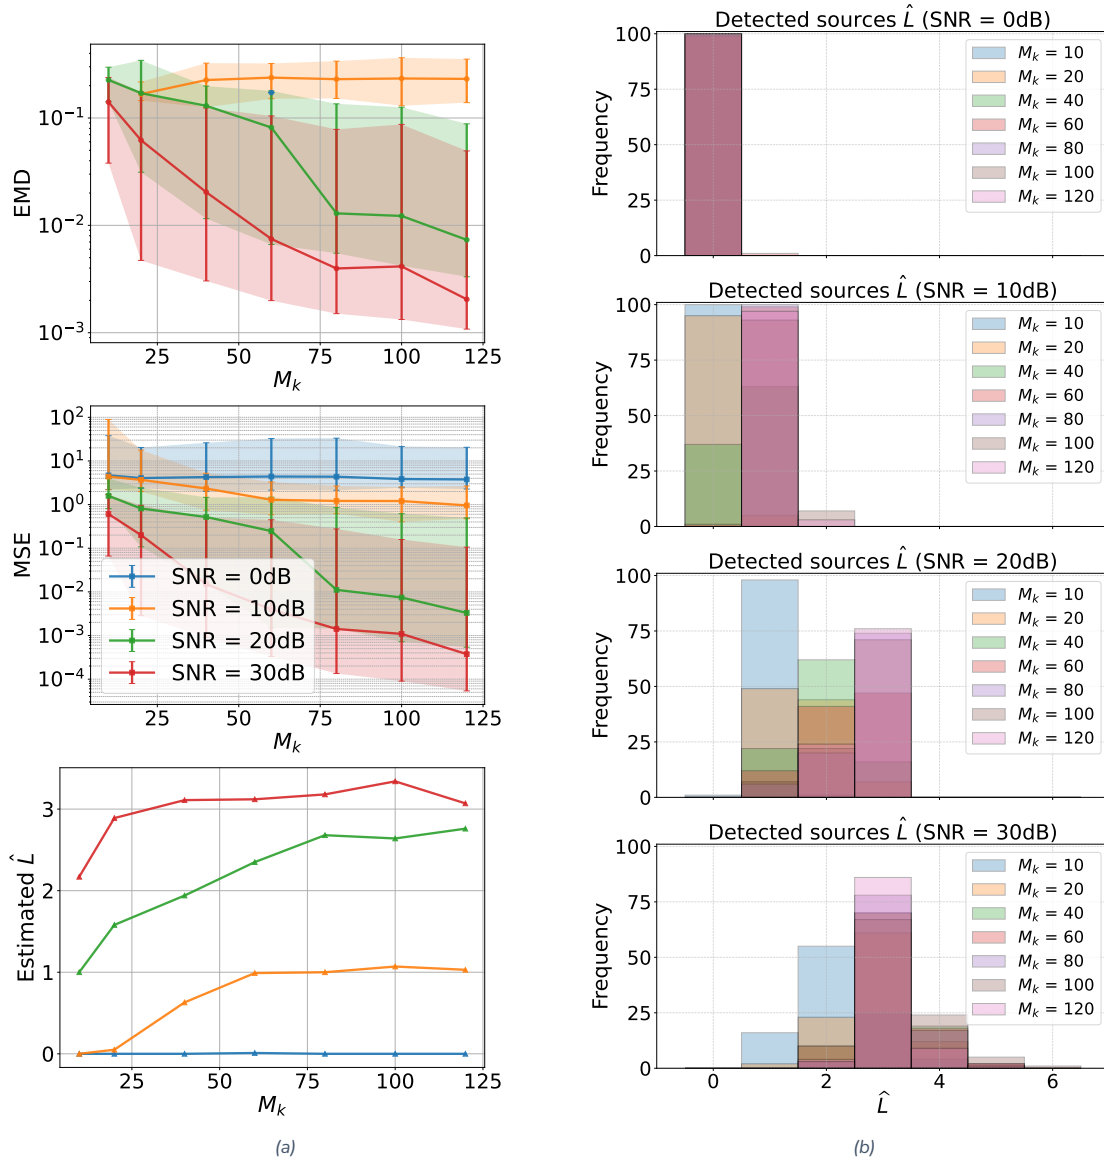





Figure 9. (a) median and 90% confidence for EMD and MSE performance, as well averaged number of detected sources \hat{L} versus M_k for different SNR. (b) Histograms of \hat{L} for different numbers of measurements and SNR values.

Adaptive source localization

Below, we discuss an extension of the distributed estimation scheme to an online, adaptive setting, allowing asynchronous streams of data to be collected, as agents explore the environment. This is particularly relevant in practice, when sensor data comes from multiple agents in streams, where not all digital samples are useful concentration values: some concentrations are too low to be reliably detected. The proposed approach thus permits more flexibility by incorporating only informative measurements when they become available, bypassing the need to store all measurements, reducing memory requirements, and effectively “decoupling” sensing and inference.





To make the algorithm adaptive, we note that both $\bar{\Phi}[t] \in \mathbb{R}^{N \times N}$ and $\bar{\mathbf{y}}[t] \in \mathbb{R}^N$ can be computed not only cooperatively (over the network), but also adaptively (over time), as measurement samples become available. Indeed, it is easy to show that both $\bar{\Phi}[t]$ and $\bar{\mathbf{y}}[t]$ can be computed as

$$\bar{\Phi}[t] = \bar{\Phi}[t-1] + \lambda_\xi \sum_{k \in K} \phi_k[t] \phi_k[t]^\top, \quad (73)$$

$$\bar{\mathbf{y}}[t] = \bar{\mathbf{y}}[t-1] + \lambda_\xi \sum_{k \in K} \phi_k[t] z_k[t]. \quad (74)$$

In other words, once a new measurement at time step t becomes available to any agent, a network only needs to cooperatively compute two correction terms

$$\Delta_\phi[t] = \lambda_\xi \sum_{k \in K} \phi_k[t] \phi_k[t]^\top, \text{ and } \delta_z[t] = \lambda_\xi \sum_{k \in K} \phi_k[t] z_k[t],$$

using e.g., an averaged consensus. After this each agent can “locally” update $\bar{\Phi}[t] = \bar{\Phi}[t-1] + \Delta_\phi[t]$ and $\bar{\mathbf{y}}[t] = \bar{\mathbf{y}}[t-1] + \delta_z[t]$ from (73) and (74). Given the latter, $\mathbf{D}[t]$ and $\mathbf{d}[t]$ are evaluated, followed by the Distributed Super-resolution Gas Source Localization (DSR-GSL) algorithm [55].


Let us stress that the update iterations of the DSR-GSL estimation can be run in parallel to data acquisition, essentially as an independent process. They are based exclusively on current values of $\mathbf{D}[t]$ and $\mathbf{d}[t]$, which can be computed asynchronously via averaged consensus as concentration samples are being collected.

In the original DSR-GSL formulation [54, 55], the algorithm assumes that all data is available to the network prior to the estimation. This implicitly assumes a fixed number of “active” sources with non-zero release rate. Consequently, the algorithm is initialized by selecting some (large) initial location set, which is “sparsified” as the algorithm proceeds – an approach we will term a top-down strategy (see [54, 55] for more details). Yet when the data arrives sequentially, new sources can be detected. To account for this, we need an alternative strategy that can include new components as new data arrives. To this end, we propose a cooperative *bottom-up* approach, largely inspired by orthogonal matching pursuit [59] – a greedy signal recovery method. The new strategy incrementally attempts to add new potential sources to optimally match novel observations, while removing irrelevant sources. As we will show, the resulting strategy also depends *only* on the variables $\mathbf{D}[t]$ and $\mathbf{d}[t]$, thus allowing for an adaptive, decentralized processing approach.

Assume that $\bar{\Phi}[t]$ and $\bar{\mathbf{y}}[t]$ are available to each agent; also let $\mathbf{g}(\boldsymbol{\theta}) \triangleq [G(x_1, \boldsymbol{\theta}), \dots, G(x_N, \boldsymbol{\theta})]^\top$ denote a Green’s function response for the source located at $\boldsymbol{\theta}$. Define now an agent’s residuum signal $\mathbf{r}_k[t] \triangleq z_k[t] - \Phi_k[\hat{\Theta}] \hat{\mathbf{w}}$, $k \in K$, as an error at time t between collected measurements and reconstructed response based on current estimates of $\hat{\Theta}$ and $\hat{\mathbf{w}}$. Similarly to (67) we define

$$\bar{\mathbf{r}}[t] \triangleq \sum_{k \in K} \lambda_\xi \Phi_k[t]^\top \mathbf{r}_k[t] = \bar{\mathbf{y}}[t] - \bar{\Phi}[t] \mathbf{G}(\hat{\Theta}) \hat{\mathbf{w}}. \quad (75)$$

Note that all quantities on RHS of (75) are known to each agent after the consensus; thus $\bar{\mathbf{r}}[t]$ can be computed locally. To add/initialize a new (potential) source at location $\boldsymbol{\theta}^*$ an agent can use the residuum $\bar{\mathbf{r}}[t]$ instead of actual measurements that are distributed over the network. Indeed, we recognize that the value $\|\mathbf{d}[t]\|^2$ measures how well $\mathbf{G}(\hat{\Theta})$ aligns, or rather correlates, with $\bar{\mathbf{y}}[t]$ – measurements projected onto N -dimensional model space. Thus, $|\mathbf{g}(\boldsymbol{\theta}^*)^\top \bar{\mathbf{r}}[t]|^2$ will measure how well a potential source at location $\boldsymbol{\theta}^*$ aligns with the current residuum. This alignment is “measured” with SBL based on the estimate of the corresponding sparsity parameter α^* : if $\alpha^* = \infty$,



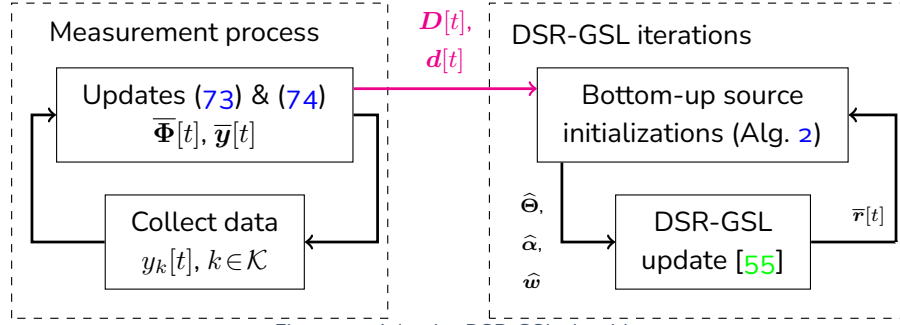




Figure 10. Adaptive DSR-GSL algorithm.

the source is dropped; otherwise, a potential source is included in the model and is updated with other detected sources following the DSR-GSL algorithm. The key steps of this procedure are summarized in Algorithm 1. Naturally, for $t = 0$ we let $\bar{\mathbf{r}}[0] = \bar{\mathbf{y}}[0]$ and the number of estimated

Algorithm 2 One-source Bottom-Up initialization, agent $k \in \mathcal{K}$

- 1: Set: $\vartheta_l, l = 1, \dots, \tilde{L}$; Residual $\bar{\mathbf{r}}[t]$ from (75)
 - 2: Solve: $\theta^* = \underset{\vartheta}{\operatorname{argmax}} |g(\vartheta)^\top \bar{\mathbf{r}}[t]|^2$ # Best fitting source
 - 3: $D_{\text{new}} \leftarrow g(\theta^*)^\top \bar{\Phi}[t] g(\theta^*)$, $d_{\text{new}} \leftarrow g(\theta^*)^\top \bar{\mathbf{r}}[t]$
 - 4: **if** $\frac{d_{\text{new}}^2}{D_{\text{new}}} > 1$ **then**
 - 5: # Add a detected source at location θ^*
 - 6: $\hat{\Theta} \leftarrow \hat{\Theta} \cup \theta^*$
 - 7: $\alpha^* \leftarrow \left[\left(\frac{d_{\text{new}}}{D_{\text{new}}} \right)^2 - \frac{1}{D_{\text{new}}} \right]^{-1}$
 - 8: $\hat{\mathbf{w}} \leftarrow \left(D[t] + \hat{\Gamma}^{-1} \right)^{-1} d[t]$
 - 9: $\hat{L} \leftarrow \hat{L} + 1$
 - 10: **else**
 - 11: # Source is not detected, continue.
 - 12: **end if**
-

sources to $\hat{L} = 0$. As new measurements become available, in line 2 we select the source location that best fits the residual data – the observed remaining concentration, with other sources “canceled”. The quantities in line 3 are scalar counterparts of $D[t]$ and $d[t]$ computed for a single Green’s function centered at the best-matching location θ^* . Line 4 mirrors that in DSR-GSL algorithm for support estimation (see also [54, 55]). It is this line that determines if a new source is detected: if the condition fails, SBL refrains from adding a source at θ^* to the model, and the initialization stops. Otherwise, the newly detected source is included in the model. Subsequently, each agent evaluates the new response $\mathbf{f} = \mathbf{G}(\hat{\Theta})\hat{\mathbf{w}}$ and updates its local residual $\mathbf{r}_k[t]$. At this stage, new sources can be added to the model. We summarize the key steps of the new adaptive DSR-GSL (aDSR-GSL) algorithm in Fig. 10. As we see, the measurement and parameter estimation are two separate steps that can be run in parallel, with the latter requiring only current values of $D[t]$ and $d[t]$ updated adaptively.



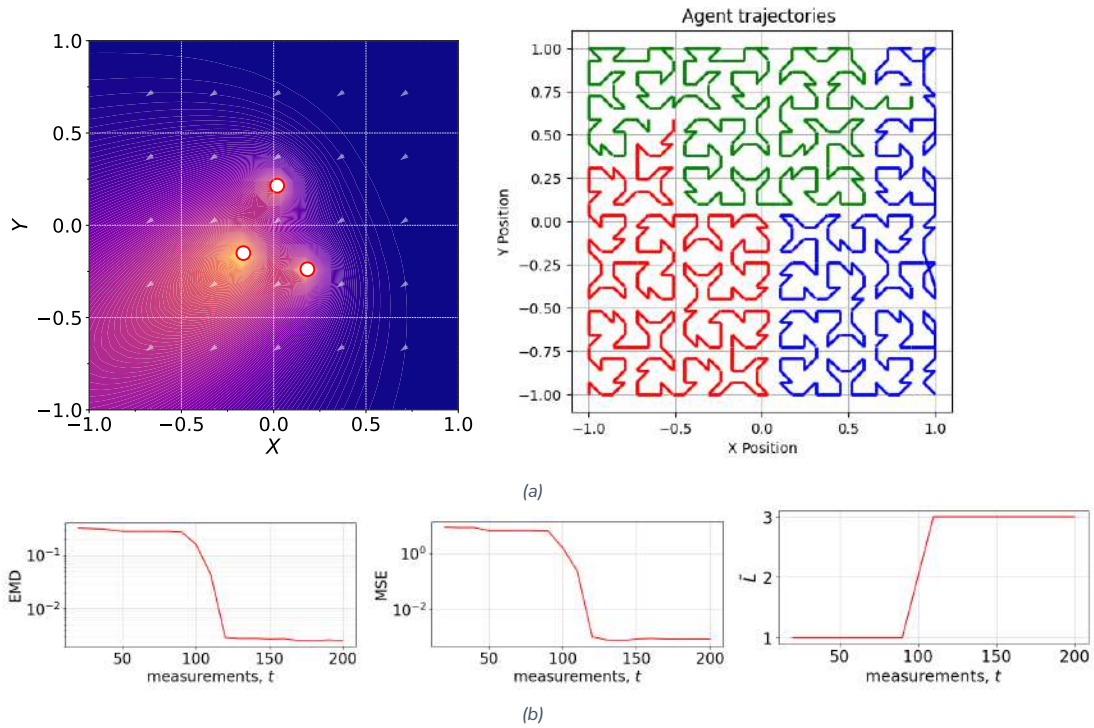





Figure 11. (a, left) f_{true} concentration for $L_{\text{true}} = 3$ sources marked by dots, (a, right) agent sampling trajectories; (b) Estimated EMD, MSE, and the number of sources \hat{L} versus number of measurements per agent.

Simulation results for the adaptive inference algorithm

Let us now illustrate the algorithm's performance with a few simulations. Similarly to the previous case we set $\Omega = [-1, 1] \times [-1, 1]$, $\kappa = 0.03$. Performance is measured with MSE criterion $\frac{1}{N} \|\hat{f} - f_{\text{true}}\|^2$ between an estimate and true concentration, as well as Earth Mover's Distance [60]. We begin by showing a single run of the algorithm. To this end we partition Ω into $N = 900$ discrete cells and set $L_{\text{true}} = 3$. The true source locations are randomly distributed around the origin, with release rates set to 1. The components of the wind v are uniformly drawn from the interval $[-1, 1]$. For this run we also set SNR, defined as $10 \log_{10} \left(\frac{\lambda \xi}{N} \|f_{\text{true}}\|^2 \right)$, to 25dB. The corresponding ground truth concentration f_{true} is shown in Fig. 11a. For the estimation, we use a network of $K = 3$ fully connected agents, each sequentially collecting samples uniformly spaced along a Hilbert curve trajectory [62] within Ω (see also Fig. 11a). We kick-start the algorithm by first collecting 20 initial samples per agent, followed by Distributed Super-resolution Gas Source Localization iterations; then 10 samples are collected before a new Distributed Super-resolution Gas Source Localization update cycle begins.

