Annual Reviews in Control 61 (2026) 101054

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

Full Length Article

A perspective and review on adaptive mission planning and hazard monitoring during emergencies with autonomous unmanned systems

Daniele Giovanni Gioia [a], Katharina Glock [b], Jens Kahlen [a],*

[a] German Aerospace Center (DLR), Institute for the Protection of Terrestrial Infrastructures, Sankt Augustin, Germany
[b] FZI Research Center for Information Technology, Karlsruhe, Germany

ARTICLE INFO

Keywords:
Unmanned systems
Information gathering
Situation awareness
Inverse problems
Hazard monitoring

ABSTRACT

Environmental emergency scenarios often involve dynamic hazards whose progression is difficult to predict due to inherent uncertainties and time dependencies. Although mathematical models exist to approximate these dynamics, they are often incomplete or imprecise. Consequently, real-time data are crucial for refining understanding of the situation, enhancing situation awareness and improving response strategies.

Recent advances in autonomous unmanned systems, such as drones or tracked robots, have enabled rapid, on-the-ground data collection. These systems often offer cost-effective and flexible means to continuously update models, allowing for adaptive monitoring that can track evolving hazards. However, the variety of systems, application domains, and methodological approaches spread across various research domains presents a challenge for future research.

This work provides a perspective and structured synthesis of techniques for integrating in situ observations into adaptive models for situation awareness, framed through an overarching sequential decision-making view and a control hierarchy for autonomous unmanned systems. Specifically, it focuses on recurring design principles in common emergency application contexts such as wildfires, CBRN, and flood and storm systems, and in adaptive information gathering and data fusion methods. Common principles that have emerged largely independently across domains are highlighted and key challenges that originate in this heterogeneity of approaches are identified. Particular attention is given to their transferability to real-world applications. Therefore, this study organizes fragmented research fields, offering a unifying set of concepts, limitations, and future perspectives for adaptive monitoring of dynamic environmental hazards with unmanned systems in emergency scenarios.

1. Introduction

In environmental emergencies characterized by, e.g., wildfires, industrial accidents, or floods, responders often face significant challenges due to a limited understanding of the affected area, infrastructure, and availability of critical resources. This lack of information can impede effective response efforts, as decision makers may not fully comprehend the current situation, thus complicating interventions like evacuations, medical assistance, or containment measures. Therefore, one urgent task during the response phase is to gather information to improve *situation awareness*, providing the best possible instruments for effective decision making of emergency response units.

A first step for situation awareness usually concerns a data-driven static description of an affected area, thus requiring a *mapping* phase. However, some emergencies such as chemical leakages, floods, and hurricanes, are generated by inherently dynamic hazards that evolve over time. Dealing with the explicit time dependence and dynamic representations, the nomenclature typically changes to *tracking* and *monitoring* rather than simple mapping, thus monitoring shifting conditions, such as flood progression or wildfire spread. Further compounding time dependencies to already extensively affected areas, monitoring activities often necessitate considerable resources, thus requiring efficient decision-making to allocate them.

For many environmental emergencies, the dynamics of the hazard follow reasonably understood physical principles. Common examples are the fluid dynamics governing atmospheric dispersion of contaminants or wildfires burning according to vegetation and weather. In principle, a priori knowledge of these underlying processes allows for models and simulations, providing valuable insights. In fact, by simulating the potential paths and intensities of an event, decision-makers can provide better responses, anticipate risks, and plan interventions.

---

* Corresponding author.
E-mail addresses: daniele.gioia@dlr.de (D.G. Gioia), kglock@fzi.de (K. Glock), jens.kahlen@dlr.de (J. Kahlen).

https://doi.org/10.1016/j.arcontrol.2026.101054
Received 22 August 2025; Received in revised form 2 February 2026; Accepted 27 April 2026

Having a clear overview of the present situation and, ideally, being able to generate quick but reliable forecasts can be the difference between containing the event early and allowing it to escalate.

Nevertheless, the reliability of a priori knowledge is usually limited. Dynamics are sensitive to, e.g., wind direction, temperature, and local obstacles. Relying on a single initial guess on the boundary and initial conditions can be misleading. For this reason, information gathered in situ and, preferably, in real time can be merged into the existing model, improving the reliability. By leveraging spatial and temporal correlations, observations may help estimate the distribution of hazardous factors and predict their future development, *adapting* the model whenever real measurements become available.

In general, observations from satellites are a precious source for accurate and comprehensive information gathering. For emergency response, well-known public international programs and commercial companies can provide satellite-based rapid mapping services with varying resolution, also enhancing the product with additional data (Denis et al., 2016). The most prominent example for this is the mapping service of the *Emergency Management Service* (EMS) of Copernicus (Copernicus, 2025). Examples of growing commercial companies are Maxar, active during the 2025 fires in California (Maxar Technologies, 2025), or BlackSky (BlackSky, 2023). In spite of that, even if, together with a continuously increasing resolution, satellite revisit frequencies are improving, aiming for *near* real-time monitoring, observations are subject to the limitations of such rigid technologies. Practical examples are coordination for complex access procedures, cost-effectiveness, communication latency, and the impossibility of immediate data collection.

Possible alternatives are unmanned systems. In fact, during emergencies, they can serve multiple roles such as search and rescue, rubble removal, structural inspection, and many others. A history of real-world applications during disasters is well documented, mainly due to their ability to operate in hazardous environments without risking human lives (Murphy et al., 2016). For real-time reconnaissance, unmanned systems can be used as mobile sensor carriers, gathering real-time critical data about the evolving emergency in a finer spatial resolution than offered by satellite-based services. Typically, these robots have been directed by human operators, but human-agent prototypes have shown how algorithms for automatic task allocation can be coupled with human teams (Bailon-Ruiz et al., 2022; Ramchurn, Huynh, et al., 2016). In parallel, more comprehensive and completely *Autonomous Unmanned Systems* (AUS) have begun to emerge.

Given the limited time and resources during an emergency, unmanned systems cannot regularly sweep an entire disaster zone; even more so in the case of time-dependent measurements. It follows that sensor measurements are often sparse (and noisy), covering a small proportion of the physical scenario over a reduced time horizon. Measurement locations and timing should be chosen intelligently to maximize information gain, ideally organizing the mission by generating a movement policy for the sampling selection process and exploiting both real-time sensor readings and *a priori* information such as known building layouts, digital elevation maps, vegetation distributions, and other relevant characteristics to the specific emergency type. Basically, leveraging on domain, context, and process knowledge, as new data arrive, the system should adaptively refine its assumptions about the state of the environment and identify the most valuable points for subsequent measurements. The objective function of the movement is to maximize the usefulness of the information on the current situation. Furthermore, models might be abstracted to not only be the goal, but also the exploration tool: collecting valuable knowledge with adaptive mechanisms that improve awareness and decisions when integrated directly or using optimized policies.

The practical implementation of AUS as mobile sensor carriers to gather information typically involves three interlinked activities: collecting measurement data; integrating these real-time sensor readings with existing models of the environment; and making autonomous decisions about where to sample next. Each of these individual phases is potentially a complex task comprising numerous technical layers that need to be addressed in order to achieve true autonomy. Different research communities have studied one of these activities or combinations thereof with numerous different goals, research domain-specific approaches, formalisms, and nomenclatures. Therefore, it is challenging to gain an overview of the research that has been conducted in these different areas and to analyze which of the developed approaches handles the dynamic, uncertain nature of the surroundings while adapting to the specific characteristics of each case study.

1.1. Related works

Several survey works cover aspects that overlap with the focus of this study. We use a set of representative overviews as reference points to frame the scope and position our contribution. Table 1 summarizes these surveys across robotics, control, environmental monitoring, and emergency management, indicating the domains and applications in which unmanned systems are deployed and the type of phenomenon that unmanned platforms are tasked with observing. The table is meant to highlight relevant research directions.

Most of the available studies exploring practical applications focus on specific domains, and many are limited to a specific type of mobile system. A promising attempt to organize the literature streams on movement and mission planning (adaptive and static) has been (Glock, 2020). In the context of AUS for *mapping* during chemical contamination emergencies, the work addresses the literature. However, it focuses on *Unmanned Aerial Vehicles* (UAV) and then develops solutions for static mapping, with no temporal dimension on the process. Still considering chemical scenarios, Hutchinson (2019) instead focuses more on using dispersion models to exploit knowledge of the physical process responsible for emergencies, with static and mobile sensors. Part of the work is also available in Hutchinson et al. (2017). Focusing on Autonomous Underwater Vehicles (AUV), Hwang et al. (2019) give a general overview of mission planning for adaptive sampling, also discussing some applications for chemical hazards. Autonomous UAVs for information gathering during wildfires are discussed in Bailon-Ruiz and Lacroix (2020), presented as a central emerging topic, and, most importantly, considering whole operative systems rather than being restricted to the UAVs' embedded components. More in general, the role of artificial intelligence and UAVs for wildfire is explored in Boroujeni et al. (2024). These application-specific studies are supplemented by works focusing on environmental monitoring with various systems (e.g., Asadzadeh et al., 2022; Holbach et al., 2023 and Iqbal et al., 2023). Additionally, several research streams focus on largely domain-independent methods, instead addressing the question of how mobile systems operate in unknown environments (e.g., Atyabi et al., 2018; Campbell et al., 2020). The survey closest related to this study is the work presented by Popović et al. (2024), which provides a comprehensive overview of adaptive informative path planning, with a particular focus on learning-based methods for robotics. In contrast, prioritizing emergency contexts, this work adopts a more interdisciplinary focus, addressing a broader spectrum of solution approaches and extending the application scope to a number of applications for hazard monitoring during emergencies. This entails a stronger emphasis on integrating domain-specific models, decision-making frameworks, and the practical challenges of deploying unmanned systems in real-world dynamic and uncertain environments.

1.2. Research questions and outline

Accordingly, to advance interdisciplinary work in hazard monitoring during emergencies, this work presents a comprehensive overview that extends beyond domain boundaries and examines the following research questions:

Table 1
Scope of selected overviews.

Reference	Systems	Domain	Context(s)	Emergency-focused
Asadzadeh et al. (2022)	UAVs	Monitoring in the petroleum industry	Oil and gas	Partially
Atyabi et al. (2018)	AUVs, UAVs	Domain-independent mission planning	n.a.	No
Bailon-Ruiz and Lacroix (2020)	UAVs	Wildfire mapping and monitoring	Fire front and thermal information	Yes
Boroujeni et al. (2024)	UAVs	Wildfire management, pre- and post-disaster management	Fire front and thermal information	Yes
Campbell et al. (2020)	Mobile robots	Domain-independent path planning	n.a.	No
Chen and Huang (2019)	Mobile robots	Odor source localization	Chemical substances, esp. airborne	No
Glock (2020)	UAVs	Contamination mapping	Airborne substances	Yes
Holbach et al. (2023)	n.a.	Cyclone observation	n.a.	No
Hutchinson (2019)	UAVs	Search and source term estimation	Airborne substances	Yes
Hutchinson et al. (2017)	Static and mobile sensors	Source term estimation	Airborne substances	No
Hwang et al. (2019)	AUVs	Sampling, mapping and monitoring of maritime phenomena	Different physical, biological and chemical substances	Partially
Iqbal et al. (2023)	UAVs	Flood mapping and monitoring	n.a.	Yes
Kowadlo and Russell (2008)	Mobile robots	Odor localization, search and identification	Chemical substances, esp. airborne	No
Popović et al. (2024)	Autonomous mobile robots	Domain-agnostic information gathering	n.a.	No
Quero and Martinez-Carranza (2025)	UAVs	Search-and-rescue	Local and static targets	Partially
Sung et al. (2023)	Mobile robots	Environmental monitoring	Different physical, biological and chemical substances	No

- Considering the diverse formalism across different communities on AUSs, what are the fundamental concepts of movement planning for information gathering? On the other hand, what are the principal strategies to integrate data into dynamic models? How have these two research streams been meaningfully connected?
- In what emergency contexts have AUS been used for improving situation awareness when environmental dynamic hazards are involved? How was their movement planned? How do they compare regarding similarities, differences, and current limitations?

To answer the research questions, Section 2 starts from the data integration and situation awareness perspective, briefly introducing how models are actually adapted during dynamic phenomena. Then, unmanned systems are discussed, presenting an interdisciplinary meta-review on adaptive movement planning. General unifying paradigms are presented with a focus on policies for information gathering. In Section 3, this work presents studies that address situation awareness with autonomous information gathering using unmanned systems for environmental hazard monitoring, also comparing different approaches on similar scenarios and abstracting their key factors. Section 4 offers an overview of related applications and research fields. While these areas fall outside the main scope of this work, they can help foster a broader understanding of the interdisciplinary landscape of emergency response. Lastly, discussions and conclusive remarks are presented in Section 5. Table 2 indicates recurrent acronyms.

2. Interdisciplinary meta-review on data integration and adaptive movement planning

Several research fields, each with a specialized yet similar scope, have studied how to *control*, and more specifically *move*, autonomous devices to collect and use context-related information. Alongside and sometimes intersecting, data assimilation and machine learning have investigated how to adapt dynamic representations of a system by incorporating new data, especially under uncertainty and/or with a limited amount of observations.

In this section, the first part concerns methodologies and paradigms to integrate observations into dynamic models and summarizes conceptual frameworks for sequential decision making that offer a frame of reference for the specific applications considered in this work. Then, from the unmanned systems point of view, this section gives an overview of the most pertinent research fields concerning movement and mission planning, discussing its relation to the connected literature on mapping, tracking, and spatial coverage. Eventually, data integration, tracking, and movement planning for information gathering are addressed from a joint perspective, highlighting the logic behind adaptive decisions to improve situation awareness for dynamic environmental hazards. Nevertheless, a clear introduction of terminology is necessary first.

2.1. Defining dynamic and adaptive aspects in hazard monitoring

In the existing literature, which is spread across a wide range of applications and research communities, terminology and conceptual frameworks are often used inconsistently, making it difficult to establish a clear-cut taxonomy. Consequently, the literature does not share a common understanding of what constitutes *dynamic* aspects in hazard monitoring. For example, the term may refer to the phenomenon itself, to the belief about the phenomenon, and/or to the decision-making policy, even in case of static hazards. This lack of conceptual clarity may pose challenges for researchers, particularly when comparing approaches or developing new methodologies. To address this, this work aims to establish a common understanding to capture nuances in the existing literature better.

From an emergency viewpoint, all scenarios are, in some way, evolving over days and weeks, but some might evolve much more

Table 2
Acronyms and abbreviations.

