

# Real-Time Autonomous Systems for Tracking and Responding to Uncontrolled Fires in the Ambient Environment: A Review

Anonymous authors

Paper under double-blind review

## Abstract

Fire-induced air pollution—originating from wildfires and industrial fires—poses a rising threat to public health and environmental systems. These episodic but increasingly frequent events release hazardous mixtures of particulate matter and gases, often overwhelming existing monitoring and response infrastructures. Traditional approaches to air quality sensing, health risk modelling, and emergency coordination are limited in spatial resolution, real-time responsiveness, and system integration. This literature review investigates how artificial intelligence (AI) and autonomous systems can address these limitations by enabling more adaptive, predictive, and interconnected fire pollution management strategies. Using a structured thematic synthesis, the review analyses 128 papers across four domains: (1) risks and impacts of fire-induced air pollution, (2) real-time autonomous systems for sensing, forecasting, and simulation, (3) AI-enhanced health risk modelling, and (4) governance and policy frameworks. Key findings reveal strong potential for UAV-based plume tracking, multi-agent learning systems, and data-driven health forecasting, but also highlight persistent gaps in regulatory readiness, system interoperability, and equity. The review argues for a coordinated, AI-centric framework to improve environmental sensing, health protection, and governance in fire-prone contexts.

## 1 Introduction

In recent years, uncontrolled fires in the ambient environment have emerged as a pressing global challenge, driven by their rising frequency and severity. Both wildfires and urban-industrial fires are occurring at increasingly destructive scales, associated with climate change, poor land management, and the consequences of increasing urbanisation. Between 2003 and 2023, the frequency of the most extreme wildfires more than doubled, with recent seasons ranking among the most severe ever recorded Cunningham et al. (2024). These fires emit high concentrations of fine particulate matter (PM<sub>2.5</sub> and PM<sub>10</sub>), gases, and volatile organic compounds (VOCs), that individually and in combination pose significant risks to public health, ecosystems, and societal infrastructure Efimova & Rukavishnikov (2021); Gao et al. (2023). The impacts are not confined to immediate burn zones—plumes can travel hundreds of kilometres and lead to degraded air quality across regions. For example, stubble burning in Southeast Asia routinely affects air quality across multiple countries, while the 2005 Buncefield oil depot fire in the UK produced a plume that was tracked across northern Europe Vautard et al. (2007). Often the plumes disproportionately affect vulnerable populations Griffiths et al. (2025).

Despite increased concern, current air quality monitoring practice for emergency response purposes remains inadequate where predictions about pollutant concentrations are needed to provide health protection advice. Ground-based sensors routinely lack the spatial density to detect localised peaks, while satellite systems suffer from latency and limited vertical resolution Efimova & Rukavishnikov (2021); Naqvi et al. (2023). Dispersion modelling is used to provide the basis for health protection advice but frequently underestimates pollutant concentrations in critical downwind areas Białowicz et al. (2021). Thus, significantly elevated pollutant emissions remain poorly characterised in real-time Griffiths et al. (2022). This makes timely public

health risk assessments and health protection interventions difficult—particularly the response phase during fast-evolving events such as wildfires and industrial fires (including landfill fires). Prospective public health implications are substantial, with short-term spikes in  $PM_{2.5}$  linked to respiratory distress, cardiovascular events, and increased mortality, with associated demand for clinical services Hughes et al. (2024); Moore et al. (2023). While some frameworks such as US EPA Air Quality Indices differentiate risk based on pre-existing health conditions, existing health risk models may not fully capture acute exposure windows or individual-level vulnerabilities, potentially leading to under-informed decision-making during emergencies.

### Conceptual overview of current and envisioned fire-plume monitoring and response pipelines

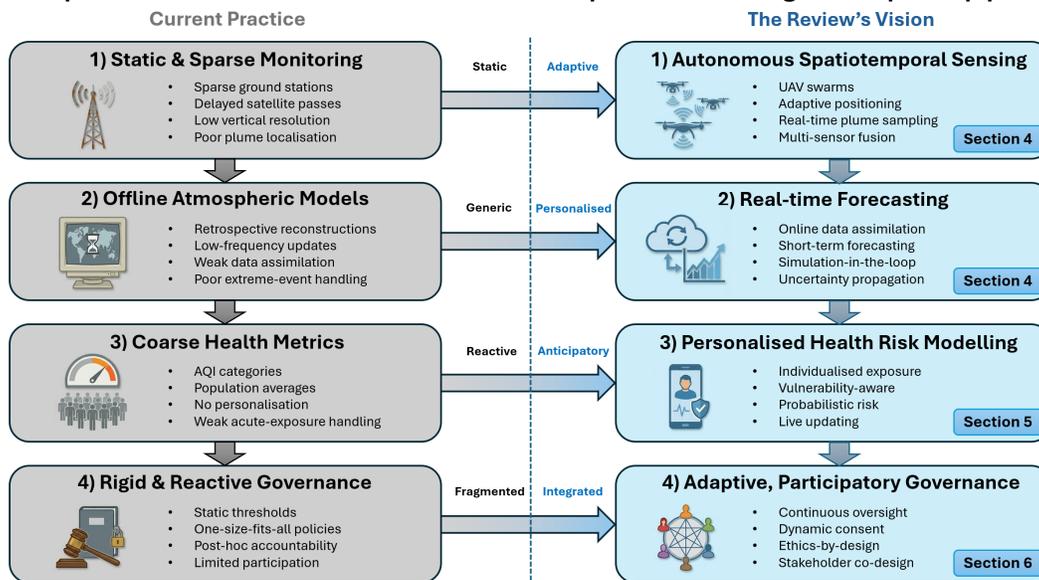


Figure 1: Conceptual contrast between current fire-plume monitoring and response pipelines (left) and the integrated, real-time, human-centred system envisioned in this review (right). Existing approaches are characterised by fragmented sensing, retrospective modelling, population-averaged health metrics, and rigid governance structures. In contrast, the proposed paradigm integrates autonomous spatiotemporal sensing, online forecasting and simulation, personalised health risk modelling, and adaptive, participatory governance into a tightly coupled pipeline. The central arrows highlight the conceptual shift from static to adaptive, generic to personalised, reactive to anticipatory, and fragmented to integrated systems.

While significant progress has been made in individual fields—such as satellite sensing, air quality modelling, and health impact research—these advances often remain siloed, with limited integration across disciplines. Few comprehensive efforts exist to integrate technological developments in autonomous environmental monitoring with epidemiological modelling and public policy. This review addresses that gap. It seeks to synthesise knowledge across environmental science, autonomous systems, artificial intelligence, and public health in the context of uncontrolled fires in the ambient environment. The goal is to assess the current state of the field, identify critical limitations, and explore how emerging technologies—particularly real-time autonomous platforms and AI-enhanced analytics—can reshape our capacity to monitor and mitigate plumes from uncontrolled fire events.

The scope of this review includes both wildfire and industrial uncontrolled fires to review their environmental, chemical, and health impacts. It covers current and emerging technologies for real-time tracking of fire plumes, including unmanned aerial vehicles (UAVs), multi-sensor fusion, and reinforcement learning algorithms. It also evaluates the evolution of health risk modelling approaches, from traditional exposure-response curves to personalised, dynamic risk assessments informed by AI. Finally, it considers the governance and policy frameworks that shape the deployment of these technologies, and the challenges involved in integrating them into public health systems.

## 1.1 Limitations of Existing Reviews and Research Gaps

A growing body of literature has examined the environmental, chemical, and health impacts of fire-induced air pollution, as well as the development of sensing technologies and atmospheric dispersion models. Several well-established reviews have synthesised evidence on the toxicological and epidemiological effects of wildfire smoke exposure, including impacts on respiratory and cardiovascular health and population-level mortality (Reid et al., 2016; Johnston et al., 2012; Black et al., 2017). In parallel, other reviews have focused on specific technical components of air quality monitoring, such as satellite-based remote sensing of surface pollutants (Martin, 2008) and atmospheric transport and dispersion modelling frameworks (Stein et al., 2015).

More recently, a separate body of review literature has emerged around the use of artificial intelligence and machine learning for air pollution monitoring and prediction. These include systematic reviews of machine learning and Internet of Things (IoT) approaches for outdoor air quality monitoring (Gryech et al., 2024), interpretable machine learning methods for pollutant forecasting (Houdou et al., 2024), and deep learning architectures for air quality prediction (Zhang et al., 2024b). Collectively, these studies demonstrate rapid methodological progress in data-driven environmental modelling and sensing infrastructures.

However, these strands of research remain largely fragmented. Existing reviews typically examine environmental sensing and atmospheric modelling, health impacts of smoke exposure, or AI-based pollution prediction in isolation, without addressing how these components could be integrated into a unified, real-time response framework for uncontrolled fire events. In particular, current reviews do not jointly consider how autonomous sensing platforms, adaptive plume tracking, spatiotemporal forecasting, health risk modelling, and governance mechanisms could be co-designed and operationalised as part of a coherent socio-technical system.

To our knowledge, no existing review systematically synthesises these domains to examine how real-time autonomous systems and AI-enhanced analytics can support end-to-end plume monitoring, exposure assessment, and public health decision-making under the extreme dynamics of wildfire and industrial fire events. This fragmentation limits both scientific understanding and the practical deployment of next-generation response infrastructures. Addressing this gap is essential for developing fire management systems that are not only more accurate and timely, but also more equitable, transparent, and actionable.

## 1.2 Contributions of This Review

This review makes the following key contributions:

- We provide a cross-disciplinary synthesis of research on fire-induced air pollution, autonomous sensing platforms, AI-based plume forecasting, and health risk modelling, unifying perspectives that are typically studied in isolation.
- We critically evaluate the limitations of existing monitoring and modelling paradigms, particularly their inability to capture acute exposure dynamics and support real-time public health interventions.
- We survey emerging autonomous and AI-enabled systems for plume tracking, multi-agent coordination, and spatiotemporal forecasting, highlighting their potential to transform emergency response.
- We examine how advances in health risk modelling can be integrated with real-time environmental data to support personalised and context-aware decision-making.
- We analyse the governance, regulatory, and ethical challenges associated with deploying these technologies at scale, with particular attention to issues of trust, equity, and accountability.

The remainder of the review is organised as follows: Section 3 explores the risks and impacts of fire-induced air pollution, including environmental hazards, emission profiles, and toxicological effects. Section 4 focuses on real-time autonomous monitoring, forecasting, and simulation systems, and their potential for high-resolution plume tracking. Section 5 examines the landscape of health risk modelling and emerging data-driven approaches. Section 6 discusses governance challenges, regulatory gaps, and future policy directions.

Together, these sections provide a cross-disciplinary synthesis aimed at informing the development of more responsive, equitable, and intelligent air pollution management systems in an era of intensifying fire risks.

## 2 Methodology

This literature review adopts a structured thematic synthesis approach to evaluate interdisciplinary research at the intersection of artificial intelligence (AI), fire-induced air pollution, environmental monitoring, health risk modelling, and governance. The core objective is to examine how AI-enabled systems can enhance real-time monitoring, plume forecasting, and decision-making related to airborne plumes from uncontrolled fires and the associated risks.

### 2.1 Search Strategy

Because of its broad cross-disciplinary indexing and accessibility to recent advances in AI, autonomous systems, and environmental science, Google Scholar was used for the literature search. A thematic, section-based approach was adopted, whereby keyword searches were conducted separately for each major review domain: (1) Risks and Impacts, (2) Real-Time Autonomous Systems, (3) Health Risk Modelling, and (4) Governance.

Papers were systematically organised by thematic domain to enable iterative refinement of search terms and cross-referencing between topics. Table 1 summarises the search strategy by thematic domain, including representative keywords and the number of papers selected for each area. In total, over 30 distinct keyword queries were used, producing a high-volume result set from which 128 papers were ultimately selected based on relevance and coverage.

Table 1: Literature Search Strategy by Thematic Domain

Thematic Domain	Representative Keywords	Papers Selected
Risks and Impacts (Section 3)	“air pollution fire”, “wildfire pollution”, “industrial fires pollution”, “industrial fire air pollution UK”, “frequency of wildfires”	23
Real-Time Autonomous Systems (Section 4)	“autonomous vehicles environmental monitoring”, “drone plume tracking”, “UAV tracking fire”, “multi-agent reinforcement learning”, “MARL drones”, “UAV tracking MARL”, “spatiotemporal forecasting”, “plume prediction”, “transformer air quality”, “wildfire simulation”, “CFD fire smoke plume”	84*
Health Risk Modelling (Section 5)	“AI health risk modelling”, “AI health risk modelling evacuation”, “AI health risk modelling gas plume evacuation”, “fire plume citizen notification”	15
Governance (Section 6)	“Governance Frameworks AI Public Health”, “Community Engagement Public Trust AI”, “Bias Equity Environmental AI”, “Coordination Failures Emergency Systems AI”	29
<b>Total unique papers</b>	<b>30+ keyword combinations</b>	<b>128</b>

\*Includes papers on autonomous platforms (19), reinforcement learning (17), forecasting (32), and simulation (16), with some overlap between subtopics.

### 2.2 Inclusion Criteria

Papers were included if they contributed meaningfully to one or more of the review’s central themes, even if they were not explicitly focused on fire events or AI technologies. Included papers addressed:

- Air pollution from uncontrolled fires, including wildfires and industrial fires.
- Autonomous systems, environmental sensing, and AI-based decision tools, even if applied to ambient air quality or environmental monitoring.
- Health risk modelling methods (including exposure-response functions and guideline threshold values) applicable to short-term or episodic pollution events.
- Governance, regulatory, and community frameworks related to environmental technologies or emergency response.

Papers were excluded only if they were clearly unrelated to any of the thematic areas described in Section 2.3, duplicated other work, or lacked academic or technical rigour.

### 2.3 Thematic Categorisation

Selected papers were organised into four major thematic domains, each representing a core pillar of the review:

1. **Risks and Impacts of Fire-Induced Air Pollution:** This section addresses the sources, chemical characteristics (emission profiles and plume composition), and health impacts from airborne pollutants.
2. **Real-Time Autonomous Systems:** This domain includes four subtopics: (i) types of autonomous platforms (e.g., UAVs, UGVs); (ii) sensing technologies for air quality and plume detection; (iii) AI-based tracking and forecasting systems; and (iv) simulation environments, such as CFD and multi-agent platforms used for fire dynamics and plume dispersion modelling.
3. **Health Risk Modelling:** Studies in this area include traditional and AI-enhanced health impact modelling, short-term exposure forecasting, and decision-support systems for risk communication and evacuation.
4. **Governance and Policy:** This category includes literature on public trust, regulatory frameworks, coordination failures, ethical issues, and equity in the use of AI and autonomous systems for environmental or health-related applications.

### 2.4 Potential Limitations

This review does not follow formal systematic review protocols and instead provides a curated, high-coverage synthesis focused on emerging and interdisciplinary literature. Non-English sources and grey literature were generally excluded. Non-peer-reviewed materials were included only when technically rigorous and directly relevant to the review themes. Some specialised technical subfields (e.g., hardware engineering or combustion chemistry) were only included when their relevance to AI-driven environmental sensing or health modelling was specific and made a meaningful contribution to the review objectives. The emphasis remains on AI as a unifying tool across environmental data acquisition, predictive analytics, and risk governance in the context of fire-related air pollution.

## 3 Risks and Impacts of Fire-Induced Air Pollution

Uncontrolled fires generate complex and rapidly evolving air pollution hazards that extend far beyond the immediate burn zone. These hazards are not only environmental, but also chemical, toxicological, and socio-economic in nature, with impacts that vary spatially, temporally, and across population groups. This section establishes the problem space that motivates the need for next-generation monitoring and response systems. We first examine the environmental drivers and geographic patterns of episodic fires, then characterise the composition and transport of resulting pollutant plumes. We subsequently review the associated health risks, socio-economic burdens, and existing monitoring and management practices, highlighting their limitations and the growing need for more adaptive, real-time, and integrated approaches.

### 3.1 Environmental Hazards and Episodic Fires

The scale of environmental challenges—such as climate change, land mismanagement, and industrial expansion—has intensified the frequency and severity of episodic fires, including wildfires and urban fires. These fires contribute significantly to transient air pollution events, posing risks to public health and ecosystems. Their unpredictable and sporadic nature makes real-time monitoring and timely modelling difficult.