For initialization, we generate $\tilde{L} = 900$ initial source locations ϑ_l , $l = 1, \dots, \tilde{L}$, using Vogel's algorithm [63] to cover the area around the origin, where true sources are expected. This set is used in line 1 of Alg. 2. The Alg. 2 is repeated a maximum of 5 times, i.e., per measurement cycle, up to 5 new sources can be added. The corresponding estimation results after $t = 200$ steps are shown in Fig. 11b. As it can be seen, after about 100 measurements (per agent), both EMD and MSE drop, indicating correct support estimation; also \hat{L} converges to exactly 3 sources. It should be





stressed, however, that these detection results are obtained using adjustment of the SBL pruning condition as used in [64]. Standard SBL pruning criterion is not sufficient to remove estimation artifacts, in particular due to source location estimation, which is characteristic to SBL for nonlinear models – an effect that was reported in e.g., [65, 66]. Due to space constraints, we, however, omit the corresponding detailed discussion of this effect here.

Next, we vary SNR between 0dB and 30dB and compute EMD, MSE, as well as \hat{L} versus t over different source locations and wind directions. To this end, we use the same setting as before, but average the results over 100 independent runs. The corresponding results are shown in Fig. 12.






Figure 12. (top, middle) Solid lines shows median EMD and MSE performance vs. number of measurements per agent for aDSR-GSL; (bottom) averaged number of detected sources \hat{L} vs. number of measurements. Shaded regions show 90% confidence interval, color indicates the same SNR for all plots. Dashed lines show the non-adaptive DSR-GSL performance.

We will also compare the results to the non-adaptive DSR-GSL algorithm [55] (dashed lines in Fig12), i.e., when the whole measurement batch is used for processing with top-down initialization. As we see, aDSR-GSL generally detects fewer sources as its non-adaptive counterpart that utilizes the whole data set at once; for SNR=0dB aDSR-GSL makes 0 detections (and thus makes EMD computation impossible). With growing SNR, however, aDSR-GSL approaches the performance of DSR-GSL, with a significant improvement after about 125 samples per agent. This is about 30% of the whole exploration area and corresponds roughly to the proportion of Ω with increased concentration, i.e., with more informative samples. This highlights the need for agents to follow more informative strategies, rather than focusing on coverage, especially for





advection-dominated processes.

3.4. Experimental campaign in Montiferru

3.4.1. Montiferru Data Collection

A great amount of effort in the stated period was focused on developing the framework to collect enough data to both validate and feed the model. For instance, the lack of either trace gas or locally collected wind data makes the historical trials unsatisfactory for testing any models, hence why PDM-tech-03 was designated as a research component. For this reason, we present a section that relates to the work to collect and assess real data on TEMA's first trial in Sardinia, taken in June 2025.


3.4.1.1. 18.o6 Wood Pile Burn Setup and Wind Conditions

On 18.o6, the prescribed burns consisted of burns of wood piles. In agreement with other scientific partners in TEMA, the prescribed burn would follow the behavior of a hypothetical fire. For this reason, piles of fuel would be ignited with the help of starter fuel in a sequence that corresponded to the observed wind direction. DLR-KN placed static sensor boxes, collectively known in TEMA as the Ground Sensor Network, to collect fixed data on the scene. Due to the limited amount of working static sensor boxes on hand, the locations of these boxes were determined on site with knowledge of the wind direction. Fig. 13 shows the placement of the sensors and the wood piles burned on 18.o6. Additionally, see Fig. 14 for images of the wood pile burns.



Figure 13. Prescribed Burn Setup on 18.o6.

This plot shows the physical setup of the prescribed burns on the first day of the Sardinia Trial. The RGB image in the background is a photogrammetric reconstruction of the domain using a DJI drone on site, before the burns occurred. The pale orange dots represent the locations of the static Ground Sensor network DLR-KN implemented. The light pink areas represent the locations of burning wood piles. The background is an OpenStreetMap tile for reference.






Figure 14. Wood Piles Burned During Sardinia Trial.

The left image shows an unburned wood pile, roughly 10 × 3m in size. The right shows a burning wood pile, collected from a DJI-M31 drone, and shows the crew of RAS involved in ensuring it burns safely. Note that though the fire is burning visibly in the image, there is little visible smoke.





Figure 15. Hardware for Ground Truth Collection.

The left represents the Ground Sensor units with trace gas sensing capabilities and anemometer. The right unit represents a Holybro X500 drone with a PCB that has trace gas sensors.

We now discuss the information DLR-KN collected from the burns themselves during this trial. First, we discuss collected wind information from the ground sensors, seen in Fig. 15. Wind data were filtered with a triangular filter of window 60 seconds to lessen the effects of turbulence on the more desired prevailing winds. The results are shown in wind rose plots in Fig. 16.⁸ We can see that the wind data collected on the 18.06 shows a relatively uniform distribution with respect to wind direction, primarily coming from the north. Some differences in the distribution exist, possibly due to slight errors in readings on sensor orientation. Additionally, the northernmost sensor experienced slightly higher wind speeds than the wind sensor placed on the south. This could be explained by differences in topography and vegetation around the sensors.

⁸One Ground Sensor unfortunately experienced a fault in wiring during the experiments and was not able to record wind data.








Figure 16. Observed Winds During 18.06 Prescribed Burn Period.
 The locations of the wind sensors are represented by red dots and the wind rose graphs fanning the points show both (1) the direction from which the wind came and (2) relative frequencies of these winds and their magnitudes. The wind direction is consistent between the sensors (with slight deviation due to errors in placement).

DLR-KN also flew drones in the region during a portion of the prescribed burn. The drone was flown several times through the smoke plume of a burning wood pile. The data themselves are discussed in a later section, when they are compared against the smoke model PDM-03.

3.4.2. 19.06 Simulated Burn Setup and Ground Truth Collection

On 19.06.25, RAS performed a regional prescribed burn. RAS ignited wood piles as before, in a relative sequence corresponding to wind direction. However, while the wood piles burned, RAS ignited other local shrubs and grasses in the region, again with fuels, to simulate a broader area burned. This traded precision in knowing the exact burn location for more realistic spreads of the burned fuel, since a real forest fire does not consist of concentrated fuel piles but a more dispersed burning area. The burnt region consisted primarily of local shrubs and grasses and had few trees, and so would correspond to an ignited grassland on the periphery of a forest fire (with significant additional trace gases present due to the igniter).







Figure 17. Prescribed Burn Setup on 19.06 (Wood Piles and Grassland Burn).

This plot shows the physical setup of the prescribed burns on the second day of the Sardinia Trial. The RGB image in the background is a photogrammetric reconstruction of the domain using a DJI drone on site, after the burns occurred. The pale orange dots represent the locations of the static Ground Sensor network DLR-KN implemented. The dark red dot represents an auxiliary anemometer, which collected additional wind data to compensate for the lack of data from one sensor. The orange areas represent the locations of burning wood piles. Additionally, the blue areas represent an estimate of the burnt regions. The background is an OpenStreetMap tile for reference.

As we can see in Fig. 17, the burnt region spans a significantly larger area than the one spanned by the burning wood piles. The burnt area was estimated by comparing domain images from before and after the burn and modeled in QGIS. One of the sensors was placed within the burnt area, but due to the controlled nature of the burn, it was safe from destruction.

We now discuss the data collected during the trial. Figure 18 describes the wind data collected on the 19.06. As before, we can see that the distribution of the winds matches between the sensors; however, there is more discrepancy between the observed winds. The extra auxiliary anemometer (northernmost sensor) has a more notably different wind direction distribution compared to the other sensors. This discrepancy is explained by the fact that the auxiliary anemometer was not able to record wind data for the last hour of observation, during which there was a shift in direction. Additionally, we see differences in wind speeds measured in the northeast compared to the southwest. This is best explained by topographic effects. The hill on the west side sloped exactly in the direction of the prevailing wind.







Figure 18. Observed Winds During 19.06 Prescribed Burn Period. The locations of the wind sensors are represented by red dots, and the wind rose graphs fanning from the points show both (1) the direction from which the wind came and (2) relative frequencies of these winds and their magnitudes. The third wind station present in the northernmost location is an auxiliary ultrasonic anemometer and was not connected to DLR-KN's GroundSensor Network.

During the prescribed burn, DLR-KN flew the sensing drone over large portions of the domain in question. Gridded trajectories were chosen to map the smoke plume over a larger region of the domain. The data and comparison to the model are present in Section 3.4.3.

3.4.3. Smoke Model Evaluation


3.4.3.1. Model Definition


The smoke model used is a time-varying model

$$M(\mathbf{x}) = \sum_i k_i G_{\mathbf{v}(t), \kappa}(\mathbf{x}, \theta), \tag{76}$$

where $G_{\mathbf{v}, \kappa}$ represents the 2D Green's function for the advection diffusion equation in (9) with constant wind vector \mathbf{v} and diffusion coefficient κ . Here, k_i are the source weights, to be explained below. The diffusion coefficient of $5\text{m}^2\text{s}^{-1}$ we used was chosen to correspond to diffusion in turbulent conditions with high mixing.

For the wind coefficient $\mathbf{v}(t)$ we computed time-varying prevailing winds from onsite measurements. First, we convolved a triangular window function of window size 5 minutes with the





wind's x and y components. This window computes a weighted average of the wind experienced in a centered 5-minute window around each time sample, which removes the turbulent component. Then, to compute the domain's prevailing wind, we simply average the three wind measurements to get an estimate of the prevailing wind over the domain. As we saw in Figs. 16, 18, the measured wind distributions are reasonably consistent with a spatially constant wind.

The model weights k_i in (76) were chosen to be 1 if the source was a wood pile, and $0.00125A_i$, where A_i is the area of the burnt region in square meters if the region was burnt grass, as in the regional prescribed burn on 19.06. This corresponds to a 50-fold difference in the contribution from a wood pile and burning grass if areas of both are kept equal. This number was chosen to represent that a wood pile burns with much more energy than the same patch of grass.

3.4.3.2. Model Evaluation Methods


To assess the model's performance, we compare the model with data collected by the X500 with trace gas sensors. The smoke model of (76) is compared to the voltage of the Semeatech 4-series electrochemical sensor measuring the plume gas carbon monoxide (CO), which showed the highest responsiveness to the gas plume.

The drone data were preprocessed to remove noise. This was achieved by convolving trace gas data with a 15-second triangular window. Some passes through the plume were done at much higher altitudes away from the plume and are excluded from the data. Finally, data were normalized to match peaks in the signal. In the following sections, we compare these trace gas signals from 18.06 and 19.06 with the model, in both the spatial and the time domain.

In the following section, in addition to noting the MSE of the model over the drone flight, we also compare the effectiveness of the model at capturing the presence of a plume. We track peaks in the trace gas signal of 20mV of width 9 to 30 seconds as a detection of a plume, and compare the model to it.

3.4.3.3. Model Evaluation 18.06 - Wood Pile Burn

The drone was flown on 18.06 during the second of the wood pile burns, when the first pile was not burning. Figure 19 shows an example of the estimated distribution of the smoke model at this time, which displays the location of the burning wood pile.






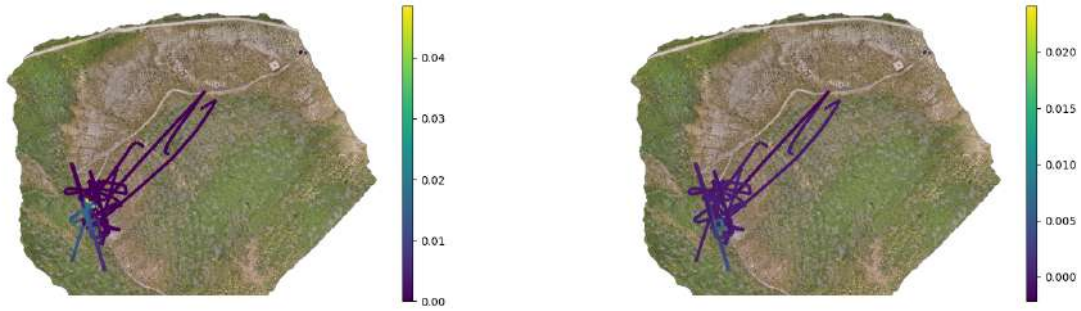



Figure 19. 18.06 Smoke Model Evaluation.

The Smoke Model evaluated at a sample time of drone data collection, with the red dot corresponding to the location of the active wood pile. The wind value corresponds to the prevailing wind calculated over the time of collection. Values lower than 0.001 are not displayed on the plot for clarity to show the problem domain. Note that the northernmost woodpile is not included in this model, due to the fact that it had already burned by the time the drone collected data.

Figure 20 shows the spatial distributions of the drone-collected data and compares them to the model predictions over the path. We can see that the model and drone predictions agree that there is a high concentration of smoke around the location of the drone. We can see that the smoke model predicts higher concentration values close to the wood pile at slightly more northern locations. We believe this is due to the altitude of the drone sampling path. At 20m above the burning wood pile, it is possible that the effective source location is shifted slightly windwards due to the fact that the smoke has been carried by the wind for some time. This is currently not reflected in the model, which is 2D.





(a) Smoke Model

(b) Drone Trace Gas Measurement


Figure 20. 18.o6 Model v. Drone Trace Gas, Spatial Representation

The left image is the model prediction. The right shows trace gas data collected by the human-flown drone. Both models show the (expected, actual) conditions overlaid on the physical location of the drone at the timestamp of the drone path.

We now observe the time series results, present in Fig. 21. We can see that the drone data matches the prediction for all time periods of flight, save one. This is reflected by the fact that all 5 peaks in the drone signal, marked in green, are reflected by the presence peaks in the smoke model. We can see that the heights of the peaks do not match the heights of the peaks of the smoke model at some times.

Additionally, we see that between 9:15 and 9:20 UTC, there is an interval where the model indicates the presence of a plume, and the trace gas sensor registered no noticeable change. This corresponds in location to the drone flight, very close to the source, slightly southwest of the source.

The MSE of the model is 0.0113, representing an 89% MSE.








Figure 21. 18.06 Drone and Time Series Evaluation.

The simulated smoke field compared to the observed measurements from the drone. Green vertical lines correspond to a plume detection by the drone. Gaps in the data reflect times the drone was not in the air or exceeded a certain altitude, which affects the voltages of the trace gas sensor.

3.4.3.4. Model Evaluation 19.06 - Prescribed Burn

We now discuss the model evaluation during the prescribed burn. Due to the large and diffuse nature of the prescribed burn, the drone path covered a much larger area and focused on two key sectors. In Fig. 22 we see an estimate of what the smoke distribution over the domain would have looked like - more diffuse over the whole region, and with many emitters. Unfortunately, the exact time and locations of these grass burns were not able to be recorded with the equipment TEMA had on site, so the model does not reflect the changes in source intensities.





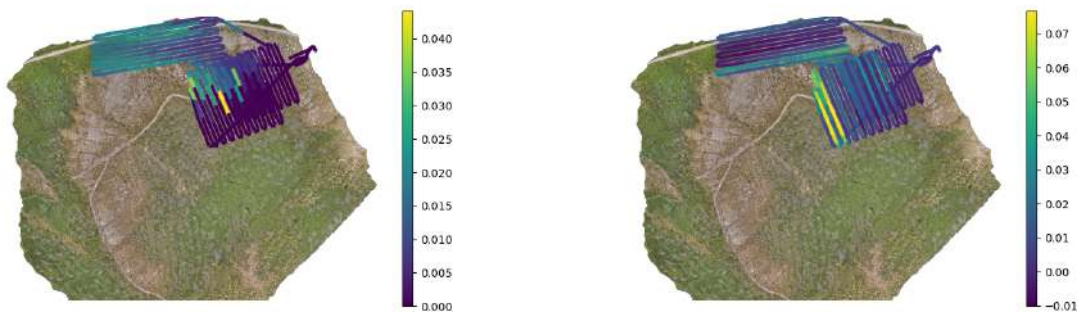


Figure 22. 19.06 Smoke Model Evaluation.

The Smoke Model evaluated at a sample time of drone data collection, with the red dots corresponding to the location of the active wood piles and orange dots the grass burnt. The wind value corresponds to the (average) prevailing wind calculated over the whole time of collection. Values lower than 0.001 are not displayed on the plot for clarity to show the problem domain.

Fig. 23 displays the gridded drone path and shows model and measured distributions of the plume. We can see that the model and measurements agree that the smoke intensity is strongest closest to the wood piles. However, we see that the peak is observed farther south, not windward, in the drone measurement. We believe that the sensor response rate caused it to exhibit high voltages due to previously observing a high value, and these phantom peaks are not physical.





(a) Smoke Model

(b) Drone Trace Gas Measurement

Figure 23. 19.06 Model v. Drone Trace Gas, Spatial Representation

The left image is the model prediction. The right shows trace gas data collected by the human-flown drone. Both models show the (expected, actual) conditions overlaid on the physical location of the drone at the timestamp of the drone path.





We now discuss the timeseries data in Fig. 24. We can see that peaks in the drone data are again accompanied by peaks in the smoke model. In fact, of the 27 peaks in the drone signal, only 4 are unexplained by a high model value. Of note here is a region between 9:21-9:26 UTC. This corresponds to the line connecting the northeast-most corners of the two grids in the drone path.

In comparison to the data on 18.06, we can also see at times that the voltage from the sensor may drift to be negative. We believe this is due to some cross-sensitivities of the sensor with water vapor/humidity in the plume.