Acronym	Meaning
AUS	Autonomous Unmanned System
AUV	Autonomous Underwater Vehicle
UAV	Unmanned Aerial Vehicle
ASV	Autonomous Surface Vehicle
AGV	Autonomous Ground Vehicle
AD	Advection–Diffusion
IPP	Informative Path Planning
MDP	Markov Decision Process
MPC	Model Predictive Control
S&R	Search and Rescue
STE	Source Term Estimation
CBRN	Chemical–Biological–Radiological–Nuclear

quickly than others. Therefore, whether or not this study considers an application to focus on *dynamic* hazards depends on the users' needs regarding updated information. Specifically, the term is used here to refer to applications in which *information on time-dependent aspects of the phenomenon or its evolution is relevant to a decision maker*. Note that this excludes a measurable criterion in terms of temporal resolution or update frequency. Instead, the unifying criterion is whether an application accounts for temporal aspects.

From the unmanned systems and information gathering perspective, the term *adaptive* has a somewhat more common understanding, referring, in a very broad sense, to any decision strategy for data-collection that depends on observations made during the mission itself, even if it does not explicitly update a model or reacts to a forecast of the future progression of the surveyed phenomenon. This view highlights the close relationship between adaptiveness and sequential decision-making, where actions are chosen in response to observed states or measurements over time (see Section 2.2.3). However, it also means that despite a diverse literature on adaptive methods, they often do not translate to dynamic processes.

It is worth emphasizing that this leaves overlaps and edge cases. For instance, a series of 'snapshot' data may provide sufficient information to the decision maker, even though the methods for achieving these snapshots are not, themselves, adaptive. Conversely, adaptive strategies can be applied to steady-state phenomena. As this work focuses on adaptive movement planning, these edge cases are included when the methodology naturally connects to the monitoring of environmental hazards.

2.2. Integrating data into dynamic models

Models representing environmental hazards vary substantially between the applications included in this study, with some approaches not relying on an explicit environment model at all. Frequently used approaches are briefly touched on in this section, before discussing methods for updating these models more in depth.

2.2.1. Overview of modeling approaches

Common modeling approaches are *Gaussian Processes* (GPs) (Williams & Rasmussen, 2006); widely used for continuous spatial modeling and generalizable as *Spatio-temporal GPs* to explicitly consider time dependencies. Approximate versions exist in the form of *Sparse GPs*, which decrease computational effort, especially for large-scale applications. More generally, *probabilistic* methods, with *Bayesian* approaches providing a principled way to fuse priors and data, allow modeling of spatiotemporal relationships in complex processes and naturally lend themselves to environmental-hazard mapping by yielding calibrated uncertainty under noisy observations. In fact, *Machine Learning* modeling solutions for complex relations that can be applied to dynamic phenomena are increasingly available (Price et al., 2025).

Complementary to this are physical and process-based models that originate primarily in environmental sciences and seek to provide a detailed model of the underlying process. As such, they rely more on direct approximations of physical laws, e.g., in the form of dispersion or hydrological models drawing on fluid or atmospheric dynamics. Naturally, hybrid models are increasingly popular, combining these two research streams. For instance, ML methods may be enhanced by explicitly incorporating information about physical dependencies, or ML may provide a surrogate model for complex physics-based simulations.

2.2.2. Data integration and model update

Given a model that represents a dynamic hazard for situation awareness purposes, the next step is to incorporate incoming data. From a practical point of view, this generally means solving an *Inverse Problem*, combining observations and dynamical model outputs (Asch et al., 2016). Common research fields for integrating data are *Data Assimilation* and *Machine Learning*. In fact, connections between data assimilation and machine learning paradigms have been presented by Geer (2021) under the framework of Bayesian networks. *Data Fusion* and *Information Fusion* topics extend to cases where multiple sources of information are available and merged. Classical sub-domains span across different communities. Well-known are the frameworks based on Kalman filter extensions and more general Bayesian approaches. Specific examples with unmanned systems are the *Particle Filter* (Hutchinson, 2019) and direct Bayesian paradigms, especially useful when dealing with high nonlinearities. Other well-known tools from this family are the *Extended*, the *Unscented*, and the *Ensemble* Kalman filters. Strictly connected, variational methods (e.g., 3D-Var) are customary in meteorology and geosciences (Carrassi et al., 2018). Nevertheless, representations are not necessarily detailed and can also correspond to rough approximations of the hazard, relying more on direct Bayesian paradigms and oriented more towards macroscopic effects rather than accurate dynamics.

The high computational costs of simulation and assimilation methods are severe bottlenecks for real-time integration of accurate representations. However, machine learning (and more specifically, deep learning) based models are usually faster at inference (once trained) than purely physics-based assimilation schemes. They are, by definition, flexible approximators that can be used to generate probabilistic physical emulators under the correct modeling assumptions (Geer, 2021). New solutions are emerging and, for example, *Data Learning* mixes data assimilation and machine learning, making the most out of the two (Buizza et al., 2022), addressing computational problems and overcoming traditional strict assumptions of data assimilation strategies like linearity and normally distributed errors. Machine learning based models enable assimilation and representation in reduced spaces, also acting as a surrogate of complex physical dynamics. More specifically, the physical representation is usually compressed to a *latent* space where assimilation is more manageable. Approaches span from more established projection-based techniques to *autoencoder* with deep neural network architectures, particularly attractive for highly nonlinear problems (Cheng et al., 2023). Besides operations in latent spaces, data assimilation and machine learning are widely reciprocally combined to address specific problems arising from each of the two research streams, such as, e.g., an efficient use of highly informative priors, a detailed quantification of the uncertainty, and the learning of dynamics over systems where physics are not fully known.

Relatedly, Popović et al. (2024) have identified four machine-learning paradigms that are commonly used for data integration and model update in adaptive planning for learning environment models: supervised learning, reinforcement learning, imitation learning, and

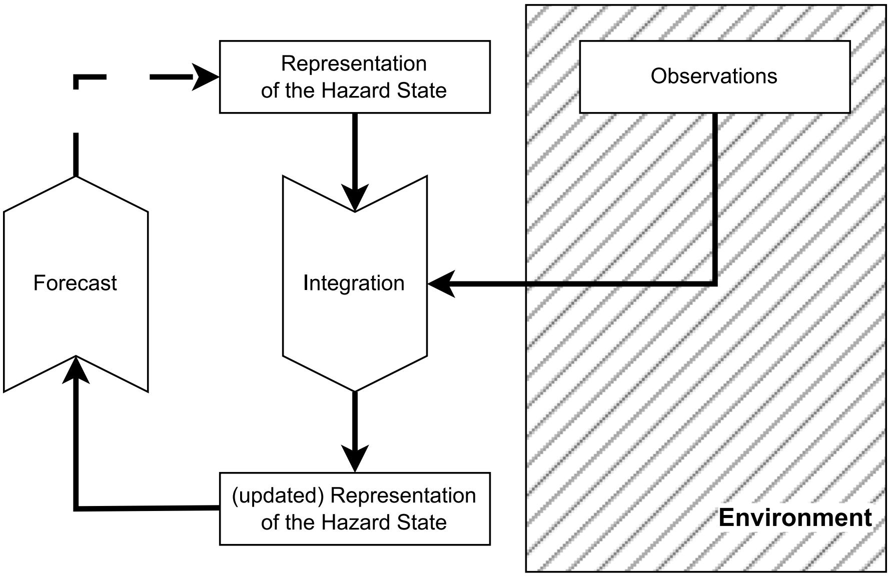

Fig. 1. General scheme for integrating observations into dynamical systems.

active learning. Moreover, the authors briefly discuss their applicability to different situational characteristics. While temporal characteristics and dedicated forecasting methods are not the primary concern of that study, the reader is referred to their work for a more in-depth discussion of potential application-agnostic learning methods.

A general scheme dealing with observations and a dynamic model that represents the hazard for situation awareness is given in Fig. 1. The representation, usually given by a model, defines the system state and can also be interpreted as a *Digital Twin* or the current *belief* about the scenario. Collected observations are, for the sake of simplicity, indicated in a single box, but generalizations with pre-processing, information fusion and different models or sources integrated together are a natural byproduct. Once observed, the framework uses the new information to update and adjust the representation of the emergency, which can then be used for forecasts and, ultimately, for informed decision-making. The process continues iteratively, progressing to the next time step, depicted as a dashed loop in Fig. 1. More formally, pre- and post-integration are in fact *forecast* and *analysis* variables in the filtering nomenclature. The term ‘integration’ is chosen to be neutral with respect to data assimilation, machine learning, data fusion, and other specific research streams. Note that steady-state dynamic processes and static data collection, like mapping applications, will not have a forecast step, only integrating new data to update the current state. A typical example is the use of *Gaussian Processes* for leveraging spatial correlations (see Glock & Meyer, 2020); the logic of which connects to Kalman filters in the case of time dependency.

2.2.3. Sequential decisions, abstractions and unifying paradigms

Sequential decision-making over dynamical systems has proven fundamental in various engineering and mathematical communities. However, driven by different problems, they have developed solutions in separate and at times isolated areas, using different nomenclatures and formalisms. Nonetheless, fundamental paradigms define the universal logic, and it is worth discussing them.

When using approximate information to make sequential decisions, uncertainty is taken into account to some degree, epistemically and/or aleatorically. Stochastic sequential decision problems are typically modeled as *Markov Decision Processes* (MDP) in the operations research and reinforcement learning communities. Moreover, when it is only possible to partially observe the state of a dynamical system (e.g., when only partial observations are available over an emergency scenario), the nomenclature refers to them as *Partially Observable Markov Decision Process* (POMDPs). There, decisions rely on an (adaptive) belief about the true state of the system, refined whenever new observations are made. Ideally, policies for information gathering should aim to improve the current belief, while simultaneously addressing other user-defined goals. Unification attempts exist to merge similar mathematical tools from different communities. For example, Powell (2022) considers POMDPs as multi-agent problems with an environment and a controlling agent, connecting them to a more general paradigm for sequential decision making.

Extending classical MDP approaches, neural networks and deep architectures have been successfully applied in sequential decision making for various practical problems. More specifically, this includes the predominant (deep) reinforcement learning based solutions in robotics and navigation problems for mobile systems (Zhu & Zhang, 2021). Due to the time, cost, and complexity required to train decision policies for autonomous systems, navigation strategy performances are evaluated and improved offline by using simulations. Various software exist to simulate environments as a playground (Kim et al., 2021). Giving some specific examples, the deprecated AirSim Simulator (Shah et al., 2017) and its successor Colosseum (LLC, 2025) are well-known high-fidelity physical simulators for robotics and autonomous systems that aim to mimic the real world to simulate how robots interact. Nevertheless, it is important to notice that, when developed online, these simulations are usually not directly embedded in the device policy. In fact, the device abstracts the characteristics of the explored use case during learning.

The explicit use of a model to control in the present while accounting for future behavior is commonly known as *Model Predictive Control* (MPC) in the control theory community. MPC consists of considering some time steps in the future to anticipate future system behavior based on current knowledge, then optimizing control inputs and only using the first control, discarding the rest, and iterating. Situations in which only partial knowledge about a system’s parameters is available and unexpected misalignments arise online have been largely explored as well under the term *Adaptive Control*. When combined with explicit models, specific architectures (also deep ones) that mix different representations and jointly use approximations of different nature and fidelity are known to be effective. For example, using truncated or simplified rollout policies in adaptive contexts to perform online adjustment given a changing system with complex dynamics has been shown to be an effective solution within the MPC domain Bertsekas (2022). It is also worth mentioning that, even if the MPC logic is widely applied on sequential problems of any kind (Bertsekas, 2022), much of the MPC

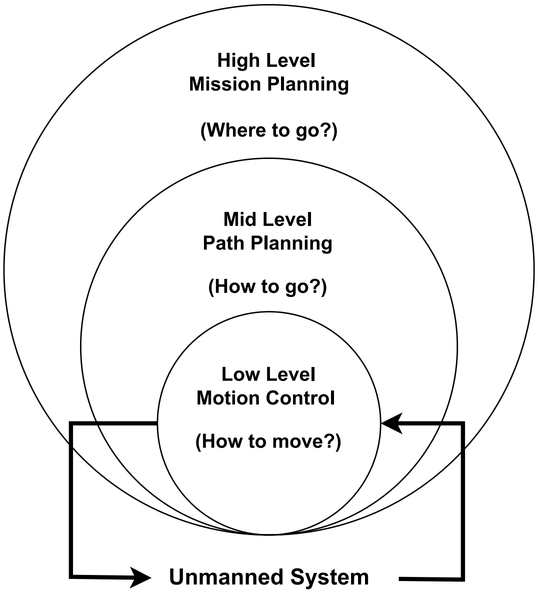

Fig. 2. Control hierarchy for autonomous unmanned systems.

community has historically been concerned with low-level controls, as further explained in Section 2.3, where a control hierarchy is discussed. In this context, learning the specifics of a dynamical model while performing adaptive control is defined as *system identification* (Ljung, 2010). This field is enormous and has a solid legacy. Nevertheless, it aims to accurately learn the effects of input signals on a controlled system, where such a system is usually the device itself or its components rather than the environment in which it operates (Hoffer et al., 2014).

The research about adaptive control is very active (Bertsekas, 2024) and is tightly connected with reinforcement learning. In fact, unifying frameworks under the umbrella of *Abstract Dynamic Programming* (Bertsekas, 2024) and even more general sequential optimization paradigms (Powell, 2022) are beneficial to connect similar ideas arising from different scientific communities, operating at varying levels of system control and with different problems in mind. More specifically, in Powell (2022), the family of sequential optimization strategies relying on an explicit model-based evaluation of the effect of a here-and-now decision into the future is unified under the term *Direct Look-Ahead (Policies)*, but three other state-dependent families of policies are introduced (i.e., *Policy Function Approximations*, *Cost Function Approximations* and *Value Function Approximations*) to abstract even further the many other intersecting research streams discussed and others beyond the scope of this work.

2.3. Movement planning for autonomous unmanned systems

Focusing on AUSs and to efficiently introduce different perspectives on their control, Fig. 2 gives a simplification with a three-tiered hierarchy, where each part focuses on a distinct aspect of decision-making and execution. The highest level, *Mission Planning* (“Where to go?”) usually defines the overall objective or strategic goals, adapting to dynamic priorities. The mid-level, *Path Planning* (“How to go?”), translates these high-level objectives into a feasible trajectory by connecting way-points, considering physical constraints, and obstacle avoidance. Finally, at the lowest level, *Motion Control* (“How to move?”) directly communicates with the device and ensures precise execution of the planned trajectory by converting control inputs into physical motion and maintaining, e.g., stability. While all three levels work in coordination, movement planning for information gathering mostly concerns the highest level or, at the very most, the middle one.

Besides the very general categorization of Fig. 2, both the literature and the real-world execution refer to mission, path, and motion planning with a blurred boundary. Furthermore, specific problems in mapping and tracking with robotics (and, consequently, unmanned systems) define standalone research streams and nomenclatures. Therefore, the next sections provide additional context and intermediate concepts that progressively link these related strands to adaptive mission planning for information gathering.

2.3.1. Movement, path and mission planning

Most of the theoretical works dealing with dynamic emergencies only consider high-level control commands, disregarding any navigation-related complexity. Since most practical studies use existing application-independent implementations when needed, a detailed treatment of low-level navigation issues is omitted, with emphasis placed instead on emergencies and models of dynamic environmental hazards.