#### 3.1.1 Wildfires

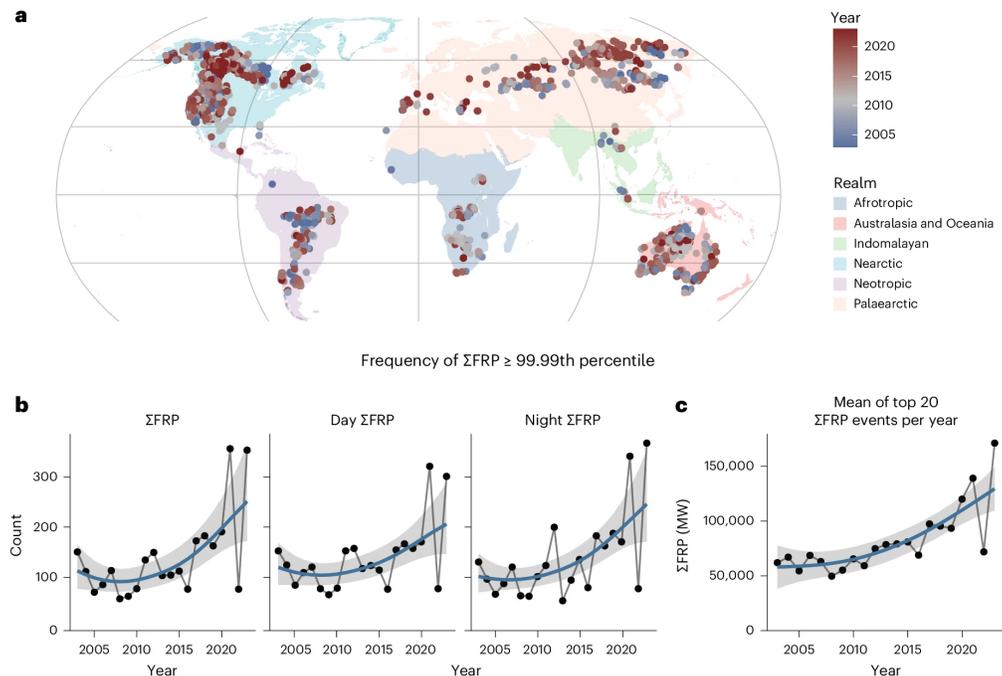


Figure 2: **Reprinted from** Cunningham et al. (2024, Fig. 1). Trends in extreme wildfire events (2003–2023) based on MODIS Fire Radiative Power (FRP) data: **(a)** Geographic distribution of extreme events ( $\geq 99.99$ th percentile of daily  $\Sigma\text{FRP}$ ), **(b)** Frequency doubling trend with 95% CI, **(c)** Intensity increase in top 20 annual events. Note: 2023 data excludes December.

Wildfires are among the most frequent and devastating uncontrolled fire types. Although these wildfires are unpredictable, their causes and effects can be predicted by studying past events, with studies identifying climate change to be the single greatest cause behind their increasing frequency in recent years. This notion is supported by Efimova & Rukavishnikov (2021), who assessed the smoke pollution caused by wildfires in the Baikal region of Russia, revealing that climate change has intensified wildfires, leading to prolonged fire seasons and increased air pollution. These regional observations are now corroborated by global satellite data: Cunningham et al. (2024) demonstrate a 2.2-fold increase in the frequency of the most extreme wildfires (top 0.01% by intensity) from 2003 to 2023 (Figure 2b), with the last seven years containing the six most extreme fire seasons—a pattern consistent with accelerating climate impacts.

Gao et al. (2023) state that wildfire frequency is projected to increase by up to 50% globally by 2100 as a result of climate change, based on global bioclimatic modelling of fire regimes under future climate scenarios (Kelley et al., 2022) and UNEP assessments of landscape fire risks (Popescu et al., 2022). This projection aligns with the observed biome-specific trends in Figure 2a, where temperate conifer forests (including California) saw an 11.1-fold rise in extreme fires—the highest of any biome. Similarly, Hughes et al. (2024) report that 61% of countries experienced increased wildfires between 2018–2021 compared to 2001–2004, while Graham et al. (2020) note a UK government projection of 30%–50% higher wildfire risk by 2080. The consistency between these regional studies and global MODIS data underscores climate change as the dominant driver of intensifying fire regimes.

As a result of climate change, wildfire activity has intensified in many regions in recent years. Many studies have identified North America, Australia, and the Amazon rainforest as notable hotspots for wildfire disasters (Gao et al., 2024; Jaiswal et al., 2022; Yue et al., 2024), with events such as the 2019–2020 Australian bushfires and 2020 California wildfires among the most extensively documented (Chen et al., 2021a). Figure 2a quantifies this geographic concentration: the Nearctic (North America) and Australasia/Oceania realms experienced 60% of extreme events despite covering only 20% of global burnable land area—a disparity reflecting climate-biome interactions. California in particular has seen a large increase in the frequency of wildfires. A study by Naqvi et al. (2023) also focuses on the 2020 California wildfire season, stating that it was one of the worst on record, with a significant increase in the number of fires, reporting 9,917 fire incidents, and burning a total of 1.77 million hectares of land. This aligns with the sharp upward trend in Figure 2c, where the average intensity of California’s top 20 annual fires increased 2.3-fold over the study period.

Another study by Boaggio et al. (2022) mentions that the 2020 fire season saw widespread smoke covering the West Coast for weeks, verifying the claims of Naqvi et al. (2023). The same study by Boaggio et al. (2022) also uses another example of California as a hotspot for wildfire disasters, mentioning that incident activity significantly increased with over 3.5 million hectares burned in 2017 and 2018 alone. Although this verifies California as a natural wildfire hotspot, it is not the only location that commonly hosts wildfire disasters. Canada is also a highly affected location, with a study by Moore et al. (2023) revealing that in 2019 alone, 4,000 wildfire incidents, burning over 1.8 million hectares of land, took place. Cunningham et al.’s (2024) identification of boreal forests as the second-most vulnerable biome (7.3-fold increase in extreme fires) explains Canada’s escalating fire activity, linking it to climate-driven drying of northern peatlands and conifer stands. A limitation of these studies is the lack of comparison between the causes and effects of wildfires in multiple regions, as they all focus on a singular geographic location. Despite this focus, the consensus seems to be clear that even though wildfires occur all around the world, there are areas which are most predisposed to such incidents.

### 3.1.2 Urban Fires

Urban fires, comprised of industrial and residential sources, are among the most common and damaging uncontrolled fire types. Often caused by anthropogenic activities or errors, they are unpredictable and because of their close proximity to populations may cause injury and mortality. To understand their distribution and scope, data from the National Fire Protection Association (NFPA) provides a granular classification of fire incidents by property type in the United States. As shown in Figure 3, NFPA data from 2018–2022 reveals significant variations in fire frequency across different property categories. Outside or special property locations (including open land, beaches, campsites, highways, streets, parking areas, bridges, landfills, railroads, and construction sites) experience the highest number of incidents at 619,554 annually, followed by residential properties with 552,240 incidents per year (National Fire Protection Association, 2023).

While urban uncontrolled fires are typically smaller in scale than wildfires, their environmental and health impacts can be just as severe due to the materials burnt and their proximity to densely populated areas.

Since 2009, across England, Scotland and Wales, the globally novel Air Quality Cell (AQC) has responded to *significant* major air pollution incidents that could pose a threat to public health taking account of the duration and severity of the burn, and the proximity of a population. Between 2009 and 2016, 76 AQC incidents were recorded across the three nations (Griffiths et al., 2022).

Griffiths et al. (2018) highlighted that industrial incidents were most likely at waste/resource management facilities near populations and would involve materials recovered from urban waste streams, such as plastics, tyres, and chemical waste, that when burnt, lead to a complex of chemical species that can have a negative impact on human health (Griffiths et al., 2022). Using back trajectory modelling for a tyre fire in an urban area, Griffiths et al. (2025) have indicated that potentially thousands of residents could have been exposed to hazardous concentrations of  $PM_{10}$ .

Policies related to the use of urban land and the spatial relationship between residential and industrial uses exacerbate these risks.

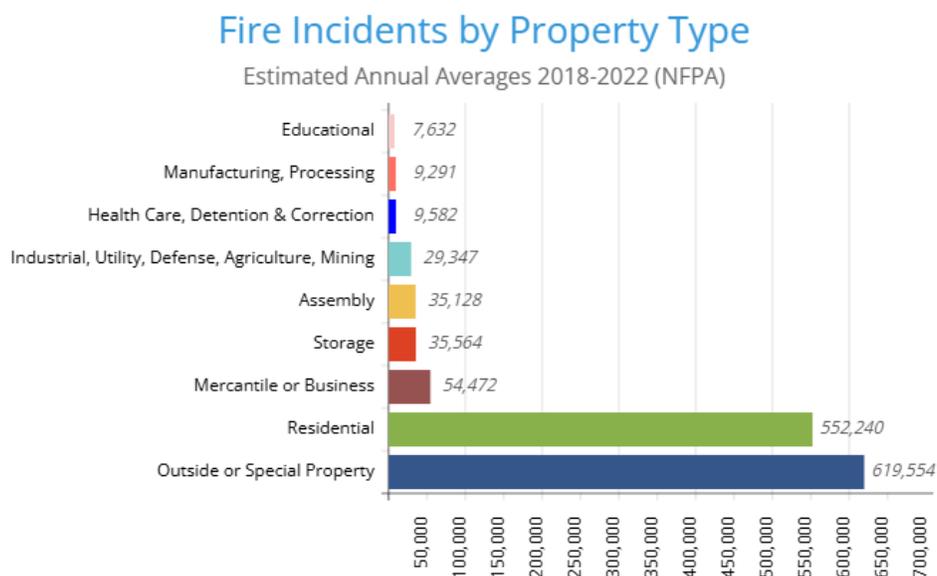


Figure 3: Fire Incidents by Property Type in the United States: Estimated Annual Averages 2018–2022 (National Fire Protection Association, 2023).

Similar trends are observed globally. In Poland, Białowicz et al. (2021) reported that 2018 saw 79 large landfill fires involving mixed municipal and industrial waste (including general waste, plastics, tyres, recyclable materials, and construction debris), with very large fires increasing by 2.4 times and big fires by 2.6 times compared to the 2010–2017 average. This surge underscores how poor waste management and rising waste generation rates contribute to environmental degradation and fire risk, in part caused by increased urban populations. Urban sprawl and the encroachment of residential areas into industrial zones further amplify these dangers, exposing communities to complex products of combustion, and soil and water contamination. These trends highlight systemic failures in waste management regulation and land-use planning, underscoring the need for policy reform.

As illustrated in Figure 3, industrial, utility, defence, agricultural, and mining fires (29,347 incidents) are less frequent than residential fires in the U.S. but often involve hazardous materials, leading to more severe environmental and health consequences (National Fire Protection Association, 2023). Meanwhile, mercantile or business properties experience 54,472 incidents annually, while storage (35,564), assembly (35,128), healthcare and detention facilities (9,582), manufacturing (9,291), and educational institutions (7,632) each contribute to the urban fire landscape (National Fire Protection Association, 2023).

Although industrial fires tend to be the most severe, residential fires remain the most widespread. Alharbi et al. (2021) note that over 3 million residential fires occur globally each year, causing significant damage, fatalities, and environmental harm—particularly in low-income areas with substandard building materials and inadequate fire safety measures. The intersection of urbanisation, industrial activity, and climate change (e.g., prolonged droughts increasing fire risk in cities such as California) further intensifies these challenges, demanding stronger regulatory frameworks and sustainable urban planning to mitigate future disasters.

### 3.2 Air Pollutant Plumes from Uncontrolled Fire

Wildfires and industrial fires are significant sources of airborne pollutants, emitting a complex mixture of particulate matter (PM), toxic gases, and organic compounds. Multiple studies highlight the release of  $PM_{2.5}$  and  $PM_{10}$  at extreme concentrations, particularly in industrial fires, where Griffiths et al. (2018) recorded peaks of  $>6,500 \mu\text{g}/\text{m}^3$  for  $PM_{10}$  and  $>650 \mu\text{g}/\text{m}^3$  for  $PM_{2.5}$ . Similarly, landfill fires (Białowicz et al., 2021) contribute substantially to  $PM_{10}$ , with short-term spikes exceeding  $100 \mu\text{g}/\text{m}^3$ , while urban wildfires (Boaggio

et al., 2022) reintroduce phased-out pollutants such as lead at concentrations 40 times higher than baseline levels, contrary to the likely conceptualisation of a ‘clean burn’. Wildfires, increasingly intensified by climate change (Yue et al., 2024), emit not only CO<sub>2</sub> and black carbon but also ozone precursors (NO<sub>x</sub> and volatile organic compounds), with boreal fires showing rising emissions despite reductions in burned area.

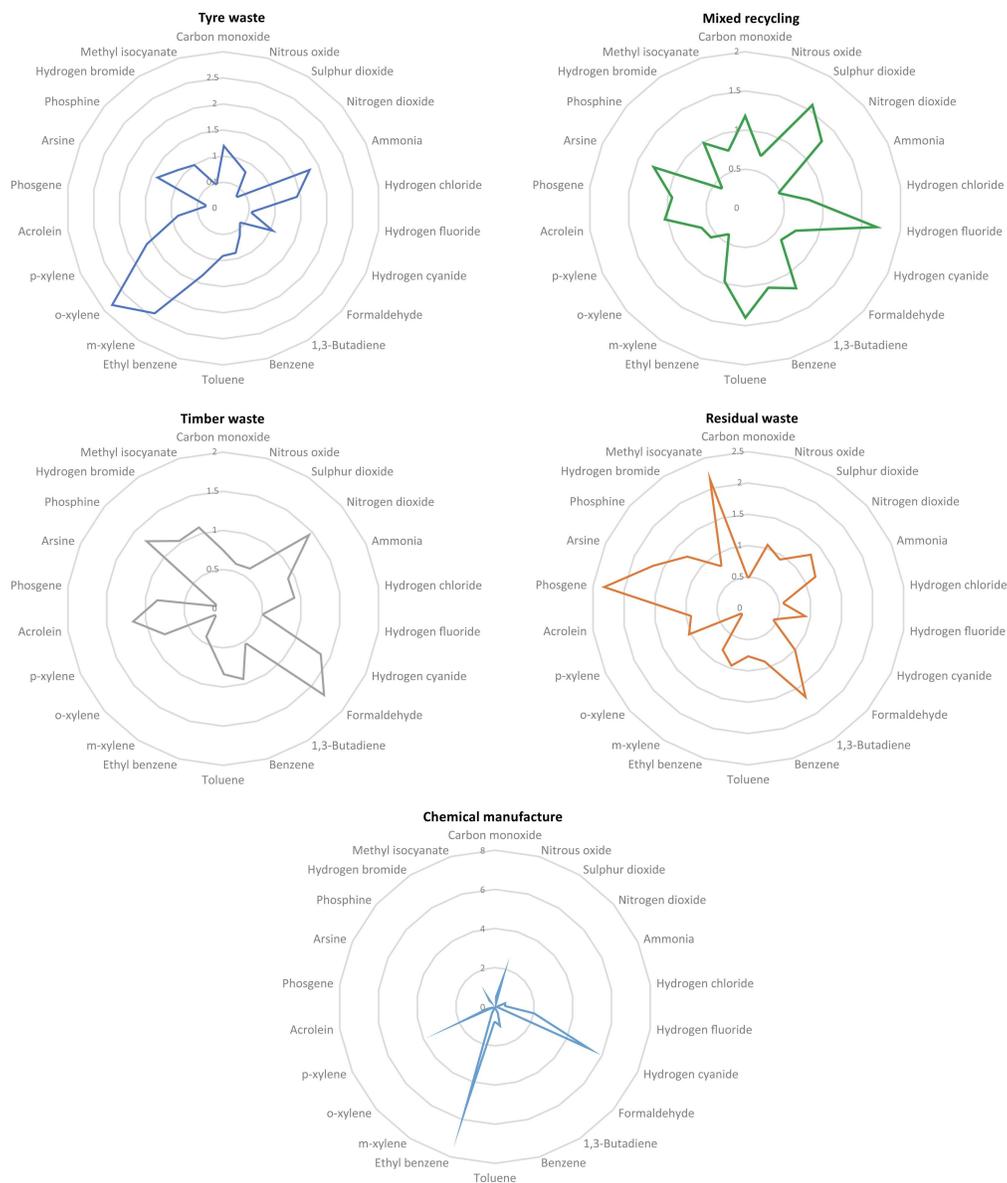


Figure 4: **Reprinted from** Griffiths et al. (2022, Fig. 6), showing relative emission factors for industrial fires by waste type. Tyre fires exhibit dominant xylene and ammonia emissions, while mixed recyclables show broader Volatile Organic Compound plumes (ethyl benzene, 1,3-butadiene) and inorganic species (HBr, HF). Timber combustion produces formaldehyde and nitrogen-derived volatiles, contrasting with residual waste’s uniform distribution of chlorinated compounds (COCl<sub>2</sub>) and isocyanates (C<sub>2</sub>H<sub>3</sub>NO).

The chemical composition of fire emissions varies by fuel type and combustion conditions, as clearly demonstrated in Figure 4 by Griffiths et al. (2022). Industrial and residential fires (Alharbi et al., 2021) release hydrogen cyanide (HCN), benzene, and sulphur dioxide (SO<sub>2</sub>), while wildfires produce formaldehyde and polycyclic aromatic hydrocarbons (PAHs) (Chen et al., 2021b). Notably, urban wildfires—such as California’s Camp Fire—generate metal-laden PM<sub>2.5</sub> from incinerated infrastructure, complicating air pollution profiles

with legacy contaminants like lead (Boaggio et al., 2022). However, some studies, e.g., Gao et al. (2023), rely on modelled PM<sub>2.5</sub> data, which may underestimate localised peaks due to sparse ground monitoring. A critical gap remains in quantifying real-time emissions from high-intensity fires, as most measurements, e.g., FTIR spectroscopy in Griffiths et al. (2022), capture only stabilised plumes, missing initial combustion phases where pollutant ratios may differ.