The MSE of the model was .109, corresponding to a 76% MSE over the time interval measured.





Figure 24. 19.06 Drone and Time Series Evaluation. The simulated smoke field compared to the observed measurements from the drone. Green vertical lines correspond to a plume detection (peak) by the drone. Gaps in the data reflect times the drone was not in the air or exceeded a certain altitude, which affects the voltages of the trace gas sensor.


3.4.3.5. Discussion

DLR-KN's smoke model maps the smoke distribution in both prescribed burn events. Some discrepancies exist between the model and measurements, but the model describes the behavior of the smoke to a reasonable degree.

Overall, we can see that PDM-03 did not match signal heights in some instances on both 18.06 and 19.06. This, in combination with the prediction of a signal where there was none, hurt its performance with respect to MSE. However, the model was able to explain 27 of the 31 total peaks in the drone signal.

These discrepancies are partly due to the fact that the electrochemical sensor used to sample gas distribution, while one of the best sensing technologies available at its form factor, still may stay low- or high-valued after observation of a plume, due to the sensor's response time. Combined





with fast drone motion, we cannot be certain that the data were indeed a ground truth. However, they would indicate the presence of trace gases in cases where the smoke still does not register visually.

3.5. Fulfillment of TEMA KPIs

The Dispersion Model Update rate KPI (OB₁) was to have the Model/map update of the plume and smoke concentration distribution every 10 s. PDM-o₃ in its current state fulfills this. Current model evaluations with the analytical solution run in less than 1 second. However, DLR-KN also undertook research work with finite element simulations using the FEniCS library tuned to the Montiferru domain, see 25 for an example of the simulation mesh with 32,000 points and 64,000 cells. Regardless, these models also output results in one second for a given time, meaning the KPI is still fulfilled if the model changes. The finite element model of smoke concentration remains future research work.






Figure 25. Finite Element Mesh for Domain Representation of the Montiferru region.


The Smoke Plume and Concentration Accuracy KPI (OB₂) was to have a Mean Squared Error of the estimated spatial phenomenon under controlled conditions of less than 15%. This KPI was demonstrated in simulations studies under controlled conditions (see [54, 55, 67]). However, in uncontrolled studies in Sardinia trial this KPI was not demonstrated. Due to the strong dependency on the wind direction, as well as model mismatch between the actual dispersion and theoretical models we were able to achieve MSE errors only between 75% and 89% described in section 3.4.3.5 against drone-collected trace gas data. This clearly shows the need for non-stationary model approaches, real-time wind incorporation, as well as appropriate sensor data pre-processing. However, when the metric is instead adjusted to a detection problem, we can see that the smoke model reflects the collected data distribution in a reasonable sense – it explains 27 of the 31 peaks in the drone-acquired data, thus performing well in full 87% of cases.


Finally, we discuss the automation and temporal resolution KPIs (OC₂, OC₃): No interaction required by human operators for mapping plume and concentration distribution, with near-real-time updates of the model. This KPI is fulfilled. Moreover, in [67] we developed a theoretical formulation for information-driven path planing of drones to improve data collection – a foundation for infotaxis and autonomous drone navigation. All inputs, wood pile locations, and wind are





present in the TEMA platform and require no extra input from an operator to be displayed. PDM-03, being still in development as a research component, was not included during the Sardinia Trial in real time. However, this capability will be demonstrated during the next in-person trial. Further research work remains to update the model in real-time with additional estimates of what is on fire and with what emission rates, for more accurate models under real-time burn conditions.





4. Wildfire behavior modeling

4.1. State of the art and TEMA's contribution

Currently, most advanced commercial wildfire simulators are built on well-established physical models, integrating detailed information such as fuel models, topography (DEM), fuel moistures, forecast weather, etc. In any case, many of these inputs are difficult to obtain or are not as accurate as this type of simulation requires.

Wildfire Analyst FireSim produces fire-spread and impact projections within seconds and performs large-scale parallel simulations to deliver robust probabilistic analyses that are quite useful for decision-making processes in natural disaster management. Despite being one of the most powerful fire simulators at a global scale, it still lacks integrations that allow for real-time data assimilation. While it is true that efforts are being made in the field of research to include real-time or near-real-time observations in wildfire simulators [68, 69], commercial and fully operational tools already in use in natural disaster management do not yet have this feature or do not have an integrated data collection component process.

During the current reporting period, the technology developed in the TEMA project has enabled FireSim to ingest real-time weather data collected in the vicinity of the fire to enhance the weather forecast model. Moreover, it has developed capabilities to analyze and extract control points from flame front observations, enabling accurate calculation of rate of spread adjustment factors for better calibration of simulations.

These upgrades have been carried out through two different approaches:


1. Improvement of weather forecast model.


To compute the simulations of forest fire behaviour, the simulator core requires certain weather information as input. The system stores this meteorological data, which the fire simulation will use to compute the simulations. These weather variables are: temperature, relative humidity, wind speed and direction, and solar radiation, if possible. This weather forecast data is downloaded daily from the Global Forecast System (GFS) (<https://www.ncei.noaa.gov/products/weather-climate-models/global-forecast>) through FireSim's weather API.

During the Sardinian trial, sensors in the field recorded real-time weather data from the Smoke and wind modelling engine component. This data was sent to the Digital Enabler module and was converted into a JSON file as FireSim input, where the GFS forecast model was upgraded with the actual information. The weather forecast model was modified by using an Inverse Distance Weighting (IDW) interpolation method with different exponential factors in the spatial and temporal dimensions. This interpolation was done independently for each weather variable so that different sensors providing different variables could be used in different locations.

2. Fire spread calibration.

The Information Fusion component processes all available information from satellite images, drone images, and other information sources to generate fire probability maps. These maps are filtered over a certain probability threshold to obtain those pixels representing the most likely





fire presence. Once the map has been generated and sent to the Digital Enabler, a GeoJSON file with fire presence positions is generated and used to feed the new simulations.

These points generate a collection of control points with the actual fire's arrival time. The technique is based on minimizing the error between the simulated fire growth and the actual fire by identifying the optimal rate of spread adjustment variables using a least squares approach.

In Sardinia, FireSim extracted the real location of the fire fronts, ingested them as control point, and reran a new and calibrated simulation. Chapter 4.6 shows the differences between the standard and the calibrated simulations calculated with the Sardinian trial inputs.

4.2. Summary of advances in the last period

During the M1-M18 period, only the development of the methodology and conceptual basis of the technology was worked on. During this period, the work has been focused on:


- Improvement of the weather forecasting model through the integration of real-time weather data from field sensors measured by the PDM-tech-o3 component.
- Calibration of simulations based on the actual location of fire fronts in real time, generated from fire probability maps (Maps4Fire) from the Information Fusion component (PDM-tech-o5).


In order to perform fire behaviour simulations, FireSim requires a series of variables as inputs, such as:

- Fuel models: the type and structure of the vegetation, with its associated parameters such as fuel load, extinction moisture, fuel bed depth, etc.
- Canopy characteristics.
- Live and dead fuel moisture.
- Digital elevation model.
- Ignition point(s).

All these basic inputs are already embedded in the FireSim api for the considered scenarios. In addition, to simulate the fire spread, weather forecast information is mandatory to understand fire behaviour. This data is obtained and processed from the GFS dataset (<https://www.ncei.noaa.gov/products/weather-climate-models/global-forecast>), a National Centers for Environmental Prediction (NCEP) weather forecast model that generates data for dozens of atmospheric and land-soil variables, including temperatures, winds, precipitation, soil moisture, and atmospheric ozone concentration. The system couples four separate models (atmosphere, ocean model, land/soil model, and sea ice) that work together to accurately depict weather conditions. FireSim processes this information and converts it into a stack of layers in NetCDF format for its use in the simulator.

This should make it clear that, no matter how good a simulation model is, if the inputs are not accurate, the simulation outputs will not be either. At this point, fuels are layers generated based on remote sensing techniques and expert review that may not fully reflect reality, and the weather forecast is only a forecast model. Furthermore, determining the moisture content of both living and dead fuels requires complex models that are not always developed or trained for the areas to be simulated.





In the context of the TEMA project, FireSim first calculates a standard simulation, using only its scenario data, ignition points, and weather forecast model. Then, when other inputs from real-time information become available, it calculates a calibrated simulation to better adjust the simulation to the actual fire behavior conditions. To do this, and following the KPIs established for the FireSim component, R&D work is the focus of the following sections.

Furthermore, integration of the wildfire prediction modeling software within the TEMA platform continued during this current review period.

4.3. Enhancement of the weather forecast model with real-time data

To improve the weather forecast model, real-time measurements from ground sensors (PDM-tech-03) are used to modify the forecast model data. To do this, more weight is given to real data during the first hours of simulation, with increasing weight given to forecast data as we move further away from the start of the simulation.

Once this functionality was developed, FireSim is able to:

- Modify the existing NetCDF wind field based on real measurements.
- Allow any number of stations (ground sensors) at every time step (every 1 hour).
- Allow stations to change in number and/or location at every time step.
- Allow to define the weight of measurements with respect to weather forecast data (WRF).

This is carried out by the new weather model through:

- Computing the effect of an individual simulation on the wind field.
- Interpolating the individual contributions of each station using an Inverse Distance Weight algorithm (IDW).

The weight (relevance) of the new measurement with respect to the existing weather forecast data (WRF) is computed based on the distance. The closer the measurement to a grid point (i,j), the greater the effect. This weight is subjective and is controlled by two parameters, alpha and beta. The equations used in this computation are shown below:

$$d_{i,j} = \sqrt{(sts_x - x_{i,j})^2 + (sts_y - y_{i,j})^2} \quad (77)$$


$$V_{x_{i,j}} = (1 - w_{i,j}) * WRFv_{x_{i,j}} + w_{i,j} * sts_{vx} \quad (78)$$


$$V_{y_{i,j}} = (1 - w_{i,j}) * WRFv_{y_{i,j}} + w_{i,j} * sts_{vy} \quad (79)$$

$$w_{i,j} = \frac{1}{1 + (\alpha d_{i,j})^\beta} \quad (80)$$

Where:

- sts_x and sts_y are the coordinates (x, y) of a station.
- $x_{i,j}$ and $y_{i,j}$ are the coordinates (x, y) of a given cell (i, j) in the underlying raster.
- $d_{i,j}$ is the distance between the station and the cell (i, j).
- $w_{i,j}$ is the weighting factor applied to each cell (i, j) to compute the final wind speed.





- $V_{x_{i,j}}$ and $V_{y_{i,j}}$ are the wind speeds in the x and y directions at cell (i, j).
- $WRFvx_{i,j}$ and $WRFvy_{i,j}$ are the wind speeds in the x and y directions predicted by the WRF model at cell (i, j).
- sts_{vx} and sts_{vy} are the wind speeds in the x and y directions at the station location.
- α and β are parameters controlling how the stations relevance decays with distance.

Figure 26 shows the wind field calculation for a single station.

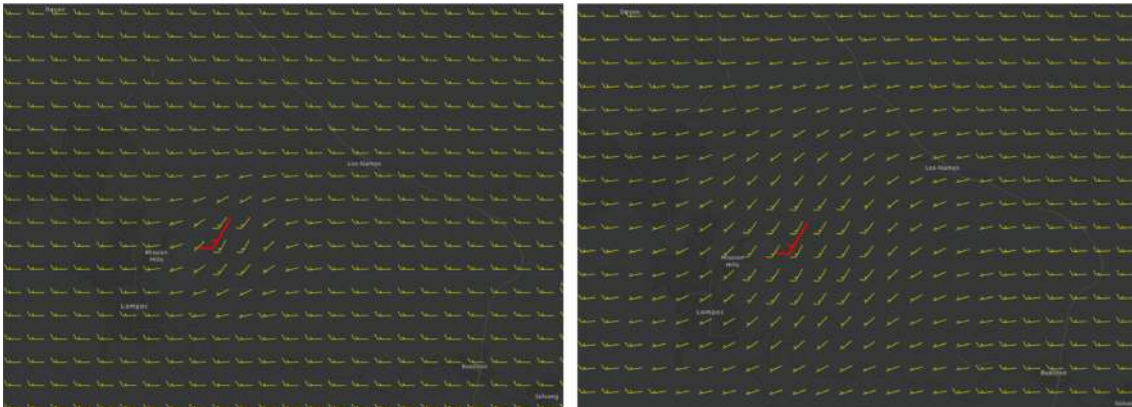


Figure 26. Effect of a single wind measurement (red) on a constant horizontal flow (yellow) using different weights parameter alpha and beta. Main wind from the west in the example.

For cases where we have more field measurement stations, as is usually the case, and knowing the effect that each station produces independently on each grid point, an IDW at each grid point (i,j) to obtain the final wind speed is performed as follows:

$$V_x = \frac{\sum_{k=1}^n w_{i,j}^k * V_{x_{i,j}}^k}{\sum_{k=1}^n w_{i,j}^k} \quad (81)$$


$$V_y = \frac{\sum_{k=1}^n w_{i,j}^k * V_{y_{i,j}}^k}{\sum_{k=1}^n w_{i,j}^k} \quad (82)$$

$$w_{i,j}^k = \frac{1}{d_{i,j}^{\theta}} \quad (83)$$

Where:

- k is a parameter that associates the k-th station.
- θ is the decay factor of Inverse Distance Weighting. It usually has a value of 2.
- V_x and V_y are the final wind components after weighing the WRF and all the stations that enter the model.

Final wind field considering several stations is shown in Fig. 27:








Figure 27. Effect of three different wind stations (red) on a constant horizontal flux (yellow). Main wind from the west in the example.

4.4. Calibrated simulations through actual fire fronts locations

In addition to improving the weather model, the functionality developed in the TEMA project also improves simulation results by taking into account the actual fire behaviour. This helps to mitigate possible deviations from reality that may arise from inaccurate input data. To do so, the Information Fusion component (PDM-tech-05) processes all available information from satellite and drone images, and other information sources to generate fire probability maps (Maps4Fire). These maps are filtered over a certain probability threshold to obtain those pixels representing the most likely fire presence. Afterwards, this raster is resampled to meet the spatial resolution of the FireSim scenario. Once this final raster is generated, the pixel centroids are calculated and exported as another input for FireSim. Hence, a GeoJSON file with fire presence positions is generated and used to feed the new simulations. These points generate a collection of control points by assuming that the actual fire arrival time is known, see Fig. 28. The technique is based on minimizing the error between the simulated fire growth and the actual fire by identifying the optimal rate of spread adjustment parameters using a least squares approach.








Figure 29. Selection of control points within the standard simulation perimeter.


With these features implemented, FireSim extracts the locations of the new fire locations, ingests them as adjustment points in addition to the real-time weather information, and reruns new and calibrated simulations each time the new current information reaches the Digital Enabler component (TEMA technology SV-tech-02).

4.5. Live Fuel Moisture Content (LFMC)

The Live Fuel Moisture Content (LFMC) is a basic parameter in fire propagation. The higher the LFMC, the more energy will be needed to dry out the fuels so that they can combust. On the other hand, this data requires complex models and training techniques that are beyond the scope of the TEMA project. Because of this, it was decided to set these values for herbaceous and woody LFMC at summer conditions, when most fires occur.

On the other hand, the grass was still green during the RAS trial, and the humidity conditions were not quite the same as those set in FireSim, resulting in a larger simulation than would be expected given the environmental conditions. Because of this, it was decided to set herbaceous and woody LFMC as parameters within the alert, so that, if these values need to be modified based on expert criteria, this would be possible before running the mission.





4.6. Fulfillment of TEMA KPIs

Tecnosylva KPIs assigned in the TEMA project were:

1. Simulation results accuracy - adjustment mode (TV: Calibration of fire simulations in real-time through field fire observations. Increase the simulation results accuracy with an error less than 15%).
2. Weather inputs accuracy (TV: Increase the fire simulation results accuracy by 10% over SoA through the use of weather inputs calibration).

Both KPIs (OB2) regarding the enhancement of the weather model and calibration of fire spread through field fire observations have been fully achieved. In the absence of real cases, historical fires could be analyzed using information about their actual weather conditions and the location of the fire fronts, although this information is not always available, nor would the technologies developed for TEMA be used.

An example of this achievement was shown in the historical case of the Montiferru fire, when meteorological information on the event and the position of the flame fronts at different times was used, collected in a technical report provided by RAS. By comparing the standard and the adjusted simulations with the real fire spread, the simulation adjusted with the technology developed by TEMA matched the burned perimeter almost perfectly, while the standard simulation barely reached half of the actual fire perimeter, see Fig. 30.






Figure 30. On the left, a standard simulation using only the weather forecast and scenario data. On the right, the same simulation is enhancing the weather forecast with actual weather information during the fire and extracting fire fronts' locations as control points from the RAS report.

An opposite example occurred during the actual RAS trial in Sardinia, conducted in June 2025, when, due to the LFM parameters set in FireSim (see Section 4.5), the simulation propagated much further than expected, knowing that the moisture conditions of the fuels present were not those of the summer conditions preset in FireSim.





Despite this, the adjusted simulation that considered real-time weather conditions and the positions of the fictitious flame fronts placed on different piles of wood that were lit to simulate the possible real spread of a natural fire significantly reduced the spread of the simulation. This shows how, even when there are errors in basic parameters such as LFMC, calculating adjusted simulations with TEMA technology greatly increases the accuracy of FireSim simulations by including real-time information. Figure 31 shows the obtained results.