The first step for navigation in unknown environments usually concerns self-localization and map building. In particular, the *Simultaneous Localization and Mapping* (SLAM) problem deals with the incremental construction of a map of the surroundings, while self-localizing the device. Sub-families of the topic, e.g., its *active* form (Cadena et al., 2016) (see also *Planning* (SPLAM) form Leung et al., 2006), adaptively make decisions and could be considered as a close relative to the focus of this work, but with an antithetical and preliminary purpose. There, a decision-making process is involved to optimize the movement of autonomous systems in order to estimate the state of the system, intended as the position of the device within an environment and the (usually static) map of the surroundings. The commonly adopted *static-world* assumption (Cadena et al., 2016) for SLAM problems is the outcome of a different focus where not the dynamics of phenomena in the environment, but rather those of the observer are of interest. Nevertheless, exploration–exploitation balance and information-gain are fundamental components as well.

Deciding how to move mobile robots mainly relates to *path planning* in the literature. In particular, such a problem considers the identification of a collision-free path of a robot from a specific source to a defined end-point, with goals dependent on metrics of the path itself (e.g., minimizing the cumulative travel time). Several reviews of common techniques are available, e.g., Campbell et al. (2020) for mobile robots in general, and Jones et al. (2023) with emphasis on the additional complexities of UAVs. Approaches dealing with known or partially unknown dynamic environments do exist (Jones et al., 2023), usually focusing on the dynamics of objects and obstacles. What is modeled besides devices and obstacles is usually restricted to the physical properties of the surroundings and does not coincide with what is causing specific environmental conditions.

From a more general perspective and according to Fig. 2, path planning may be interpreted as a problem embedded within *mission planning*: concerning, e.g., resource allocation, goal re-planning or task management. Here, situation awareness and cognition are central components of the actual autonomous unmanned system (Atyabi et al., 2018); able to independently re-prioritize mission objectives and, ideally, to adapt to unfolding dynamic scenarios. In such a context, information gain over the system of the environment may be the goal of the planning, commonly referred to as *Informative Path Planning* (see also Section 4.1).

When planning the movement of an unmanned system for emergency response, situation awareness is both a goal and a necessity, as the unmanned system should follow the emergency dynamics in real-time and information gathering is pivotal for developing effective counter-measures within the emergency scenario. In this direction, embedding a decision-making process that exploits a priori knowledge about the phenomena and is enhanced by collected observations may improve the utility of data collected for situation awareness. Dedicated mechanisms for mission planning in this context are discussed in more depth in the following two sections.

2.3.2. From mapping to tracking

Having a static or dynamic representation of the environment as a goal can make a difference in both methodological and operational terms. As previously mentioned, gathering information to learn a data-driven static model of the physical environment of operation of a device is usually referred to as *Mapping*. Conversely, interest concerning the dynamics of a phenomenon is mentioned as *Tracking* or *Monitoring*. From the application side, the separation is not so clear-cut and, rather than the temporal characteristics considered, the goal of the application itself may determine the nomenclature.

More in the direction of mapping a phenomenon occurring in the environment rather than the environment itself and more specifically in the case of emergencies, *Rapid Mapping* relates to the provisioning of geographical data, historically mainly relying on satellite data and aerial photographs, and recently expanding to other sources (e.g., Lu et al., 2011). As an extension, *Emergency Rapid Mapping* refers to the problem of reconnaissance in the aftermath of a disaster, leveraging dedicated, mobile sensor platforms, spatial interdependencies and interpolation techniques. A data-driven example is given in Glock and Meyer (2020) for pre-determined planning during contaminant dispersion, then also investigated with online adaptive planning strategies in Glock (2020). More generally, when information is gathered by leveraging on the spatial proximity of the visited locations, *Spatial Coverage* problems, largely explored in the operations research fields as extensions of the vehicle routing problem, are natural connections. A comprehensive unifying view is given in Glock and Meyer (2023). It is also worth noting that Stampa et al. (2021) stress how *Assessment Mapping*, where plenty of time is available for data acquisition and processing, is very different from emergency mapping; as the latter requires increased processing capabilities and a higher level of operative autonomy, without increasing the cognitive load of emergency operators. In fact, even if, according to their analysis, assessment mapping has market-ready, fully functioning solutions, the same was not true for emergency mapping, especially in complex physical three-dimensional domains.

When moving to tracking for dynamic environmental hazards, reacting to changing environmental conditions and leveraging a priori knowledge about temporal dependencies is valuable for effective information gathering. Adaptive strategies solving sequential decision problems, where the movement follows a policy that reacts to the current *state of the system*, are an active research focus. In particular, the movement planning for an autonomous unmanned system may follow pre-planned movement policies. However, when dealing with dynamic scenarios and/or uncertainty, adaptiveness with respect to the current knowledge is especially valuable. Nevertheless, a first source of confusion can be what exactly is understood as the state of the system. In fact, different problems require different abstractions, entailing different definitions. In robotics, the state of the system traditionally includes the position of the robot, specific mechanical elements or, for example, some specification of the object towards which a device is to interact (Kurniawati, 2022). Alternatively, a useful unifying definition of the state of the system that fits the various dynamic contexts hereto discussed is given by Powell (2019). Namely, the state of the system *"contains all the information that is necessary and sufficient to model our system from time t onward"*. Nevertheless, the ‘system’ itself may refer to the components of a device, to the conditions of the environment where it operates or to other context-specific definitions. More general definitions and remarks are given in Powell (2022).

2.3.3. Modeling informativeness

As practical applications typically face constraints prohibiting a complete assessment of the area in question, the balance between exploration and exploitation is often crucial, as applications seek to achieve both a reasonably complete overview and the refinement of data on high-value areas. This directly connects to the concept of *cost/reward function*. Several approaches referenced in Section 2.2 can support techniques that seek to achieve such a balanced outcome by assessing the *informativeness* of the additional observations to be integrated. Specifically, probabilistic methods can provide quantitative information about the uncertainty of their prediction, then exploited in *Active Sensing* and *Active Learning* applications, as summarized by Popović et al. (2024). Examples typically rely on information-theoretic measures such as *Mutual Information*, which aim to accelerate model convergence by prioritizing high-impact data points, as their main objective.

Modeling the underlying process in more detail may also lead to task-specific formulations for modeling informativeness, e.g., by using cost functions that support *Value of Information* measures tailored to the mission objectives. Implicit models and heuristics also exist, considering, for example, information about gradients (e.g., in wildfire front mapping) or modeling spatial coverage as a proxy for the collected information. Not all models require such an indication of informativeness, although the idea of valuing potential options against one another, based on previously collected data, is typically prevalent.

Lastly, it is worth noting that mobile sensor platforms, as well as the setting in emergency surveillance, come with additional requirements and restrictions that forbid a straightforward application of application-agnostic approaches. Most importantly, movement restrictions mean that valuable information may simply be impractical to obtain. This is why adaptive movement planning extends these methods in decision-making frameworks that consider these additional aspects.

2.4. Adaptive movement planning for information gathering

Fig. 3 represents an abstraction of the adaptive movement problem for an observational agent (the unmanned system) where the most general sequential decision making paradigm introduced by Powell (2022) is followed. More specifically, the state of the system is separated into: the current knowledge about some selected features of the hazard (which refers to the current *belief* about the state of the hazard); the physical constraints of the devices and some additional information necessary to make decisions. New observations are collected according to a *policy* that decides where to go to observe. Such a policy is optimized according to an *objective function*, usually related to a quantification of the information gain for the specific problem at hand. This decision is the *action* of the device. Observations are used to update the *belief*, integrating collected data. Other varying quantities (e.g., position of the devices or battery) are also updated accordingly, potentially depending on other exogenous factors, and the function responsible for the update is commonly known as *transition function*, mainly within the reinforcement learning community. The abstraction is quite general, allowing for different policies, state of the system descriptions, objective functions for which the policy is optimized, transition function components, and exogenous information generated. Nevertheless, exceptions to such a simple architecture are the rule in practical applications.

The ideal system would make decisions in real-time, with loops working at every small bit of information collected and using a very accurate model to which to supply data with low latency. However, this is not the case in the real world. Such a sequential problem might be solved off-line and off-board; pre-determining plans that are updated occasionally, collecting observations received in batches and only then solving seemingly accurate inverse problems for integration. Even when an integration of the data is fully implemented on board or on a synchronized central system, the selected features of the hazard can vary in complexity and level of abstraction, being not coincident with comprehensive descriptions better suited for situation awareness. For an unmanned system, it may be both impossible and useless to have a complex representation and accurate integration on board, even when a more comprehensive model is theoretically available. In general, the available knowledge about the hazard is instrumental to achieve a specific goal, in terms of information gathering, by developing a

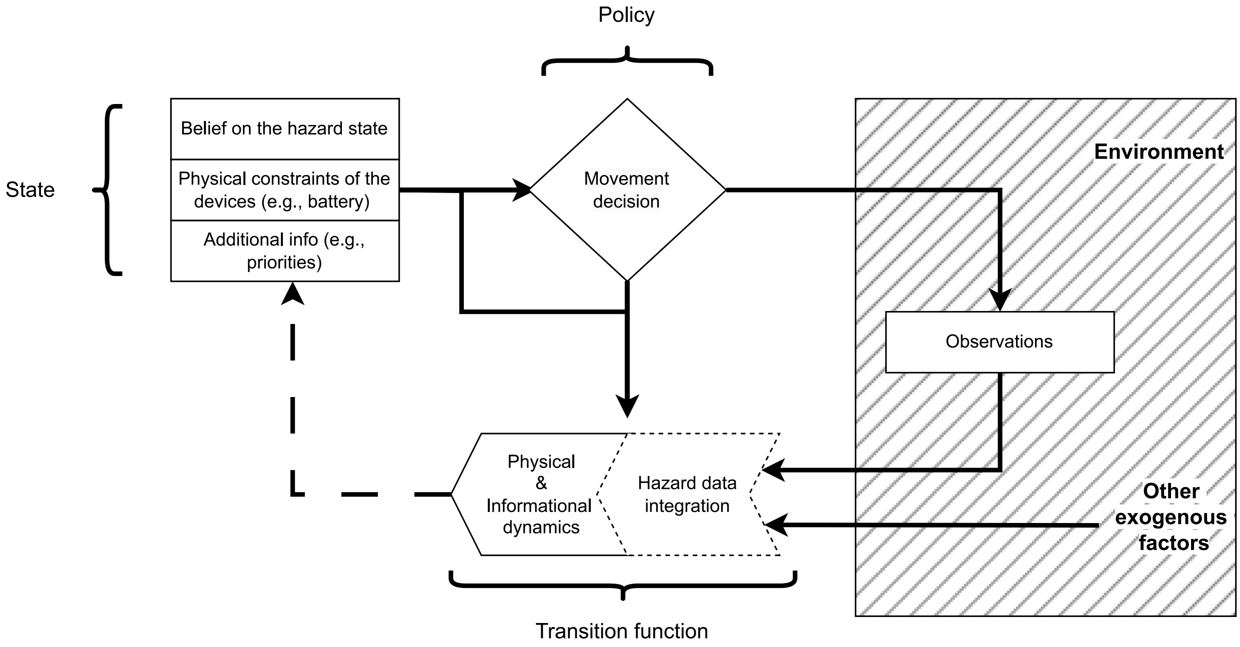

Fig. 3. The adaptive movement problem for an observational agent.

good policy and optimizing it (e.g., selecting a specific deep learning architecture and training with a correct *reward*). Various sequential decision frameworks do not consider explicit integration of the observation at all, as an iterative update of the current belief is often confined to POMDPs studies and various case studies collect data only for asynchronous and/or in batches integration. The movement can also be adapted according to rather simple policies. For example, AUS may provide data for assimilation purposes while relying on very little data of the emergency to make decisions (e.g., following the gradient of a concentration of a chemical).

Not only computational constraints, but also limited communication between devices and navigation complexities may make detailed and real-time integration infeasible. Usually, multiple agents are involved, pursuing a common goal, but relying on scattered knowledge and different policies. In particular, the concept of decisional autonomy for multi-device systems has been studied from centralized to distributed configurations, being evaluated according to the available capabilities of the selected architecture (Gancet et al., 2005). For this reason, along with the general provided paradigm, the next section provides a context-driven presentation of the most common problems for dynamic environmental hazard monitoring and their respective practical and effective solutions.

3. Emergency contexts and case studies

In various emergency contexts, AUSs have gathered information for monitoring dynamic environmental hazards, using very different policies, models, sensors and also time scales. Given that monitoring ongoing emergencies with in-situ measurements is the central topic of this work, even if situation awareness can be connected to every part of a crisis, the focus of this section is mostly on the *response* stage of an emergency. In fact, at least two other canonical phases are usually mentioned in the literature: preparedness and mitigation in pre-crisis; and recovery, after the crisis (Lettieri et al., 2009). Applications on such stages will not be discussed, thus not including the aftermath of instantaneous stressors like, e.g., an earthquake or a bomb, but rather focusing on hazards sufficiently dispersed over time, allowing for the use of some preliminary knowledge about their dynamics to deploy an adaptive movement policy. Furthermore, since the aim is where to move and how to use observations while monitoring hazards, solutions where data are collected according to a full exploration of the environment and/or a movement policy completely independent of what is observed are not discussed, but only briefly mentioned if key to progress in context-specific practical testing.

In the following, the three most prominent application contexts are discussed. Specifically, *Chemical–Biological–Radiological–Nuclear* (CBRN) scenarios are covered in Section 3.1; wildfires in Section 3.2; and floods and storm systems in Section 3.3. For each context, specific strategies on movement planning for information gathering with autonomous unmanned systems are explored. Table 3 summarizes the different research streams to provide an overview of the case studied, discussed thereafter.

3.1. Chemical–biological–radiological–nuclear

CBRN covers a wide range of natural and man-made disasters: from the accidental or deliberate release of hazardous agents to ecological disasters such as volcanoes. Emergency monitoring applications are typically concerned with identifying the source and/or the dispersion of contaminants, which are often airborne. Environmental models representing these hazards typically rely on diffusion and/or advection mechanisms that may also be modeled using stochastic models. Major sources of uncertainty lie in the source location and chemical behavior of the substance given environmental factors such as wind speed or direction.

A wide number of unmanned systems and sensors have been used in CBRN scenarios; each offering distinct advantages (and disadvantages) with regards to the emergency, type of hazard and operational setting. For example, a useful overview on thermal infrared cameras, optical imaging cameras and various laser-based systems is given in Asadzadeh et al. (2022) for UAVs for environmental monitoring and, specifically, the petroleum industry. For underwater applications with AUV, Hwang et al. (2019) discuss detection systems for hydrocarbons. Examples focusing on airborne contaminants and concentration readings by conductometric sensors and photo-ionization detectors are given in Hutchinson (2019) for AGVs and UAVs, then also tested empirically for adaptive movement planning during gas leakages. However, given the extensive range of sensors and systems available, an exhaustive listing would exceed the scope of this work. Consequently, the focus will rather concern the use and value of information for

Table 3
Overview of prominent application contexts.