### 3.2.1 Fire Toxicological Profile and Human Health Risks

Fire-induced air pollution poses both acute and chronic health risks. Short-term exposure to elevated PM levels triggers inflammatory responses in the respiratory tract (Griffiths et al., 2018), with industrial fire plumes inducing cardiovascular stress even in healthy adults (Deary & Griffiths, 2024). Vulnerable populations experience exacerbated effects; children exposed to wildfire PM<sub>2.5</sub> exhibit 13% higher asthma-related emergency visits during smoke events (Moore et al., 2023), while prenatal exposure correlates with gestational hypertension and preterm birth (Hughes et al., 2024).

The toxicological profile varies by emission source, as demonstrated by the distinct chemical signatures in Figure 4. Alongside the expected particulates, industrial fires release hydrogen cyanide and phosphine exceeding acute exposure guidelines (Griffiths et al., 2022)—particularly from tyre combustion (xylenes, ammonia) and chemical waste (phosgene, methyl isocyanate) shown in Figure 4. These contrast with wildfire PAHs (Chen et al., 2021b) and metal-laden PM<sub>2.5</sub> (Boaggio et al., 2022) that drive oxidative stress through different pathways. Longitudinal studies associate wildfire PM<sub>2.5</sub> with increased cancer mortality (0.4% per 10  $\mu\text{g}/\text{m}^3$  increment) (Gao et al., 2024), particularly lung (1.1%) and upper aerodigestive cancers (2.7%). Firefighters demonstrate elevated risks of chronic obstructive pulmonary disease and neurological disorders from volatile organic compound exposure (Alharbi et al., 2021).

Emerging evidence suggests systemic effects beyond cardiopulmonary outcomes. A 1  $\mu\text{g}/\text{m}^3$  increase in annual wildfire PM<sub>2.5</sub> is associated with 2.0% higher suicide rates in rural areas (Molitor et al., 2023), while ocular surface damage manifests through tear film disruption (Jaiswal et al., 2022). The 2020 California wildfires demonstrated compounded mortality risks with COVID-19, where PM<sub>2.5</sub> levels  $>240 \mu\text{g}/\text{m}^3$  amplified respiratory susceptibility (Naqvi et al., 2023). Critical knowledge gaps persist regarding dose-response relationships for ultrafine particles ( $\leq 100 \text{ nm}$ ), for which there are no health exposure guidelines, which remain understudied in wildfire smoke (Chen et al., 2021b), long-term neurocognitive impacts of metal inhalation (Boaggio et al., 2022), and efficacy of respiratory protection during extreme smoke events ( $>500 \mu\text{g}/\text{m}^3$ ) (Hughes et al., 2024).

### 3.2.2 Socioeconomic Costs and Inequality of Fire

Fire-related air pollution imposes substantial economic burdens across multiple sectors. Industrial fires in the UK result in £21.1 million in economic costs per major incident due to mortality and lost productivity (Graham et al., 2020). In the U.S., wildfire smoke exposure generates \$11 billion in annual short-term health costs and \$76 billion in long-term health burdens (Hagler et al., 2021), with the Government expending \$3 billion annually on wildfire suppression (Burke et al., 2021). These costs disproportionately affect vulnerable populations, as low-income communities near industrial sites face compounded risks due to inadequate fire safety measures (Griffiths et al., 2018).

Occupational impacts reveal significant productivity losses. Firefighters experience elevated cancer risks from chronic exposure to benzene and other VOCs (Alharbi et al., 2021), while rural agricultural workers lose productivity during smoke events due to the need to avoid outdoor exposure (Molitor et al., 2023). The 2020 California wildfires demonstrated compounded economic impacts, where PM<sub>2.5</sub> levels exceeding 240  $\mu\text{g}/\text{m}^3$  coincided with increased healthcare expenditures during the COVID-19 pandemic (Naqvi et al., 2023).

Spatial inequities further exacerbate these economic consequences. Rural communities experience 2.0% higher suicide rates per  $\mu\text{g}/\text{m}^3$  increase in wildfire PM<sub>2.5</sub> due to limited mental health infrastructure (Molitor et al., 2023), while urban industrial zones face 40-fold increases in lead exposure during infrastructure fires (Boaggio et al., 2022). Current regulatory frameworks inadequately address three critical gaps: transboundary smoke transport (Burke et al., 2021), acute exposure windows during fire ignition (Griffiths et al., 2022), and cumulative impacts on environmental justice communities (Yue et al., 2024).

### 3.2.3 Current Methods for Monitoring, Modelling and Managing Plumes

The monitoring of fire emissions and their associated plumes is essential for protecting environmental and public health. In the United Kingdom, the Air Quality in Major Incidents (AQinMI) service can be deployed to an uncontrolled fire and provide ‘live’ plume species monitoring results to AQC’s to support public health risk assessments. Monitoring teams have a standard range of indicative portable monitors but can also call mobile laboratories with reference standard equipment available. Routinely deployed are light scattering PM monitors (OSIRIS) (Griffiths et al., 2018) and Fourier Transform Infrared spectroscopy instruments (GasMet) (Griffiths et al., 2022) to detect pollutants such as particulate matter (PM<sub>10</sub>, PM<sub>2.5</sub>), carbon monoxide (CO), sulphur dioxide (SO<sub>2</sub>), and volatile organic compounds (VOCs) (Griffiths et al., 2022).

Globally, approaches vary based on regional needs and capabilities. In the United Kingdom, the Meteorological Office provides a CHEMET (CHEMical METeorological) service to responders that can predict the direction of the plume and concentrations of pollutants to responders. Further, the AQinMI system can incorporate dispersion modelling to identify where monitoring can take place and estimate the movement and impact of fire plumes, helping authorities issue public health guidance such as sheltering advisories (Deary & Griffiths, 2024). In the United States, the Environmental Protection Agency (EPA) employs the Air Quality System (AQS) and Hazard Mapping System (HMS), which are supported by NOAA’s VIIRS satellite data to track wildfire smoke (Boaggio et al., 2022; Burke et al., 2021); Australia integrates MODIS satellite data with ground-based sensors for bushfire monitoring (Holm et al., 2021; Jaiswal et al., 2022); and in Saudi Arabia, portable detectors are used to assess firefighter exposure to toxic gases such as hydrogen cyanide and benzene (Alharbi et al., 2021). These sensing strategies form the foundational layer for understanding plume behaviour and pollutant composition.

Once data is gathered, analysis and modelling techniques are used to predict plume dispersion and pollutant exposure. During major wildfires, such as the 2018 Saddleworth Moor fire, the UK augmented its ground-based monitoring efforts with regional air quality modelling to assess PM<sub>2.5</sub> exposure, which was subsequently associated with significant health impacts (Graham et al., 2020). Similarly, Poland—facing frequent landfill fires—has applied the HYSPLIT model to simulate pollutant dispersion (Białowicz et al., 2021). These models, often informed by both ground sensors and satellite inputs, provide critical spatial and temporal estimates of air quality deterioration during fire events.

Connecting air quality measurements to tangible health outcomes is a key component of fire monitoring systems. In the UK, data from the AQinMI framework has been used to support shelter-in-place advisories based on pollutant concentration thresholds (Deary & Griffiths, 2024). The 2018 Saddleworth Moor fire exemplifies this reasoning process, where air quality modelling linked elevated PM<sub>2.5</sub> levels to an estimated 28 excess deaths (Graham et al., 2020). These cases demonstrate how monitoring data, when properly analysed, can inform understanding of short-term and long-term health risks, especially for vulnerable populations. By contextualising pollutant levels with epidemiological evidence, public health authorities and incident responders can make more informed decisions for protecting health during fire incidents—the purpose of the AQC in the UK.

The ultimate goal of fire monitoring systems is to inform timely interventions that mitigate health risks and environmental damage. The AQinMI framework serves as a model for real-time response, supporting rapid deployment of monitoring equipment and public health advisories such as shelter-in-place instructions (Griffiths et al., 2018; Deary & Griffiths, 2024). In Australia, where bushfires are frequent and severe, public health recommendations include the use of N95 respirators during high-exposure events to reduce inhalation of harmful particles (Holm et al., 2021; Jaiswal et al., 2022). While national-scale monitoring systems may not yet be fully developed in countries like Saudi Arabia, localised efforts to protect firefighters from toxic gas exposure represent important forms of intervention (Alharbi et al., 2021). Moving forward, technological advancements such as low-cost sensors, machine learning analytics, and integrated air quality platforms offer promising avenues for enhancing both the speed and precision of fire-related health interventions (Yue et al., 2024; Chen et al., 2021a; Ebi et al., 2021).

### 3.2.4 Limitations of Existing Processes and the Need for AI-Driven Solutions

Current fire and plume monitoring systems exhibit multiple critical limitations that hinder effective emergency response and public health protection. The United Kingdom’s Air Quality in Major Incidents (AQinMI) service demonstrates these challenges through well-documented constraints (Griffiths et al., 2018). Monitoring activities are restricted to locations outside the densest parts of a plume to protect personnel, potentially missing peak plume concentrations in the most affected zones (Griffiths et al., 2025). These limitations extend globally across different monitoring contexts.

In Poland, landfill fire pollution monitoring relies heavily on HYSPLIT dispersion modelling to compensate for sparse ground-based monitoring networks in rural areas (Białowicz et al., 2021). Fundamental spatial-temporal resolution gaps persist across monitoring approaches, with ground-based sensor networks providing insufficient spatial density in remote regions (Efimova & Rukavishnikov, 2021), while satellite systems like MODIS and VIIRS offer broader coverage but suffer from latency periods of hours to days (Yue et al., 2024). This became particularly problematic during California’s 2020 wildfire season when stationary monitors failed to capture acute exposure peaks due to rapidly shifting plumes (Naqvi et al., 2023). Operational constraints further exacerbate these technical limitations. The manual positioning of portable monitors creates dangerous delays in hazardous environments (Griffiths et al., 2022), while firefighters globally face significant exposure risks during responses due to limited real-time gas monitoring capabilities (Alharbi et al., 2021).

The AQinMI service also relies on dispersion modelling to identify where monitoring can take place and estimate plume behaviour and population exposure. However, these models have shown limitations in real-world scenarios. During the Saddleworth Moor fire, HYSPLIT predictions underestimated observed PM<sub>2.5</sub> concentrations in some downwind communities (Graham et al., 2020). Similarly, U.S. wildfire smoke management systems struggle with transboundary impacts, as smoke plumes regularly cross state and national borders, overwhelming air quality frameworks designed for local pollution sources (Burke et al., 2021; Hager et al., 2021). Recent advances in AI-driven approaches show promise for addressing these persistent challenges. Machine learning integrations of multi-source data (satellite, ground sensors) have demonstrated improved plume trajectory predictions compared to conventional models, though performance metrics vary significantly by study design (Yue et al., 2024).

Unlike specific organic and inorganic species where short duration exposure guidelines exist (AEGL and ERPGs), airborne particulate exposure guidelines are for ambient exposures, routinely therefore for 24-h durations. This lack of short-term exposure guidelines complicates acute public health decision-making during incidents (Deary & Griffiths, 2021). Even advanced health monitoring systems, such as those tracking paediatric asthma emergencies, often lack individual-level exposure data during wildfire events (Moore et al., 2023). These gaps in short-term exposure thresholds and real-time personal exposure data hinder the ability to translate plume behaviour into actionable health intelligence. AI-enhanced tools have the potential to bridge this gap by integrating environmental and clinical data streams for more effective reasoning and risk assessment.

Resource limitations prevent comprehensive coverage of multiple simultaneous major incidents, as evidenced during the 2018 Saddleworth Moor fire (Graham et al., 2020). Traditional dispersion models show particular limitations in urban environments, where localised terrain and structures can, for example, cause a canyon effect and alter plume dispersal patterns. Unmanned aerial systems are being actively tested for early fire detection, with performance heavily dependent on sensor specifications and environmental conditions (Boaggio et al., 2022). Autonomous platforms, particularly drone-based systems, may eventually reduce human risk during plume sampling operations while overcoming current spatial and temporal monitoring limitations, though scalability and regulatory challenges remain significant barriers (Alharbi et al., 2021). These technological developments create a compelling case for next-generation monitoring systems that combine AI-enhanced analytics with autonomous deployment capabilities.

## 4 Real-time Autonomous Systems

The limitations of existing fire plume monitoring and modelling pipelines, as outlined in Section 3, motivate the need for a new generation of systems that are not only more spatially and temporally resolved, but also adaptive, predictive, and resilient under extreme conditions. Real-time autonomous systems offer such a paradigm shift. Rather than relying on sparse static sensors, delayed satellite observations, and manually configured dispersion models, these systems integrate mobile sensing platforms, AI-driven control, spatiotemporal forecasting, and high-fidelity simulation to enable continuous situational awareness and proactive decision-making.

This section reviews the technological foundations of real-time autonomous fire plume monitoring and response. We first examine the range of autonomous vehicle platforms used for environmental sensing and plume sampling, highlighting their sensing capabilities, operational trade-offs, and coordination strategies. We then discuss how reinforcement learning and multi-agent systems enable adaptive exploration, cooperative tracking, and robust control in highly dynamic environments. Building on this, we review emerging forecasting methods that predict plume evolution across space and time, allowing systems to anticipate exposure risks rather than merely reacting to them. Finally, we examine the role of simulation environments in training, validating, and stress-testing these systems before deployment. Together, these components form an integrated sensing–prediction–decision pipeline that underpins next-generation fire response infrastructures.

### 4.1 Autonomous Vehicle Platforms for Environmental Monitoring

#### 4.1.1 Vehicle Types

Recent research has identified several categories of autonomous vehicles (AVs) that are increasingly being adapted for environmental monitoring applications, including plume tracking and ecological surveys. Ground-based systems, particularly *Unmanned Ground Vehicles (UGVs)* and autonomous cars, have been extensively studied for their sensor fusion capabilities (Ignatious et al., 2022; Parekh et al., 2022). As demonstrated by Ignatious et al. (2022), these systems typically employ LiDAR, cameras, and RADAR for obstacle detection and environmental perception, though Vargas et al. (2021) show their effectiveness is limited by weather conditions and terrain accessibility. Parekh et al. (2022) highlight how UGVs in agricultural settings utilise multi-sensor arrays for precision monitoring, while Ignatious et al. (2022) emphasise the importance of real-time processing in ground AVs for dynamic environment interaction.

*Unmanned Aerial Vehicles (UAVs)* have emerged as particularly valuable for large-scale environmental monitoring due to their mobility and aerial perspective. Bathla et al. (2022) and Chamara et al. (2022) demonstrate that drones equipped with multispectral cameras and LiDAR are being deployed for crop health assessment and forest monitoring. Besson et al. (2022) further show how UAVs with hyperspectral sensors can be adapted for gas composition analysis, though Vargas et al. (2021) caution that their operational limitations in adverse weather remain a significant challenge. Chamara et al. (2022) reveal that 12% of surveyed agricultural IoT implementations now utilise UAVs for high-resolution spatial data collection, demonstrating their growing adoption.

Emerging *hybrid systems* that combine multiple platforms (e.g., UAV-UGV teams) are gaining attention for comprehensive environmental monitoring. Bathla et al. (2022) present case studies of truck-drone collaborations for large-area surveillance, and Besson et al. (2022) propose conceptual UAV-AUV networks for marine-terrestrial interface monitoring. These systems leverage Vehicle-to-Everything (V2X) communication and Multi-Agent Reinforcement Learning (MARL) for coordinated operation (Bathla et al., 2022), though Chamara et al. (2022) note that standardisation and interoperability remain active research challenges.

Across all platforms, sensor fusion (particularly LiDAR-camera-RADAR combinations) has been identified as critical for robust environmental perception (Ignatious et al., 2022; Vargas et al., 2021). However, as noted by Vargas et al. (2021), weather-induced sensor degradation and real-time processing demands continue to limit reliability in dynamic monitoring scenarios such as wildfire plume tracking. Recent advances in edge computing (Chamara et al., 2022) and 5G-enabled coordination (Besson et al., 2022) are addressing these limitations, pointing towards more resilient autonomous monitoring systems in the future.

### 4.1.2 Addressing Plume Tracking with Drones

The evolution of Unmanned Aerial Vehicles (UAVs) for fire-induced plume monitoring represents a technological convergence that significantly advances our capacity to observe and analyse atmospheric contaminants. Building upon the sensor fusion frameworks established in ground autonomous vehicles (Ignatious et al., 2022; Parekh et al., 2022), modern drone systems have overcome the spatial-temporal resolution gaps that plagued traditional monitoring networks (Efimova & Rukavishnikov, 2021; Yue et al., 2024).

The unique advantages of UAVs become particularly evident when considering the dynamic nature of wildfire emissions. Unlike industrial plumes, fire-induced plumes exhibit rapid spatial-temporal variability due to shifting fire fronts and pyroconvective activity—characteristics that render traditional monitoring inadequate. UAVs provide a feasible solution for: (1) proximal sampling near active fire fronts where ground access is dangerous (Alharbi et al., 2021); (2) vertical profiling of injection heights crucial for transport modelling (Burke et al., 2021); and (3) real-time identification of toxic gas hotspots that pose acute risks (Moore et al., 2023). The limitations of ground-based monitoring during wildfire events—particularly their failure to capture acute exposure peaks due to rapidly shifting plumes—have been well-documented (Naqvi et al., 2023). UAV systems like the Kolibri platform (Joíca et al., 2022) demonstrate how aerial measurements can overcome these gaps, with optical particle counters achieving  $0.1 \mu\text{m}$  resolution to track particulate evolution throughout the plume column. This capability to resolve vertical and horizontal variability is transforming our understanding of wildfire smoke dispersion and its health impacts.