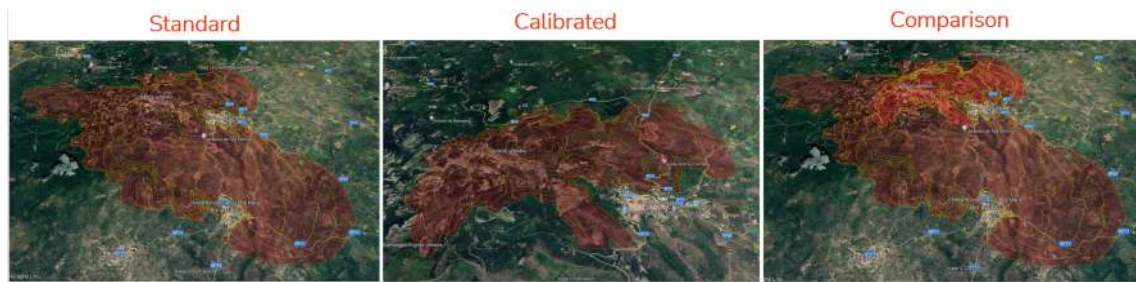




Figure 31. Comparison between standard and adjusted simulation using RAS trial inputs.

The achievement of both objectives will be verified again in the following practical cases of the Fire Mission in the coming year.





5. Near real-time flood modeling


5.1. Introduction

In emergency situations, decision makers and first responders rely on real-time or near-real-time predictions to implement protective measures. Therefore, accurate and fast hydrodynamic modeling and flood forecasting are crucial for immediate emergency response and long-term risk assessment and urban planning. To achieve near-real-time flood simulations, research is performed to accelerate the 3Di Water Management software by aiming to reduce the computational cost without compromising on accuracy. 3Di is specifically a suitable candidate for this investigation, as it already benefits from a modeling method that allows for a strong increase in accuracy for coarse grid-based simulations, while keeping the computational cost affordable.

5.2. State of the art and TEMA's contribution

In hydrodynamic modeling, it is conventional to cover the domain of interest with a grid of computational cells. In 3Di, a so-called staggered grid is used. This means that water levels and volumes are computed in cell centers, and velocities and discharges are computed at cell edges. This offers multiple advantages, including maintaining orthogonal flow direction, simple discretization for the equations, and accurate estimations for shear stresses. One can refer to the report D4.1-Report on phenomenon prediction and information fusion for an elaborate explanation of the fundamental physics laws behind 3Di.

Furthermore, 3Di is a subgrid-based model [70, 71]. The subgrid system is a modeling approach designed to account for small-scale bathymetry features and their effect on the large scale. This technique is based on the principle that variations in bed level are generally greater than variations in water level. Traditional shallow water models divide the domain into grid cells, each representing averaged properties of the terrain and flow. However, if the grid resolution is too coarse to capture small-scale changes in bathymetry, significant errors can occur in predicting the behavior of the flow and water storage. The subgrid method addresses this by embedding detailed topographic information within each coarse grid cell. This detailed information, namely bed elevation, roughness, vegetation, and infiltration rates, is typically derived from high-resolution digital maps. As such, 3Di uses two grids: a coarse computational grid and an underlying subgrid with higher resolution (illustrated in Fig. 32). This allows the model to represent small-scale variations in bathymetry and wetting-drying behavior while keeping the computational cost reasonable.







Figure 32. Illustration of the grid systems used in 3Di; coarser grid (green) as computational grid and finer grid (grey) as subgrid






Figure 33. Illustration of a computational grid including a quadtree grid refinement

On the other hand, in most model domains, there are regions where the flow dynamics are highly complex or for which more elaborate results are desired. In such cases, a finer spatial resolution is required to adequately capture the flow behavior. Given that the computational cost is quadratically proportional to the total number of computational cells, 3Di provides the option of local grid refinement, allowing higher resolution only in areas of interest. Using a so-called quadtree refinement method, grid cells are recursively subdivided into four smaller cells, allowing a hierarchical structure where cell sizes vary smoothly across the domain [72]. The method enforces a constraint that the edge length of a refined cell cannot be smaller than half of that of its neighboring cells (illustrated in Fig. 33).

Building on the state-of-the-art of hydrodynamic modeling concepts, NS has conducted exten-





sive research within WP4 on a promising alternative to model disconnected flow regions with subgrid-scale features and coarse-grid 2D models: Clone Cells (grid cloning). This approach enables more precise identification and prediction of evolving flood phenomena without compromising on calculation cost.

5.3. Summary of advances in the last period

During the M19-M36 period, the development of Clone Cells started, whereas before (M1-M18), efforts were mainly focused on the development of the methodology and conceptual basis of the technology.


The integration of Clone Cells in hydrodynamic modeling is promising. Accurate hydrodynamic simulations depend very much on the resolution of input data. However, enhanced input data requires more detailed analysis and usually results in higher computational cost. Specifically, in TEMA, detailed information about disconnected flow regions can be obtained during a flood event by drone images or sensor data. When new or updated information is included in the scenario input data of the simulation software, 3Di does not want to compromise much on calculation cost by adding more detail to the schematization of the real-time event. If, during floods, due to collapsed infrastructure, such as buildings or bridges, areas in the flood region become disconnected, clone cells add detail to the simulation with only a minimal increase in computational cost.


Besides, 3Di calculates a preliminary flood simulation upon receiving the alert and an Area Of Interest (AOI), using basic information, such as the low-resolution digital elevation model (DEM) from Copernicus and weather forecast information (ICON EU-forecast). When real-time information becomes available, 3Di reruns the simulation with updated scenario information with enhanced input data. However, when updated (field) data is unavailable or delayed, 3Di still calculates new simulation output at desired timestep outputs for the end-users, based on the altered weather forecast that also progresses through time. The advances in the hydrodynamic modeling with clone cells have been summarized in Section 5.4.

Furthermore, integration of the flood modeling software continued during this current review period. Updated expectations of real-time flood updates were retrieved from end-users and implemented as such on the TEMA platform.

5.4. Hydrodynamic modeling with Clone Cells

Accurate hydrodynamic models are crucial for reliable flood forecasting. For forecasting floods, it is especially important to accurately handle complex flow areas such as bifurcations, narrow barriers, and adjacent waterways, where flow regions are separated. Using a coarse grid, where a computational cell contains disconnected flow regions with different water levels and/or different flow directions, fails to capture essential hydrodynamic behaviors. This is due to the assumption that the flow variables are uniform per computational cell. Applying local grid refinement,





although efficient, does increase the computational cost, particularly when many refinement levels are required. Moreover, it introduces additional numerical diffusion due to grids of different resolutions. A new method, the so-called clone cells (grid cloning) [73], is investigated to allow for accurate and efficient modeling of disconnected flow regions.

The clone cell method provides users with a means to align their computational grid accurately with the underlying landscape without the need for grid refinements. It allows for flow separation within the length scale of a single grid cell, such as bifurcated creek branches, closely spaced waterways, and barriers, improving model accuracy and flexibility with a limited increase in computational cost. This offers an interesting trade-off between computational efficiency and numerical precision.

The subgrid method allows for partially wet domains within a computational cell. This flexible active domain can have an arbitrary shape, adjusting to complex landscapes and water levels. Using this existing quality of the subgrid method, it becomes possible to define two computational cells occupying the same spatial location (i.e., they are each others clones). The clones can be conceptualized as two overlapping cells with complementary active domains (as illustrated in Fig. 34). The division between the active domains of these cells is specified by the user with a line called a cutline. Each cutline is assigned a crest level, representing the elevation at which flow exchange might occur. An advanced algorithm is implemented to generate a computational grid allowing for both quadtree refinements and clone cells. The algorithm needs to determine active domains, active neighbors, and detect possible flow exchange between clone cells. Moreover, a flow formulation is developed to calculate the (possible) flow dynamics between the two clones.





Figure 34. a) a computational cell without grid cloning, b) a computational cell with grid cloning, c) active domain of each clone cell

In a nutshell, clone cells promise an advancement in hydrodynamic modeling, enabling improved accuracy in coarse-grid simulations with only a minimal increase in computational cost (see Fig. 35).



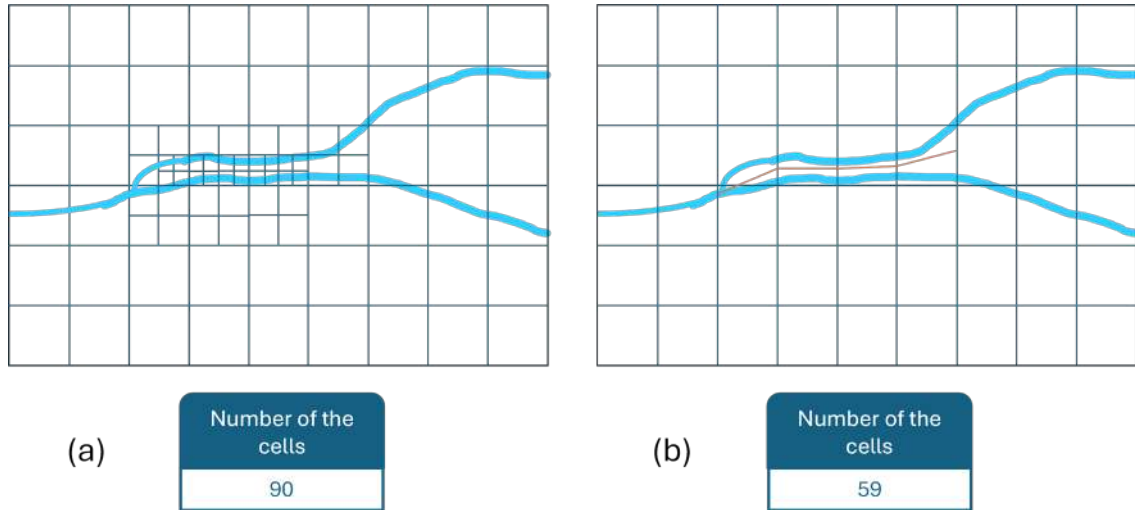



Figure 35. Total number of computational cells using: a) quadtree grid refinement, b) clone cells

Before applying the methods to the cases of the TEMA project, the new method is formally tested on ideal test cases. The test case below (Fig. 36) presents a channel in which the water flows from left to right. In the middle of the channel, an obstacle is introduced that locally narrows the width of the channel. At the location of the obstacle, a cutline generates clone cells, the active domains of which are indicated by colored fragments. The crest level assigned to the cutline is sufficiently high to prevent the flow exchange between the clone cells within each quadtree cell. The result shown in Fig. 37 illustrates the direction of the flow and the corresponding discharge gradient. It can be observed that the flow is hindered at the location of the obstacle and is consequently redirected towards the narrow opening. This test case serves to verify the implementations of the clone cells method and represents a simplified equivalent of more complex and realistic scenarios where elevated roads, railroads, levees, or dikes are involved.




Figure 36. Grid system of the test case (the yellow line represents the cutline and the colored fragments the clone cells)






Figure 37. The discharge and water level map of the test case





5.5. Fulfillment of TEMA KPIs


The following KPIs were linked to flood modeling in the TEMA project:

- OB1: Increase model-based prediction responsiveness/speed for evolving phenomena
- OB2: Increase model-based prediction accuracy for evolving phenomena
- OC2: Increase situational awareness in NDM
- OC4: Prototype a proof-of-concept TEMA system for NDM in flash floods and regional floods

All KPIs for NS are met in the TEMA project. Both OB₁ and OB₂ required the development of clone cells, whilst OC₂ and OC₄ needed improvements in the integration and automation of workflows from 3Di software to the TEMA platform. The real-time Application Programming Interface (API) of 3Di provides a solid foundation for enhancing real-time updates and improving precise phenomenon prediction in TEMA Task T_{4.1}.

For objective OB₁, the target was a **real-time update rate**, tailored to the needs of the end-users. Currently, when new information becomes available in end-user cases, the model and simulations can be updated within half an hour. Exact times differ depending on the type of the new data set. In case an improved DEM (Digital Elevation Model) becomes available, both the model and simulations are to be updated. In case a new weather forecast provides improved data on the precipitation event or wind event, only the simulations need to be updated. The clone cells method developed aims at improving the accuracy, without increasing computational cost, as explained in Section 5.4.

Objective OB₂ has a target where the accuracy of flood simulation results is increased. This should result in a target of a **MSE of less than 15%**. The clone cell developments are a special technique that focuses on dealing with obstacles and bifurcations in the flow. So the increase in accuracy strongly depends on the measure of performance. We can further optimize the Clone Cells method to allow for faster and more accurate results. Locally, we can increase the accuracy from totally wrong to good. An example is shown in Figure 38. The flooding extend is computed with more accuracy, whilst still using coarser (and therefore fewer) computational cells. This can result in an MSE of much less than 15%.








Figure 38. Overview of computational grid with a flooded domain in Blue and in Green, an obstacle that should block the flow. Left) Original approach, having refinements in combination with a method that blocks the flow on cell edges. Right) New setup with clone cells (marked cells), where active domains can be defined accurately.

Objective OC2 was achieved by automation of workflows, implementing systematic model updates, and reducing the post-processing time of simulation results. The real-time updates of flood results contribute to the situational awareness of first responders, firefighters, or emergency response teams.

Objective OC4 is focused on the validation of the TEMA system in experiments, where a successful experiment completion is foreseen. Currently, we have run several historical test cases (Ahrtal use-case during the Hackathon in Munich (April 2025) and the DMALLIAN use-case (October 2025) successfully. Figure 39 shows four snapshots of the flood model results in the TEMA platform.

Next up is the historical use case of the Ahrtal region in November 2025. The former results showed promising options to improve the generation of models for future speed-up of model generation.



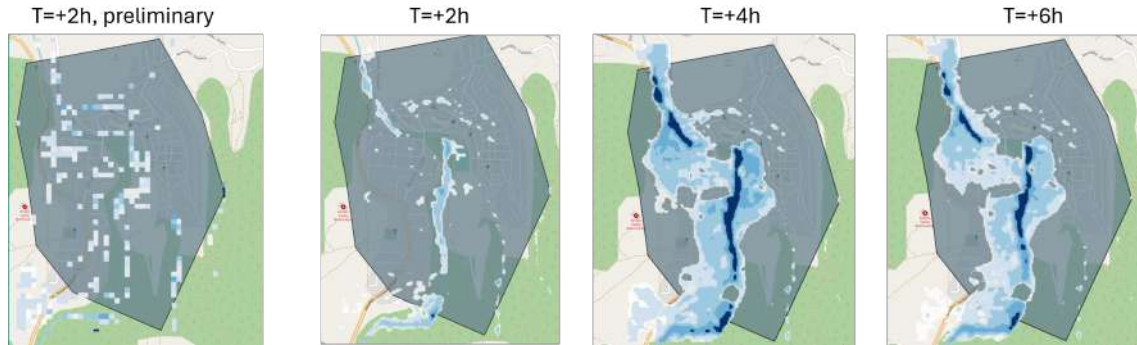





Figure 39. Snapshots of flood simulation results in the Mantoudi region. The left image is the simulation results at 2 hours after the start of the flood based on preliminary (basic) data. The second to fourth images are flood simulation results at respectively 2 hours, 4 hours, and 6 hours after the start of the flood, based on enhanced data from the TEMA platform.





6. Disaster prediction from user-generated data sources


6.1. Introduction


As per the project's necessity to connect the work carried out across tasks, some of the research started in Task T3.3 (Geo-social media analysis) and Task T4.3 (Information fusion) is closely connected to Task T4.1 (Precise phenomenon prediction). Accurate and timely prediction of natural disasters remains a critical challenge, especially for rare or rapidly evolving events. Leveraging user-generated data from multiple sources offers promising opportunities to capture near-real-time signals that traditional monitoring systems may miss. Phenomenon prediction from user-generated data is therefore critical to ensure the real-world utility of the TEMA outputs for end users.

6.2. State of the art and TEMA's contribution

Recent research has demonstrated that user-generated content constitutes a critical resource for natural disaster detection and prediction, offering complementary, near-real-time information beyond traditional sensor networks. Early approaches relied on keyword filtering of geo-referenced Twitter data, which enabled the rapid identification of earthquake-related impacts and even operational integration into systems such as Global Disaster Alert and Coordination System (GDACS) [74, 75]. Subsequent work advanced towards machine learning methods, ranging from topic models [76] and similarity joins incorporating spatiotemporal context [77] to deep learning architectures such as Convolutional Neural Network (CNN) and Bidirectional Encoder Representations from Transformers (BERT) for flood and wildfire detection [78, 79, 80, 81]. More recent studies highlight the importance of multimodal integration, combining text with images, spatiotemporal features, or even weather data for improved detection and monitoring [82, 83, 84, 85]. Beyond social media, other crowd-sourced platforms such as Global Database of Events, Language and Tone (GDELT) news streams [86], have shown significant potential for disaster detection, albeit with persistent challenges in data quality and geographic bias.