Context	Research stream	Sensors (examples)	Common paradigms
Chemical–Biological–Radiological–Nuclear	Odor localization	Chemical sensors and anemometers	Reactive searches (chemotaxis, anemotaxis); belief-based search with entropy/information criteria (infotaxis)
	Source term estimation	Chemical sensors and anemometers	Model-based inverse problems using transport models (e.g., plume/puff, partial-differential-equation models); sequential Bayesian estimation; entropy/information-driven routing (entrotaxis)
	Boundary tracking/contour mapping	Visual or remote-sensing imagery	Contour/level-set representations; perimeter following and coverage-style control; cooperative tracking with centralized or decentralized coordination
	Volcanic/radioactive plumes	Dedicated hazard sensors + imagery	Sampling designed for assimilation into dispersion models
Wildfires	Monitoring (front/perimeter)	Visual and thermal imagery	Grid-based belief maps; coverage and tracking policies; embedded firefront simulators for sequential re-planning
	Suppression & monitoring	Visual and thermal imagery	Coupled sensing–action planning: belief-aware tasking and routing; criteria combining information value and expected suppression effectiveness
Flood & Storm Systems	Flood/river-state estimation	Imagery and/or water depth and surface-velocity measurements	Hydrologic/hydraulic models and inverse modeling
	Flood extent boundary tracking	Visual/thermal imagery	Boundary following as a geometric control problem; coordination laws for multi-vehicle fleets
	Storm systems	Dropsondes	Targeted observations for better forecasts with long-horizon missions planned under airspace/regulatory constraints, typically with humans-in-the-loop.

planning the movement of an AUS rather than on the technicality behind its collection.

An initial introduction is necessary on *olfactory* systems and *Odor Localization*, then naturally moving towards more complex representations with *Source Term Estimation* and time-dependent models, eventually discussing more general tracking strategies and other types of sensors focusing on pointwise estimates and boundaries of affected areas. Each problem addresses a distinct facet of inferring critical information, considering different representations and adaptive movement policies from sparse and/or noisy measurements.

3.1.1. *From odor localization to source term estimation*

Many works in CBRN scenarios use robotic platforms in a fully autonomous way for *Odor Localization* problems, where the aim is to find the source of volatile dispersed substances in e.g., atmospheric or aquatic environments. Arguably, finding the source of a dispersion is the first step to understanding its dynamics, perhaps also using this information in terms of an initial condition for simulations. Unlike the two application contexts in the remainder of this section, this covers the least dynamic of the discussed emergency scenarios, being the source location usually static, with contaminants being released continuously and distributions often only considered in *steady state*.

This is an old and well-researched problem and a direct example of an inverse problem. In fact, considering the size of the literature, numerous reviews are available like (Kowadlo & Russell, 2008) and more recently (Chen & Huang, 2019). Alternatively, Hwang et al. (2019) consider underwater problems with AUV. Usually, in academic experiments, the focus is on devices with electronic noses in indoor settings and steady state assumptions on wind direction and intensity. Historically, early works dealt with movement policies like gradient-based solutions (*chemotaxis*), upwind directions (*anemotaxis*) and bio-inspired algorithms, thus embedding on the device (and its policy) only local features of the hazard. Since the introduction of *infotaxis* in Vergassola et al. (2007), many studies and applications followed on probabilistic and map-based algorithms, explicitly incorporating the information content of the measurement within the mission planning problem and taking into account sparse and turbulent readings. In fact, these contributions present a probabilistic approach for the state of the system (represented by the probability distribution of the unknown source), dynamically integrating in a Bayesian fashion the collected measurements and following a policy that reduces the *entropy* of the estimated distribution.

When source localization is aided by additional information about the generating dispersion process, the problem nomenclature gradually transforms to *Source Term Estimation* (STE). STE aims for better performances by including a more general representation of the dynamics of the hazard. However, this typically comes at the cost of increased computational effort. Furthermore, even if the localization is usually the primary objective, AUS collects information to improve a model that depends on additional parameters related to the overall dynamics, thus aiding the localization by exploiting specific assumptions while providing improved situation awareness as collateral. Therefore, the AUS does not just move towards the source.

A notable example of movement policy tested and developed on STE problems is *entrotaxis* (Hutchinson et al., 2018), extensively explored theoretically and practically in Hutchinson (2019). Such a strategy is still connected to entropy and its introduction was oriented towards tracking multiple parameters characterizing the dispersion process. An organized presentation of STE problems and solutions with static and mobile sensors is given in Hutchinson et al. (2017), where it is discussed how, relying on different atmospheric transport models used as a representation of the emergency, the spread of hazardous material is predicted.

Given their simplicity and generality, parametrical *Gaussian plume* (or *puff*, for instantaneous releases) are often the representation choice, even though their accuracy can be limited. Many pure-theoretical contributions are available. For example, an early work on a simplified case study on a flat, obstacles-free, environment is Christopoulos and Roumeliotis (2005), where a mobile robot estimates the parameters (i.e., location, time of explosion, gas release mass and diffusivity) over a simple setting with constant wind and diffusivity hypotheses, deciding its movements by using the trace of the expected covariance matrix of the estimated parameters with respect to candidate points. With a similar Gaussian solution, but with additional care on the low-level motion constraints of specific devices, Euler and von Stryk (2017) simulate a swarm of quadrotor UAVs and use an MPC approach with a decentralized cooperative control that maximizes information gain with respect to the assumed parameterization. To reduce the gap between theory and real-world, they simulate the UAVs and the adaptive problem with ROS and Gazebo. However, all theoretical works have limitations. Successful practical tests for steady state problems (sequentially adapting the uncertain parameters of a Gaussian plume) have been made in wind tunnels (Rahbar et al., 2019) as well as uncontrolled outdoor settings in Hutchinson (2019) by using a *Particle Filter*. Other important real-world experiments are given in Reggente (2014), where Gaussian kernels are the main tools. Lastly, even if the movement of the device is pre-planned, valuable discussions on experiments regarding the effect of empirical wind conditions and other important factors during assimilation are given in Darynova et al. (2023).

More complex dispersion models are common choices as well. Still in steady state conditions, empirical examples for source location mapping in non-convex domains, are, e.g., Khodayi-mehr et al. (2019) and Jin et al. (2023). In Khodayi-mehr et al. (2019), a mobile robot considers the *Advection-Diffusion* (AD) equation within the decision process and a variational approach over a reduced order model is used for assimilating data in the embedded representation. Target points are selected by following a policy that computes the amount of information that a measurement may provide about the source related parameters. Alternatively, as an example of deep architectures for surrogate representations, Jin et al. (2023) uses *Convolutional Neural Networks* (CNN) for representing the gas dispersion.

In addition to single devices, interest naturally extends to solutions for odor localization problems that account for swarms of unmanned systems. For example, in Wiedemann et al. (2021), a swarm of unmanned tracked robots is moved by relying on a stationary AD model, stochastically relaxed by assuming Gaussian distributed deviations from the deterministic AD equation. By directly observing the concentration of the gas, a Bayesian estimation is sequentially performed, with a policy selecting locations with the highest variance. From a more theoretical perspective, but tracking considering explicit time dependency, Hinsen et al. (2023) introduce policies based on *potential fields* across the spatial domain, also taking into account collision avoidance.

Lastly, control methods that blend model dynamics and observations focusing on state estimation error with classic control theory are also often explored. An interesting simulated example that focuses on the direct estimation of the concentrations but follow a moving source as collateral is Egorova et al. (2016), where a UAV moves by following the maximum state-estimation error for a Luenberger observer with a 3D-AD-model. A conceptual connection on data assimilation techniques is possible with methods using a feedback term that directly corrects differential-equation-based dynamics, i.e., *Nudging* (Asch et al., 2016).

3.1.2. Boundary tracking

Often associated with slower-moving or longer-lasting contaminants in denser media (e.g., water), another important family of problems concerns *Boundary Tracking*. In fact, following a dynamic perimeter of a contaminated region, rather than locating a source or accurately quantifying the concentration distribution, is a common problem for CBRN emergency contexts. As a traditional example for dispersion of contaminants on water surfaces, and concerning disasters like large oil spills, an operative solution in environmental monitoring for detection and rapid mapping is the *Synthetic Aperture Radar* (SAR) imagery, provided by, e.g., the Sentinel-1 (European Space Agency, 2025) satellite. However, direct mapping and tracking of the hazard dispersion boundary (i.e., the contour of the concentration field) with unmanned systems is an active research field.

Unlike olfactory systems for odor localization, more studies assume visual data. Many theoretical contributions are available, often tested on simulations, dealing with swarms of cooperative devices with centralized or decentralized control and full or partial/no shared information between devices. However, frameworks frequently do not explicitly use the physics of the dispersion, neither to develop a movement policy nor to explicitly solve an inverse problem. In fact, boundary tracking often only focuses on the geometry of the problem.

To provide some examples, starting from pure mapping, in Odonkor et al. (2019) the movement is planned according to a particle swarm fashion with UAVs on oil spill scenarios. As is common in such studies, knowledge is local and decentralized, sporadically communicated between devices. Regarding the tracking of a moving boundary, again assuming sea surface oil spills and a swarm of devices, Pashna et al. (2020) use a fuzzy movement policy. Both studies move devices so that the identification of the boundary of a geometric shape is the goal. Concerning approaches where a quantitative concentration reading is used within an assimilation framework, Wang et al. (2019) is a theoretical example where a Luenberger observer is assumed for a dynamic pollutant plume over water following an AD model. A sensing and a (boundary) tracking ASV follow an ad-hoc collaborative movement policy to collect data.

It is worth noting that boundary tracking is a problem concerning all three emergency contexts presented in this work. In fact, the logic behind specific algorithms tested or developed for CBRN scenarios can be generalized to other contexts. Nevertheless, beyond pure contour mapping, the models related to the specific hazard dynamics differ. Practically speaking, if CBRN scenarios usually focus on AD equations, floods in Section 3.3 may rather use *Shallow Water Equations* (SWE), while wildfires in Section 3.2 often have ad-hoc front propagation simulations (e.g., based on Rothermel's model Rothermel, 1972).

3.1.3. Volcanoes, radioactive clouds and other examples

For volcanic plumes, examples of autonomous devices used for information gathering are Schellenberg et al. (2021) and, successively, Rolland et al. (2024). Tested on simulation, but based on data collected by the same research group at Volcán de Fuego in Guatemala, a UAV first autonomously localizes the volcanic plume using a YOLO model (Rolland et al., 2024). Then it follows a movement that adapts to the anemometric conditions. Specifically, the goal is to sample the particle size distribution in two locations according to a policy developed in Schellenberg et al. (2021); then to be used for assimilation on dispersion models.

Considering instead radioactive cloud scenarios with stochastic wind forces rather than gases, an interesting example is Wu et al. (2016), where a Geiger counter is assumed as a carried sensor. The movement of the unmanned system is directly connected to the goal of the first responders, as the movement follows the *value of information* for their specific task, thus assessing a value beyond the pure modeling accuracy. Specifically, a set of possible current states of the radioactive clouds is sampled, simulating the dynamics only for such a subset and assessing the value of information with regard to first responders' actions. The overall emergency response framework is presented in Ramchurn, Wu, et al. (2016) and they define a prototype of a functional emergency response system. In particular, the radioactive cloud is based on the Smoluchowski diffusion equation and integration/predictions of the dynamics of the stressor are performed using a latent force model and computed using an extended Kalman filter to deal with its stochastic behavior.

Lastly, for chemicals, some works assume the source location as known (e.g., after an odor localization step) or, for example, disregard it by focusing on the concentration on every point of a discretization of the domain while, e.g., considering instantaneous releases. In general, but even more for airborne dispersed gases, inferring the pointwise concentration of a puff is very complex. Several works deal with the pointwise tracking problem in simplified simulations (e.g., Euler et al., 2015; Gioia et al., 2024), but most are far away from reasonable assumption for field tests. A more practical example is Kato et al.

(2017), where, using an AUV for underwater oil tracking, the location of the seepage and initial vertical water current are assumed to be surveyed beforehand, so as to optimize the AUV movement for data assimilation purposes. Unfortunately, field-tests were interrupted due to fatal faults on the underwater mass spectrometer.

At the opposite end, it is also possible to not infer the physical evolution of a chemical plume and to not solve any inverse problem, but rather to track its approximate position regardless of its source or detailed distribution. A field experiment for underwater oil droplets (simulated with air bubbles as a surrogate) with an AUV is given in Hwang et al. (2023), where the turning angle of the device is adapted using measurement of a sonar sensor. Another example, using an ASV with an ad-hoc movement policy for remaining inside an oil slick by using measurements on contaminant presence, wind and water current velocity is Rathour et al. (2015), where tests used neoprene sponge rubber as oil surrogate.

3.1.4. Further considerations

As mentioned in various studies, and particularly for information gathering applications (Popović et al., 2024), common limitations of unmanned systems like localization uncertainty, processing time, sensors reliability and limited communication bandwidth (Stampa et al., 2021) naturally apply to fully autonomous devices and emergencies. In general, there is no guarantee of performance when moving from simulated environments to real emergency applications, where unexpected behaviors are the norm. Nevertheless, the first steps for functional systems require simulation and verification.

Many preliminary contributions focusing on information gathering with AUS assume devices as massless points to be routed. However, it is important to consider their physical dimensions and structural constraints. Open source software development kits like *Robot Operating System* (ROS) and simulation environments like Gazebo (Macenski et al., 2022) are valuable tools in between high abstractions and real working framework prototypes, especially as there exist frameworks for integrating these tools to facilitate moving from simulated to real-world environments and to coordinate sensor data fusion. As an example, an integrated ROS-based system for detecting nuclear sources has been tested in robotics challenges, with data and implemented software published afterwards (Schwaiger et al., 2024). It is also worth mentioning practical examples of available ad-hoc extensions of such software. For gas dispersion simulations on complex 3D scenarios for mobile olfactory devices, a working package based on ROS is Monroy et al. (2017). Alternatively, an odor simulation plug-in is available on Webots (The Wikibooks Community, 2025). In fact, *Twin-model* approaches are often considered to test offline and validate solutions, where the test-bed simulation and the device's embedded model act as the "truth" and the belief, respectively.

When moving from simulated studies to practical situations, variability in the environment conditions, possible interference and unmodeled dynamics must be taken into account. For olfactory systems, an overview of common problems for affordable sensors and devices is given in Hutchinson (2019). Of particular note, as their usage continues to expand rapidly, is the use of UAV (Stampa et al., 2021). When used for olfaction, it is currently the subject of active research understanding how the device affects the measurement (e.g., due to the movement of rotors) and how they can effectively smell (see, e.g., Fig. 4 from Deutsches Zentrum für Luft- und Raumfahrt, 2024).