The MUST platform’s deployment at the Taichung Power Plant demonstrated this transformation, where its quantum cascade laser system achieved 1 Hz VOC measurements with sub-ppb accuracy, revealing plume-induced temperature inversions that explained persistent discrepancies in HYSPLIT model predictions (Chen et al., 2023; Graham et al., 2020). This capability directly addresses the AQinMI service’s limitation of peripheral monitoring (Griffiths et al., 2018) by enabling in-plume measurements that detected 60 ppm  $\text{CH}_4$  hotspots invisible to ground stations (Gålfalk et al., 2021), while simultaneously validating the sensor fusion approaches proposed by Bathla et al. (2022) for intelligent automation systems.

Sensor miniaturisation breakthroughs have enabled unprecedented payload configurations on UAV platforms. The CLaDS system’s achievement of  $\pm 1$  m methane localisation accuracy using drone-mounted retroreflectors (Soskind et al., 2023) exemplifies how aerial platforms can overcome the rural coverage gaps identified in Polish landfill fire monitoring (Bihalowicz et al., 2021). Similarly, the Kolibri system’s optical particle counters provide  $0.1 \mu\text{m}$   $\text{PM}_{2.5}$  resolution (Joíca et al., 2022), delivering the granularity needed to resolve acute exposure peaks during wildfire events that stationary monitors consistently missed (Naqvi et al., 2023). These advancements operationalise the sensor fusion principles outlined by Ignatious et al. (2022) while confronting the weather vulnerability challenges documented by Vargas et al. (2021), particularly the 50% LiDAR visibility reduction in heavy rain and camera performance degradation in fog. The PULSAR platform’s reconfigurable design embodies this progress, combining LiDAR with  $640 \times 512$  IR thermal cameras in a system capable of rapid mode switching between quadcopter and octocopter configurations to optimise either payload capacity (33 kg) or flight endurance (28 minutes) (Perikleous et al., 2024), fulfilling the hybrid system potential foreseen in ecological monitoring research (Besson et al., 2022).

Operational deployments have yielded transformative insights into plume dynamics and system requirements. The Taichung Power Plant studies not only revealed temperature inversion effects but also identified unique VOC signatures (high ethane/ethene ratios) that distinguished power plant emissions from urban pollution sources (Chen et al., 2023), providing critical data to refine the transboundary pollution models that have challenged U.S. wildfire smoke management (Hagler et al., 2021). Similarly, wastewater plant surveys achieved  $<5\%$   $\text{CH}_4$  flux uncertainty without ground measurements using mass balance methods (Gålfalk et al., 2021), demonstrating an alternative to the dispersion modelling approaches that struggled during the Saddleworth Moor fire (Graham et al., 2020). IoT-enabled swarms have further expanded capabilities, dynamically adjusting formations to track chlorine plume boundaries in real-time (Seiber et al., 2018), thereby overcoming the manual positioning delays that hampered emergency responses (Griffiths et al., 2022). These advancements align with the predictive frameworks developed by Seraj et al. (2022) using Adaptive Extended Kalman Filters and FARSITE fire models, while Hu et al.’s (2022) fault-tolerant navigation systems ensure continuous operation even during hardware failures.

The integration of these technological advancements faces persistent challenges that mirror those predicted in comprehensive autonomous vehicle reviews (Parekh et al., 2022). Weather vulnerabilities remain particularly problematic, with LiDAR performance degradation in precipitation and camera scattering in fog (Vargas et al., 2021) requiring mitigation through all-weather RADAR integration and polarised optical systems (Jońca et al., 2022). Energy constraints continue to limit operational durations, as evidenced by the PULSAR’s 28-minute endurance (Perikleous et al., 2024), though solutions are emerging through hydrogen fuel cell prototypes and the energy management strategies proposed in agricultural IoT systems (Chamara et al., 2022). Regulatory barriers to beyond visual line-of-sight (BVLOS) operations (Lelis et al., 2024) are gradually being addressed through compliance with emerging standards like ASTM F3411-19a for UTM integration (Bathla et al., 2022). These challenges are being systematically overcome through innovations such as Wu et al.’s (2024) loss minimisation framework for optimised drone deployment and Lelis et al.’s (2024) three-layer data architecture for standardised spatiotemporal fusion.

The convergence of these technologies creates a robust monitoring paradigm that extends from sensor hardware to public health integration. SegNet’s implementation of segmented CNN architecture achieves 98.2% wildfire detection accuracy with 240 ms/image processing (Jonnalagadda & Hashim, 2024), fulfilling the real-time processing requirements anticipated in autonomous vehicle research (Parekh et al., 2022). When combined with the plume tracking capabilities demonstrated by Assenine et al.’s (2023) cooperative deep reinforcement learning framework and the health impact models developed by Moore et al. (2023), these systems now bridge the critical gap between atmospheric measurements and actionable public health interventions. This comprehensive approach addresses the exposure risks faced by firefighters (Alharbi et al., 2021) while overcoming the resource limitations that prevented multi-incident coverage during the 2018 Saddleworth Moor fire (Graham et al., 2020). As these technologies mature, they fulfil the potential envisioned in ecological monitoring research (Besson et al., 2022), creating an integrated solution that spans from individual sensor nodes to system-wide response coordination, fundamentally transforming our capacity to monitor and mitigate the impacts of fire-induced plumes.

## 4.2 Monitoring

**Overview** Monitoring dynamic, complex environments like fire-induced plumes demands rapid, adaptive decision-making beyond traditional control methods. Reinforcement Learning (RL) and Multi-Agent Reinforcement Learning (MARL) offer powerful frameworks to enable autonomous unmanned aerial vehicles (UAVs) to learn from experience, adapt to environmental changes, and coordinate efficiently without relying on centralised control. In this section, we present the foundational principles of RL and MARL, outline their relevance to UAV-based plume tracking, and review key algorithmic strategies designed to address the unique challenges of this mission-critical application.

### 4.2.1 Reinforcement Learning

Reinforcement Learning (RL) is a type of machine learning where an agent learns to make sequences of decisions by interacting with an environment. This agent receives rewards or penalties based on its actions and aims to maximise its total reward over time.

In simplest terms, RL is the process of an agent learning what actions to take in an environment, at any possible state, through the process of trial and error.

The standard model used in reinforcement learning to define sequential decision-making processes is the *Markov decision process (MDP)*.

**Definition 1 (Markov decision process (Albrecht et al., 2024))** *A finite Markov decision process (MDP) consists of:*

- A finite set of states  $S$ , with a subset of terminal states  $\bar{S} \subseteq S$ .
- A finite set of actions  $A$ .
- A reward function  $\mathcal{R} : S \times A \times S \rightarrow \mathbb{R}$ .

- A state transition probability function  $\mathcal{T} : S \times A \times S \rightarrow [0, 1]$  such that

$$\forall s \in S, a \in A : \sum_{s' \in S} \mathcal{T}(s, a, s') = 1. \quad (1)$$

- An initial state distribution  $\mu : S \rightarrow [0, 1]$  such that

$$\sum_{s \in S} \mu(s) = 1 \quad \text{and} \quad \forall s \in \bar{S} : \mu(s) = 0. \quad (2)$$

In the context of **fire-induced plume tracking using UAVs**, these MDP components can be interpreted as follows:

- **States**  $S$ : The position of the UAV relative to the plume (e.g., coordinates, altitude, and local plume intensity).
- **Actions**  $A$ : Movement commands available to the UAV (e.g., move north, south, ascend, descend, hover).
- **Reward function**  $\mathcal{R}$ : A function that provides higher rewards when the UAV approaches the plume centre or effectively follows the plume path.
- **Transition function**  $\mathcal{T}$ : Probability of transitioning between states, factoring in environmental disturbances such as wind.
- **Initial distribution**  $\mu$ : The starting position of the UAV, usually near but outside the plume.

An MDP starts with the UAV in an initial state  $s^0 \in S$ , sampled according to  $\mu$ . At each time step  $t$ , the UAV observes its current state  $s^t \in S$  and selects an action  $a^t \in A$  according to a policy  $\pi(a^t | s^t)$ , which specifies the probability of taking action  $a^t$  given  $s^t$ .

Given  $s^t$  and  $a^t$ , the UAV transitions to the next state  $s^{t+1} \in S$  according to  $\mathcal{T}(s^t, a^t, s^{t+1})$ , and receives a reward

$$r^t = \mathcal{R}(s^t, a^t, s^{t+1}).$$

This process continues until the UAV reaches a terminal state (e.g., it exits the plume area or the battery is depleted) or a maximum number of time steps is reached.

Each independent mission is referred to as an **episode**.

**Key Concepts in Reinforcement Learning** The RL framework comprises several core components. The *agent* is the decision-maker—in our context, the UAV tracking the fire plume—while the *environment* encompasses everything the agent interacts with, including atmospheric conditions, plume dynamics, and wind patterns. At each moment, the agent observes a *state*  $s$  describing its current situation (e.g., position and sensor readings) and selects an *action*  $a$  from available options (e.g., move north, ascend, or hover). Following each action, the agent receives a scalar *reward*  $r$  indicating the immediate benefit of that action, such as higher values for approaching regions of greater plume concentration.

The agent’s behaviour is governed by a *policy*  $\pi$ , which maps states to actions and represents the learned strategy for plume tracking. To evaluate long-term performance, RL employs two key functions: the *value function*  $V(s)$ , which estimates the expected cumulative reward from a given state under the current policy, and the *Q-function*  $Q(s, a)$ , which estimates the expected cumulative reward of taking a particular action from a specific state and following the policy thereafter.

**Agent-Environment Interaction Loop** The agent-environment interaction cycle is shown in Figure 5, depicting how the agent selects actions, receives feedback, and transitions through states over time.

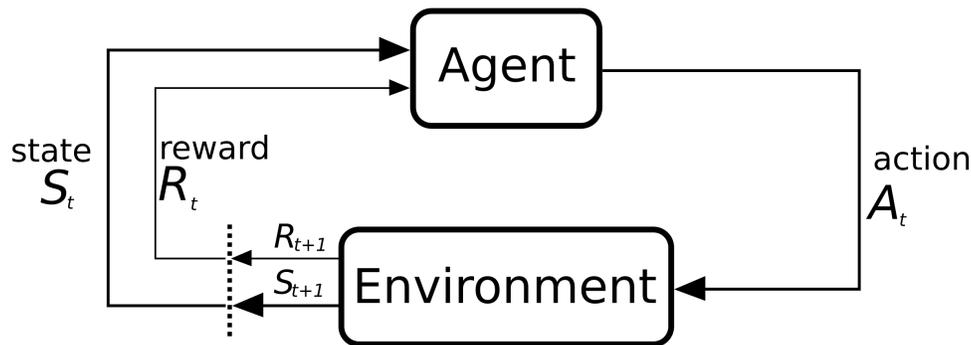


Figure 5: **Adapted from** Sutton & Barto (2018). Agent-environment interaction loop in reinforcement learning.

#### 4.2.2 Fundamental Concepts in Reinforcement Learning and Multi-Agent Reinforcement Learning

Before discussing specific algorithmic approaches for UAV plume tracking, we review key reinforcement learning (RL) and multi-agent reinforcement learning (MARL) concepts that underpin the design and application of these methods.

**Exploration vs. Exploitation** One of the core challenges in RL is the trade-off between exploration and exploitation. *Exploration* involves the agent taking actions that may not yield immediate rewards but help discover potentially better strategies in the long term—for example, a UAV might deliberately investigate a less concentrated region of a plume to locate new high-intensity areas. In contrast, *exploitation* involves leveraging the agent’s current knowledge to select actions known to yield high rewards, such as a UAV consistently following the strongest detected plume signals. Effective learning requires balancing both behaviours, especially in dynamic environments such as fire-induced plumes where conditions can change rapidly.

**On-Policy vs. Off-Policy Learning** RL algorithms can be categorised based on how they use collected experience. *On-policy methods* learn and update a policy using trajectories generated by the current version of the policy itself, prioritising consistency between behaviour and learning; examples include Policy Gradient methods, Proximal Policy Optimisation (PPO) (Schulman et al., 2017), and Multi-Agent PPO (MAPPO) (Yu et al., 2022). *Off-policy methods*, by contrast, learn a target policy using experience collected by different (possibly older or exploratory) policies, allowing more efficient data reuse through techniques like experience replay; examples include Deep Q-Networks (DQN) (Mnih et al., 2015), Soft Actor-Critic (SAC) (Haarnoja et al., 2018), and QMIX (Rashid et al., 2020). For UAV plume tracking, off-policy methods can provide sample efficiency, while on-policy methods can offer stable learning in highly non-stationary environments.

**Value-Based, Policy-Based, and Actor-Critic Methods** RL algorithms also differ in how they represent and optimise agent behaviour. *Value-based methods* estimate the value of actions or states without explicitly maintaining a policy, with the agent selecting actions that maximise estimated value; Q-learning (Kaelbling et al., 1996) is a classic example, where the agent learns a Q-function  $Q(s, a)$  predicting expected cumulative rewards. *Policy-based methods* directly learn a policy  $\pi(a | s)$  mapping states to actions by optimising expected rewards through gradient ascent, naturally handling continuous action spaces and stochastic policies. *Actor-critic methods* combine both approaches by maintaining an actor (the policy) and a critic (a value estimator); the critic evaluates the actor’s actions, and the actor improves based on this feedback, balancing stable value estimation with flexible policy improvement. In MARL, actor-critic meth-

ods such as COMA (Foerster et al., 2018) and MAPPO are widely used to address partial observability and multi-agent credit assignment challenges.

**Cooperative, Competitive, and Mixed Multi-Agent Settings** MARL problems can involve different interaction types among agents. In *cooperative* settings, all agents share a common reward function and work together towards a shared goal; fire plume tracking with multiple UAVs is inherently cooperative, as all UAVs aim to maximise plume coverage whilst minimising energy usage. *Competitive* settings involve agents with opposing objectives, such as zero-sum games, and while less common in plume tracking, remain relevant in adversarial scenarios such as avoiding malicious drones. *Mixed* settings feature some agents cooperating while others compete, a complexity that appears in general surveillance and resource allocation problems but is less central to cooperative plume tracking. The cooperative nature of UAV plume tracking motivates the use of algorithms specifically designed for decentralised cooperation and efficient joint exploration.

### 4.2.3 Multi-Agent Reinforcement Learning

Many real-world problems, including plume tracking, require multiple cooperating agents. Multi-Agent Reinforcement Learning (MARL) extends traditional RL to environments where multiple agents interact, learn, and collaborate to achieve shared or individual goals. In the context of **fire-induced plume tracking using UAVs**, MARL enables a fleet of drones to cooperatively map and follow dynamic plumes through decentralised decision-making, communication, and shared rewards. Figure 6 illustrates the core MARL framework, where agents (UAVs) observe the environment, take individual actions, and jointly modify the environment state.

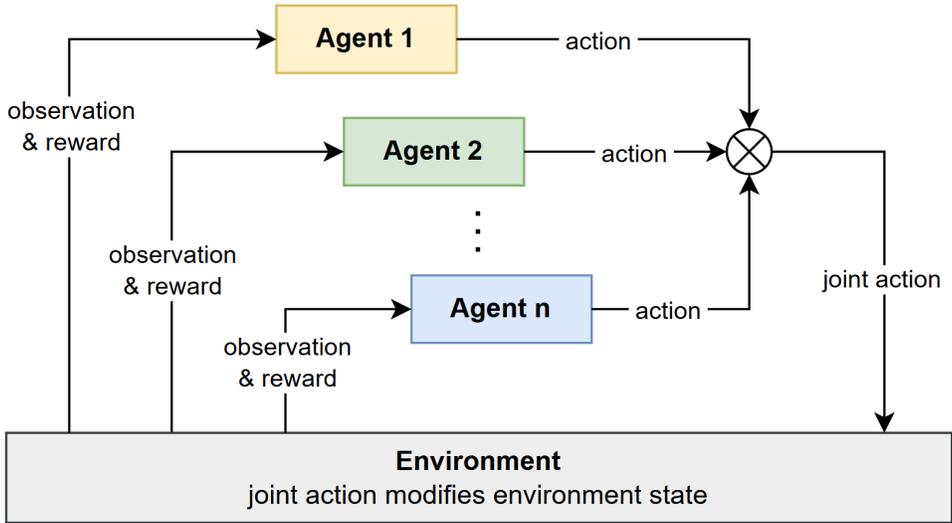


Figure 6: **MARL framework for UAV plume tracking.** Agents observe partial states, take actions, and jointly influence the environment. Adapted from Albrecht et al. (2024).

**Challenges in MARL** Multi-Agent Reinforcement Learning (MARL) introduces fundamental complexities beyond single-agent reinforcement learning, making its direct application to UAV plume tracking both challenging and fascinating (Canese et al., 2021; Orr & Dutta, 2023). These challenges must be carefully considered when designing cooperative multi-UAV systems for real-world deployments.