Building on this state-of-the-art, recent work by IT:U has significantly advanced disaster detection and prediction using user-generated and open data sources. First, IT:U conducted the first large-scale, geographically explicit assessment of Bluesky as a platform for disaster monitoring. This study demonstrated that user-generated posts on Bluesky provide meaningful temporal and spatial signals for major disaster events such as the 2024 Central Europe floods and the 2025 Southern California wildfires. The results showed strong temporal alignment with real-world event progression and revealed characteristic emotional and spatial dynamics of disaster-related discourse, suggesting that Bluesky holds high potential as a complementary source for real-time situational awareness. Based on these findings, IT:U built a methodology for natural disaster prediction that integrates Bluesky social media data, global news reports from GDELT, and meteorological observations. This integration bridges the gap between human-sensed information and environmental sensor data, creating a unified method for multi-source disaster analysis. The






approach leverages large-scale Foundation Models (FMs) capable of contextual reasoning and long-sequence processing [87], which enables the joint interpretation of textual, spatial, and temporal signals across modalities. In contrast to previous works that primarily rely on statistical anomaly detection or single-source data, the FM-based framework performs geospatial reasoning directly over multimodal inputs. Compared to statistical hotspot and weather anomaly baselines, the FM-based approach achieved higher macro F1-scores (up to 0.66), improved ROC-AUC (up to 0.69), and increased recall for disaster events. The advancement marks a major step towards automated, and near-real-time disaster monitoring and prediction within TEMA, directly contributing to the scope of Task T4.1.


6.3. Summary of advances in the last period

During the last period, IT:U has substantially advanced disaster observation and prediction using user-generated data by systematically integrating multiple heterogeneous sources, including Bluesky posts, GDELT news, and weather observations. Key achievements include the first large-scale, geographically explicit assessment of Bluesky for disaster monitoring, demonstrating strong temporal alignment with real-world events and meaningful emotional and spatial patterns across two major disasters (2024 Central Europe floods, 2025 Southern California wildfires). Beyond descriptive analyses, we developed a multimodal prediction framework leveraging large-scale FMs with in-context learning and Automatic Prompt Optimization (APO), achieving improved macro F1-scores, ROC-AUC, and event recall compared to statistical baselines and naïve multimodal ensembles. These advances establish a robust foundation for near-real-time, spatially explicit disaster detection and prediction, addressing the objectives of TEMA Task T4.1 with user-generated data.

6.4. Disaster observation with user-generated data

Within TEMA Task T4.1, significant progress has been made by IT:U in advancing disaster observation through user-generated data. The effort addresses the urgent need for viable alternatives to X (formerly Twitter) in academic research, as the platform has severely restricted data accessibility for researchers since May 2023 [88]. Therefore, IT:U has systematically investigated Bluesky as a potential data source for geographically informed situational awareness in disaster management and has submitted the corresponding research paper to the *Journal of Computational Science*. The work covers two major events: the 2024 Central Europe floods and the 2025 Southern California wildfires. Using a custom multilingual crawling and analysis pipeline, IT:U collected over 676,337 posts and applied geoparsing, semantic classification, and emotion detection to evaluate their spatiotemporal and thematic characteristics. The results demonstrate that Bluesky provides large volumes of disaster-related content with strong temporal alignment to the actual progression of events, and even shows potential for early warning and thus prediction. These findings represent a concrete step forward for Task T4.1, as they highlight the feasibility of integrating emerging social media platforms into phenomenon modelling pipelines for near-real-time disaster detection and prediction.





In addition to temporal alignment, our analysis revealed important spatial and emotional dynamics of disaster discourse on Bluesky. For the Central Europe floods, disaster-related posting activity increased gradually and peaked after the onset of heavy rainfall, while for the California wildfires, social media activity rose sharply and closely matched the number of active fires identified by National Aeronautics and Space Administration (NASA) Fire Information for Resource Management System (FIRMS). Emotion analysis further confirmed shifts in public sentiment, with fear and anger dominating during peak disaster phases. On the spatial dimension, a clear urban bias was observed, with higher post densities in large population centers such as Los Angeles or Vienna, but with strong overlap between disaster-related posts and impacted regions (Fig. 40). These results underscore both the opportunities and limitations of Bluesky for disaster monitoring. While coverage is still uneven across regions with lower platform penetration, its real-time Application Programming Interface (API) provides a solid foundation for enhancing precise phenomenon prediction in Task T4.1.

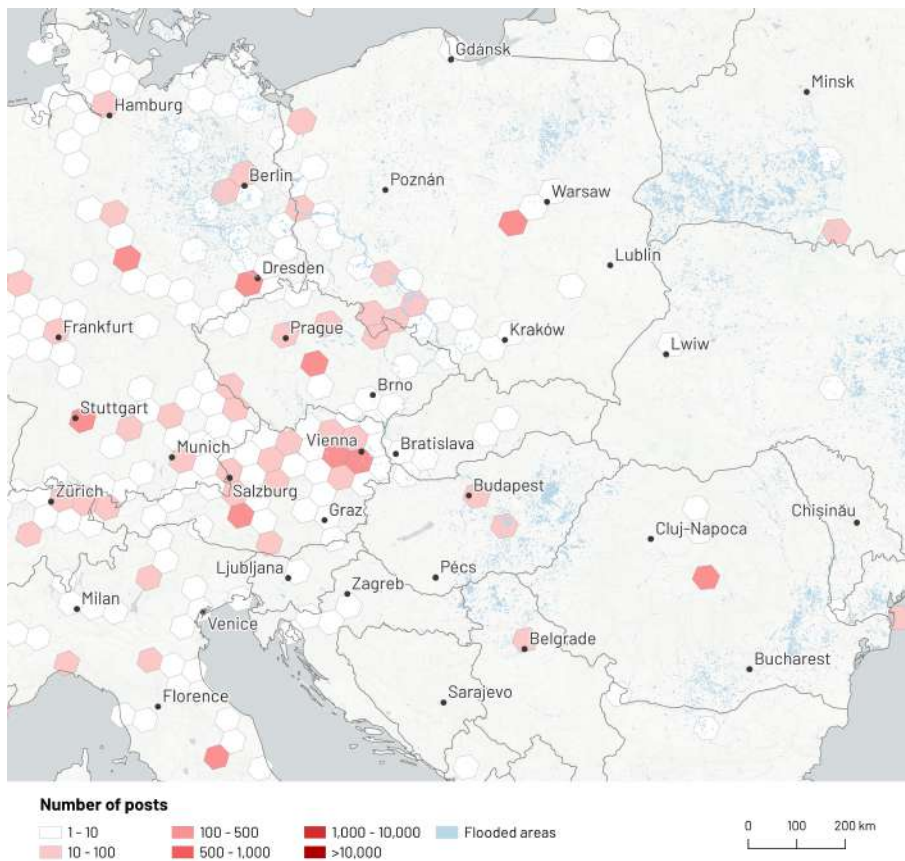




Figure 40. Relation between post count and presence of flooding on H3 grid (level 4). Regions without posts were excluded for legibility.






6.5. Disaster prediction using user-generated data

In follow-up work to IT:U's efforts described, disaster prediction was formulated as a binary classification problem over H3 grid cells – hexagonal spatial units provided by the open-source H3 geospatial indexing system developed by Uber [89]. The system divides the Earth's surface into hexagonal cells at multiple resolutions, enabling consistent spatial aggregation and analysis. The cell-based problem formulation allows for quantitative evaluation using precision, recall, macro F1-score, and ROC-AUC. For the experiments, two recent natural disasters were selected: the September 2024 floods in Central Europe and the January 2025 wildfires in Southern California. Statistical hotspot detection with Getis-Ord G_i^* [90] and weather anomaly detection served as baseline approaches. However, we found a severe limitation. For the 2024 Central Europe floods, for instance, hotspot detection with GDELT news articles achieved the highest macro F1-score of 0.51 and ROC-AUC of 0.67 but captured only 10% of actual flood events. Bluesky-based hotspot detection underperformed (M-F1: 0.45; class1 recall: 0.06), and weather anomaly detection was only slightly better (M-F1: 0.50; class1 recall: 0.29). For the 2025 Southern California wildfires, a statistical ensemble combining Bluesky and GDELT data reached a macro F1-score of 0.62, with class 1 recall of 0.37 and precision of 0.21, demonstrating that naïve multimodal combinations partially mitigate but do not fully resolve these limitations.

To overcome these constraints, IT:U implemented multimodal geospatial reasoning using large-scale FMs with in-context learning (Fig. 41). By simultaneously incorporating social media posts from Bluesky, GDELT news, and weather observations per H3 cell, the FM-based approach exploits spatial and temporal context, interdependencies across modalities, and near-real-time signals. APO further enhanced performance. Applying APO to the Qwen3 FM, macro F1-scores increased by 0.04 and improved ROC-AUC from 0.65 to 0.69 (Table 1). When combined in an ensemble with statistical methods, the FM-based predictions achieved higher recall for disaster events than statistical baselines alone, while maintaining balanced precision. Ablation studies confirmed the additive value of each data modality, and Jaccard index analyses across methods highlighted substantial overlap in detected events while revealing complementary contributions from social media, news, and weather signals. These results demonstrated that multi-source event prediction – and therefore also detection – with geospatial reasoning not only improves capture rates for rare disaster occurrences but also provides a robust framework for integrating diverse, user-generated datasets in near-real-time, directly addressing the objectives of Task T4.1 with user-generated data. The work is currently being finalised in collaboration with DLR-DFD.

6.6. Fulfillment of TEMA KPIs

By integrating social media, news, and weather data in near-real time, the methodology enhances the accuracy (OA2) of extreme data analysis algorithms, as reflected in improved event detection metrics (macro F1-score, recall, precision) compared to baseline statistical methods. The use of large foundation models with in-context learning and automatic prompt optimisation supports responsiveness and speed (OA3, OB1), enabling timely predictions per H3 cell. Furthermore, multimodal, geospatial reasoning improves the trustworthiness (OA1) of predictions through interpretable, modality-specific contributions, while the frameworks design contributes to model-based prediction accuracy and responsiveness (OB2, OB1) for evolving phenomena



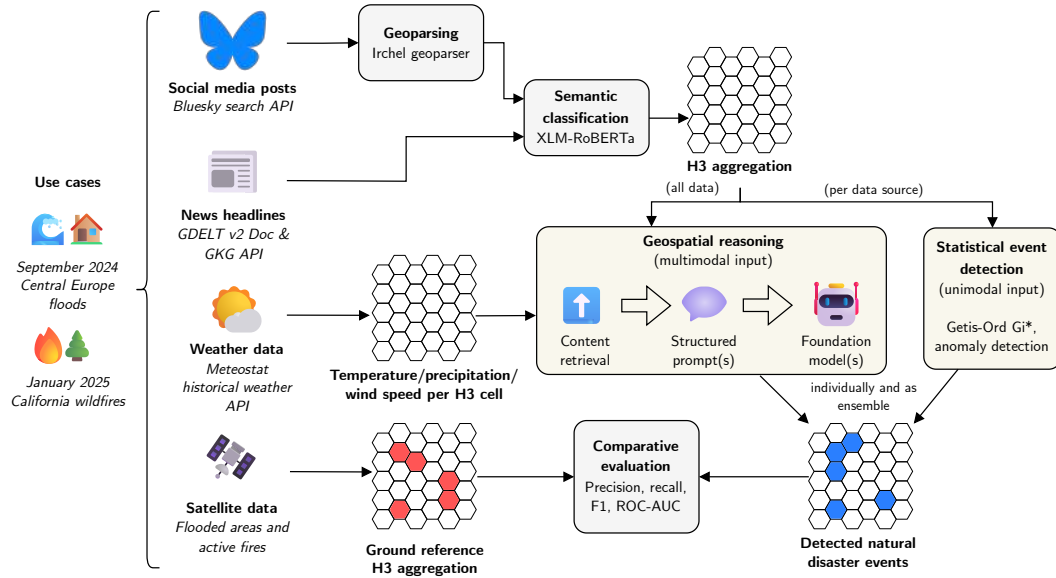





Figure 41. Overview of our systematic evaluation of disaster prediction methodologies, including the FM-based methods.

Table 1. Average model performance across the two disaster events on the human-engineered base prompt versus the optimised prompts after APO with zero-shot learning. All results were computed on a subset of 40 samples per disaster (20 disaster-positive and 20 disaster-negative) to minimise inference time and costs. The best scores are marked in bold, the second-best scores are italic.

Model	Base prompt			Optimised prompt		
	M-F1	ROC-AUC	Invalid	M-F1	ROC-AUC	Invalid
GPT-4.1-nano	0.56 (±0.04)	0.64 (±0.03)	0.00 (±0.00)	0.61 (±0.02)	0.67 (±0.02)	0.00 (±0.00)
GPT-5-nano	0.66 (±0.04)	0.66 (±0.05)	0.00 (±0.00)	0.65 (±0.03)	0.72 (±0.05)	0.00 (±0.00)
Gemma3	0.59 (±0.01)	0.68 (±0.01)	23.60 (±0.42)	0.65 (±0.02)	0.69 (±0.02)	25.40 (±0.56)
Llama3.1	0.31 (±0.18)	1.00 (±0.00)	76.60 (±0.27)	0.55 (±0.04)	0.73 (±0.03)	18.12 (±1.42)
Qwen2.5	0.45 (±0.00)	0.56 (±0.02)	0.00 (±0.00)	0.48 (±0.00)	0.59 (±0.02)	0.00 (±0.00)
Qwen3:14B	0.56 (±0.00)	0.66 (±0.01)	0.00 (±0.00)	0.60 (±0.01)	0.66 (±0.02)	0.00 (±0.00)
Qwen3:8B	0.61 (±0.01)	0.65 (±0.01)	0.00 (±0.00)	0.66 (±0.01)	0.68 (±0.01)	0.00 (±0.00)
Deepseek-R1	<i>infinite thought loop, no output</i>					

such as floods and wildfires.





7. Response Planning

7.1. Introduction

The objective of drone response planning is to efficiently assign trajectories to each available drone in order to optimally complete the Natural Disaster Management (NDM) mission. It is an essential functionality to enable the operation and success of the TEMA technologies. The trajectories are computed and provided to the human operator for validation. Once validated, they are sent to the drones for execution, enabling autonomous operation with a human in the loop. Response planning has been developed in Task 4.1 “Response Planning and Recommendations” starting in M19 and finishing in M36.


This section provides the theoretical and technical framework for the *PDM-tech-04: Drone Planning* technology within the TEMA project. The drone planning component of the TEMA architecture is responsible for generating optimal multi-UAV mission plans for disaster response operations, specifically addressing flood and fire scenarios. The response planning component transforms probabilistic disaster information into executable UAV trajectories while considering operational constraints, resource limitations, and mission objectives.


The response drone planning approach is built on established search theory principles, adapted for actual multi-UAV systems, and incorporates real-world operational requirements and integration with TEMA technologies. This section details the mathematical foundation, algorithmic implementation, system architecture, and validation procedures that constitute the complete response planning solution.

7.2. State of the art and TEMA’s contribution

Due to the development in drones in the last decades, they have become an ideal tool for search missions not only due elimination of human safety risks in hazardous environments but also because of their versatility and easy deployment. Several works [91, 92, 93] have already tapped into the research of using drones for search missions in a systematic manner.

When multiple robots need to cooperate, a coordination strategy is required. These strategies can vary from fully centralized approaches to robot swarms, each characterized by the capabilities of the robots, as well as the decision-making authority of the system. A centralized approach consists of robots that simply execute instructions delivered from a single entity through a low-latency, high-throughput communication channel. Such a centralized system can optimally utilize its resources as it has a global knowledge of the environment [94]. On the other hand, a robotic swarm is usually composed of simple robots with limited sensing, computation, and communication capabilities. Its operation usually relies on simple control and behaviour policies, the aggregate results of which achieve the overall goal of the system [95]. In general, such an approach is usually resilient to failures of sensors, communications, and even of robots. In addition, given that the decision and control computations are performed locally, based on local observations, such systems have good scalability. Despite the resilience and robustness of the robot team, heavily decentralized systems cannot provide any guarantees on the quality of the task execution.





In natural disaster scenarios, drones must coordinate for disaster monitoring and for performing tasks such as search and rescue. Search mathematical foundations can be found in the Theory of Search [96, 97, 98, 99], a branch of Operations Research dedicated to the study of optimal search in different scenarios. Several authors extended and refined the theory to more general scenarios [100, 101]. Many reviews and books have been published with extended revision on the state of the Theory of Search [102, 103, 104]. This theory has been successfully adapted and applied in real search missions. It was used to recover the H-Bombs lost on the southern coast of Spain in 1966 after a B-52 military aircraft and a tanker crashed [105]. It was also applied to develop the Search and Rescue Optimal Planning System (SAROPS) [106] software that the U.S. Coast Guard uses for real missions. This software was used to find the Air France Flight 447, an Airbus 330-200 with 228 passengers and crew, that disappeared over the South Atlantic during a night flight from Rio de Janeiro, Brazil, to Paris, France on 1 June 2009 [107].


A recent paper [108] makes use of the well-known Traveling Salesman Problem (TSP) and coverage path planning to compute the optimal path to find a certain number of lost targets after an earthquake. However, it lacks a solid probabilistic foundation to systematically assign priority levels to each target and often uses arbitrary metrics. Another paper [109] presents a route planning algorithm for maritime search and rescue missions with moving targets. It provides a sophisticated and precise probabilistic model that addresses many characteristics of these kinds of dynamic missions. It also develops the corresponding multi-UAV path planning. However, the application is limited to maritime scenarios and lacks any consideration of the physical maneuvering and autonomy of the devices used.


A wide range of existing methods has been developed to coordinate multiple drones in surveillance and monitoring tasks. In TEMA, the drone response planning method generates optimal near-real-time trajectories for Natural Disaster (ND) surveillance within the Area of Interest (AOI). While the mathematical foundation draws from search theory principles originally developed for search and rescue applications [110], our approach adapts these concepts specifically for disaster surveillance tasks, including monitoring flood propagation, detecting wildfire hotpots, identifying re-ignition zones, and assessing disaster development. Instead of employing systematic area coverage (e.g., zig-zag patterns), our method in TEMA optimizes drone trajectories to maximize information gain from probabilistic disaster models. This information-theoretic formulation enables dynamic prioritization of surveillance areas based on real-time disaster evolution.