Lastly, concerning datasets, differently from visual ones (e.g., for volcanic plumes Rolland et al., 2024 or for water surface oil spills De Kerf et al., 2024) usually intended for detection, collections of concentrations over a complete spatio-temporal domain require expensive set-ups. The Fusion Field Trail 2007 (FFT-07), usually cited as reported in Storwold (2007), consists of accurate distributed measurements for several dynamic dispersion settings and it is particularly worth mentioning as it has been used for various studies on CFD models. A more exhaustive description of FFT-07 is available in Kumar et al. (2017) and it belongs to the Defense Threat Reduction Agency (DTRA). Specific for odor localization, Hinsen et al. (2024) have collected measurements from a low-speed wind tunnel, making them available in the form of a challenge for researchers. Clearly, the given list is not a comprehensive one, but illustrative. Furthermore, new datasets are often released, given the advancements in artificial intelligence applications.

3.2. Wildfires

Wildfires, even though they may be human-ignited, are typically driven by environmental factors such as drought, wind, topography and vegetation. The primary factor in the evolution of a disaster is the potentially rapid wind-driven spread, which is also the major concern of the emergency services. Note that secondary CBRN hazards may be associated with wildfires due to the ensuing air pollution and toxic smoke. When referring to wildfire monitoring, the main focus concerns the position of the fire front (connected with the boundary tracking problem in CBRN), its *Rate of Spread* (RoS) and/or the maximum height of the flames (Merino et al., 2012).

Satellite measurements are the most known operative solutions and, beside the role of the EMS of Copernicus, an additional example from private companies comes from Google. They are actively focusing on machine learning solutions to detect, track and simulate wildfires, providing solutions like their Wildfire Boundary Tracking, Wildfire Simulation, FireBench, and FireSat (Google Research, 2025). Unmanned systems are a natural alternative and are currently widely investigated.

An extensive presentation of the current research towards the use of artificial intelligence techniques for wildfire modeling, detection, classification, segmentation and monitoring is provided by Boroujeni et al. (2024). UAV-based solutions are thoroughly discussed as, even if of different types (i.e., fixed-wings, rotary-wings, etc.), the interest for wildfire contexts mainly focuses on aerial devices. Furthermore, Boroujeni et al. (2024) give a comprehensive review of various device characteristics and utilized sensors that, in contrast to CBRN contexts, are predominantly intended for computer vision based approaches. Connected, but concentrating on the autonomy of UAVs for information gathering during wildfires, the study of Bailon-Ruiz and Lacroix (2020) analyze situation awareness, decision making and collaboration ability for AUS. Mapping is regarded as the most common mission, but time-dependent monitoring is common as well. Furthermore, given the extension in space and time of wildfires, fleets of UAVs are identified in Bailon-Ruiz and Lacroix (2020) as likely the most useful configuration to gain situation awareness in such a context.

In the following, Section 3.2.1 presents some developed frameworks and strategies for wildfire monitoring where AUS adaptively move for information gathering. However, besides information gathering, and differently from most of the CBRN applications, a second sequential problem is also studied with wildfires: suppressing the fire in selected zones. Several studies consider the two decision processes jointly and some examples are discussed in Section 3.2.2. Additional insights on the AUS operations are given in Section 3.2.3.

3.2.1. Monitoring

A renowned study with field tests for wildfire monitoring is Merino et al. (2012). In such a work, a team of heterogeneous UAVs is maneuvered with a central decisional system, following the architecture presented in Gancet et al. (2005). Besides standard sensors for navigation (i.e., GPS receivers and inertial measurement units), infrared and visual cameras are used for data collection. The integration of new data is processed with a discrete Bayes filter over a rectangular grid representation of the hazard, where a probability of fire and complete exhaustion is assigned for each cell.

A more recent example of a complete framework for wildfire monitoring with UAV is presented and tested on mixed reality in Bailon-Ruiz et al. (2022). The aim is to provide information on the perimeter to first responders. Therefore, UAVs movements must maximize information

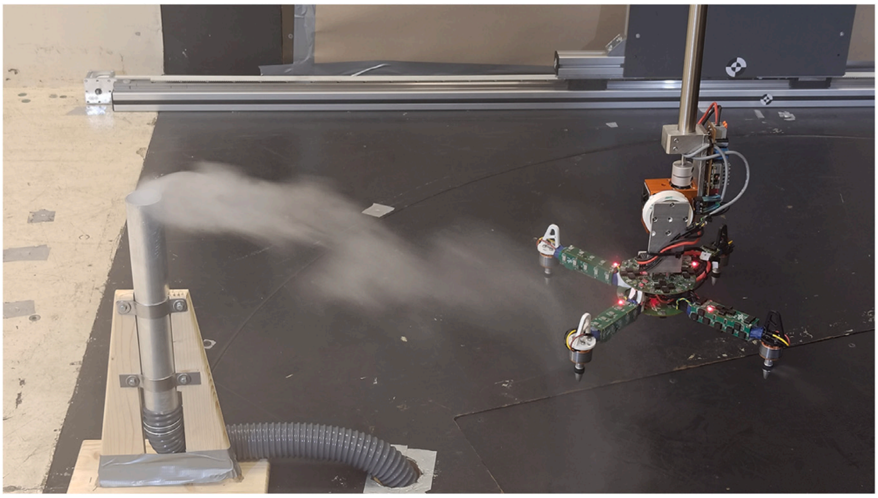

Fig. 4. Recording gas concentration with quadrotor systems. Photo: DLR, CC BY-NC-ND 3.0.

gain with respect to such a goal. The mission plan is not updated in real-time on board, but a spatial coverage problem (see, e.g., Glock & Meyer, 2023) that takes into account the dynamics of the fire and the physical constraints of the selected device (fixed-wing) is solved using a metaheuristic. Collected data are integrated in a propagation model by interpolation of the points observed within the front and the system naturally allows for sequential re-planning during the wildfire evolution. Therefore, adaptiveness is intended between batches of observations and not online and on board.

More theoretical studies are numerous, especially considering the vast amount of contributions focusing on general-purpose multi-device coordination and the growing number of model-free formation control policies developed with deep architectures. To mention some examples in wildfires with sequential decisions in real-time, Shobeiry et al. (2021) use twin-model experiments with DEVS-FIRE simulator as ground-truth and a simple linear propagation of the front as an embedded model, formulating the problem as a POMDP and practically implementing a Kalman filter for data integration. The mission is planned according to an MPC-like approach that minimizes the trace of the covariance matrix of the Kalman filter. Alternatively, another example without embedded complex models but simplified assumptions to allow for on-board real-time decisions is Pham et al. (2020). In particular, they consider a coverage and tracking distributed control policy with respect to an objective function made of a coverage cost dependent on the field of view of the camera-equipped UAVs, additionally weighted according to the fire heat intensity (assumed to be measurable quantitatively) to give importance to the front region. At a given time step, the desired position of a UAV follows the negative gradient of the coverage-tracking objective function and is further processed by a potential field control component to avoid collisions and ease formations. Tests are made using a simplified version of the well-known modeling tool FARSITE (Finney, 1998). The simplified version of this model is also used more explicitly within movement policies proposed in the literature. For example, in Seraj and Gombolay (2020), an adaptive extended Kalman filter is developed to integrate data and move a team of UAVs according to the corresponding uncertainty of the fire front according to the filter, a formation control component and, more importantly, a focus on areas where humans are operating; thus aligning the goal to a more human-centered problems rather than just coverage or tracking. Notably, some studies for UAV-based monitoring focus more on the data integration problem rather than the mission planning one, (Lin et al., 2019) being an example of using FARSITE and developing a Kalman filter with theoretical guarantees.

High-level movement policies based on deep architectures are also common and examples are Julian and Kochenderfer (2019) and Viseras et al. (2021). More specifically, in Julian and Kochenderfer (2019), two policies are presented: the first one directly reacting to observations of burning zones with no explicit representation of the wildfire besides the immediate local information. The other, investigating the value of an embedded belief made of a two-layers cell-based domain discretization with a first level concerning a binary value for burning zones and the second counting the time steps since the last observation of the zone. The study of Viseras et al. (2021) builds upon (Julian & Kochenderfer, 2019), implementing additional multi-agent deep policies focusing on cooperative learning approaches.

3.2.2. *Suppression & monitoring*

During wildfires, agents can directly affect the hazard by suppressing the fire in selected zones. This is a sequential decision problem on its own and the current understanding of the hazard dynamics is of great value for selecting where to operate. Furthermore, the state of the system for the observational system is directly altered by operational decisions, thus influencing the mission planning. The suppression problem relates to *Dynamic Resource Allocation* problems and Bertsimas et al. (2017) is a numerical example tackling it. In particular, they sequentially decide where to assign suppression teams, aiming to minimize the cost related to burning zones. An MPC-based and *Monte Carlo Tree Search* (MCTS) suppression policy are compared; dealing with uncertainty by simulating, for each decision step, the stochastic system, in the MCTS case, and using deterministic linear approximation in the MPC approach. The propagation is described by a boolean variable that follows stochastic dynamics, representing whether a specific zone is burning, and a fuel level variable of the same zone.

The AUS-based suppression and information gathering problems have been investigated together in the literature. A simulated example is Thul and Powell (2023), where authors consider a multi-agent system with a drone and a helicopter collaborating to extinguish a fire. The drone sequentially gathers information and the helicopter exploits the current belief about the wildfire to decide where to autonomously operate next. They use a set of difference equations as an embedded model of the fire fuel over different zones, equipped with a Bayesian structure with multinomial and Bernoulli distributions to take into account the dynamical adaptivity of the belief about the fire dynamics. They assume to observe the fire intensity directly, retrieving as data the burned fuel per observed zone. The goal is to mitigate the damage, intended as the cost of burned fuel. The gathered information is directly used by another agent (the helicopter), having its own policy, state of the system, objective function, transition function, but shared goal. Another example is Chen et al. (2024), where the authors introduce a general multi-UAV framework for wildfire monitoring and suppression that adapts by assessing the presence of fire using vision sensors like thermal imaging cameras and aims to minimize the affected area. Integration of observations is made with a direct Bayesian approach over a domain discretized in cells. Different from Thul and Powell (2023), where agents with specific characteristics have different roles, here each UAV has both the movement and the suppressing operation action. The abstract formulation of the scheduler allows for any stochastic embedded model that simulates the fire spreading. Devices are moved using an MPC policy based on the information gain of a location (with entropy-based evaluations) and the operational gain of a suppression (leveraging on the current understanding of the hazard).

3.2.3. Further considerations

Given the importance of computer vision and, more in general, artificial intelligence based systems for monitoring and tracking wildfires, data are fundamental to train and test practical solutions. Available visual datasets concerning fire and smoke have been cataloged and reviewed by Boroujeni et al. (2025). Such datasets have also been of particular value when considering the increasing interest in deep architectures. In fact, deep learning solutions for detection, segmentation and classification are fundamental. Nevertheless, it is also worth mentioning that, particularly for mission planning for monitoring applications, model-based abstractions for data integration and policies are equally important.

The value of data and simulations is even greater when considering the complexity and required resources of field tests. As an example, in the experiments of Merino et al. (2012), eighty firemen were required to allow for semi-operative tests. Besides the simulation frameworks mentioned in 3.1.4, Robotarium (Pickem et al., 2017) is a useful framework to remotely perform initial tests on real robotic hardware; especially useful for UAVs fleet experiments and already used in, e.g., Seraj and Gombolay (2020).

For practical applications with multiple UAVs, recent publications have put forward a system that integrates autonomous agent implementations for cooperative agents based on a Belief–Desire–Intention (BDI) model in a ROS-based communication scheme (de Oliveira Silvestre et al., 2025, 2023). While evaluated in simulated wildfire monitoring scenarios, the corresponding mission planning library is, by its nature, application-independent. Similar to the approaches discussed in Section 3.1.4, these publications start bridging the gap between theoretical works and real-world implementations.

In conclusion, autonomous wildfire monitoring with UAVs is, as of today, a developing technology that also has to face hardware complexities like effective communications within the UAVs fleet and navigation–perception errors and interference.

3.3. Floods & storm systems

The last family of problems investigated in this study concerns floods and storm systems. Mainly driven by weather and climate aspects and influenced by specific geographical and hydrological features of the affected area, they include disaster scenarios of vastly different scales. Dedicated modeling tools exist that can incorporate hydrological aspects, rainfall as well as terrain features, etc. (e.g., the United States Environmental Protection Agency Storm Water Management Model). Especially in the case of flash floods and heavy rainfall, the primary concern is the sudden onset and potentially rapid evolution. Consequently, the use of unmanned systems in monitoring, mapping and detection of floods has received wide attention. Focusing on UAVs, the bibliometric review of Iqbal et al. (2023) shows how contributions related to flood management applications are increasing over all the phases of such emergencies.

A necessary mention regards river flow monitoring. In fact, waterway measurements can be useful for situation awareness and a practical proof-of-concept demonstrating progress towards the use of adaptive movement policies of AUS is given by Dille et al. (2025). The problem is a mapping one, as adaptivity concerns accuracy metrics for surface velocity distribution rather than dynamic changes during the mission. Nevertheless, initial autonomous field tests with adaptive planning strategies for dwell time and way-point sequences are presented. Furthermore, it merits attention that the machine-generated way-points were communicated to human pilots during field tests for safety and regulatory ease, thus mimicking the fully autonomous behavior. In fact, regulatory constraints are often neglected in theoretical studies, but are fundamental constraints during field tests.

Closer to the topic treated in this study, various studies offer a preliminary theoretical perspective, disregarding some real-world complexities but providing initial insights. The adaptive use of measurement of water velocity together with water depth is discussed for dynamic flooded urban scenarios in Exadaktylos et al. (2022), where a UAV is assumed to monitor a flood event on a theoretical case study over the city center of Nicosia in Cyprus. A simplified flooding model is used for sequential observational way-point selection, while adjusting the rainfall intensity. Water depth is also considered in Garg and Jha (2024), where deep reinforcement learning policies for monitoring areas of high interest with a swarm of UAVs are trained using a dynamic hydrological model and a reward based on: a collision penalty, a dispersion factor and a critical level dependent on population density and water level. Other examples of movement policies based on deep-architectures trained on flood simulations have been studied in Baldazo et al. (2019), where fixed wings UAVs track the progression of floods over random generated terrain following the assumed collapse of a dam.

Boundary tracking problems have also been investigated for floods; following similar geometrical paradigms and assumptions of the aforementioned application contexts. For example, simple control strategies simulated in ROS and Gazebo, where UAVs automatically adapt their distances within a fleet of devices are given in Bai et al. (2020). They assume groups of unmanned systems to track the mid-point of the land–flood regions with a vision-based algorithm, thus following the boundary, simultaneously covering the inner flooded regions.