The first major challenge is *non-stationarity*. In MARL, each agent’s policy evolves as learning progresses, causing the environment to appear non-stationary from any individual agent’s perspective (Canese et al., 2021). Unlike single-agent RL, where environmental dynamics are typically fixed, an agent in MARL must adapt not only to environmental uncertainties such as wind and fire behaviour, but also to the shifting behaviours of teammates. For UAV plume tracking, this necessitates adaptive policies that can cope with evolving group dynamics, fluctuating plume structures, and unexpected disturbances. Methods such as

centralised critics and shared learning signals, as employed in COMA (Canese et al., 2021), help stabilise learning under these conditions.

*Scalability* presents another significant obstacle. As the number of agents increases, the joint action space grows exponentially ( $|A|^n$  for  $n$  agents), making centralised action planning computationally intractable—a phenomenon known as the “curse of dimensionality” that severely impacts both learning efficiency and convergence speed. Techniques like QMIX (Xia et al., 2021) mitigate this challenge by learning factored value functions that approximate the true joint Q-function through monotonic mixing, enabling decentralised decision-making whilst preserving coordination.

*Partial observability* further complicates MARL in realistic plume tracking scenarios, where UAVs possess only local and often noisy observations from onboard gas concentration sensors and limited-range cameras. No single UAV has access to the full global state of the plume or peer positions at all times, requiring agents to infer hidden environmental dynamics or coordinate without direct communication. Methods like MADDPG (Jiandong et al., 2021; Lowe et al., 2017) and MAPPO (Su & Qian, 2023) address this limitation by introducing communication protocols, belief models, or shared centralised critics during training.

Finally, *credit assignment* poses difficulties in cooperative settings where agents receive a shared global reward, making it challenging to determine each agent’s individual contribution to overall success. Poor credit assignment can lead to suboptimal or even adversarial behaviours, such as UAVs avoiding risky but necessary plume sectors. Counterfactual baseline methods like COMA (Canese et al., 2021) and reciprocal reward shaping (Zhou et al., 2022) improve credit assignment by estimating what would have happened had an individual agent acted differently, or by incentivising supportive actions among neighbouring UAVs.

**MARL Approaches for UAV Plume Tracking** To effectively manage these challenges in UAV plume tracking tasks, a number of specialised MARL frameworks and architectures have been developed. These approaches tailor fundamental MARL ideas to the unique requirements of decentralised cooperation, energy efficiency, and robustness under environmental uncertainty.

*Centralised Training with Decentralised Execution (CTDE)* paradigms are especially suited for UAV missions because they allow agents to exploit full environment knowledge during training—such as access to global plume dynamics and full team observations—while operating independently at deployment. QMIX (Xia et al., 2021) exemplifies this approach by learning individual agent Q-functions conditioned on local observations, which are then combined by a learned mixing network that enforces monotonicity with respect to individual utilities. This structure enables highly scalable coordination without requiring agents to explicitly model teammates, and has been successfully applied to problems resembling plume tracking, such as dynamic target pursuit (Hou et al., 2023). COMA (Canese et al., 2021) offers an alternative CTDE approach by implementing a centralised critic that computes counterfactual baselines for each agent’s action, allowing fine-grained credit assignment. In plume tracking tasks, COMA can reward UAVs not only for following dense plume areas but also for positioning themselves to improve the team’s collective sensing coverage.

*Communication-augmented MARL* methods address the partial observability and dynamic environments characteristic of plume tracking by augmenting agent policies with learned communication channels. CommNet (Jung et al., 2021; Shin et al., 2019) introduces differentiable communication by sharing mean-aggregated embeddings of local observations among agents during decision-making, reducing redundant exploration of the same plume regions and improving coverage efficiency. The BRNN-Actor-Critic approach (Jiandong et al., 2021) uses bidirectional recurrent neural networks (Schuster & Paliwal, 1997) to model communication implicitly, allowing agents to share internal hidden states across time steps and enhancing synchronisation for continuous control tasks such as smooth UAV plume tracking in turbulent winds.

When explicit centralised training or communication is infeasible, *decentralised collaboration* strategies emphasise implicit coordination through shaped incentives or limited message passing. Reciprocal reward methods (Zhou et al., 2022) encourage cooperation by modifying agents’ rewards based on their positive influence on nearby agents’ success; in plume-tracking settings, UAVs receiving reciprocal rewards are incentivised to avoid congesting sectors already adequately monitored by neighbours. Sparse-interaction MARL (Darwin & Tamba, 2023) introduces selective collaboration by dynamically switching between independent policies and

group equilibrium policies, enabling robust exploration in sparse regions of the plume whilst maintaining collective behaviour when clustering is necessary.

**UAV-Specific Applications** Beyond general MARL strategies, specific applications and algorithmic adaptations have emerged for plume tracking and wildfire monitoring. For *plume coverage optimisation*, algorithms such as MUSAC (Xia et al., 2021) prioritise maximising spatial information entropy (SIE) across the coverage area. Higher SIE indicates more uniform and comprehensive mapping of the plume, preventing UAVs from redundantly sampling the same high-concentration regions whilst balancing information gain against battery conservation.

*Dynamic task allocation* frameworks address the need for flexible coordination as conditions evolve. The Cooperative Dynamic Task Allocation (CDTA) framework (Liu et al., 2022) facilitates flexible sector division among UAVs through a proposer-responder mechanism that enables rapid adaptation when new plume branches emerge or when UAVs drop out due to battery depletion, ensuring continuous high-fidelity tracking without explicit global supervision.

For *path planning* in complex environments, hybrid approaches such as SAC with AIT\* (Zhao et al., 2024) combine RL with sampling-based motion planning. Soft Actor-Critic (SAC) provides entropy-regularised exploration strategies, while AIT\* (Strub & Gammell, 2020) (Asymptotically-Optimal Incremental Sampling) guarantees asymptotically optimal, collision-free paths. This dual-layered approach is particularly useful in cluttered fire environments where UAVs must navigate complex airspace whilst tracking plumes.

**Key Insights** Synthesising lessons from recent literature and practical deployments, several important insights emerge regarding MARL for UAV plume tracking. With respect to *algorithm selection*, CTDE frameworks such as QMIX and COMA consistently outperform purely decentralised or fully centralised alternatives, effectively balancing the richness of centralised knowledge during training with the robustness and scalability of decentralised execution in the field (Su & Qian, 2023).

*Reward design* proves equally critical: global objectives like maximising cumulative plume coverage must be complemented by local penalties for behaviours that hinder mission goals, such as energy wastage, inter-UAV collisions, or redundant sampling (Hou et al., 2023). Multi-tiered reward structures often lead to more reliable emergent coordination.

Finally, *realism in simulation and training* has emerged as essential for successful deployment. Recent works highlight the importance of physics-based simulation environments such as Gazebo and Webots, along with consideration of operational constraints including limited battery life, sensor noise, and wind effects (Orr & Dutta, 2023; Xia et al., 2021). Oversimplified training environments produce brittle policies that fail when transferred to real-world fire scenarios.

**Summary** Reinforcement Learning and Multi-Agent Reinforcement Learning provide a versatile and scalable toolkit for monitoring fire-induced plumes with autonomous UAV fleets. By framing plume tracking as a sequential decision-making problem under uncertainty, RL and MARL approaches enable UAVs to continuously adapt, explore, and collaborate in dynamic, partially observable environments. Whilst challenges such as non-stationarity, scalability, and credit assignment complicate learning, specialised algorithmic frameworks—ranging from CTDE strategies like QMIX and COMA to communication-augmented and decentralised collaboration methods—offer promising solutions. As research advances, integrating realistic environmental models and carefully engineered reward structures will be essential for deploying robust, reliable monitoring systems in real-world wildfire scenarios.

### 4.3 Forecasting

Fire-induced plume forecasting involves predicting the spatiotemporal evolution of combustion byproducts such as carbon dioxide ( $\text{CO}_2$ ), fine particulate matter ( $\text{PM}_{2.5}$ ), and volatile organic compounds (VOCs) as they disperse through the atmosphere. Hybrid physical/data-driven methods show promise for this task, as demonstrated by Li et al. (2024) for hydrogen safety applications and Liu et al. (2024) for geologic  $\text{CO}_2$  storage, though fire-specific adaptations remain necessary. Current implementations face limitations in

dimensionality and data resolution—whilst Wang et al. (2024) achieved 1 km horizontal resolution for  $\text{PM}_{2.5}$  using satellite-based methods, vertical resolution often remains limited to column-integrated values.

Accurate forecasting provides operational value across three key domains:

- **Early warning systems:** Yu et al. (2023) showed transformer-based models can predict  $\text{PM}_{2.5}$  with RMSE of  $6.92 \mu\text{g}/\text{m}^3$ , demonstrating improved accuracy during wildfire events.
- **Emergency response:** Transfer learning methods from Shi et al. (2023) reduced computational costs by 72% for gas leak forecasting, though fire plumes introduce additional complexities like pyrolysis and buoyancy.
- **Public health protection:** Adaptive graph networks developed by Liu et al. (2023) improved exposure risk assessments.

Key challenges in fire-induced plume forecasting include:

- **Dynamic atmospheric conditions:** Zhang & Zhao (2021) demonstrated that uncertainties in wind fields propagate exponentially, but fire plumes lack equivalent high-resolution input data.
- **Topographic complexity:** Urban geometries induce micro-scale atmospheric circulations that complicate dispersion modelling (Asahi et al., 2023).
- **Emission source uncertainty:** Bayesian frameworks can quantify prediction uncertainty in dispersion modelling (Shi et al., 2021).

Most existing systems operate in two dimensions, whether using ground-based networks (Peralta et al., 2022) or satellite imaging (Rouet-Leduc & Hulbert, 2024). Fully 3D forecasting remains experimental, with frameworks like the CFD-ML model from Li et al. (2024) showing potential but requiring validation for fire scenarios. Forecast horizons vary widely, from 15-minute nowcasting models (Shi et al., 2023) to decadal-scale  $\text{CO}_2$  dispersion projections (Fan et al., 2024), highlighting trade-offs between temporal resolution and forecast duration.

### 4.3.1 Spatiotemporal Forecasting Models

Current approaches to fire-induced plume forecasting can be broadly categorised into three paradigms, as illustrated in Figure 7: physical models grounded in fluid dynamics, data-driven methods leveraging machine learning, and hybrid approaches that combine both.

**Physical Models** Whilst Li et al. (2024) achieved high precision ( $\text{MSE}: 4.60 \times 10^{-5}$ ) in controlled hydrogen plume simulations, these methods face additional challenges in wildfire scenarios including turbulent buoyancy and pyrolysis. The computational requirements are substantial—Asahi et al. (2023) needed 750 LBM (Lattice Boltzmann Method) simulations to model urban dispersion. This creates a fundamental trade-off where increasing spatial resolution improves accuracy but exacerbates computational costs, limiting utility for emergency response.

**Data-Driven Methods** Machine learning models demonstrate strong adaptability to real-world conditions. The graph network from Liu et al. (2023) demonstrated improved accuracy ( $R^2 = 0.94$ ) through dynamic spatial attention mechanisms. However, these gains often come at the cost of physical interpretability, as shown by Fan et al. (2024) through integrated gradients analysis. A key research question is whether architectures can maintain the performance of models like the transformer from Yu et al. (2023) ( $\text{RMSE}: 6.92 \mu\text{g}/\text{m}^3$ ) whilst preserving explainability.

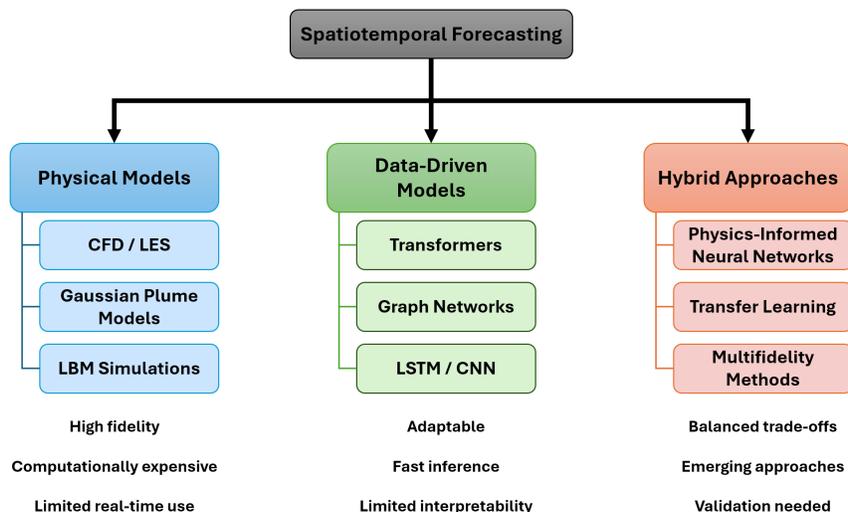


Figure 7: Taxonomy of spatiotemporal forecasting approaches for fire-induced plumes, showing the three main paradigms and representative methods within each.

**Hybrid Approaches** Combining physical and data-driven paradigms has yielded mixed results. The physics-informed neural networks from Zhang & Zhao (2021) achieved wind field reconstruction with RMSE as low as 0.263 m/s, but required 17-hour training times that hinder real-time deployment. Transfer learning frameworks like that from Shi et al. (2023) reduced CFD runs by 72%, but rely heavily on synthetic training data which may not generalise to diverse fire scenarios.

**Future Directions** Four key research directions could address current limitations. First, *3D forecasting* remains underdeveloped: current models focus primarily on 2D horizontal dispersion (Li et al., 2024), whilst vertical plume dynamics such as smoke injection height remain understudied despite their importance for long-range transport. Second, *real-time execution* at scale presents ongoing challenges; whilst Shi et al. (2023) achieved 13 ms inference times, scaling these approaches to regional forecasts whilst maintaining sub-minute latency remains difficult. Third, achieving *field-deployable accuracy* requires bridging the gap between high-resolution urban predictions (Asahi et al., 2023) and the computational constraints of remote deployment, potentially through innovations like quantised neural networks. Fourth, *edge computing integration* offers promising opportunities, with adaptive graph networks (Liu et al., 2023) showing potential for drone swarm implementations with onboard processing. Together, these directions could lead to more resilient and responsive forecasting systems that integrate effectively with real-time autonomous platforms.

#### 4.3.2 Data Sources and Preprocessing for Forecasting

The accuracy of fire-induced plume forecasts depends fundamentally on available data sources and preprocessing methods. Whilst multi-source data integration has advanced significantly, important limitations remain in observational capabilities and their alignment with modelling requirements.

**Remote Sensing Capabilities** Current remote sensing platforms offer complementary strengths and weaknesses for plume monitoring. Earlier-generation instruments like MODIS and VIIRS provide daily global coverage but lack spatial precision for localised dispersion modelling. Geostationary satellites such as Himawari-8 provide hourly AOD data for  $\text{PM}_{2.5}$  estimation (Wang et al., 2024), whilst polar-orbiting platforms like Sentinel-2 have 5-day revisit cycles that may miss critical plume evolution phases (Rouet-Leduc & Hulbert, 2024). Although these systems achieve 1 km resolution for aerosol optical depth and 20 m for methane detection (Khirwar & Narang, 2024), their column-integrated measurements cannot resolve

vertical stratification crucial for understanding smoke injection heights. This limitation necessitates reliance on reanalysis data like ERA5 for vertical profiling, introducing potential errors into forecasting models.

**Ground and UAV-based Monitoring** Ground stations and UAV networks help address satellite resolution limitations but face different constraints. Fixed monitoring stations used in adaptive graph networks (Liu et al., 2023) suffer from spatial sparsity—for example, Beijing’s network of 35 stations provides only one sensor per 50 km<sup>2</sup>, insufficient for neighbourhood-scale exposure assessment. UAVs could provide intermediate-scale coverage, but their use in operational forecasting remains limited due to endurance constraints. The lack of UAV-derived datasets in published research highlights an important gap between monitoring capabilities and predictive applications.

**Meteorological Data Challenges** Atmospheric modelling of fire-plume interactions involves inherent uncertainties. Whilst Zhang & Zhao (2021) achieved 0.263 m/s RMSE in wind field reconstruction using physics-informed neural networks, this precision requires LIDAR data rarely available during wildfires. Operational systems typically substitute coarser reanalysis data, such as ERA5’s 31 km resolution wind fields (Yu et al., 2023), creating scale mismatches between atmospheric forcing and plume dispersion. These issues are particularly significant for fire-generated pyroconvection, where Asahi et al. (2023) showed that sub-kilometre wind variations can alter concentration patterns by 300%.

**Preprocessing Requirements** Current preprocessing methods face challenges in balancing competing priorities. Z-score normalisation effectively handles sensor variability but discards measurement confidence intervals (Wang et al., 2024), whereas Bayesian frameworks preserve uncertainty propagation but require substantially greater computational resources (Shi et al., 2021). This trade-off becomes particularly problematic for real-time systems aiming for 13 ms inference times (Shi et al., 2023). Future preprocessing approaches should integrate more closely with end-to-end learning systems, following examples like the Dartboard Spatial MSA from Liang et al. (2023) that compensates for imperfect inputs through architectural innovation. As forecasting moves towards real-time autonomous deployment, preprocessing must evolve from passive data cleaning to an active interface between sensing systems and models.