The developed method advances the state of the art in drone planning for disaster response as follows:

- **Mathematically Rigorous Framework:** We extended optimal search theory with information-theoretic principles, using normalized entropy to systematically quantify information gain and prioritize search areas, overcoming the heuristic approaches in existing works.
- **Multi-Disaster Applicability:** Our approach is validated for both flood and fire scenarios, providing a generalizable solution beyond domain-specific limitations in current literature.
- **Computational Efficiency:** The developed planning technology generates optimal multi-UAV missions within practical time-frames (20 seconds for 100 km^2 areas).
- **System Integration:** The drone planning technology is fully integrated within the TEMA platform architecture, enabling seamless data exchange with information fusion technology (PDM-tech-05) and supporting real-time mission adaptation

The developed drone planning technology (PDM-tech-04) bridges theoretical optimality with practical implementation, providing a field-ready planning solution that addresses the complex requirements of multi-UAV coordination in disaster response scenarios.





7.3. Response planning requirements

The drone response planning is designed to address specific end-user requirements identified in TEMA Deliverable D2.1 "Report on TEMA Requirements". Based on the End User requirements in Tables 10 and 11 (pages 80-87), the drone planning technology must address the end-user requirements described below.

7.3.1. Non-Functional Requirements

EU-RQ-NF-03: Information provision as soon as possible: The drone planning technology generates optimal drone trajectories within 5 minutes of disaster detection, enabling rapid area assessment and immediate information gathering. The mathematical optimization framework ensures maximum information gain per unit time, providing critical situational awareness to first responders in the earliest phases of disaster response.

7.3.2. Functional Requirements

EU-RQ-FUNC-03: Monitor the development, the size of the affected area: The drone planning technology generates optimal coverage paths that systematically scan the entire affected area, providing comprehensive data for mapping the exact extent and development of floods or fires. The probability-based approach ensures efficient monitoring of disaster progression.

EU-RQ-FUNC-05: Reveal fires as they start: The drone planning technology allocates optimal search effort, prioritizing areas with the highest probability of fire occurrence, enabling early detection through systematic UAV surveillance before fires escalate.


EU-RQ-FUNC-06: Teams involved: The multi-UAV coordination system within the drone planning technology provides clear task allocation and assignment visibility, showing which drones are assigned to which subareas and their specific missions.

EU-RQ-FUNC-07: Geolocation of people who are in danger of life: The drone planning technology optimizes search patterns to locate people in danger, with detection capabilities that can identify human presence in affected areas and provide precise geolocation data for rescue operations.

EU-RQ-FUNC-08: Model of fire propagation: The drone planning technology integrates with fire propagation models (PDM-tech-01) through the Information Fusion technology (PDM-tech-05) to dynamically adjust drone surveillance patterns based on predicted fire spread, ensuring optimal data collection for model validation and updating.

EU-RQ-FUNC-09: Monitor the area of interest for fire revival: The drone planning technology supports continuous monitoring missions and can generate follow-up surveillance plans to detect fire revival in previously affected areas, with optimized patterns for hotspot detection.

EU-RQ-FUNC-10: Response planning for extinguishing fires: The drone planning technology provides optimal surveillance routes that support firefighting operations by identifying critical areas, monitoring fire fronts, and assessing resource deployment effectiveness.







Figure 42. Probability map example for the historical flood case of Ahrtal in Germany. The blue intensity represents the probability of flood presence, with darker blues indicating higher probabilities. This visualization demonstrates the spatial heterogeneity of disaster likelihood across the affected region. (The squares on the grid do not match the actual pixels.)

EU-RQ-FUNC-12: Flood propagation modelling: The drone planning technology coordinates drone fleets to collect data that feeds into and validates flood propagation models (PDM-tech-02), with adaptive planning to focus on areas of predicted flood expansion.


7.4. Probability maps as the foundation for response planning



The response planning technology addresses the fundamental challenge that natural disaster events (ND), whether floods or wildfires, are spatially extensive phenomena that cannot be adequately represented as point events. Unlike localized incidents, these disasters occupy a subregion with varying intensities and probabilities of presence, requiring a sophisticated mathematical representation that can handle spatial uncertainty and correlation.

Our planning pipeline, therefore, builds on *PDM-tech-05: Information Fusion*, which provides *probability maps* as primary inputs. A probability map is a geo-referenced raster whose pixels encode the probability (in $[0,1]$) that the event is present at the corresponding geographic location (see Fig. 42).


If a pixel value is 0.6, this means that there is a 60% chance of that pixel containing either flood or fire. Let us then consider probability maps as depicted in Eq. 84

$$\mathcal{P}_M = \{M \in \mathcal{M}_{n_x, n_y} \mid 0 \leq M_{ij} \leq 1; \forall i, j\} \quad (84)$$





(a) Difference in indexing interpretation between the raster and the matrix that represents a probability map.



(b)

Figure 43. Mathematical basis of information-driven planning.

where \mathcal{M}_{n_x, n_y} denotes the space of $n_x \times n_y$ matrices and M_{ij} is the probability at cell (i, j) . In our indexing convention, columns index correspond to the horizontal direction (easting/longitude) and rows to the vertical direction (northing/latitude). This convention aligns with the standard image processing and numerical analysis practices, facilitating straightforward software implementation while maintaining mathematical clarity ($f(x_i, y_j) \leftrightarrow f_{ij}$), cf. Fig. 43a.

7.5. Multi-target search theory and information quantification


Building upon optimal search theory [97, 98], our approach of drone planning extends traditional formulations to address multi-target disaster scenarios. For n_T targets with probability density functions $p_t: A \rightarrow \mathbb{R}^+$, the probability of exactly k targets in cell c is depicted in Eq. 85:

$$P^{(k)}(c) = \sum_{\substack{T_n \subset T \\ |T_n|=k}} \prod_{t \in T_n} P_t(c) \prod_{t' \in T \setminus T_n} [1 - P_{t'}(c)] \quad (85)$$

Now, consider that a certain probability map $P \in \mathcal{M}_{n_x, n_y}$ for either flood or fire is given. By visiting or scanning regions of the map, we can obtain local information on the status of the catastrophe in that region. Consider that we have a drone that has a detector (in the context of the TEMA project, we consider regular cameras that are managed by the technology SV-tech-01) and that can be sent to certain points of interest, where it can take measurements with enough precision to correctly determine the status at that location. The amount of information that we gain can be estimated in information units (such as bits). As the "random" variable that models the presence of flood or fire is binary, we compute what we call the information or the normalized entropy matrix S as:

$$S_{ij} = H(P_{ij}) = -\frac{1}{\log 2} [P_{ij} \log(P_{ij}) + (1 - P_{ij}) \log(1 - P_{ij})] \quad \forall i \in [1, n_x], \forall j \in [1, n_y]. \quad (86)$$

Each component of S measures how much information units we gain after completely determining the status of position (i, j) (see Figs. 43b and 44). In Fig. 43b, we plot the dependency of








Figure 44. Complete planning pipeline showing transformation of probability maps through information quantification, search effort optimization, area decomposition, multi-UAV coordination, and trajectory generation. Each stage maintains mathematical optimality while addressing practical constraints.


S_{ij} on P_{ij} . As neither the case where a catastrophe is present or not should be special over the other, the entropy value is symmetric to the transformation $P \leftrightarrow 1 - P$. The special point $P = 1/2$ is where we are most uncertain about the local status. This means that a scan of that pixel will provide the maximum information and, as such, a maximum in S is reached. When the status is known with very good precision, that is, we are close to either $P = 0$ or $P = 1$, no information can be gathered, and a minimum is reached.

As in initial states of a catastrophe, we have almost no knowledge about its status, the information maps are expected to be almost flat around $P = 1/2$ over the entire region of interest. This is not too realistic for several reasons, the first one being that outside the area, neither flood nor fire should be considered. Therefore, the probabilities are exactly 0 outside the rasterized area. Even if the selected area does not contain the entirety of the catastrophe, there is an abrupt change in probability within the boundaries. Somehow, the boundaries must be selected carefully to ensure that everything is correct. Another issue is that usually, alerts are activated due to local detectors. The probability of the catastrophe being located in the vicinity of triggered detectors should be higher than in regions that are further apart.

In any case, the first steps are critical as no new information is available to better define the probability maps. For this, a general reconnaissance flight mission is proposed. It should cover much of the area close to the detectors, and if several are triggered within a short period of time, the ones with higher certainty should be prioritized. This can be done with a priority-aware TSP-based model.

However, it is important to note that the drone planning technology operates on probability maps provided by the Information Fusion component (PDM-tech-05), which already incorporate sensor characteristics, detection probabilities, and Bayesian updating of the ND likelihoods. Given these probability maps, the planning algorithm defines an information value matrix M that quantifies the potential information gain across the operational area.





7.6. From single-look updates to an information footprint

Consider that the probability of disaster at a certain position of space \mathbf{bx} is $P(\mathbf{bx})$. If the drone u is at position \mathbf{bx} , the probability of detecting the disaster at \mathbf{bx}_c if it is present at that location is $P_D^{(u)}(|\mathbf{bx} - \mathbf{bx}_c|)$. Notice that we already consider that it depends only on the distance between the detector position and the scanned region (isotropy). Then, after the measurement, the probability of disaster after no detection is depicted in Eq. 87.

$$P'(\mathbf{bx}_c|\text{ND};\mathbf{bx}) = \frac{1 - P_D^{(u)}(|\mathbf{bx} - \mathbf{bx}_c|)}{1 - P_D^{(u)}(|\mathbf{bx} - \mathbf{bx}_c|)P(\mathbf{bx}_c)} P(\mathbf{bx}_c) \quad (87)$$

The information gained after the measurement is depicted in Eq. 88:

$$\text{IG}^{(u)}(\mathbf{bx}, \mathbf{bx}_c) = H(P') - H(P), \quad (88)$$

and the *information footprint* (or M -matrix) for a UAV positioned over cell (i, j) aggregates contributions over all cells (k, l) is defined as depicted in Eq. 89:

$$M_{ij}^{(u)} = \sum_{k,l} \text{IG}^{(u)}(\mathbf{bx}_{ij}, \mathbf{bx}_{kl}), \quad (89)$$


with \mathbf{bx}_{ij} the center of cell (i, j) . The computation in Eq. 89 resembles a convolution of S , but the nonlinearity in Eq. 88 distinguishes it. This expression could also be extended to take into account the directionality of the drones' trajectories by eliminating the isotropy assumption.


At this point, we know where each of the drones has more information to collect based on position, environmental status, and the specific detection law $P_D^{(u)}$. At this point, we can consider several options to compute the actual trajectories.

The most obvious one is to find an algorithm that computes the routes that maximize data collection, that is, the set of an ordered list of (i, j) pairs that overfly the most amount of information. This could be done using a TSP approach and solving it by using an approximate solution. For that, we could define each cell as a node $n \equiv (i, j)$ in a planar graph (it can be drawn in a 2D plane without any intersection). Each of these nodes has a weight $W_n^{(u)}$. Then, the following cost function can be used:

$$F(\{p_n^{(u)}\}) = \sum_n \sum_u p_n^{(u)} W_n^{(u)}, \quad (90)$$

where $p_n^{(u)}$ is a binary variable (free to optimize) that is active if and only if the drone u visits n . For more details, see [110]. This approach is limited because the sweeping width would depend only on the original pixel size and cannot be changed.





7.7. Search-effort unification via an exponential detection law

Since the detection function depends on each individual drone, we could try to make it universal by introducing the search effort density σ . It represents how dense the search must be at each location. Although the geometrical interpretation differs for each drone, the value of σ is shared between all of them. The only assumption is that the functional dependency of P_D with σ is the exact same for all detection laws in the mission. This is something that can be assumed since the vast majority of cases fall under the “exponential law” case, as shown in Eq. 91.

$$P_D(\mathbf{b}x, \sigma(\mathbf{b}x)) = 1 - e^{-\sigma(\mathbf{b}x)} \quad (91)$$

For sensor u with range shaping $f(|\mathbf{b}x - \mathbf{b}x_c|)$, we write:

$$P_D^{(u)}(|\mathbf{b}x - \mathbf{b}x_c|) = f(|\mathbf{b}x - \mathbf{b}x_c|)(1 - e^{-\sigma(\mathbf{b}x)}), \quad (92)$$

which is already independent of u . This allows us to define a unique and global $M(\sigma)$ matrix. Now the problem is to compute the routes. For that, we first need to compute the optimal search effort density σ distribution. Consider that the total amount of information gathered in all cells scanned is depicted in Eq. 93:

$$I[\sigma] = \sum_{i,j} M_{ij}(\sigma_{ij}), \quad (93)$$

which is a function on σ . To solve this, we can take the limit to the continuum and use the Euler-Lagrange equations:

$$I[\sigma] \Delta A = \sum_{i,j} M_{ij}(\sigma_{ij}) \Delta A \stackrel{\Delta A \rightarrow 0}{\approx} \int M(\mathbf{b}x, \sigma(\mathbf{b}x)) dA \quad (94)$$


If we consider that the total amount of search effort $\Phi_T = \sum_{i,j} \sigma_{ij} \Delta A \approx \int \sigma dA$ is limited, the Lagrangian is depicted in Eq. 95:


$$\mathcal{L}(\mathbf{b}x, \sigma(\mathbf{b}x)) = M(\mathbf{b}x, \sigma(\mathbf{b}x)) - \lambda \sigma(\mathbf{b}x), \quad (95)$$

where λ is the Lagrange multiplier for the limited search effort constraint. One important nuance is that $\sigma \geq 0$. To add this constraint, we need to do it explicitly by computing the EL equation for the variable $z = \sigma^2$ and reverting back to this change of variable before solving the actual equation.

7.8. Practical multi-UAV planning implementation

The optimal search effort density σ^* obtained from the optimization is typically non-uniform across the area. We translate this mathematical optimum into practical mission plans through a hierarchical decomposition process:





Priority-based area decomposition: The optimal σ^* distribution is processed to extract contour levels $z_1 > z_2 > \dots > z_n$, defining priority bands as depicted in Eq. 96.

$$R_i = \{c \in \mathcal{D}(A) | z_i \leq \sigma^*(c) \leq z_{i-1}\}, \quad (96)$$

We balance the workload such that each region R_i contains comparable search effort $\phi \approx \Phi_T/n$, ensuring efficient resource distribution while respecting priority ordering.

Multi-UAV coordination and routing: We model the planning problem as a graph where vertices represent subareas and the drone base, with edges carrying travel costs d_{uij} that account for distance, altitude changes, and turning penalties. We employ a Travel Salesman Problem (TSP) formulation [110] with cost function as depicted in Eq. 97

$$F(\{z, \theta\}) = \sum_{u=1}^{n_S} \left[\sum_{i,j \in \mathcal{S}(A)} d_{uij} z_{uij} - K \sum_{k \in \mathcal{S}(A)} \theta_{uk} P_k \right], \quad (97)$$

where travel costs d_{uij} incorporate energy consumption models [111]. This formulation ensures all subareas are visited exactly once by one drone, avoiding redundant coverage while respecting priority ordering and operational constraints.

Coverage path planning with adaptive resolution: Within each priority region R_i , we compute detailed coverage paths that adapt to the priority level:

- Compute average search effort density: $\langle \sigma \rangle_i = \frac{1}{\sum_{c \in R_i} \Delta A} \sum_{c \in R_i} \sigma(c) \Delta A$
- Determine optimal track spacing based on sensor characteristics:

$$W_i = \frac{R_S}{\langle \sigma \rangle_i}, \quad W_i \in [W_{\min}, W_{\max}], \quad (98)$$


where R_S represents the sensor detection range or acuity


- Apply Energy-aware multi-drone Coverage Path Planning (mCPP) algorithms [112] that minimize turns and optimize energy consumption.

7.9. Integration within the TEMA platform

The drone planning technology (PDM-tech-04) is tightly integrated within the broader TEMA technological platform, as demonstrated by the complete processing pipeline shown in Fig. 44 and the resulting trajectory examples in Fig. 45:

- **PDM-tech-05 (Information Fusion):** Provides continuously updated probability maps based on multi-source data fusion, including satellite imagery, ground sensors, and social media analysis.
- **SV-tech-01 (Drone Data Acquisition):** Executes the generated trajectories and streams real-time sensor data back for continuous probability map updates.
- **TFA-tech-05/06 (Detection/Segmentation):** Provides real-time detection capabilities that feed back into the information fusion technology (PDM-tech-05).
- **PDM-tech-01/02 (FireSim/3D hydrodynamic):** Provides real-time simulation results for the monitored ND event that feed back into the information fusion technology (PDM-tech-05).





- **Visualization Systems (SV-tech-04/06/07):** Displays planning results, mission progress, and situational awareness information to operators.




Figure 45. Final optimized drone trajectories demonstrate efficient coverage of priority areas while respecting operational constraints. The paths show systematic coverage of high-priority regions with efficient transitions between subareas and safe separation between multiple UAVs.

For practical deployment, the system uses JSON-based configuration files for mission parameters and REST APIs for integration with UAV autopilot systems and ground control stations. This architecture supports rapid deployment and adaptation to evolving mission requirements.


In operational scenarios, particularly during the early stages of an event when probability maps may be relatively flat ($P \approx 1/2$ over much of the area), the system can use the M matrix (from Equation 89) directly as a proxy for σ to bootstrap the planning pipeline. This approach enables rapid generation of initial reconnaissance missions that prioritize areas near triggered detectors, with higher-certainty alerts receiving immediate attention.