Besides direct measurements from unmanned devices, sensors can be dropped from aerial devices for in-situ measurements (i.e., *dropsondes*). Where to drop them is a problem connected to the movement planning of the device and an example for monitoring floods by inverse modeling is discussed in Abdelkader et al. (2013) and then in Abdelkader et al. (2014). Even if observations are not integrated in the trajectory update within the presented simulations, the planning problem for the unmanned system addressed in Abdelkader et al. (2014) uses the current belief with an ensemble of models, polynomial chaos expansion and least squares to minimize the uncertainty on the estimated parameters.

It is also worth mentioning that the use of unmanned systems and adaptive sampling strategies for storm systems monitoring in purely forecast problems is a very active research topic. Specifically, in-situ observations during the development of storms like tropical cyclones are critical for data assimilation into meteorological numerical models and real-time assessment. For such purposes, the use of autonomous systems is increasingly common, together with advanced sampling strategies for maximizing the impact on forecast models (Holbach et al., 2023). A practical example moving towards an adaptive selection of observational locations given a simulated model is the NOAA's mission Sensing Hazard with Operational Unmanned Technology (SHOUT) (Wick et al., 2020). Focusing on tropical cyclones and winter storms, a Global Hawk drone from NASA has been used for direct observations and for releasing dropsondes for vertical readings, collecting data according to their usefulness to improve forecasts by using ensemble-based sensitivity algorithms. If observations are assimilated in the models used for mission planning, the UAV follows an adaptive policy. However, developing emergencies like tropical cyclones usually have way longer time frames and spatial domains than, e.g., a local gas leakage. Therefore, even if adaptive to the current belief on the dynamic process, missions are planned days before the flights and have humans-in-the-loop, possibly changing with short notice due to, e.g., air traffic control requirements and external technical cooperation. Simpler unmanned systems like buoyancy-driven balloons have also been investigated as adaptive observer for storm systems monitoring in, e.g., Floriano et al. (2025), where a policy based on a neural network model predictive control method is considered for maximizing coverage of the area of interest and managing intermittent communications.

4. Related approaches and generalized mechanisms

So far, the focus of this study has been on the core contexts of dynamic hazard monitoring. However, there are numerous related topics and methodological approaches that, although outside the primary scope of this work, naturally connect to research in emergency applications. Most importantly, this includes lines of research that apply similar concepts and mechanisms that can also be subsumed under the framework introduced in Section 2.4, namely informative path planning, search and rescue and optimal sensor placement.

Due to their generality, these approaches may offer promising insights from an interdisciplinary point of view, and readers may benefit from a delineation between the different disciplines. The approaches summarized below all belong to the same broad conceptual class, which is interested in obtaining information on a phenomenon with spatial aspects subject to sensing restrictions, and have a somewhat broader scope instead of focusing on environmental monitoring. However, they come from different academic traditions, leading to different nomenclature, a largely separated literature and emphasizing different aspects in their solution techniques. Future research would likely benefit from a more interdisciplinary perspective, fusing the advances originating from diverse fields.

4.1. Informative path planning

Informative Path Planning (IPP) explicitly aims to design trajectories that maximize the acquisition of valuable information about a target environment within a given resource budget. This is typically achieved by developing paths that minimize uncertainty about the environment or specific parameters of interest. Specifically, probabilistic models (e.g., Gaussian processes) and information-theoretic metrics, such as mutual information, often serve as a basis for quantifying the value of information gained and constructing reward functions that guide optimal control policies (Sung et al., 2023). These measures can be integrated into sampling-based planners, such as rapidly exploring random tree (RRT) algorithms, which explicitly account for environmental uncertainty by simulating multiple extensions to the search tree under varying conditions. When information is used to update the paths online, the nomenclature refers to *Adaptive* Informative Path Planning (AIPP), for which (Popović et al., 2024) provides a comprehensive review of applications and solution approaches. Note that neither IPP nor AIPP specifically focuses on dynamically evolving environments; instead, even adaptive applications tend to be at least in a steady state, with few exceptions (e.g., Meliou et al., 2007).

As (A)IPP is primarily directed towards efficient data collection via mobile sensor systems, it is closely related to hazard monitoring as it offers strategies for information gathering from a generalized point of view. Research in this domain is predominantly driven by the robotics community, with a traditionally vast range of possible applications that range from environmental monitoring to target detection, the latter being more closely related to search and rescue (see Section 4.2). Interestingly, even though closely related from a methodological point of view, the overlap between the literature on AIPP and the research domains covered in this study is surprisingly sparse. Probably this is not due to a fundamental inapplicability or discrepancy in methodology. Instead, literature on AIPP tends to focus on environmental mapping from a more general point of view, and is largely driven by academia with a particular emphasis on formalizing information gain and deriving efficient strategies and (provable) performance guarantees. Studies on hazard monitoring originate in a larger number of domains and consider different stakeholders and objectives, leading to more heterogeneous approaches that rarely, if ever, directly use the methods or terminology introduced in AIPP literature. Nonetheless, the shared conceptual and methodological approaches are promising for future research that might leverage insights from different domains.

4.2. Search and rescue

One of the most widely studied uses for autonomous systems in emergencies is *Search and Rescue* (S&R), where AUSs provide rapid assessment and victim localization capabilities. For instance, UAVs equipped with high-resolution imagery and thermal sensors are crucial for locating missing persons in challenging conditions. AGVs complement aerial surveillance by accessing areas that may be contaminated, unstable, or otherwise unsafe for human crews. A recent survey of UAVs in this domain is provided in Quero and Martinez-Carranza (2025), although algorithmic approaches for mission planning are only of secondary interest.

In addition to nonadaptive strategies, e.g., coverage path planning approaches such as lawnmower-based strategies or preplanned paths (Scherer et al., 2015), learning mechanisms play an increasing role in path and mission planning in S&R applications. For instance, belief maps of where victims might be can be updated iteratively, and adaptive policies may adjust to new information on environmental conditions such as blocked routes. Moreover, SLAM techniques (see Section 2.3.1) enable robot movements within unknown environments, supplemented by, e.g., deep learning strategies for victim and object recognition or for exploring cluttered environments (Niroui et al., 2019).

Similar to AIPP and emergency monitoring, S&R is concerned with navigating one or several agents through an environment to gather information. All three domains rely on probabilistic models of the physical environment, though in S&R this typically includes both a target-location distribution and a representation of traversal constraints (terrain, hazards), whereas in AIPP it is often a spatial or spatiotemporal field model. The shared challenge of balancing exploration and exploitation in the face of tight budget constraints means that the underlying ideas behind the applied approaches are typically very closely related, with both lines of research applying sequential decision-making frameworks to ingest and adapt to new information.

Despite these similarities, the underlying objectives and model assumptions often differ. S&R focuses on maximizing the probability of detecting targets, usually in hazardous conditions, whereas AIPP applications emphasize the estimation of spatiotemporal phenomena. Moreover, many S&R applications assume targets to be stationary, which means that adaptive methods focus on adjusting to newly obtained data but not to an inherently dynamic environment. While there are fundamental similarities in how uncertainty about the *physical* environment is modeled, obstacles and traversal constraints are much more emphasized in the literature on S&R as opposed to AIPP or emergency monitoring. The most promising overlap is related to *active perception*, which, in the context of S&R, is an umbrella term referring to techniques that enable agents to choose sensing inputs that improve overall mission outcome, e.g., by moving to other sensing locations but also by modifying sensor configurations with or without physical movement (Queralta et al., 2020).

4.3. Optimal sensor placement and design of experiments

A brief mention of *Optimal Sensor Placement* (OSP) and the broader fields of *Design of Experiments* (DoE) concludes this section; which can be seen as a more generalized problem of choosing sensor locations subject to constraints to obtain as much information about a phenomenon as possible. Such methods aim to minimize uncertainty using data-driven analysis that reflects environmental and sensor constraints, generally assuming static sensors.

The primary objective of OSP is to identify sensor locations that yield the most informative data, thereby minimizing monitoring costs while simultaneously meeting predefined performance requirements that may be related to locations where direct observations may not be possible. Planning objectives can be, e.g., coverage (i.e., percentage of the monitored area) as a proxy for the obtained information, target detection probability, expected uncertainty reduction or robustness. Simpler strategies may also apply algorithms such as Voronoi coverage control that rely on spatial features exclusively to ensure a stable distribution of sensors. When information-theoretic criteria (e.g., mutual information) are used, these strategies bear close resemblance to those discussed in Section 2.3.3. For static placements, these strategies correspond to submodular optimization, i.e., the optimization of a submodular objective function. For these algorithms, Krause et al. (2008) have shown that a simple greedy algorithm yields a near-optimal sensor set.

When routing considerations come into effect, the problem becomes more closely related to informative path planning. In fact, AIPP, S&R and hazard monitoring can be seen as dynamic extensions of the (static) OSP problem, with several placement problems being solved sequentially. From this point of view, approaches in OSP and DoE may offer insights into designing new policies that exploit, e.g., strong performance guarantees from sensor placement strategies, or translate multi-sensor placement strategies to multi-agent cooperative approaches. For example, recent approaches have discussed translating optimality criteria from OSP to AIPP (Ott et al., 2024).

5. Conclusions, limitations, and future directions

The use of autonomous unmanned systems for hazard monitoring during emergencies holds valuable potential and can benefit from efficient mission planning strategies to enhance information gathering. Although this work does not aim to list comprehensively every article or give detailed specifications for all the endless technical possibilities across the diverse application contexts and theoretical frameworks discussed, it highlights an interdisciplinary and unified perspective on the challenges, common problems, and interconnections between various applications and strategies in adaptive mission planning for autonomous unmanned systems and dynamic environmental hazard monitoring during emergencies. The formalization of sequential frameworks and control policies spans multiple research domains, yet connects logical foundations based on emerging unifying frameworks in control theory and machine learning. Moreover, since the integration of data within adaptive policies also varies across different fields, shared connections are emphasized from an interdisciplinary perspective for developing effective solutions. In fact, some common problems, such as boundary tracking and fleet coordination, offer generalized paradigms applicable across various contexts like wildfires, floods, and CBRN scenarios. Conversely, other methods leverage idiosyncratic features, specialized sensors, and domain-specific algorithms to address unique challenges and dynamics.

Also, limitations overlap between different contexts. For example, observers are often detached from their physical constraints, being treated as idealized point masses without inspection of their real-world complexities like localization and control uncertainties. Other aspects, such as energy conservation, are rarely integrated into these approaches due to the added complexity and contradictory requirements. Similarly, inherent complications of sensors like noise, calibration errors, or environmental interference are often overlooked, leading to simplified models that, while providing useful insights, may require additional care if deployed under operational stress. Regulatory constraints, weather variability, and the unpredictability of natural hazards further complicate robust deployments. For instance, adverse weather conditions can degrade sensor performance, disrupt communication, or limit flight duration, while regulatory frameworks often restrict airspace access, particularly in densely populated or critical infrastructure zones.

A core issue lies in balancing sensitive domain needs with current technical limitations, particularly in achieving full autonomy and robust performance in unpredictable and dynamic conditions. Many studies limit themselves to simulations, missions in controlled environments, or theoretical performance guarantees. Results obtained with models ready to run on industry-standard hardware in real-world scenarios are rare, and as the ground truth is often unknown, their performance is difficult to assess. This is particularly relevant in the case of computationally expensive methods and models that require extensive model design phases, such as reinforcement learning models, which often remain limited to simulation studies.

Many existing studies in this field are relatively small in scale, rely on unpublished benchmarks, or differ substantially in terms of application, AUS types, and experimental setup. As a result, there is currently no standardized benchmarking framework or simulation environment. This lack of consistency systematically limits both the comparability and further development of algorithms. No means exist for consistently evaluating related approaches, neither within nor across the discussed application scenarios. To address this, dedicated frameworks should provide appropriate simulation environments (e.g., Gazebo), utilize standardized communication protocols such as ROS, and allow for the uniform specification of vehicle and sensor properties. Such an approach would enable the currently fragmented literature to become more comparable and thus foster more systematic progress in the field. These challenges underscore the need for context-aware systems that balance technical capabilities with operational feasibility. Nevertheless, the topic is attracting scientific interest, and it is likely to be increasingly helpful for improving response capabilities during emergencies.

Despite the many advances made in the field during the last years, to date, fully functioning solutions for truly autonomous monitoring of environmental hazards during emergencies are rarely achieved in real emergency operations; most deployed solutions remain human-supervised and are limited to tightly constrained operating conditions. Numerous research groups from multiple domains, such as robotics, control theory, machine learning/reinforcement learning and operations research, actively address the key subtasks required for such systems, including physical-environment modeling, collision-free navigation of AUSs, assimilating sensor measurements into environmental models, and autonomously selecting future sensing locations. A growing body of literature is connecting relevant approaches from different domains. However, the overlap in literature and coordination of the different research areas towards a common goal has yet to gain even more traction.

The establishment of a unified understanding, nomenclature and terminology would certainly facilitate such collaborations. In this direction, a promising development of recent years can be seen in efforts to develop general frameworks for sequential decision-making, which explicitly capture the interplay between actions, evolving system state, and information acquisition; central to context-aware autonomous hazard monitoring. Nevertheless, the diversity of emergency hazards implies that no single method can be optimal across contexts. Method and tool depend on both the hazard characteristics and the available technology. Therefore, future solutions to fully autonomous monitoring of environmental hazards will likely exhibit a diversity of approaches, while hopefully relying on a set of joint core principles.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

We gratefully acknowledge the DLR Institute of Communications and Navigation for their permission to use their images in this article.

Data availability

No data was used for the research described in the article.

References

Abdelkader, M., Shaqura, M., Claudel, C. G., & Gueaieb, W. (2013). A UAV based system for real time flash flood monitoring in desert environments using Lagrangian microsensors. In *2013 international conference on unmanned aircraft systems* (pp. 25–34). http://dx.doi.org/10.1109/ICUAS.2013.6564670.

Abdelkader, M., Shaqura, M., Ghommem, M., Collier, N., Calo, V., & Claudel, C. (2014). Optimal multi-agent path planning for fast inverse modeling in UAV-based flood sensing applications. In *2014 international conference on unmanned aircraft systems* (pp. 64–71). http://dx.doi.org/10.1109/ICUAS.2014.6842239.

Asadzadeh, S., de Oliveira, W. J., & de Souza Filho, C. R. (2022). UAV-based remote sensing for the petroleum industry and environmental monitoring: State-of-the-art and perspectives. *Journal of Petroleum Science and Engineering*, *208*, Article 109633. http://dx.doi.org/10.1016/j.petrol.2021.109633.

Asch, M., Bocquet, M., & Nodet, M. (2016). *Data assimilation: methods, algorithms, and applications*. SIAM.

Atyabi, A., MahmoudZadeh, S., & Nefti-Meziani, S. (2018). Current advancements on autonomous mission planning and management systems: An AUV and UAV perspective. *Annual Reviews in Control*, *46*, 196–215. http://dx.doi.org/10.1016/j.arcontrol.2018.07.002.