### 4.3.3 Real-time Forecasting Systems and Deployment

The implementation of fire-induced plume forecasting systems in real-world scenarios reveals significant challenges in transitioning from theoretical models to practical deployments. Whilst studies such as Li et al. (2024) and Shi et al. (2023) demonstrate promising results in controlled environments, several critical limitations emerge when considering operational deployment.

**Latency and Accuracy Trade-offs** Current forecasting systems can be categorised into two approaches with distinct limitations. Cloud-based solutions like that proposed by Liang et al. (2023) achieve accurate 72-hour predictions but may face latency challenges in real-time deployment, making them unsuitable for rapid response scenarios. In contrast, edge computing implementations such as Shi et al. (2023) achieve 13 ms response times but sacrifice accuracy by relying on simplified Gaussian plume models. Hybrid approaches like Yu et al. (2023) attempt to balance these trade-offs, but still face challenges with false positives during sudden wind changes.

**Data Integration Challenges** Operational forecasting systems must integrate data from multiple sources with varying characteristics. Regarding *temporal resolution*, satellite revisit cycles (5+ days for Sentinel-2 (Rouet-Leduc & Hulbert, 2024)) differ significantly from UAV sampling rates (minutes) and fixed sensor frequencies (seconds). *Spatial resolution* poses additional constraints, as MODIS sensors provide only 1 km resolution (Wang et al., 2024), insufficient for capturing small-scale eddies affecting plume dispersion. Furthermore, most systems face limitations in *physical modelling*, lacking real-time pyrolysis data and relying instead on constant emission assumptions. As demonstrated by Asahi et al. (2023), these mismatches can lead to 300% concentration errors in urban environments. The absence of standardised data fusion protocols further complicates system integration.

**Practical UAV Limitations** Whilst UAVs are frequently proposed as mobile sensing platforms in studies like Li et al. (2024) and Liu et al. (2024), several practical constraints hinder their adoption. These include airspace restrictions during active fire events, limited battery life (typically less than 30 minutes for drones carrying gas sensors), payload constraints (under 500 g for most commercial UAVs), and the lack of standardised protocols for hazardous operations. These limitations explain the absence of real UAV field data in current research, despite their theoretical potential for plume monitoring.

**Operational Implementation Challenges** A disconnect exists between forecasting system outputs and practical decision-making needs. Whilst probabilistic frameworks like Shi et al. (2021) provide rigorous uncertainty quantification, their complexity exceeds the requirements of most incident command systems. Conversely, simplified binary alerts as used by Asahi et al. (2023) discard potentially critical uncertainty information. This mismatch between technical capabilities and operational needs contributes to the limited real-world adoption of advanced forecasting systems.

Future progress will require coordinated development of both technical solutions and operational frameworks. The layer-wise prediction approach demonstrated by Fan et al. (2024) represents one promising direction, but comprehensive solutions must address the entire system from sensor networks to decision support tools. Without such holistic development, even the most sophisticated forecasting models risk remaining academic exercises rather than practical tools for emergency response.

#### 4.3.4 Evaluation Metrics and Uncertainty Quantification

Current approaches to validating fire-induced plume forecasting systems face significant challenges in bridging the gap between technical metrics and operational requirements. Traditional evaluation methods often fail to capture the complex decision-making needs of environmental emergencies.

The current landscape of performance metrics reveals important limitations. Whilst Li et al. (2024) report  $R^2$  values approaching 0.98 for hydrogen plume forecasting, such aggregate measures can mask critical spatial variations in prediction quality. Asahi et al. (2023) demonstrate through  $\text{FAC}_2$  analysis that only 64% of predictions fall within a factor of 2 of ground truth, despite high  $R^2$  values, limiting their practical reliability. This highlights a fundamental challenge in current evaluation frameworks—the mismatch between measurable statistics and operational requirements. Yu et al. (2023) further show how conventional MAE/RMSE metrics may overlook crucial timing errors in peak concentration predictions.

Current approaches to uncertainty quantification present two distinct methodologies with complementary strengths and weaknesses. The Bayesian framework developed by Shi et al. (2021) provides rigorous probabilistic bounds through Monte Carlo sampling and variational inference, but requires  $17\times$  more computational resources than deterministic models. In contrast, the ensemble methods employed by Fan et al. (2024) generated 108 simulation scenarios covering geological and operational variability, though with potentially reduced statistical rigour.

The practical communication of uncertainty information remains a significant challenge for operational deployment. Liang et al. (2023) observed that despite having advanced probabilistic capabilities, many systems ultimately revert to simpler binary alerts in practice. Similarly, the three-tier confidence system developed by Shi et al. (2023) achieved better adoption than more complex alternatives, suggesting practical limits to how much uncertainty information emergency responders can effectively utilise during crises.

Moving forward, the field requires evaluation frameworks that better address the needs of emergency response. Future developments should focus on metrics that directly quantify decision-critical errors rather than statistical abstractions, whilst maintaining a balance between computational feasibility and information value. Most importantly, forecasting systems must be evaluated not just by their technical sophistication, but by their ability to support effective decision-making during actual emergencies. Achieving this will require sustained collaboration between model developers and operational end-users to ensure evaluation approaches reflect real-world requirements whilst maintaining scientific rigour.

## 4.4 Simulation

### 4.4.1 The Role and Rationale for Simulation

Simulation has emerged as a foundational tool in wildfire plume research, addressing critical gaps in both data availability and system development. Real-world data on plume dispersion is inherently limited by safety concerns, logistical constraints, and temporal sparsity during fast-evolving wildfire events. As Brucker et al. (2022) highlight, simulation experiments help overcome these challenges by enabling controlled exploration of burn dynamics and their environmental impacts, particularly when collecting in situ data is infeasible or hazardous.

Moreover, wildfire behaviour is shaped by complex fire-atmosphere interactions, making accurate forecasting and autonomous system training heavily dependent on dynamic simulation environments. Moisseeva & Stull (2021) use a coupled WRF-SFIRE large-eddy simulation (LES) approach to generate synthetic plume data, showing that energy balance parameterisations can achieve reasonable predictive performance at low computational cost. Similarly, Baggio et al. (2022) demonstrate the importance of simulating plume dynamics in tandem with meteorological feedback using a Meso-NH and ForeFire coupled model. These studies reinforce the value of simulation not only for understanding plume rise mechanics, but also for supporting real-time emergency response and autonomous decision-making.

The need for simulation also stems from the lack of long-term, spatially dense monitoring networks, especially in urban and complex terrain environments. As Barlow (2014) discusses in the context of the urban boundary layer, surface heterogeneity and turbulence complicate direct measurements, requiring model-driven interpolation and forecasting tools. This becomes especially relevant when considering drone deployment in urban-wildland interface regions, where localised meteorology influences both plume transport and UAV flight behaviour.

Sun et al. (2023) further emphasise simulation’s role in risk assessment, fire spread forecasting, and preparedness strategy development. Simulation models allow researchers to explore variable conditions, assess response strategies, and guide decision support systems. Mateus et al. (2024) show how qualitative and quantitative validation techniques using smoke visualisation and CFD help verify model fidelity, making simulation a powerful bridge between controlled experimentation and predictive modelling.

Finally, Ke et al. (2021) underscore the importance of plume simulation in global atmospheric modelling, where plume rise height influences smoke transport and pollutant lifetime. Their development of a global dataset and parameterisation scheme reinforces the simulation’s critical role in climate-scale modelling as well as short-term operational forecasting.

Taken together, these studies illustrate the indispensable role of simulation in generating realistic data for forecasting algorithms, validating plume behaviour under diverse conditions, and training autonomous UAV systems for intelligent fire response.

### 4.4.2 Key Parameters and the Necessity of 3D Modelling

Realistic simulation of wildfire smoke plumes requires careful attention to the physical, chemical, and environmental parameters that govern their evolution. Core variables such as wind speed, atmospheric stability, fuel heat release, injection height, and turbulence intensity strongly influence plume structure and trajectory. As highlighted by Tenti & Ferrero (2024), even variations in the drag coefficient—used in plume rise models—can significantly affect the predicted plume height and pollutant concentration fields. Their study, based on prescribed fire campaigns, emphasises that accurate representation of these parameters is essential for meaningful dispersion modelling.

Plume height itself is a particularly critical variable, as it governs the vertical extent of smoke injection into the atmosphere and determines subsequent passive pollutant transport. Ke et al. (2021) show that parameterising plume rise accurately is vital for calculating the lifetime and dispersal range of smoke particles in climate models. This importance extends to shorter timescales, where fine-grained 3D structure influences forecasting and UAV-based sensing strategies. Similarly, Moisseeva & Stull (2021) demonstrate that crosswind-integrated injection height can be approximated using simple energy balance models derived

from LES-based simulations, providing a balance between physical realism and computational efficiency. We note that with uncontrolled fires, their plumes are released to atmosphere at ground level.

Dimensionality is a key concern. Whilst 2D or Gaussian-based models offer computational efficiency, they are insufficient for capturing the vertical dynamics and spatial heterogeneity present in real plumes. Clements et al. (2024) compare a Large Eddy Simulation (LES) implemented in OpenFOAM with a Gaussian Plume Model (GPM) implemented in ADMS. Their findings show that whilst both models can reproduce observed  $\text{PM}_{2.5}$  concentration profiles, the LES approach provides greater spatial accuracy—particularly in vertical plume structure—albeit at a higher computational cost. This accuracy is critical for simulating UAV sensor exposure across different altitudes and for informing control strategies that depend on dynamic vertical gradients.

Three-dimensional modelling also becomes essential when considering heterogeneous environments and atmospheric chemistry. Mateus et al. (2024) validate a CFD model for thermal plumes using real-scale smoke visualisation experiments, revealing strong agreement in vertical velocity and temperature profiles. This form of validation is particularly valuable for ensuring UAV agents can rely on realistic sensory input in simulation. Li et al. (2022) further show that wind speed and platform motion significantly influence plume rise behaviour, underlining the need to capture dynamic boundary conditions in 3D.

Moreover, Wang et al. (2021) use a chemistry-coupled LES model to examine chemical heterogeneity within fire plumes, finding that photochemical reactions are highly stratified, with active oxidation at the plume edges and different reaction pathways in the core. Such insights would be lost in models lacking vertical and cross-sectional resolution. Finally, Barlow (2014) underscores how surface heterogeneity and the urban roughness sublayer complicate the structure of the atmospheric boundary layer, reinforcing the need for spatially resolved, high-fidelity simulations in built environments.

Taken together, these studies demonstrate that high-resolution 3D simulations—whilst computationally intensive—are indispensable for capturing the complex, multiscale dynamics of smoke plumes, validating forecasting models, and informing UAV decision-making in real-world operational contexts.

#### 4.4.3 Existing Simulation Tools and Their Capabilities

A wide range of simulation tools have been developed for modelling wildfire behaviour and plume dispersion, each with varying degrees of physical fidelity, computational cost, and application specificity. Table 2 summarises the key characteristics of prominent tools discussed in this section. These models can be broadly categorised into empirical methods, dynamic physical simulators, and high-fidelity computational fluid dynamics (CFD) frameworks.

Table 2: Comparison of wildfire and plume simulation tools. Fidelity refers to physical accuracy; speed indicates computational efficiency; 3D indicates full three-dimensional modelling capability; plume indicates explicit smoke/pollutant dispersion modelling. Classifications based on descriptions in the cited literature; see referenced studies for detailed specifications.

<b>Tool</b>	<b>Type</b>	<b>Fidelity</b>	<b>Speed</b>	<b>3D</b>	<b>Plume</b>
PHOENIX Rapidfire	Empirical	Low	Fast	No	Limited
FlamMap / Farsite	Dynamic	Medium	Medium	No	No
Prometheus	Dynamic	Medium	Medium	No	No
WRF-SFIRE	Coupled	Med-High	Slow	Yes	Yes
Meso-NH + ForeFire	Coupled	Med-High	Slow	Yes	Yes
WFDS	Full CFD	High	Very Slow	Yes	Yes
FIRETEC	Full CFD	High	Very Slow	Yes	Yes
OpenFOAM (LES)	CFD	Very High	Very Slow	Yes	Yes
ADMS (GPM)	Gaussian	Low	Very Fast	Limited	Yes
HYSPLIT	Lagrangian	Medium	Fast	Yes	Yes

Singh et al. (2025) provide a comprehensive review of wildfire simulators used in Australia, including PHOENIX Rapidfire, SPARK, AUSTRALIS, REDEYE, and IGNITE. These tools range from empirical

models based on historical correlations to dynamic physical simulators that integrate atmospheric conditions and terrain data. Empirical models offer computational efficiency but are limited in their ability to capture complex plume behaviours, whilst dynamic tools are more physically grounded but require detailed environmental input and longer runtimes.

Ghodrat et al. (2022) present a broader review of simulation software used in wildfire and wind-fire interaction modelling, including Prometheus, Phoenix, FireStation, HFire, FlamMap, WRF-Fire, FIRETEC, WFDS, and FIRESTAR. Each of these platforms has distinct strengths: for example, WFDS and FIRETEC are full-physics CFD tools capable of modelling fire spread and smoke transport in three dimensions, whilst tools like FlamMap and Farsite focus more on terrain-based fire progression. The authors provide a critical comparison of their capabilities and limitations, emphasising the need for more efficient, accurate, and flexible modelling tools.

Multifidelity simulation approaches offer a promising compromise between physical realism and computational cost. Valero et al. (2021) demonstrate that control variates and multilevel Monte Carlo strategies can accelerate full-physics fire simulations by up to two orders of magnitude, whilst still capturing key behaviours such as fire rate of spread and environmental sensitivity. Their use of the Wildland-Urban Interface Fire Dynamics Simulator (WFDS) highlights the role of high-fidelity CFD in modelling coupled fire and smoke dynamics.

Several tools have been adapted or extended specifically for plume modelling. Baggio et al. (2022) combine Meso-NH (an atmospheric model) with ForeFire (a fire spread model) to simulate plume dispersion with meteorological feedback effects. Moisseeva & Stull (2021) use WRF-SFIRE in LES mode to generate synthetic datasets for smoke plume rise. Clements et al. (2024) compare two widely used dispersion models—LES (OpenFOAM) and GPM (ADMS)—highlighting their respective strengths in accuracy and speed. Similarly, Li et al. (2022) employ CFD to simulate plume behaviour from moving ships, proposing simplifications to reduce computational demand.

Some models have also been applied to industrial and built environments. John et al. (2021) evaluate the STAR-CCM+ CFD tool for simulating fire and smoke propagation in nuclear facility scenarios, finding that ventilation settings significantly affect smoke mixing and stratification. Mateus et al. (2024) validate CFD models using coloured smoke visualisation in large indoor spaces, offering insights into qualitative and quantitative consistency between measured and simulated plume behaviour.

Salis et al. (2021) demonstrate how simulation outputs can be combined with land use, ignition data, and weather records to assess wildfire exposure and transmission in real geographic contexts. Whilst their focus is not on plume dynamics per se, it illustrates how simulation platforms can be embedded into decision-support frameworks for risk management.

Collectively, these tools span a spectrum from rapid, empirical predictors to fine-grained, three-dimensional CFD models. Their selection depends heavily on the intended application—whether for real-time forecasting, autonomous control training, or policy planning—and on the trade-offs between physical realism, spatial resolution, and computational feasibility.

#### 4.4.4 Challenges and Integration with Forecasting and Control

Despite significant advances in simulation methods for wildfire plumes, several critical challenges remain—particularly when integrating these simulations with forecasting algorithms and autonomous control systems. One key limitation is the trade-off between physical fidelity and computational efficiency. High-resolution models such as Large Eddy Simulation (LES) provide excellent spatial detail, but are often too computationally demanding for real-time applications or extensive training datasets. Clements et al. (2024) compare LES and Gaussian plume models, showing that whilst LES offers greater accuracy, it comes at the cost of efficiency—limiting its applicability for rapid-response systems or continuous UAV control feedback loops.

This concern is echoed by Moisseeva & Stull (2021), who propose an energy balance parameterisation for plume rise that achieves reasonable accuracy with minimal computational burden. Their approach illustrates a broader need for hybrid or reduced-order modelling strategies that can be used in autonomous system training without sacrificing too much realism. Similarly, Valero et al. (2021) demonstrate how multifidelity

methods, such as multilevel Monte Carlo and control variates, can retain the predictive capabilities of high-fidelity models whilst reducing simulation time by orders of magnitude. These approaches also enable uncertainty quantification, which is crucial when forecasting plume spread in operational contexts.

Another major challenge lies in the co-simulation of fire dynamics and atmospheric feedback. Baggio et al. (2022) address this by coupling Meso-NH with ForeFire to produce smoke dispersion forecasts that adapt to changing meteorological conditions. Their framework supports faster-than-real-time predictions, making it particularly relevant for emergency response and mission planning. However, such tightly coupled systems are still rare, and their integration with autonomous UAV control architectures remains largely unexplored.

The chemical complexity of plumes also introduces challenges for integration with AI-based forecasting systems. Wang et al. (2021) highlight how photochemical and dark chemical reactions evolve spatially within plumes, often requiring high spatial resolution to model accurately. This poses difficulties for machine learning models trained on low-resolution or homogenised data, which may miss critical gradients or overgeneralise.