The complete processing pipeline, as illustrated in Fig. 44, transforms probability maps through multiple stages of mathematical optimization and practical constraint handling to generate missions that are both theoretically optimal and operationally feasible, fully leveraging the integrated TEMA technological framework for effective disaster response. The final output, exemplified in Fig. 45, demonstrates the system’s capability to produce coordinated multi-UAV missions that efficiently address the complex requirements of disaster response operations.

7.10. Fulfillment of TEMA KPIs

The following KPIs were linked to drone response planning in the TEMA project:

- OB1: Increase model-based prediction responsiveness/speed for evolving phenomena
- OB2: Increase model-based prediction accuracy for evolving phenomena
- OC1: Reduce latency in NDM
- OC4: Prototype a proof-of-concept TEMA system for NDM in flash floods and regional floods






All KPIs have been fully met for USE in the TEMA project.


Although KPIs OB₁ and OB₂ refer to model-based prediction modeling speed (OB₁) and accuracy (OB₂), drone images play a very relevant role in extracting close and high-quality images and measurements that feed the probabilistic model that is used to speed up and refine these model-based predictions. The drone response planning tool efficiently assigns tasks and trajectories to each available drone in order to optimally complete the NDM mission. Different NDM missions have been considered in the drone response planning, including: surveillance (e.g., for detecting fires as they start), monitoring (e.g., to gather high-quality images of the ND), and search (e.g., to detect and locate people at risk).

Drone response planning has been implemented by using efficient and scalable algorithms that strongly help reduce the response time, helping to fulfill OB₁. In addition, drone response planning enables fulfilling complex missions that enable gathering high-quality images from a short distance of the ND, enabling the extraction of accurate ND measurements by other components of the TEMA architecture and their integration through the **Information Fusion** component to keep updated the actual status of the ND, which is crucial for the accuracy of NDM predictions.

Although KPI OC₁ refers mainly to satellite-based NDM, drone images play a very relevant role in ND monitoring and management in general. Drones' response planning is flexible and can be configured to cover an area of interest with a very high frequency (e.g., one coverage mission every hour). This can strongly reduce latency in NDM missions such as surveillance, ND monitoring, people search missions, or wildfire burnt area measurement. This high frequency fulfills the Target Values (TV) specified in OC₁. Of course, this depends on the size of the area of interest, and a suitable dimensioning of the number of drones in the fleet should be performed to take into account the size of the specific area of interest affected by the ND.

Finally, KPI OC₄ is also fully met for USE in the TEMA project. The drone response planning technology (PDM-tech-04) is integrated in the TEMA architecture and has been used in the TEMA trials.






8. Conclusion


Deliverable D4.3 “Report on Phenomenon Prediction and Response Planning” is the third deliverable of WP4 of the TEMA project. The main objective of TEMA WP4 “Prediction and decision-making” (M13-M36) is to exploit the outputs of WP3 analysis methods to enrich and accelerate response to emergencies/disasters, via novel approaches to fast modeling engines for the TEMA use-cases, automated response planning for making optimal sensor placement recommendations, and decision support services for remote sensing. This document reports the final research results of Tasks T4.1 “Precise phenomenon prediction” and T4.2 “Response planning and recommendations” carried out between M19-M36. These developments are interlinked with WP3, WP5, and WP6. This document serves as a summary of the main research outputs and serves as a reference point for researchers summarizing the technical challenges.

Major technical advancements achieved during this reporting period include:

In Phenomenon Prediction (task T4.1):

- **Wildfire Smoke Dynamics Modeling (DLR-KN):** Development of a novel sparse Bayesian learning (SBL) framework for gas/smoke source localization using stationary advection-diffusion equations with Green’s function solutions. The method enables grid-less source localization with arbitrary spatial placement, distributed multi-agent implementation, and theoretical performance characterization through Cramer-Rao bounds. Experimental validation was conducted during the Montiferru campaign in Sardinia (June 2025). With the research and development performed in M19-M36, the research in the field of wildfire smoke dynamics modeling has been completed, fulfilling the specified TEMA objectives and assigned KPIs.
- **Wildfire Behavior Modeling (TYSL):** Enhancement of the FireSim simulator with real-time data assimilation capabilities. Implemented integration of field sensor measurements (PDM-tech-03) to improve weather forecasting models using inverse distance weighting, and developed calibration methods for fire spread using actual fire front locations extracted from fire probability maps (Maps4Fire from the Information Fusion component, PDM-tech-05). Validation against the historical Montiferru fire demonstrated significant accuracy improvements. With the research and development performed in M19-M36, the research in the field of wildfire behavior modeling has been completed, fulfilling the specified TEMA objectives and assigned KPIs.
- **Near Real-time Flood Modeling (NST):** Research and development of the “Clone Cells” method within the 3Di hydrodynamic model to accurately represent disconnected flow regions with minimal computational overhead. This approach allows modeling of subgrid-scale flow separations (e.g., behind barriers or in bifurcations) without requiring fine grid resolution, enabling faster and more accurate flood predictions. With the research and development performed in M19-M36, the research in the field of near real-time flood modeling has been completed, fulfilling the specified TEMA objectives and assigned KPIs.
- **Disaster Prediction from User-Generated Data (IT:U, linked to Tasks T3.3 and T4.3):** Created a multimodal prediction framework integrating Bluesky social media data, GDELT news streams, and weather observations. Leveraged foundation models with automatic prompt optimization (APO) for geospatial reasoning, achieving improved performance over statistical hotspot detection methods in both the 2024 Central Europe floods and 2025 Southern California wildfires. With the research and development performed in M19-M36, the research in the field of disaster prediction from user-generated data has been com-







pleted, fulfilling the specified TEMA objectives and assigned KPIs.

In Response Planning (task T4.2):

- **Optimal Drone Routing (USE):** Developed mathematical frameworks for UAV path planning based on probability maps of disaster phenomena. Formulated the problem using information theory (normalized entropy) to maximize information gain, with approaches including Traveling Salesman Problem (TSP) formulations and search effort density optimization for determining optimal reconnaissance trajectories. The method has been developed and integrated into the TEMA architecture as technology PDM-tech-04. With this research and development, the research in the field of drone response planning has been completed, fulfilling the specified TEMA objectives and assigned KPIs.


In addition, all developed technologies have been successfully integrated into the TEMA platform and have demonstrated fulfillment of relevant TEMA Key Performance Indicators (KPIs), including improvements in prediction accuracy (Objective OA₂), responsiveness (OA₃), and situational awareness (OC₂). Validation through historical case studies (Montiferru fire, Ahrtal and Mantoudi floods) and the Sardinia experimental campaign (June 2025) confirmed the operational readiness and practical effectiveness of the developed approaches for natural disaster management. These research advances provide significant capabilities for precise phenomenon prediction and optimal response planning in wildfire and flood emergency scenarios, representing substantial progress beyond state-of-the-art approaches in real-time natural disaster management systems.







References


- [1] A. L. Westerling et al., “Warming and earlier spring increase western us forest wildfire activity,” *Science*, vol. 313, no. 5789, pp. 940–943, 2006.
- [2] J. T. Abatzoglou and A. P. Williams, “Impact of anthropogenic climate change on wildfire across western us forests,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 42, pp. 11770–11775, 2016.
- [3] C. B. Clements et al., “Observing the dynamics of wildland grass fires: Firefluxa field validation experiment,” *Bulletin of the American Meteorological Society*, vol. 87, no. 11, pp. 1369–1382, 2006.
- [4] W. Khalaf, C. Pace, M. Gaudioso, et al., “Gas detection via machine learning,” *Int. J. Comput. Electr. Autom. Control Inf. Eng.*, vol. 2, no. 1, pp. 61–65, 2008.
- [5] M. Reggente and A. J. Lilienthal, “Using local wind information for gas distribution mapping in outdoor environments with a mobile robot,” in *SENSORS, 2009 IEEE*, pp. 1715–1720, 2009.
- [6] D. Badawi, I. Bassi, S. Ozev, and A. E. Cetin, “Deep-learning-based gas leak source localization from sparse sensor data,” *IEEE Sensors Journal*, vol. 22, no. 21, pp. 20999–21008, 2022.
- [7] A. S. A. Yeon, A. Zakaria, S. M. M. S. Zakaria, R. Visvanathan, K. Kamarudin, and L. M. Kamarudin, “Gas source localization via mobile robot with gas distribution mapping and deep neural network,” in *2022 2nd International Conference on Electronic and Electrical Engineering and Intelligent System (ICE3IS)*, pp. 120–124, 2022.
- [8] H. Kim, M. Park, C. W. Kim, and D. Shin, “Source localization for hazardous material release in an outdoor chemical plant via a combination of lstm-rnn and cfd simulation,” *Computers & Chemical Engineering*, vol. 125, pp. 476–489, 2019.
- [9] G. Strang, *Computational Science and Engineering*. Wellesley-Cambridge Press, 2007.
- [10] A. Francis, S. Li, C. Griffiths, and J. Sienz, “Gas source localization and mapping with mobile robots: A review,” *Journal of Field Robotics*, pp. 1341–1373, 2022.
- [11] H. Ishida, Y. Wada, and H. Matsukura, “Chemical sensing in robotic applications: A review,” *IEEE Sensors Journal*, vol. 12, no. 11, pp. 3163–3173, 2012.
- [12] A. Lilienthal and T. Duckett, “Experimental analysis of gas-sensitive Braitenberg vehicles,” *Advanced Robotics*, vol. 18, no. 8, pp. 817–834, 2004.
- [13] A. Rutkowski, S. Edwards, M. Willis, R. Quinn, and G. Causey, “A robotic platform for testing moth-inspired plume tracking strategies,” in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, vol. 4, pp. 3319–3324, 2004.
- [14] X. xing Chen and J. Huang, “Odor source localization algorithms on mobile robots: A review and future outlook,” *Robotics and Autonomous Systems*, vol. 112, pp. 123–136, 2019.
- [15] S. Li, Y. Guo, and B. Bingham, “Multi-robot cooperative control for monitoring and tracking dynamic plumes,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 67–73, 2014.







- [16] X. Jiang and S. Li, "Plume front tracking in unknown environments by estimation and control," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 2, pp. 911–921, 2019.
- [17] X. Jiang, S. Li, B. Luo, and Q. Meng, "Source exploration for an under-actuated system: A control-theoretic paradigm," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 3, pp. 1100–1107, 2020.
- [18] Z. Shen, Z. He, S. Li, Q. Wang, and Z. Shao, "A multi-quadcopter cooperative cyber-physical system for timely air pollution localization," *ACM Trans. Embed. Comput. Syst.*, vol. 16, apr 2017.
- [19] M. Hutchinson, H. Oh, and W.-H. Chen, "A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors," *Information Fusion*, vol. 36, pp. 130–148, 2017.
- [20] J. R. Bourne, E. R. Pardyjak, and K. K. Leang, "Coordinated bayesian-based bioinspired plume source term estimation and source seeking for mobile robots," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 967–986, 2019.
- [21] M. Hutchinson, C. Liu, and W.-H. Chen, "Source term estimation of a hazardous airborne release using an unmanned aerial vehicle," *Journal of Field Robotics*, vol. 36, no. 4, pp. 797–817, 2019.
- [22] B. Ristic, D. Anglely, B. Moran, and J. L. Palmer, "Autonomous multi-robot search for a hazardous source in a turbulent environment," *Sensors*, vol. 17, no. 4, 2017.
- [23] B. Ristic, A. Skvortsov, and A. Walker, "Autonomous search for a diffusive source in an unknown structured environment," *Entropy*, vol. 16, no. 2, pp. 789–813, 2014.
- [24] C. Rhodes, C. Liu, P. Westoby, and W.-H. Chen, "Autonomous search of an airborne release in urban environments using informed tree planning," *Autonomous Robots*, vol. 47, pp. 1–18, Jan. 2023.
- [25] M. Vergassola, E. Villermaux, and B. I. Shraiman, "Infotaxis as a strategy for searching without gradients," *Nature*, vol. 445, pp. 406–409, 2007.
- [26] X.-Y. Dai, J.-Y. Wang, and Q.-H. Meng, "An infotaxis-based odor source searching strategy for a mobile robot equipped with a tdlas gas sensor," in *2019 Chinese Control Conference (CCC)*, pp. 4492–4497, 2019.
- [27] M. Hutchinson, H. Oh, and W.-H. Chen, "Entrotaxis as a strategy for autonomous search and source reconstruction in turbulent conditions," *Information Fusion*, vol. 42, pp. 179–189, 2018.
- [28] T. Wiedemann, C. Manss, D. Shutin, A. J. Lilienthal, V. Karolj, and A. Viseras, "Probabilistic modeling of gas diffusion with partial differential equations for multi-robot exploration and gas source localization," in *2017 European Conference on Mobile Robots (ECMR)*, pp. 1–7, 2017.
- [29] T. Wiedemann, C. Manss, and D. Shutin, "Multi-agent exploration of spatial dynamical processes under sparsity constraints," *Autonomous Agents and Multi-Agent Systems*, July 2017.
- [30] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *J. Machine Learning Res.*, vol. 1, pp. 211–244, June 2001.







- [31] D. Wipf, J. Palmer, and B. Rao, "Perspectives on sparse Bayesian learning," in *Proceedings of the 16th International Conference on Neural Information Processing Systems, NIPS'03*, (Cambridge, MA, USA), pp. 249–256, MIT Press, 2003.
- [32] T. Wiedemann, A. Lilienthal, and D. Shutin, "Analysis of model mismatch effects for a model-based gas source localization strategy incorporating advection knowledge," *Sensors*, Jan. 2019.
- [33] T. Wiedemann, D. Shutin, and A. Lilienthal, "Model-based gas source localization strategy for a cooperative multi-robot system - a probabilistic approach and experimental validation incorporating physical knowledge and model uncertainties," *Robotics and Autonomous Systems*, 2019.
- [34] T. Wiedemann, D. Shutin, and A. Lilienthal, "Experimental validation of domain knowledge assisted robotic exploration and source localization," in *IEEE International Conference on Autonomous Systems (ICAS)*, Oct. 2021.
- [35] T. Wiedemann, *Domain Knowledge Assisted Robotic Exploration and Source Localization*. PhD thesis, Örebro University, School of Science and Technology, 2020.
- [36] J. Nocedal and S. Wright, *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering, Springer New York, 2006.
- [37] P.-J. Chung, "A max-search approach for DOA estimation with unknown number of signals," *IEEE J. Sel. Topics in Signal Process.*, vol. 4, June 2010.
- [38] B.-S. Shin and D. Shutin, "Joint distributed traveltime and full waveform tomography for enhanced subsurface imaging in seismic networks," *IEEE Transactions on Computational Imaging*, vol. 10, pp. 600–612, 2024.
- [39] G. Korn and T. Korn, *Mathematical Handbook for Scientists and Engineers: Definitions, Theorems, and Formulas for Reference and Review*. Dover Civil and Mechanical Engineering Series, Dover Publications, 2000.
- [40] P. Kythe, *Green's Functions and Linear Differential Equations: Theory, Applications, and Computation*. Chapman & Hall/CRC Applied Mathematics & Nonlinear Science, CRC Press, 2011.
- [41] V. S. P. Ruiz, "Neural Network Surrogate Models for the Localization of Sources in Semilinear Elliptic PDEs," Master's thesis, TUM School of Computation, Information and Technology of the Technical University of Munich, Germany, 2023.
- [42] M. J. Beal, *Variational Algorithm for Approximate Bayesian Inference*. PhD thesis, University College London, 2003.
- [43] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York: Springer, 2006.
- [44] R. M. Neal and G. E. Hinton, *A view of the EM algorithm that justifies incremental, sparse, and other variants*, ch. III, pp. 355–366. Kluwer Academic Publishers, 1998.
- [45] M. E. Tipping and A. C. Faul, "Fast marginal likelihood maximisation for sparse Bayesian models," in *Proc. 9th Int. Workshop Artificial Intelligence and Statistics*, (Key West, FL, USA), Jan. 2003.







- [46] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. of the Royal Statist. Soc. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977.
- [47] R.-E. Plessix, "A review of the adjoint-state method for computing the gradient of a functional with geophysical applications," *Geophysical Journal International*, vol. 167, pp. 495–503, Nov. 2006.
- [48] S. Zhang, R. Pöhlmann, T. Wiedemann, A. Dammann, H. Wymeersch, and P. A. Hoeher, "Self-aware swarm navigation in autonomous exploration missions," *Proceedings of the IEEE*, vol. 108, no. 7, pp. 1168–1195, 2020.
- [49] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.
- [50] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD: The Johns Hopkins University Press, 3rd ed., Oct. 1996.
- [51] V. S. P. Ruiz, D. Shutin, T. Wiedemann, and P. Hinsin, "Physics-guided neural networks for distributed sparse gas source localization using poissons equation and greens function method," in *Proc. 32nd European Sig. Process. Conf. (EUSIPCO)*, (Lyon, France), August 2024. to appear.
- [52] V. P. Ruiz, P. Hinsin, T. Wiedemann, D. Shutin, and C. Christof, "Gas source localization using physics-guided neural networks," in *2024 IEEE International Symposium on Olfaction and Electronic Nose (ISOEN)*, pp. 1–3, 2024.
- [53] R. E. Hunt, "Poisson's equation, chapter 2." Web, 2002. Lecture notes.
- [54] D. Shutin, T. Wiedemann, and P. Hinsin, "Detection and estimation of gas sources with arbitrary locations based on poisson's equation," *IEEE Open Journal of Signal Processing*, vol. 5, pp. 359–373, 2024.
- [55] D. Shutin, T. Wiedemann, and P. Hinsin, "Distributed superresolution gas source localization based on poisson equation," in *2023 IEEE 9th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 371–375, 2023.
- [56] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.
- [57] L. Xiao, S. Boyd, and S. Lall, "A scheme for robust distributed sensor fusion based on average consensus," in *IPSN 2005. Fourth International Symposium on Information Processing in Sensor Networks, 2005.*, pp. 63–70, IEEE, 2005.
- [58] C. Manss, D. Shutin, and G. Leus, "Consensus based distributed sparse bayesian learning by fast marginal likelihood maximization," *IEEE Signal Processing Letters*, vol. 27, pp. 2119–2123, 2020.
- [59] D. Needell, J. Tropp, and R. Vershynin, "Greedy signal recovery review," in *2008 42nd Asilomar Conference on Signals, Systems and Computers*, pp. 1048–1050, 2008.
- [60] Y. Rubner, C. Tomasi, and L. Guibas, "A metric for distributions with applications to image databases," in *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pp. 59–66, 1998.