Bai, Y., Asami, K., Svinin, M., & Magid, E. (2020). Cooperative multi-robot control for monitoring an expanding flood area. In *2020 17th international conference on ubiquitous robots* (pp. 500–505). http://dx.doi.org/10.1109/UR49135.2020.9144931.

Bailon-Ruiz, R., Bit-Monnot, A., & Lacroix, S. (2022). Real-time wildfire monitoring with a fleet of UAVs. *Robotics and Autonomous Systems*, *152*, Article 104071. http://dx.doi.org/10.1016/j.robot.2022.104071.

Bailon-Ruiz, R., & Lacroix, S. (2020). Wildfire remote sensing with UAVs: A review from the autonomy point of view. In *2020 international conference on unmanned aircraft systems* (pp. 412–420). IEEE, http://dx.doi.org/10.1109/ICUAS48674.2020.9213986.

Baldazo, D., Parras, J., & Zazo, S. (2019). Decentralized multi-agent deep reinforcement learning in swarms of drones for flood monitoring. In *2019 27th European signal processing conference* (pp. 1–5). IEEE, http://dx.doi.org/10.23919/EUSIPCO.2019.8903067.

Bertsekas, D. P. (2022). *Lessons from AlphaZero for optimal, model predictive, and adaptive control*. Athena Scientific.

Bertsekas, D. P. (2024). Model predictive control and reinforcement learning: A unified framework based on dynamic programming. *IFAC-PapersOnLine*, *58*, 363–383. http://dx.doi.org/10.1016/j.ifacol.2024.09.056, 8th IFAC Conference on Nonlinear Model Predictive Control NMPC 2024.

Bertsimas, D., Griffith, J. D., Gupta, V., Kochenderfer, M. J., & Mišić, V. V. (2017). A comparison of Monte Carlo tree search and rolling horizon optimization for large-scale dynamic resource allocation problems. *European Journal of Operational Research*, *263*, 664–678. http://dx.doi.org/10.1016/j.ejor.2017.05.032.

BlackSky (2023). Respond to disasters faster with on-demand satellite imagery. https://blacksky.com/respond-to-disasters-faster-with-on-demand-satellite-imagery/. (Accessed 05 May 2026).

Boroujeni, S. P. H., Mehrabi, N., Afghah, F., McGrath, C. P., Bhatkar, D., Biradar, M. A., & Razi, A. (2025). Fire and smoke datasets in 20 years: An in-depth review. http://dx.doi.org/10.48550/arXiv.2503.14552, ArXiv Preprint.

Boroujeni, S. P. H., Razi, A., Khoshdel, S., Afghah, F., Coen, J. L., O'Neill, L., Fule, P., Watts, A., Kokolakis, N. M. T., & Vamvoudakis, K. G. (2024). A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management. *Information Fusion*, Article 102369. http://dx.doi.org/10.1016/j.inffus.2024.102369.

Buizza, C., Quilodrán Casas, C., Nadler, P., Mack, J., Marrone, S., Titus, Z., Le Cornec, C., Heylen, E., Dur, T., Baca Ruiz, L., Heaney, C., Díaz Lopez, J. A., Kumar, K. S., & Arcucci, R. (2022). Data learning: Integrating data assimilation and machine learning. *Journal of Computational Science*, *58*, Article 101525. http://dx.doi.org/10.1016/j.jocs.2021.101525.

Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., & Leonard, J. J. (2016). Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, *32*, 1309–1332. http://dx.doi.org/10.1109/TRO.2016.2624754.

Campbell, S., O'Mahony, N., Carvalho, A., Krpalkova, L., Riordan, D., & Walsh, J. (2020). Path planning techniques for mobile robots a review. In *2020 6th international conference on mechatronics and robotics engineering* (pp. 12–16). IEEE, http://dx.doi.org/10.1109/ICMRE49073.2020.9065187.

Carrassi, A., Bocquet, M., Bertino, L., & Evensen, G. (2018). Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *Wiley Interdisciplinary Reviews: Climate Change*, *9*, Article e535. http://dx.doi.org/10.1002/wcc.535.

Chen, X. x., & Huang, J. (2019). Odor source localization algorithms on mobile robots: A review and future outlook. *Robotics and Autonomous Systems*, *112*, 123–136. http://dx.doi.org/10.1016/j.robot.2018.11.014.

Chen, X., Xiao, Z., Cheng, Y., Hsia, C. C., Wang, H., Xu, J., Xu, S., Dang, F., Zhang, X. P., Liu, Y., & Chen, X. (2024). SOScheduler: Toward proactive and adaptive wildfire suppression via multi-UAV collaborative scheduling. *IEEE Internet of Things Journal*, http://dx.doi.org/10.1109/JIOT.2024.3389771.

Cheng, S., Quilodrán-Casas, C., Ouala, S., Farchi, A., Liu, C., Tandeo, P., Fablet, R., Lucor, D., Iooss, B., Brajard, J., Xiao, D., Janjic, T., Ding, W., Guo, Y., Carrassi, A., Bocquet, M., & Arcucci, R. (2023). Machine learning with data assimilation and uncertainty quantification for dynamical systems: A review. *IEEE/CAA Journal of Automatica Sinica*, *10*, 1361–1387. http://dx.doi.org/10.1109/JAS.2023.123537.

Christopoulos, V. N., & Roumeliotis, S. (2005). Adaptive sensing for instantaneous gas release parameter estimation. In *Proceedings of the 2005 IEEE international conference on robotics and automation* (pp. 4450–4456). IEEE, http://dx.doi.org/10.1109/ROBOT.2005.1570805.

Copernicus (2025). Copernicus emergency management service - mapping. https://mapping.emergency.copernicus.eu/. (Accessed 09 August 2025).

Darynova, Z., Blanco, B., Juery, C., Donnat, L., & Duclaux, O. (2023). Data assimilation method for quantifying controlled methane releases using a drone and ground-sensors. *Atmospheric Environment: X*, *17*, Article 100210. http://dx.doi.org/10.1016/j.aeaoa.2023.100210.

De Kerf, T., Sels, S., Samsonova, S., & Vanlanduit, S. (2024). A dataset of drone-captured, segmented images for oil spill detection in port environments. *Scientific Data*, *11*, 1180. http://dx.doi.org/10.1038/s41597-024-03993-8.

de Oliveira Silvestre, I., de Lima, B., Dias, P., Buss Becker, L., Fisher, M., Hübner, J., & de Brito, M. (2025). Enhanced agent-oriented programming for robot teams. *Engineering Applications of Artificial Intelligence*, *158*, Article 111390. http://dx.doi.org/10.1016/j.engappai.2025.111390.

de Oliveira Silvestre, I., de Lima, B., Dias, P., Buss Becker, L., Hübner, J., & de Brito, M. (2023). UAV swarm control and coordination using Jason BDI agents on top of ROS. In *International conference on practical applications of agents and multi-agent systems* (pp. 225–236). Springer, http://dx.doi.org/10.1007/978-3-031-37616-0_19.

Denis, G., de Boissezon, H., Hosford, S., Pasco, X., Montfort, B., & Ranera, F. (2016). The evolution of earth observation satellites in europe and its impact on the performance of emergency response services. *Acta Astronautica*, *127*, 619–633. http://dx.doi.org/10.1016/j.actaastro.2016.06.012.

Deutsches Zentrum für Luft- und Raumfahrt (2024). When swarms of autonomous robots sniff out hazardous gases. https://www.dlr.de/en/latest/news/2024/when-swarms-of-autonomous-robots-sniff-out-hazardous-gases. (Accessed 09 August 2025).

Dille, M., Vespignani, M., Bruce, J., & Wong, U. (2025). Advancing river flow monitoring with small uncrewed aircraft and simulation-driven development. *Frontiers in Remote Sensing*, *6*, Article 1520963. http://dx.doi.org/10.3389/frsen.2025.1520963.

Egorova, T., Gatsonis, N. A., & Demetriou, M. A. (2016). Estimation of gaseous plume concentration with an unmanned aerial vehicle. *Journal of Guidance, Control, and Dynamics*, *39*, 1314–1324. http://dx.doi.org/10.2514/1.G001453.

Euler, J., Ritter, T., Ulbrich, S., & von Stryk, O. (2015). Centralized ensemble-based trajectory planning of cooperating sensors for estimating atmospheric dispersion processes. In *International conference on dynamic data-driven environmental systems science* (pp. 322–333). Springer, http://dx.doi.org/10.1007/978-3-319-25138-7_29.

Euler, J., & von Stryk, O. (2017). Decentralized data-driven control of cooperating sensor-carrying UAVs in a multi-objective monitoring scenario. *IFAC-PapersOnLine*, *50*, 15828–15834. http://dx.doi.org/10.1016/j.ifacol.2017.08.2316, 20th IFAC World Congress.

European Space Agency (2025). Sentinel-1. https://sentiwiki.copernicus.eu/web/sentinel-1. (Accessed 09 August 2025).

Exadaktylos, S., Kolios, P., & Eliades, D. G. (2022). An integrated modelling framework for drone-based flood monitoring. In *2022 international conference on unmanned aircraft systems* (pp. 636–645). IEEE, http://dx.doi.org/10.1109/ICUAS54217.2022.9836177.

Finney, M. A. (1998). *FARSITE: Fire area simulator—model development and evaluation*. Ogden UT: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station, Revised 2004.

Floriano, B. R., Hanson, B., Bewley, T., Ishihara, J. Y., & Ferreira, H. C. (2025). A novel policy for coordinating a hurricane monitoring system using a swarm of buoyancy-controlled balloons trading off communication and coverage. *Engineering Applications of Artificial Intelligence*, *139*, Article 109495. http://dx.doi.org/10.1016/j.engappai.2024.109495.

Gancet, J., Hattenberger, G., Alami, R., & Lacroix, S. (2005). Task planning and control for a multi-UAV system: Architecture and algorithms. In *2005 IEEE/RSJ international conference on intelligent robots and systems* (pp. 1017–1022). IEEE, http://dx.doi.org/10.1109/IROS.2005.1545217.

Garg, A., & Jha, S. S. (2024). Learning continuous multi-UAV controls with directed explorations for flood area coverage. *Robotics and Autonomous Systems*, *180*, Article 104774. http://dx.doi.org/10.1016/j.robot.2024.104774.

Geer, A. J. (2021). Learning earth system models from observations: Machine learning or data assimilation? *Philosophical Transactions of the Royal Society, Series A*, *379*, Article 20200089. http://dx.doi.org/10.1098/rsta.2020.0089.

Gioia, D. G., Bonari, J., Lichte, D., & Popp, A. (2024). Sequential drone routing for data assimilation on a 2D airborne contaminant dispersion problem. In *2024 sensor data fusion: trends, solutions, applications* (pp. 1–8). http://dx.doi.org/10.1109/SDF63218.2024.10773899.

Glock, K. (2020). *Emergency rapid mapping with drones: models and solution approaches for offline and online mission planning* (Ph.D. thesis), Karlsruher Institut für Technologie (KIT), http://dx.doi.org/10.5445/IR/1000124518.

Glock, K., & Meyer, A. (2020). Mission planning for emergency rapid mapping with drones. *Transportation Science*, *54*, 534–560. http://dx.doi.org/10.1287/trsc.2019.0963.

Glock, K., & Meyer, A. (2023). Spatial coverage in routing and path planning problems. *European Journal of Operational Research*, *305*, 1–20. http://dx.doi.org/10.1016/j.ejor.2022.02.031.

Google Research (2025). Wildfires. https://sites.research.google/gr/wildfires/. (Accessed 09 August 2025).

Hinsen, P., Wiedemann, T., Fan, H., Munoz Martos, C., Scott Prieto Ruiz, V., Zhang, S., Shutin, D., & Lilienthal, A. (2024). ICASSP 2025 SP grand challenge: Gas source localization from real-world spatial in-situ concentration and wind measurements. *IEEE Dataport*, http://dx.doi.org/10.21227/x0sf-ad36.

Hinsen, P., Wiedemann, T., Shutin, D., & Lilienthal, A. J. (2023). Exploration and gas source localization in advection–diffusion processes with potential-field-controlled robotic swarms. *Sensors*, *23*, 9232. http://dx.doi.org/10.3390/s23229232.

Hoffer, N. V., Coopmans, C., Jensen, A. M., & Chen, Y. (2014). A survey and categorization of small low-cost unmanned aerial vehicle system identification. *Journal of Intelligent and Robotic Systems*, *74*, 129–145. http://dx.doi.org/10.1007/s10846-013-9931-6.

Holbach, H. M., Bousquet, O., Bucci, L., Chang, P., Cione, J., Ditchek, S., Doyle, J., Duvel, J. P., Elston, J., Goni, G., Hon, K. K., Ito, K., Jelenak, Z., Lei, X., Lumpkin, R., McMahon, C. R., Reason, C., Sanabia, E., Shay, L. K., ... Zhang, J. A. (2023). Recent advancements in aircraft and in situ observations of tropical cyclones. *Tropical Cyclone Research and Review*, *12*, 81–99. http://dx.doi.org/10.1016/j.tcrr.2023.06.001.

Hutchinson, M. (2019). *On the use of autonomous unmanned vehicles in response to hazardous atmospheric release incidents* (Ph.D. thesis), Loughborough University.

Hutchinson, M., Oh, H., & Chen, W. H. (2017). A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors. *Information Fusion*, *36*, 130–148. http://dx.doi.org/10.1016/j.inffus.2016.11.010.

Hutchinson, M., Oh, H., & Chen, W. H. (2018). Entrotaxis as a strategy for autonomous search and source reconstruction in turbulent conditions. *Information Fusion*, *42*, 179–189. http://dx.doi.org/10.1016/j.inffus.2017.10.009.

Hwang, J., Bose, N., & Fan, S. (2019). AUV adaptive sampling methods: A review. *Applied Sciences*, *9*, 3145. http://dx.doi.org/10.3390/app9153145.

Hwang, J., Bose, N., Millar, G., Bulger, C., & Nazareth, G. (2023). Bubble plume tracking using a backseat driver on an autonomous underwater vehicle. *Drones*, *7*, 635. http://dx.doi.org/10.3390/drones7100635.

Iqbal, U., Riaz, M. Z. B., Zhao, J., Barthelemy, J., & Perez, P. (2023). Drones for flood monitoring, mapping and detection: A bibliometric review. *Drones*, *7*, 32. http://dx.doi.org/10.3390/drones7010032.

Jin, W., Rahbar, F., Ercolani, C., & Martinoli, A. (2023). Towards efficient gas leak detection in built environments: Data-driven plume modeling for gas sensing robots. In *2023 IEEE international conference on robotics and automation* (pp. 7749–7755). IEEE, http://dx.doi.org/10.1109/ICRA48891.2023.10160816.

Jones, M., Djahel, S., & Welsh, K. (2023). Path-planning for unmanned aerial vehicles with environment complexity considerations: A survey. *ACM Computing Surveys*, *55*, 1–39. http://dx.doi.org/10.1145/3570723.