From the control perspective, Tzoumas et al. (2023) show that UAV swarms can achieve high wildfire detection coverage using decentralised control algorithms. Whilst their work focuses on fire detection rather than plume following, the principles of dynamic area partitioning and robustness to agent failure are highly relevant for designing autonomous sensing systems that can adapt to evolving smoke conditions. Simulation environments that accurately model plume evolution are a necessary precondition for training such control policies under realistic constraints.

Validation remains another open issue. Mateus et al. (2024) demonstrate how experimental visualisation techniques can be used to evaluate the accuracy of CFD-based thermal plume models, yet such validations are not commonly incorporated into real-time or learning-based frameworks. Ke et al. (2021) further emphasise that injection height—often a simplification or heuristic in many models—plays a dominant role in determining smoke transport and pollutant lifetime, directly affecting the quality of forecasts.

Overall, integrating plume simulation with forecasting and control requires balancing realism, efficiency, and generalisability. This involves not only improving simulation tools but also designing training protocols and feedback mechanisms that allow autonomous systems to reason under uncertainty, adapt to changing conditions, and operate safely in complex atmospheric environments.

## 5 Health Risk Modelling

While Sections 3 and 4 established how fire-induced plumes emerge, evolve, and can be tracked in real time, understanding plume dynamics alone is insufficient for effective public health response. The critical next step is translating atmospheric measurements into meaningful estimates of human harm. Health risk modelling provides this interpretive layer, transforming spatiotemporal pollutant data into predictions of morbidity, mortality, healthcare demand, and long-term population impacts.

This section reviews the methodological foundations of health risk modelling in the context of fire-induced air pollution, examining how exposure data are converted into estimates of biological response, clinical burden, and population-level risk. We first describe the core modelling pipeline—from exposure assessment through dose–response estimation to probabilistic risk quantification—before discussing how vulnerability, inequality, and population heterogeneity shape health outcomes. We then examine the limitations of existing public health protocols for acute fire events and explore how artificial intelligence enables more adaptive, personalised, and real-time risk forecasting. Together, these approaches position health modelling as the decision-making interface between autonomous sensing systems and emergency response.

Wildfires are no longer seasonal anomalies; they are becoming sustained, climate-driven threats with wide-ranging consequences. As shown in Section 3, the acute health impacts of fire-related pollutants are potentially substantial, with observed spikes in respiratory and cardiovascular morbidity during and after major fire events. Section 4 further demonstrated how spatiotemporal models can successfully track the atmospheric dynamics of these pollutants. However, understanding where smoke travels is only part of the equation; the next critical step is predicting how it will affect human health.

## 5.1 How Health Risks Are Calculated

Quantifying health risks from wildfire smoke exposure involves a multi-step modelling chain that begins with pollutant detection and ends with predictive estimates of health burden. The fundamental components of this chain—exposure assessment, dose-response estimation, and risk quantification—form the methodological foundation for health risk modelling systems. Figure 8 illustrates this modelling chain and the data sources informing each stage.

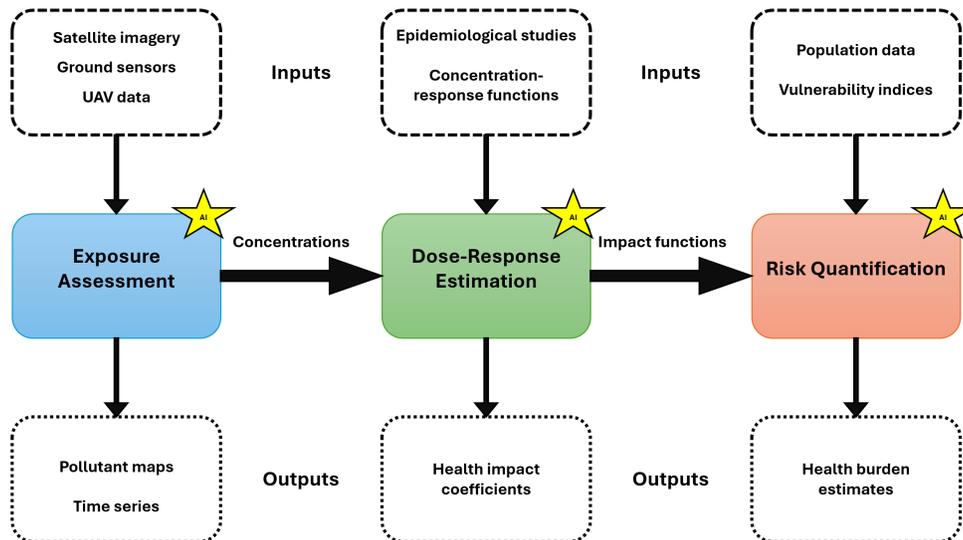


Figure 8: Health risk modelling pipeline showing the three-stage process from pollutant detection to health burden estimation, with key data inputs and emerging AI integration points.

The first stage, *exposure assessment*, focuses on identifying the concentration and spread of harmful pollutants such as PM<sub>2.5</sub>, PM<sub>10</sub>, VOCs, CO, and ozone. As established in Section 4, spatiotemporal pollutant modelling provides the geographical and temporal resolution required for precise exposure estimates. However, fire-related exposure is uniquely complex: pollutant concentrations fluctuate rapidly, vary across microclimates, and can be strongly influenced by wind and topography. Several studies have combined satellite imagery, ground-based monitoring, and machine learning to model pollutant plumes more accurately in near-real time (Aldahlawi et al., 2024; Zhao et al., 2021).

The second stage involves establishing *dose-response relationships* that translate pollutant levels into estimated biological impact. Classical models rely on concentration-response functions (CRFs) derived from epidemiological studies, often assuming linear relationships between exposure and incidence of outcomes such as asthma exacerbation, myocardial infarction, and premature mortality (Hoffmann et al., 2021). However, these relationships are frequently context-dependent; Lu et al. (2022) note that the same PM<sub>2.5</sub> exposure can yield vastly different outcomes depending on age, pre-existing conditions, and socioeconomic status. Fire-specific dose-response data remain limited, prompting increased interest in adaptive modelling techniques that can recalibrate based on local conditions and populations.

The final stage, *risk quantification*, has evolved from deterministic health burden estimations to probabilistic frameworks. Traditional approaches often apply attributable risk models or population impact fractions using static CRFs, whereas more recent innovations include Bayesian networks and Monte Carlo simulations that propagate uncertainty across the exposure-disease chain (Zhang et al., 2024a). These probabilistic models are particularly useful during fire events, where both the pollutant landscape and population movement are highly dynamic.

Whilst deterministic models provide foundational insights, they often lack the capacity to accommodate uncertainty inherent in real-world scenarios. Probabilistic frameworks such as Monte Carlo simulations and

Bayesian inference offer a more robust alternative by quantifying variability in both exposure and response. For example, Hughes et al. (2024) demonstrate the use of Bayesian models to iteratively update risk estimates based on emerging clinical and environmental data—a methodology highly applicable to wildfire contexts where real-time updates are critical. Similarly, agent-based models, such as those used in toxic gas evacuation modelling by Zhang et al. (2024a), can simulate heterogeneous individual responses to exposure events, allowing for scenario testing under variable conditions.

Emerging AI-based systems are also being integrated at this stage to capture non-linear patterns and individual-level variability. For instance, DirPred (Niu et al., 2024) uses neural attention mechanisms to model risk across longitudinal EHRs, offering the potential to personalise health impact predictions based on both environmental exposure and medical history. These probabilistic and agent-based approaches serve as a methodological bridge to more advanced AI-enabled systems, which incorporate dynamic, high-dimensional inputs for real-time individualised risk forecasting.

Despite these advances, significant challenges remain: limited temporal resolution in dose-response estimates, sparse population-specific CRFs for fire pollutants, and the need to validate AI models against real-world health outcomes. Nevertheless, these methods represent an increasingly robust toolkit for translating atmospheric data into actionable health intelligence.

## 5.2 Vulnerable Populations and Human Impact

Fire-related air pollution does not affect all populations equally (Griffiths et al., 2025). The intensity, duration, and timing of exposure—combined with individual susceptibilities—create a layered landscape of health vulnerability. Modelling health risk without accounting for these disparities can lead to misleading predictions and inequitable response strategies.

Among the most affected groups are **children**, whose developing respiratory systems and higher breathing rates make them more susceptible to inhaled pollutants. Exposure to PM<sub>2.5</sub> during wildfires has been associated with increased rates of lower respiratory infections and reduced lung development, outcomes that may persist into adulthood (Hoffmann et al., 2021; Lu et al., 2022).

**Older adults**, particularly those with pre-existing cardiovascular or respiratory conditions, experience heightened morbidity and mortality during fire events. As discussed in Section 3, emergency hospitalisations often spike for asthma, COPD, and ischaemic heart disease. de Hond et al. (2022) emphasise the importance of stratifying risk models by age and comorbidity to accurately capture these elevated risks.

**Pregnant individuals** also represent a critical population. Exposure to wildfire smoke has been associated with adverse birth outcomes such as preterm delivery and low birth weight. These impacts are thought to result from systemic inflammation and placental oxidative stress triggered by particulate matter inhalation (Albahri et al., 2023).

**First responders and frontline workers**, including firefighters and paramedics, face both acute and chronic exposure risks. Darwiesh et al. (2022) highlight the need for real-time monitoring systems tailored to occupational exposures, which often occur at pollutant concentrations far exceeding ambient levels. Long-term exposure may contribute to diminished pulmonary function, elevated cancer risk, and stress-related disorders.

Evacuation challenges further exacerbate risks for both vulnerable individuals and responders. Studies such as Munawar et al. (2022) underscore the difficulty of timely evacuation in aged care facilities during flood or fire events, pointing to the need for AI-enhanced planning tools that can optimise routing and response in real time. Similarly, Zhang et al. (2024a) use multi-agent models to simulate evacuation under toxic gas releases, demonstrating how fear, uncertainty, and communication delays can influence exposure outcomes.

**Low-income and marginalised communities** often face a compounding burden: they are more likely to reside in high-exposure zones, less likely to have access to air filtration or medical services, and more frequently excluded from real-time alerts and health risk communications. These inequities are exacerbated by structural factors such as housing quality, insurance gaps, and digital divides (Meskó & Topol, 2023).

Social vulnerability also plays a key role. Populations without access to protective infrastructure—such as HEPA filtration, sealed indoor environments, or private transportation—face disproportionately higher exposure and recovery barriers. As Darwiesh et al. (2022) note, using social media and public sentiment analysis can help identify areas of heightened concern or neglect in real time, creating opportunities for more targeted intervention.

Effective health risk models must therefore move beyond population averages and incorporate demographic, occupational, and socioeconomic dimensions. AI-enabled systems—such as DirPred—offer the potential for granular, individual-level forecasting, but only if such dimensions are explicitly integrated into model architecture and training data. Without such attention to vulnerability, technological advances may reinforce rather than reduce health disparities.

### 5.3 Existing Protocols and Gaps

Health risk assessment during wildfire events continues to rely heavily on generalised public health frameworks developed for ambient air quality management. As outlined in Section 3, existing exposure assessment techniques rely on a mix of satellite imagery, land-use regression models, and ground-based sensor networks. Whilst these tools form a technical foundation of health risk monitoring, they are typically embedded within broader public health frameworks such as Air Quality Index (AQI) and WHO guidelines, noting again that these are 24-hour based for many pollutants. Generically, these frameworks interpret raw environmental data to predict conditions that are then operationalised for ambient exposure decision-making. Chief among these, the Air Quality Index (AQI) translates pollutant concentrations—most notably  $\text{PM}_{2.5}$ ,  $\text{PM}_{10}$ ,  $\text{O}_3$ ,  $\text{CO}$ , and  $\text{NO}_2$ —into a categorical scale ranging from “Good” to “Hazardous”. The AQI is designed for clarity in public communication. As Zhao et al. (2021) note, wildfire plumes are dynamic, and pollutant concentrations can shift dramatically over short distances and timescales, often outpacing the update cycles of AQI systems.

Furthermore, the WHO’s 2021 Air Quality Guidelines (AQGs) provide globally recognised exposure thresholds, recommending significant reductions in acceptable concentrations for  $\text{PM}_{2.5}$  and  $\text{NO}_2$  compared to earlier standards (Hoffmann et al., 2021). However, these guidelines primarily reflect long-term averages (e.g., annual or daily means), making them ill-suited for evaluating acute, high-intensity exposure events such as those caused by wildfires. During extreme fire episodes, pollutant levels can exceed WHO-recommended thresholds within hours, exposing populations to levels that existing policy tools are not calibrated to address (Lu et al., 2022). An alternative approach that uses 1-hour monitoring data to predict the probability that a 24-hour standard will likely be exceeded has been suggested specifically for acute incident response (Deary & Griffiths, 2021).

Operational responses such as shelter-in-place advice, school closures, or evacuation orders are typically informed by dynamic public health risk assessments guided by exposure guidelines. These interventions often fail to incorporate individual-level risk variations such as those with pre-existing disease. Darwiesh et al. (2022) argue that such uniform guidance overlooks the heightened susceptibility of specific groups, including individuals with chronic respiratory disease, children, and the elderly. The resulting blanket policies may underprotect those most at risk or overextend resources where risk is minimal.

### 5.4 The Role of Artificial Intelligence in Health Risk Modelling

Artificial Intelligence (AI) is emerging as a pivotal tool in advancing health risk modelling during wildfire events, offering capabilities that extend beyond the static, generalised frameworks critiqued in the previous section. By integrating environmental exposure data with electronic health records, behavioural patterns, and geospatial analytics, AI enables risk prediction that is both temporally dynamic and individually responsive.

A key advancement is the shift from deterministic to probabilistic modelling. Traditional exposure-response relationships often assume linear, population-wide effects, whereas AI models can dynamically learn complex, non-linear interactions from heterogeneous datasets. For example, the DirPred model proposed by Niu et al. (2024) combines Dirichlet Process Mixture Models with attention-based neural networks to analyse

longitudinal electronic health records (EHRs). This approach not only predicts individual health risks but also generates interpretable evidence at both the cluster and local levels—allowing clinicians to understand the rationale behind each prediction.

Beyond structured clinical data, AI has also been used to leverage unstructured information such as social media signals. Darwiesh et al. (2022) demonstrate the utility of Natural Language Processing (NLP) to assess operational and reputational risks in healthcare systems during crises. Whilst their model was applied in institutional settings, the framework is adaptable to fire-related public health scenarios—for example, detecting early signals of asthma exacerbations or ER overloads through Twitter activity during wildfire events.

Privacy-conscious AI modelling has also seen recent innovation. The Spyderisk platform introduced by Carmichael et al. (2024) applies ISO 27005-compliant risk assessment in AI-enabled healthcare services, particularly for managing privacy risks in multi-stakeholder environments. Although originally developed for data donation contexts, the methodology is extensible to emergency scenarios where personal health and exposure data must be analysed in real-time whilst preserving confidentiality.

Another promising frontier is mobile and wearable data integration. AI-driven platforms can incorporate real-time biometric and geolocation data from smartphones or fitness trackers to produce hyper-personalised health risk profiles. This functionality addresses a major shortcoming of existing protocols, which typically offer uniform recommendations for heterogeneous populations. For instance, individuals with pre-existing cardiovascular conditions could receive customised alerts when  $PM_{2.5}$  thresholds exceed their personal tolerance, even if the general AQI remains “Moderate”.

Moreover, AI systems can support resource allocation during high-risk periods. Predictive models can forecast surges in emergency department visits or medication demand, allowing healthcare systems to deploy staff and supplies preemptively. As shown in studies like Lu et al. (2022), AI-based simulations enhance disaster preparedness by coupling environmental exposure data with human behavioural models.

However, these opportunities come with challenges. Meskó & Topol (2023) caution that the interpretability, fairness, and safety of AI models must be carefully evaluated before deployment. Whilst governance will be discussed in Section 6, it is worth noting here that the lack of standardised evaluation metrics and interoperability frameworks currently limits widespread adoption of AI in wildfire health risk modelling.

In summary, AI holds transformative potential for precision risk modelling in wildfire contexts. By shifting from population-level averages to individualised, real-time prediction, AI-driven systems promise a paradigm shift in how we monitor, predict, and respond to fire-related health hazards.

## 5.5 Future Directions and Recommendations

The limitations of existing health risk frameworks during wildfire events call for a shift towards predictive, adaptive, and personalised modelling approaches. As demonstrated throughout this section, AI has shown great promise in addressing these challenges. However, translating this potential into operational practice will require targeted innovation and systemic change.

One critical need is the development of standardised frameworks for evaluating AI models in health risk contexts. This includes benchmarks for interpretability, fairness, uncertainty quantification, and real-time responsiveness. Establishing these criteria will support regulatory acceptance and build public trust.

Second, future systems should prioritise integration of diverse data streams—environmental, clinical, biometric, and behavioural—in real time. The inclusion of wearable and mobile sensor data will be essential for generating individualised alerts, especially for vulnerable populations such as those with respiratory or cardiovascular conditions.

Advances in edge computing and drone-mounted sensors offer promising avenues for expanding spatial coverage in remote or under-monitored areas. When coupled with predictive AI models, such technologies could enable hyperlocal, anticipatory risk assessments that adapt to rapidly changing fire dynamics.