- [61] C. Villani, *Optimal Transport: Old and New*, vol. 338 of *Grundlehren der mathematischen Wissenschaften*. Berlin, Heidelberg: Springer-Verlag, 2009.
- [62] H. Sagan, *Space-Filling Curves*. Universitext, New York, NY: Springer, 1994.
- [63] H. Vogel, "A better way to construct the sunflower head," *Mathematical Biosciences*, vol. 44, no. 3, pp. 179–189, 1979.
- [64] D. Shutin, T. Buchgraber, S. R. Kulkarni, and H. V. Poor, "Fast variational sparse Bayesian learning with automatic relevance determination for superimposed signals," *IEEE Trans. Signal Process.*, vol. 59, pp. 6257–6261, Dec. 2011.
- [65] D. Shutin and B. H. Fleury, "Sparse variational Bayesian SAGE algorithm with application to the estimation of multipath wireless channels," *IEEE Trans. Signal Process.*, vol. 59, pp. 3609–3623, Aug. 2011.
- [66] A. Filip-Dhaubhadel and D. Shutin, "Long coherent integration in passive radar systems using super-resolution sparse bayesian learning," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 1, pp. 554–572, 2021.
- [67] D. Shutin, V. S. Prieto Ruiz, T. Wiedemann, P. Hinsén, and A. Dammann, "Cramer-rao lower bound for a gas source localization based on poisson's equation," in *2024 IEEE Workshop on Signal Processing Systems (SiPS)*, pp. 77–82, 2024.
- [68] S. Cheng, I. C. Prentice, Y. Huang, Y. Jin, Y.-K. Guo, and R. Arcucci, "Data-driven surrogate model with latent data assimilation: Application to wildfire forecasting," *Journal of Computational Physics*, vol. 464, 2022.
- [69] T. Wu, Q. Zhang, J. Zhu, L. Xu, and Y. Zhang, "Forest fire spread prediction and assimilation using the deterministic ensemble kalman filter," *Fire Technology*, vol. 61, pp. 2467–2492, 2025.
- [70] V. Casulli, "A high-resolution wetting and drying algorithm for free-surface hydrodynamics," *International Journal for Numerical Methods in Fluids*, vol. 60, pp. 391–408, 2009.
- [71] N. D. Volp, B. C. V. Prooijen, and G. S. Stelling, "A finite volume approach for shallow water flow accounting for high-resolution bathymetry and roughness data," *Water Resources Research*, vol. 49, pp. 4126–4135, 2013.
- [72] G. S. Stelling, "Quadtree flood simulations with sub-grid digital elevation models," *Proceedings of the Institution of Civil Engineers - Water Management*, vol. 165, pp. 567–580, 2012.
- [73] V. Casulli, "Computational grid, subgrid, and pixels," *Numerical Methods in Fluids*, vol. 90, pp. 140–155, 2019.
- [74] B. Stollberg and T. de Groeve, "The use of social media within the global disaster alert and coordination system (GDACS)," in *Proceedings of the 21st International Conference on World Wide Web, WWW '12 Companion*, (New York, NY, USA), pp. 703–706, Association for Computing Machinery, Apr. 2012.
- [75] S. A. Shah, S. B. Yahia, K. McBride, A. Jamil, and D. Draheim, "Twitter streaming data analytics for disaster alerts," in *2021 2nd International Informatics and Software Engineering Conference (IISEC)*, (Ankara, Turkey), pp. 1–6, IEEE, 2021.







- [76] C. Havas and B. Resch, “Portability of semantic and spatial–temporal machine learning methods to analyse social media for near-real-time disaster monitoring,” *Natural Hazards*, vol. 108, pp. 2939–2969, Sept. 2021.
- [77] X. Zhou and L. Chen, “Event detection over twitter social media streams,” *The VLDB Journal*, vol. 23, pp. 381–400, June 2014.
- [78] J. A. De Bruijn, H. De Moel, B. Jongman, M. C. De Ruiter, J. Wagemaker, and J. C. J. H. Aerts, “A global database of historic and real-time flood events based on social media,” *Scientific Data*, vol. 6, p. 311, Dec. 2019.
- [79] C. V. Pennington, R. Bossu, F. Ofli, M. Imran, U. Qazi, J. Roch, and V. J. Banks, “A near-real-time global landslide incident reporting tool demonstrator using social media and artificial intelligence,” *International Journal of Disaster Risk Reduction*, vol. 77, p. 103089, July 2022.
- [80] S. Schmidt, M. Friedemann, D. Hanny, B. Resch, T. Riedlinger, and M. Mühlbauer, “Enhancing satellite-based emergency mapping: Identifying wildfires through geo-social media analysis,” *Big Earth Data*, vol. 0, no. 0, pp. 1–23, 2025.
- [81] J. C. Pinto, H. Gonalo Oliveira, A. Cardoso, and C. Silva, “Generating Wildfire Heat Maps with Twitter and BERT,” in *Intelligent Data Engineering and Automated Learning – IDEAL 2023* (P. Quaresma, D. Camacho, H. Yin, T. Gonalves, V. Julian, and A. J. Tall3n-Ballesteros, eds.), vol. 14404, pp. 82–94, Cham: Springer Nature Switzerland, 2023.
- [82] J. Li, Y. Wang, and W. Li, “MGMP: Multimodal Graph Message Propagation Network for Event Detection,” in *MultiMedia Modeling* (B. P3r J3nsson, C. Gurrin, M.-T. Tran, D.-T. Dang-Nguyen, A. M.-C. Hu, B. Huynh Thi Thanh, and B. Huet, eds.), (Cham), pp. 141–153, Springer International Publishing, 2022.
- [83] M.-A. Kaufhold, M. Bayer, and C. Reuter, “Rapid relevance classification of social media posts in disasters and emergencies: A system and evaluation featuring active, incremental and online learning,” *Information Processing & Management*, vol. 57, p. 102132, Jan. 2020.
- [84] C. Scheele, M. Yu, and Q. Huang, “Geographic context-aware text mining: Enhance social media message classification for situational awareness by integrating spatial and temporal features,” *International Journal of Digital Earth*, vol. 14, pp. 1721–1743, Nov. 2021.
- [85] D. Hanny and B. Resch, “Multimodal Geo-Information Extraction from Social Media for Supporting Decision-Making in Disaster Management,” *AGILE: GIScience Series*, vol. 5, pp. 1–8, May 2024.
- [86] I. Owuor, H. H. Hochmair, and S. Cvetojevic, “Tracking Hurricane Dorian in GDELT and Twitter,” *AGILE: GIScience Series*, vol. 1, pp. 1–18, July 2020.
- [87] A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei, H. Lin, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Lin, K. Dang, K. Lu, K. Bao, K. Yang, L. Yu, M. Li, M. Xue, P. Zhang, Q. Zhu, R. Men, R. Lin, T. Li, T. Tang, T. Xia, X. Ren, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Wan, Y. Liu, Z. Cui, Z. Zhang, and Z. Qiu, “Qwen2.5 Technical Report,” Jan. 2025.
- [88] R. Murtfeldt, N. Alterman, I. Kahveci, and J. D. West, “RIP Twitter API: A eulogy to its vast research contributions,” 2024.
- [89] Uber, “H3: A hexagonal hierarchical geospatial indexing system,” 2023.







- [90] A. Getis and J. K. Ord, "The analysis of spatial association by use of distance statistics," *Geographical Analysis*, vol. 24, no. 3, pp. 189–206, 1992.
- [91] D. Erdos, A. Erdos, and S. E. Watkins, "An experimental UAV system for search and rescue challenge," *IEEE Aerospace and Electronic Systems Magazine*, vol. 28, no. 5, pp. 32–37, 2013.
- [92] J. Scherer, S. Yahyanejad, S. Hayat, E. Yanmaz, T. Andre, A. Khan, V. Vukadinovic, C. Bettstetter, H. Hellwagner, and B. Rinner, "An autonomous multi-UAV system for search and rescue," in *Proceedings of the First Workshop on Micro Aerial Vehicle Networks, Systems, and Applications for Civilian Use, DroNet '15*, (New York, NY, USA), p. 3338, Association for Computing Machinery, 2015.
- [93] E. T. Alotaibi, S. S. Alqefari, and A. Koubaa, "LSAR: Multi-UAV Collaboration for Search and Rescue Missions," *IEEE Access*, vol. 7, pp. 55817–55832, 2019.
- [94] I. Maza, F. Caballero, J. Capitán, J. R. Martínez-De-Dios, and A. Ollero, "Experimental results in multi-uav coordination for disaster management and civil security applications," *JOURNAL OF INTELLIGENT & ROBOTIC SYSTEMS*, vol. 60, no. 2, pp. 195–210, 2011.
- [95] J. J. Acevedo, B. n. C. Arrue, I. Maza, and A. Ollero, "Cooperative large area surveillance with a team of aerial mobile robots for long endurance missions," *J. Intell. Robotics Syst.*, vol. 70, no. 14, p. 329345, 2013.
- [96] B. O. Koopman, "The Theory of Search. I. Kinematic Bases," *Operations Research*, vol. 4, no. 3, pp. 324–346, 1956.
- [97] B. O. Koopman, "The Theory of Search. II. Target Detection," <https://doi.org/10.1287/opre.4.5.503>, vol. 4, pp. 503–531, 10 1956.
- [98] B. O. Koopman, "The Theory of Search. III. The Optimum Distribution of Searching Effort," <https://doi.org/10.1287/opre.5.5.613>, vol. 5, pp. 613–626, 10 1957.
- [99] P. M. Morse, "In Memoriam: Bernard Osgood Koopman, 1900-1981," *Operations Research*, vol. 30, no. 3, pp. viii–427, 1982.
- [100] J. de Guenin, "Optimum distribution of effort: An extension of the Koopman basic theory," *Operations Research*, vol. 9, no. 1, pp. 1–7, 1961.
- [101] J. H. Discenza and L. D. Stone, "Optimal survivor search with multiple states," *Operations Research*, vol. 29, no. 2, pp. 309–323, 1981.
- [102] J. Frost and L. Stone, "Review of Search Theory: Advances and Applications to Search and Rescue Decision Support," Tech. Rep. 00929592, United States Coast Guard, September 2001.
- [103] L. D. Stone, J. O. Royset, and A. R. Washburn, *Optimal Search for Moving Targets*, vol. 237. Springer International Publishing, 2016.
- [104] S. I. Gass and M. C. Fu, "Encyclopedia of operations research and management science," 2013.
- [105] H. R. Richardson, *Operations Analysis*. U.S. Navy, 1967.
- [106] T. M. Kratzke, L. D. Stone, and J. R. Frost, "Search and rescue optimal planning system," in *2010 13th International Conference on Information Fusion*, pp. 1–8, 2010.







- [107] L. D. Stone, C. Keller, T. L. Kratzke, and J. Strumpfer, "Search Analysis for the Location of the AF447 Underwater Wreckage," tech. rep., Metron Scientific Solutions, January 2011.
- [108] L. Yang, X. Zhang, Z. Li, L. Li, and Y. Shi, "A LODBO algorithm for multi-uav search and rescue path planning in disaster areas," *Chinese Journal of Aeronautics*, vol. 38, no. 2, p. 103301, 2025.
- [109] H. Wen, Y. Shi, S. Wang, T. Chen, P. Di, and L. Yang, "Route planning for UAVs maritime search and rescue considering the targets moving situation," *Ocean Engineering*, vol. 310, p. 118623, 2024.
- [110] A. Sojo, I. Maza, and A. Ollero, "Optimal planning for heterogeneous autonomous teams with precedence and compatibility constraints and its application on power grid inspection with Unmanned Aerial Vehicles," 2024.
- [111] L. Bauersfeld and D. Scaramuzza, "Range, endurance, and optimal speed estimates for multicopters," *IEEE Robotics and Automation Letters*, vol. 7, pp. 2953–2960, April 2022.
- [112] D. Datsko, F. Nekovar, R. Penicka, and M. Saska, "Energy-aware multi-UAV coverage mission planning with optimal speed of flight," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2893–2900, 2024.
- [113] S. Schmidt, E. Díaz Fragachan, D. Arifi, D. Hanny, and B. Resch, "Assessing the spatial accuracy of geocoding flood-related imagery using vision language models," *Spatial Information Research*, vol. 33, no. 2, p. 15, 2025.
- [114] C. Zorenböhmer, S. Schmidt, and B. Resch, "Emograce: Aspect-based emotion analysis for social media data," 2025.
- [115] M. Wieland, V. Hertel, C. Geiss, S. Martinis, and K. Lechner, "Toward Scalable Damage Assessment for Rapid Disaster Response," in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, pp. 3902–3905, July 2024. ISSN: 2153-7003.
- [116] M. Wieland, S. Schmidt, B. Resch, A. Abecker, and S. Martinis, "Fusion of geospatial information from remote sensing and social media to prioritise rapid response actions in case of floods," *Natural Hazards*, Jan. 2025.
- [117] C. Gilga, C. Hochwarter, L. Knoche, S. Schmidt, G. Ringler, M. Wieland, B. Resch, and B. Wagner, "Legal and ethical considerations for demand-driven data collection and AI-based analysis in flood response," *International Journal of Disaster Risk Reduction*, vol. 122, p. 105441, May 2025.
- [118] S. Dujardin, D. Arifi, S. Schmidt, C. Linard, and B. Resch, "Tracing online flood conversations across borders: A watershed level analysis of geo-social media topics during the 2021 european flood," 2024.
- [119] D. Hanny and B. Resch, "Multimodal geoai: An integrated spatio-temporal topic-sentiment model for the analysis of geo-social media posts for disaster management," *International Journal of Applied Earth Observation and Geoinformation*, vol. 139, p. 104540, 2025.
- [120] A. Sojo, A. Perea, M. Castell, J. Perrela, I. Maza, A. Caballero, and A. Ollero, "Multi-uav planning in search and rescue missions using optimal search effort allocation," in *2025 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 649–656, IEEE, 2025.





- [121] D. Shutin, V. S. Prieto Ruiz, and P. Hinsin, "Sparse bayesian learning for decentralized 2d smoke source localization," in *2025 IEEE International Conference on Image Processing (ICIP 2025)*, IEEE, September 2025.
- [122] D. Shutin, V. S. Prieto Ruiz, P. Hinsin, and T. Wiedemann, "Multi-agent adaptive super-resolution sparse source localization for advection-diffusion processes," in *2025 IEEE 10th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2025. to appear.





Annex A: Related publications and technical reports

No.	Title	Reference	Type
1	Enhancing satellite-based emergency mapping: Identifying wildfires through geo-social media analysis	[80]	Article
2	Assessing the spatial accuracy of geocoding flood-related imagery using Vision Language Models	[113]	Article
3	EmoGRACE: Aspect-Based Emotion Analysis for Social Media Data	[114]	Preprint
4	Toward Scalable Damage Assessment for Rapid Disaster Response	[115]	Article
5	Fusion of geospatial information from remote sensing and social media to prioritise rapid response actions in case of floods	[116]	Article
6	Legal and ethical considerations for demand-driven data collection and AI-based analysis in flood response	[117]	Article
7	Tracing online flood conversations across borders: A watershed level analysis of geo-social media topics during the 2021 European flood	[118]	Preprint
8	Multimodal GeoAI: An integrated spatio-temporal topic-sentiment model for the analysis of geo-social media posts for disaster management	[119]	Article
9	Multi-UAV Planning in Search and Rescue Missions using Optimal Search Effort Allocation	[120]	Article
10	Distributed Superresolution Gas Source Localization Based on Poisson Equation	[55]	Article
11	Detection and Estimation of Gas Sources With Arbitrary Locations Based on Poisson's Equation	[54]	Article
12	Cramer-Rao Lower Bound for a Gas Source Localization Based on Poisson's Equation	[67]	Article
13	Gas Source Localization Using Physics-Guided Neural Networks	[52]	Article
14	Physics-Guided Neural Networks for Distributed Sparse Gas Source Localization Using Poissons Equation and Greens Function Method	[51]	Article
15	Sparse Bayesian Learning for Decentralized 2D Smoke Source Localization	[121]	Article
16	Multi-Agent Adaptive Super-resolution Sparse Source Localization for Advection-Diffusion Processes	[122]	Article, to appear