Julian, K. D., & Kochenderfer, M. J. (2019). Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning. *Journal of Guidance, Control, and Dynamics*, *42*, 1768–1778. http://dx.doi.org/10.2514/1.G004106.

Kato, N., Choyekh, M., Dewantara, R., Senga, H., Chiba, H., Kobayashi, E., Yoshie, M., Tanaka, T., & Short, T. (2017). An autonomous underwater robot for tracking and monitoring of subsea plumes after oil spills and gas leaks from seafloor. *Journal of Loss Prevention in the Process Industries*, *50*, 386–396. http://dx.doi.org/10.1016/j.jlp.2017.03.006.

Khodayi-mehr, R., Aquino, W., & Zavlanos, M. M. (2019). Model-based active source identification in complex environments. *IEEE Transactions on Robotics*, *35*, 633–652. http://dx.doi.org/10.1109/TRO.2019.2894039.

Kim, T., Jang, M., & Kim, J. (2021). A survey on simulation environments for reinforcement learning. In *2021 18th international conference on ubiquitous robots* (pp. 63–67). IEEE, http://dx.doi.org/10.1109/UR52253.2021.9494694.

Kowadlo, G., & Russell, R. A. (2008). Robot odor localization: A taxonomy and survey. *The International Journal of Robotics Research*, *27*, 869–894. http://dx.doi.org/10.1177/0278364908095118.

Krause, A., McMahan, H. B., Guestrin, C., & Gupta, A. (2008). Robust submodular observation selection. *Journal of Machine Learning Research*, *9*.

Kumar, P., Singh, S. K., Ngae, P., Feiz, A. A., & Turbelin, G. (2017). Assessment of a CFD model for short-range plume dispersion: Applications to the fusion field trial 2007 (FFT-07) diffusion experiment. *Atmospheric Research*, *197*, 84–93. http://dx.doi.org/10.1016/j.atmosres.2017.06.025.

Kurniawati, H. (2022). Partially observable Markov decision processes and robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, *5*, 253–277. http://dx.doi.org/10.1146/annurev-control-042920-092451.

Lettieri, E., Masella, C., & Radaelli, G. (2009). Disaster management: Findings from a systematic review. *Disaster Prevention and Management: An International Journal*, *18*, 117–136. http://dx.doi.org/10.1108/09653560910953207.

Leung, C., Huang, S., Kwok, N., & Dissanayake, G. (2006). Planning under uncertainty using model predictive control for information gathering. *Robotics and Autonomous Systems*, *54*, 898–910. http://dx.doi.org/10.1016/j.robot.2006.05.008.

Lin, Z., Liu, H. H., & Wotton, M. (2019). Kalman filter-based large-scale wildfire monitoring with a system of UAVs. *IEEE Transactions on Industrial Electronics*, *66*, 606–615. http://dx.doi.org/10.1109/TIE.2018.2823658.

Ljung, L. (2010). Perspectives on system identification. *Annual Reviews in Control*, *34*, 1–12. http://dx.doi.org/10.1016/j.arcontrol.2009.12.001.

LLC, C. L. (2025). Welcome to Colosseum, a successor of AirSim. https://github.com/CodexLabsLLC/Colosseum. (Accessed 09 August 2025).

Lu, P., Stumpf, A., Kerle, N., & Casagli, N. (2011). Object-oriented change detection for landslide rapid mapping. *IEEE Geoscience and Remote Sensing Letters*, *8*, 701–705. http://dx.doi.org/10.1109/LGRS.2010.2101045.

Macenski, S., Foote, T., Gerkey, B., Lalancette, C., & Woodall, W. (2022). Robot operating system 2: Design, architecture, and uses in the wild. *Science Robotics*, *7*, Article eabm6074. http://dx.doi.org/10.1126/scirobotics.abm6074.

Maxar Technologies (2025). Using geospatial insights from maxar imagery to help respond to the Los Angeles Wildfires. https://blog.maxar.com/earth-intelligence/2025/using-geospatial-insights-from-maxar-imagery-to-help-respond-to-the-los-angeles-wildfires. (Accessed 09 August 2025).

Meliou, A., Krause, A., Guestrin, C., & Hellerstein, J. M. (2007). Nonmyopic informative path planning in spatio-temporal models. In *Proceedings of the 22nd national conference on artificial intelligence - volume 1* (pp. 16–17). http://dx.doi.org/10.5555/1619645.1619742.

Merino, L., Caballero, F., Martínez-de Dios, J. R., Maza, I., & Ollero, A. (2012). An unmanned aircraft system for automatic forest fire monitoring and measurement. *Journal of Intelligent and Robotic Systems*, *65*, 533–548. http://dx.doi.org/10.1007/s10846-011-9560-x.

Monroy, J., Hernandez-Bennetts, V., Fan, H., Lilienthal, A., & Gonzalez-Jimenez, J. (2017). GADEN: A 3D gas dispersion simulator for mobile robot olfaction in realistic environments. *Sensors (Switzerland)*, *17*, 1–16. http://dx.doi.org/10.3390/s17071479.

Murphy, R. R., Tadokoro, S., & Kleiner, A. (2016). Disaster robotics. *Springer Handbook of Robotics*, 1577–1604. http://dx.doi.org/10.1007/978-3-319-32552-1_60.

Niroui, F., Zhang, K., Kashino, Z., & Nejat, G. (2019). Deep reinforcement learning robot for search and rescue applications: Exploration in unknown cluttered environments. *IEEE Robotics and Automation Letters*, *4*, 610–617. http://dx.doi.org/10.1109/LRA.2019.2891991.

Odonkor, P., Ball, Z., & Chowdhury, S. (2019). Distributed operation of collaborating unmanned aerial vehicles for time-sensitive oil spill mapping. *Swarm and Evolutionary Computation*, *46*, 52–68. http://dx.doi.org/10.1016/j.swevo.2019.01.005.

Ott, J., Kochenderfer, M. J., & Boyd, S. (2024). Approximate sequential optimization for informative path planning. *Robotics and Autonomous Systems*, *182*, Article 104814. http://dx.doi.org/10.1016/j.robot.2024.104814.

Pashna, M., Yusof, R., Ismail, Z. H., Namerikawa, T., & Yazdani, S. (2020). Autonomous multi-robot tracking system for oil spills on sea surface based on hybrid fuzzy distribution and potential field approach. *Ocean Engineering*, *207*, Article 107238. http://dx.doi.org/10.1016/j.oceaneng.2020.107238.

Pham, H. X., La, H. M., Feil-Seifer, D., & Deans, M. C. (2020). A distributed control framework of multiple unmanned aerial vehicles for dynamic wildfire tracking. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *50*, 1537–1548. http://dx.doi.org/10.1109/TSMC.2018.2815988.

Pickem, D., Glotfelter, P., Wang, L., Mote, M., Ames, A., Feron, E., & Egerstedt, M. (2017). The Robotarium: A remotely accessible swarm robotics research testbed. In *2017 IEEE international conference on robotics and automation* (pp. 1699–1706). IEEE, http://dx.doi.org/10.1109/ICRA.2017.7989200.

Popović, M., Ott, J., Rückin, J., & Kochenderfer, M. J. (2024). Learning-based methods for adaptive informative path planning. *Robotics and Autonomous Systems*, *179*, Article 104727. http://dx.doi.org/10.1016/j.robot.2024.104727.

Powell, W. B. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, *275*, 795–821. http://dx.doi.org/10.1016/j.ejor.2018.07.014.

Powell, W. B. (2022). *Reinforcement learning and stochastic optimization: a unified framework for sequential decisions*. John Wiley & Sons.

Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., Lam, R., & Willson, M. (2025). Probabilistic weather forecasting with machine learning. *Nature*, *637*, 84–90. http://dx.doi.org/10.1038/s41586-024-08252-9.

Queralta, J. P., Taipalmaa, J., Can Pullinen, B., Sarker, V. K., Nguyen Gia, T., Tenhunen, H., Gabbouj, M., Raitoharju, J., & Westerlund, T. (2020). Collaborative multi-robot search and rescue: Planning, coordination, perception, and active vision. *IEEE Access*, *8*, 191617–191643. http://dx.doi.org/10.1109/ACCESS.2020.3030190.

Quero, C. O., & Martinez-Carranza, J. (2025). Unmanned aerial systems in search and rescue: A global perspective on current challenges and future applications. *International Journal of Disaster Risk Reduction*, *118*, Article 105199. http://dx.doi.org/10.1016/j.ijdrr.2025.105199.

Rahbar, F., Marjovi, A., & Martinoli, A. (2019). An algorithm for odor source localization based on source term estimation. In *2019 international conference on robotics and automation* (pp. 973–979). IEEE, http://dx.doi.org/10.1109/ICRA.2019.8793784.

Ramchurn, S. D., Huynh, T. D., Wu, F., Ikuno, Y., Flann, J., Moreau, L., Fischer, J. E., Jiang, W., Rodden, T., Simpson, E., Reece, S., Roberts, S., & Jennings, N. R. (2016). A disaster response system based on human-agent collectives. *Journal of Artificial Intelligence Research*, *57*, 661–708. http://dx.doi.org/10.1613/jair.5098.

Ramchurn, S. D., Wu, F., Jiang, W., Fischer, J. E., Reece, S., Roberts, S., Rodden, T., Greenhalgh, C., & Jennings, N. R. (2016). Human–agent collaboration for disaster response. *Autonomous Agents and Multi-Agent Systems*, *30*, 82–111. http://dx.doi.org/10.1007/s10458-015-9286-4.

Rathour, S. S., Kato, N., Tanabe, N., Senga, H., Hirai, Y., Yoshie, M., & Tanaka, T. (2015). Spilled oil autonomous tracking using autonomous sea surface vehicle. *Marine Technology Society Journal*, *49*, 102–116.

Reggente, M. (2014). *Statistical gas distribution modelling for mobile robot applications* (Ph.D. thesis), Örebro university.

Rolland, E. G., Grøntved, K. A., Christensen, A. L., Watson, M., & Richardson, T. (2024). Autonomous UAV volcanic plume sampling based on machine vision and path planning. In *2024 international conference on unmanned aircraft systems* (pp. 1064–1071). IEEE, http://dx.doi.org/10.1109/ICUAS60882.2024.10556912.

Rothermel, R. C. (1972). *A mathematical model for predicting fire spread in wildland fuels*: *Vol. 115*, Intermountain Forest & Range Experiment Station, Forest Service, US Department of Agriculture.

Schellenberg, B., Richardson, T. S., Richards, A., & Watson, M. (2021). Automated real-time volcanic plume interception for UAVs. In *AIAA scitech 2021 forum* (p. 0811). http://dx.doi.org/10.2514/6.2021-0811.

Scherer, J., Yahyanejad, S., Hayat, S., Yanmaz, E., Andre, T., Khan, A., Vukadinovic, V., Bettstetter, C., Hellwagner, H., & Rinner, B. (2015). 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* (pp. 33–38). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/2750675.2750683.

Schwaiger, S., Muster, L., Novotny, G., Schebek, M., Wöber, W., Thalhammer, S., & Böhm, C. (2024). UGV-CBRN: An unmanned ground vehicle for chemical, biological, radiological, and nuclear disaster response. http://dx.doi.org/10.48550/arXiv.2406.14385, ArXiv Preprint.

Seraj, E., & Gombolay, M. (2020). Coordinated control of UAVs for human-centered active sensing of wildfires. In *2020 American control conference* (pp. 1845–1852). IEEE, http://dx.doi.org/10.23919/ACC45564.2020.9147613.

Shah, S., Dey, D., Lovett, C., & Kapoor, A. (2017). AirSim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and service robotics*. http://dx.doi.org/10.48550/arXiv.1705.05065.

Shobeiry, P., Xin, M., Hu, X., & Chao, H. (2021). UAV path planning for wildfire tracking using partially observable Markov decision process. In *AIAA scitech 2021 forum* (p. 1677). http://dx.doi.org/10.2514/6.2021-1677.

Stampa, M., Sutorma, A., Jahn, U., Thiem, J., Wolff, C., & Röhrig, C. (2021). Maturity levels of public safety applications using unmanned aerial systems: A review. *Journal of Intelligent and Robotic Systems*, *103*, 1–15. http://dx.doi.org/10.1007/s10846-021-01462-7.

Storwold, D. (2007). Detailed test plan for the fusing sensor information from observing networks (FUSION) field trial 2007 (FFT 07).

Sung, Y., Chen, Z., Das, J., Tokekar, P., et al. (2023). A survey of decision-theoretic approaches for robotic environmental monitoring. *Foundations and Trends® in Robotics*, *11*, 225–315. http://dx.doi.org/10.1561/2300000073.

The Wikibooks Community (2025). Webots odor simulation. https://en.wikibooks.org/wiki/Webots_Odor_Simulation. (Accessed 09 August 2025).

Thul, L., & Powell, W. B. (2023). An information-collecting drone management problem for wildfire mitigation. http://dx.doi.org/10.48550/arXiv.2301.07013, ArXiv Preprint.

Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature*, *445*, 406–409. http://dx.doi.org/10.1038/nature05464.

Viseras, A., Meissner, M., & Marchal, J. (2021). Wildfire front monitoring with multiple UAVs using deep Q-learning. *IEEE Access*, http://dx.doi.org/10.1109/ACCESS.2021.3055651.

Wang, J. W., Guo, Y., Fahad, M., & Bingham, B. (2019). Dynamic plume tracking by cooperative robots. *IEEE/ASME Transactions on Mechatronics*, *24*, 609–620. http://dx.doi.org/10.1109/TMECH.2019.2892292.

Wick, G. A., Dunion, J. P., Black, P. G., Walker, J. R., Torn, R. D., Kren, A. C., Aksoy, A., Christophersen, H., Cucurull, L., Dahl, B., English, J. M., Friedman, K., Peevey, T. R., Sellwood, K., Sippel, J. A., Tallapragada, V., Taylor, J., Wang, H., Hood, R. E., & Hall, P. (2020). NOAA's Sensing Hazards with Operational Unmanned Technology (SHOUT) Experiment Observations and Forecast Impacts. *Bulletin of the American Meteorological Society*, *101*, E968–E987. http://dx.doi.org/10.1175/BAMS-D-18-0257.1.

Wiedemann, T., Shutin, D., & Lilienthal, A. J. (2021). Experimental validation of domain knowledge assisted robotic exploration and source localization. In *2021 IEEE international conference on autonomous systems* (pp. 1–5). IEEE, http://dx.doi.org/10.1109/ICAS49788.2021.9551145.

Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MIT press Cambridge, MA.

Wu, F., Ramchurn, S. D., & Chen, X. (2016). Coordinating human-UAV teams in disaster response. In *Proceedings of the 25th international joint conference on artificial intelligence* (pp. 524–530).

Zhu, K., & Zhang, T. (2021). Deep reinforcement learning based mobile robot navigation: A review. *Tsinghua Science and Technology*, *26*, 674–691. http://dx.doi.org/10.26599/TST.2021.9010012.