Finally, AI systems must be validated in real-world fire events. Only through deployment in operational settings—such as emergency departments, schools, or community health platforms—can these tools be refined and their public health value truly assessed.

## 6 Governance

The technological and methodological advances reviewed in the preceding sections—ranging from autonomous plume tracking to AI-driven health risk forecasting—cannot be safely or equitably deployed without robust governance structures. In high-stakes contexts such as wildfire response and environmental health surveillance, failures of governance can be as harmful as failures of technology. Issues of accountability, transparency, inclusion, and legal responsibility directly shape whether these systems protect vulnerable populations or exacerbate existing inequalities.

This section examines governance as a foundational layer that must co-evolve with spatiotemporal forecasting and AI-enabled public health systems. We adopt a whole-systems perspective, viewing governance not as a static regulatory constraint but as a dynamic process embedded across the entire lifecycle of data collection, modelling, deployment, and intervention. We first review pipeline-wide governance frameworks, then examine the ethical and legal foundations of AI in health contexts. We subsequently analyse the roles of transparency, accountability, trust, and equity, before exploring participatory and adaptive governance models. Finally, we outline unresolved challenges and future directions for governing AI in public and planetary health.

### 6.1 Governance Across the Pipeline

Effective governance of AI systems in public and planetary health must extend across the entire pipeline—spanning data collection, algorithmic modelling, and deployment in real-world interventions. Rather than treating governance as an afterthought, recent scholarship emphasises that it must be designed as an integral part of system development and implementation. Figure 9 illustrates this whole-systems approach, showing governance considerations at each stage of the AI lifecycle.

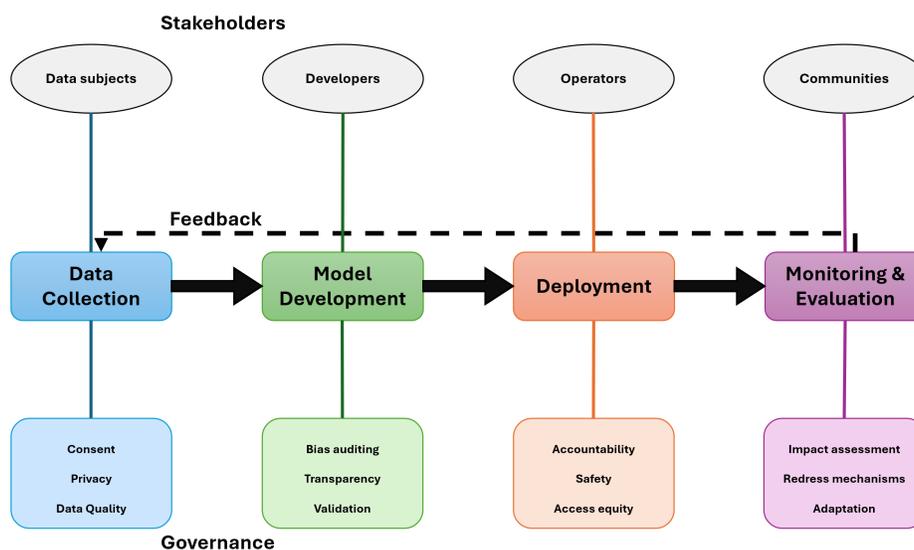


Figure 9: Governance framework across the AI pipeline, showing key considerations and stakeholder involvement at each stage from data collection to operational deployment.

Morley et al. (2022) highlight the importance of a “whole systems approach,” which incorporates governance mechanisms at each stage of the AI lifecycle, from data sourcing and preprocessing to model deployment and monitoring. This approach is echoed by De Almeida et al. (2021), who propose a regulatory framework with continuous feedback loops between technology development and regulatory oversight. Their conceptual

model includes technology-to-regulation (T2R) and regulation-to-technology (R2T) processes, acknowledging that governance must be both adaptive and anticipatory.

These governance challenges are not purely technical. Wagner et al. (2024) argue for embedding public health ethics into governance frameworks, enabling better navigation of trade-offs between individual autonomy and collective benefit. They propose that harm reduction, attention to social determinants of health, and the One Health perspective can help contextualise AI decisions within broader public and environmental systems.

Several studies emphasise the importance of stakeholder participation across the pipeline. Banerjee et al. (2022) and Bazzano et al. (2025) argue that co-designing AI systems with patients, healthcare workers, and community members not only builds trust but also ensures that the systems align with local needs and contexts. Similarly, Zhang & Zhang (2023) identify stakeholder accountability and human oversight as critical safeguards against algorithmic opacity and bias.

However, as Ulnicane et al. (2021) note, AI governance is often framed in policy documents as a balancing act between competing pressures—innovation, safety, public engagement, and market power. Whilst calls for inclusivity and transparency are common, specific mechanisms for implementation are frequently lacking (Wilson, 2022).

In sum, governing spatiotemporal forecasting AI for public health demands a multilayered framework that links regulatory tools, ethical principles, and participatory processes across the full AI pipeline. Each stage—data collection, model design, and intervention—introduces distinct risks and responsibilities that require contextualised, adaptive governance strategies.

## 6.2 Ethical and Legal Foundations

The ethical and legal foundations of AI governance are especially critical in public and planetary health applications, where the consequences of algorithmic decisions can directly impact human lives and wellbeing. Governance in this context must move beyond abstract ethical principles to include enforceable, context-sensitive legal structures and operational safeguards.

The World Health Organisation outlines six key principles for the ethical use of AI in health: autonomy, safety, transparency, accountability, equity, and sustainability (World Health Organization, 2021). These principles offer a normative baseline for AI development, particularly in high-risk environments such as health emergencies or environmental crises. However, operationalising these ideals remains challenging in practice.

Zhang & Zhang (2023) highlight that biased training data, opaque model logic, and unclear responsibility chains remain prevalent in medical AI systems. They argue for greater transparency and stronger mechanisms for responsibility attribution, proposing interventions such as algorithmic audits, standardised reporting protocols, and improved dataset governance. These proposals echo broader calls for traceability and explainability as pillars of trustworthy AI.

From a legal perspective, adaptive and anticipatory governance models are increasingly advocated. De Almeida et al. (2021) introduce the concept of regulatory-to-technology (R2T) and technology-to-regulatory (T2R) feedback loops, suggesting that governance must co-evolve with innovation. Their framework promotes dynamic interaction between regulators, developers, and users, allowing legal norms to adjust alongside technical advances.

Taeihagh (2021) further emphasises the need for regulatory pluralism—where formal laws are complemented by soft governance mechanisms such as standards, guidelines, and institutional norms. In fast-moving fields like AI for public health, where novel risks emerge faster than laws can adapt, such hybrid approaches provide necessary flexibility.

Informed consent is another ethical concern under re-evaluation. Pickering (2021) critiques traditional consent models as too rigid for fast-paced, data-intensive systems. He proposes a trust-based framework of “dynamic consent,” where ongoing negotiation and contextual understanding replace static, one-time disclosures. This model is particularly relevant in settings like wildfire response or pandemic surveillance, where population-level data are collected and acted upon in real time.

Finally, legal and ethical governance must account for global disparities. Joshi et al. (2022) point out that governance models designed for high-income countries often overlook infrastructural, institutional, and cultural barriers in low- and middle-income contexts. Their work calls for flexible, locally adapted governance strategies that uphold global ethical principles whilst responding to local realities.

In summary, ethical and legal foundations are not standalone concerns but must be deeply integrated into the design, deployment, and oversight of spatiotemporal forecasting AI systems. Ensuring fairness, transparency, and accountability requires a combination of normative frameworks, legal innovation, and continuous stakeholder engagement.

### 6.3 Transparency, Accountability, and Trust

Transparency, accountability, and trust form an interdependent triad in the governance of AI systems for public and planetary health. In the context of spatiotemporal forecasting modelling—where real-time decisions can affect large populations—these elements are not optional but foundational to safe and ethical deployment.

Zhang & Zhang (2023) emphasise that the opacity of many medical AI models poses significant challenges for oversight. Without transparent logic or accessible documentation, it becomes difficult to identify where errors occur, how decisions are made, and who should be held responsible. Their review stresses the importance of algorithmic explainability and traceability to mitigate the risks of biased or unsafe outputs. This is particularly relevant in health risk models used during wildfire episodes or pollution crises, where algorithmic misjudgements can have immediate and widespread effects.

The issue of accountability is tightly coupled with transparency. Knowles & Richards (2021) argue that AI systems should be treated as institutional actors that must operate under mechanisms of redress and auditability. They propose that accountability should not rely solely on technical validation, but also on institutional structures that enable responsibility attribution, especially in multi-actor environments such as public health responses.

These concerns are echoed in the WHO’s ethical guidance (World Health Organization, 2021), which calls for governance frameworks that clarify roles and responsibilities, particularly where automated systems influence clinical or emergency decisions. The guidance underscores that without clear accountability, trust in AI systems is likely to erode, regardless of technical performance.

Trust itself is not simply an outcome of transparency and accountability—it is also a prerequisite for effective AI deployment. Robles & Mallinson (2025) show that public trust is shaped by broader concerns over civil liberties, data privacy, and institutional integrity. Their analysis indicates that trust can be actively built through open communication, stakeholder engagement, and demonstrable fairness in system outcomes.

Pickering (2021) offers a complementary view by proposing trust-based consent models in emergency settings. He suggests that when traditional consent frameworks are impractical, systems must instead foster legitimacy through continual negotiation and clear, respectful use of public data. This is particularly applicable to spatiotemporal forecasting systems used for wildfire exposure tracking or infectious disease surveillance, where individual consent may be impossible to obtain in real-time.

Despite widespread agreement on the importance of transparency and trust, Ulnicane et al. (2021) note that concrete governance mechanisms often lag behind rhetorical commitments. They highlight that whilst national strategies frequently reference these values, few offer detailed policies for implementation. Wilson (2022) similarly criticises the tendency to subordinate participatory and transparency ideals to innovation and economic competitiveness.

Together, these studies underscore that ensuring transparency, establishing clear lines of accountability, and building trust are mutually reinforcing goals. Achieving them requires not only technical solutions but also political will, institutional commitment, and culturally attuned governance practices.

## 6.4 Inclusion, Equity, and Capacity Gaps

Ensuring that spatiotemporal forecasting AI systems equitably serve all populations requires governance approaches that explicitly address disparities in exposure, access, and institutional capacity. These concerns are especially acute in public and planetary health contexts, where the burden of environmental hazards and infrastructural limitations disproportionately affects low-resource and marginalised communities.

Section 3 of this review highlights the spatial asymmetries in exposure to wildfire smoke, noting that remote, Indigenous, and economically disadvantaged populations are often both more exposed and less protected by monitoring infrastructure. Governance strategies must therefore confront the unequal distribution of risk and the limitations of existing surveillance systems.

Zhang & Zhang (2023) identify algorithmic bias and poor data quality as key governance risks, warning that AI systems trained on skewed datasets can reinforce existing inequalities. Trehan (2025) extends this critique by highlighting how sustainability efforts in AI, whilst well-intentioned, may unintentionally marginalise vulnerable groups if energy-efficient models are prioritised over accessible or inclusive ones.

The gap between high-income and low- and middle-income countries (LMICs) is particularly pronounced. Joshi et al. (2022) model AI implementation challenges in LMIC health systems and find that barriers such as limited infrastructure, weak governance, and data privacy concerns hinder responsible deployment. Gama et al. (2022) echo these findings, noting that whilst many governance ideals align with existing public health frameworks, AI-specific requirements—such as continuous human oversight and ethical data integration—are often unfulfilled in practice.

Wagner et al. (2024) propose the One Health approach as a means of integrating human, animal, and environmental health to more equitably govern complex, systemic risks. By recognising interdependence across domains, One Health supports governance models that prioritise social determinants and structural vulnerabilities.

Inclusion also requires meaningful participation from underserved communities. Bazzano et al. (2025) argue that co-design and community engagement are not ancillary, but essential, to building AI systems that reflect lived realities. Their work shows that participatory design processes improve system relevance, enhance trust, and support more equitable outcomes.

To close capacity gaps, governance strategies must extend beyond high-level ethical principles to include investment in infrastructure, support for local expertise, and adaptive policy mechanisms that respond to on-the-ground realities. Without such mechanisms, AI risks becoming a tool that deepens rather than alleviates global health inequities.

## 6.5 Participatory and Adaptive Governance

As AI systems increasingly influence public health outcomes, governance models must evolve to reflect not only technical and ethical demands, but also the lived experiences and priorities of affected populations. Participatory and adaptive governance offer a pathway towards legitimacy, responsiveness, and resilience—especially in the complex and uncertain contexts of spatiotemporal forecasting in public health systems.

Ulnicane et al. (2021) argue that AI governance is too often dominated by top-down, state- or industry-led narratives that marginalise societal voices. They call for more inclusive frameworks that recognise the role of communities, healthcare workers, and local institutions in shaping AI outcomes. This is especially relevant in systems that collect real-time, geo-located health data, where the stakes of surveillance, consent, and intervention are deeply personal.

Banerjee et al. (2022) and Bazzano et al. (2025) provide concrete examples of participatory design in health-care AI. Their work shows that involving patients and community members early in the design process improves usability, contextual fit, and trust. Participatory governance is not limited to feedback collection; it includes co-creation, shared decision-making, and long-term engagement.

Corbett et al. (2023) emphasise the importance of moving beyond tokenistic consultation. Drawing on Arnstein's Ladder of Citizen Participation, they show that many so-called participatory processes offer

limited power to stakeholders. Instead, they advocate for models that embed community agency into the design, deployment, and monitoring of AI systems.

Adaptive governance complements participation by acknowledging that health emergencies, environmental conditions, and social needs evolve rapidly. Hassan et al. (2024) describe governance systems that are intentionally iterative—capable of incorporating feedback, revising models, and adjusting policy in response to new data. Their PACE (Participatory AI for Civic Engagement) framework integrates dynamic consent, community input, and real-time governance loops to improve system responsiveness and ethical alignment.

Viewing AI systems as institutional actors also has implications for governance. Knowles & Richards (2021) argue that as AI systems become embedded in public decision-making, they must be held to standards akin to those applied to public institutions—transparency, accountability, and legitimacy through representation.

Ultimately, participatory and adaptive governance reframes AI not as a fixed technology, but as a living process shaped by collective values, contested interests, and evolving challenges. For spatiotemporal forecasting systems in health and disaster response, such governance models are essential to balancing innovation with social responsibility.

## 6.6 Governance Challenges and Forward Directions

Despite growing attention to AI governance in public and planetary health, substantial challenges remain in operationalising ethical principles, ensuring meaningful participation, and achieving systemic resilience. As spatiotemporal forecasting AI systems become more integral to emergency response, risk prediction, and health surveillance, the urgency to confront unresolved tensions intensifies.

One persistent challenge is the tension between responsiveness and privacy. Real-time data collection, especially from mobile or sensor-based platforms, enables more accurate risk forecasting but raises concerns about surveillance, consent, and data security. As Zhang & Zhang (2023) and Pickering (2021) both note, traditional consent models are poorly suited to fast-moving crises, and trust-based alternatives require robust institutional infrastructure to function meaningfully.

Another key issue is the fragmentation of regulatory frameworks across jurisdictions. Joshi et al. (2022) and Morley et al. (2022) highlight how uneven legal systems, resource constraints, and institutional silos create implementation gaps, particularly in low- and middle-income countries. Without international coordination, efforts to standardise AI safety, fairness, and accountability remain limited in scope and impact.

Moreover, the increasing complexity of AI systems—especially those based on deep learning—poses challenges for traceability and redress. Even when harm can be observed, it is often difficult to attribute responsibility across model developers, data providers, and deploying agencies. WHO guidance (World Health Organization, 2021) and Knowles & Richards (2021) both call for clearer documentation and oversight structures, but questions remain about enforcement and jurisdiction, especially in cross-border or planetary health applications.

Looking forward, a hybrid governance model appears most promising. De Almeida et al. (2021) propose the integration of soft law principles—like ethical guidelines and standards—with formal regulatory instruments. Their model of bidirectional feedback between regulators and developers (R2T and T2R) allows governance systems to adapt as technologies evolve. This adaptability is essential for long-term sustainability.

Additionally, whole-systems approaches are gaining traction. Morley et al. (2022) advocate for governance that spans the entire AI lifecycle—from data sourcing and modelling to deployment and impact evaluation. This approach must also be multilayered: incorporating local knowledge, national oversight, and planetary coordination to manage systemic risk in a globally interconnected health landscape.

Finally, governance must cultivate legitimacy over time. As Robles & Mallinson (2025) argue, trust is not a one-off achievement but a product of consistent, transparent, and inclusive practice. Ensuring the resilience of governance systems requires sustained investment in public engagement, technical capacity, and normative alignment.

In sum, future governance of AI for public and planetary health must embrace complexity, enable feedback, and remain attuned to power dynamics and equity. Only through such adaptive, inclusive, and system-wide oversight can AI be steered towards just and sustainable outcomes.

## 7 Conclusion

Figure 10 presents the integrative framework emerging from this review, illustrating how autonomous sensing, predictive modelling, health risk assessment, and governance form interconnected components of a comprehensive fire pollution management system.

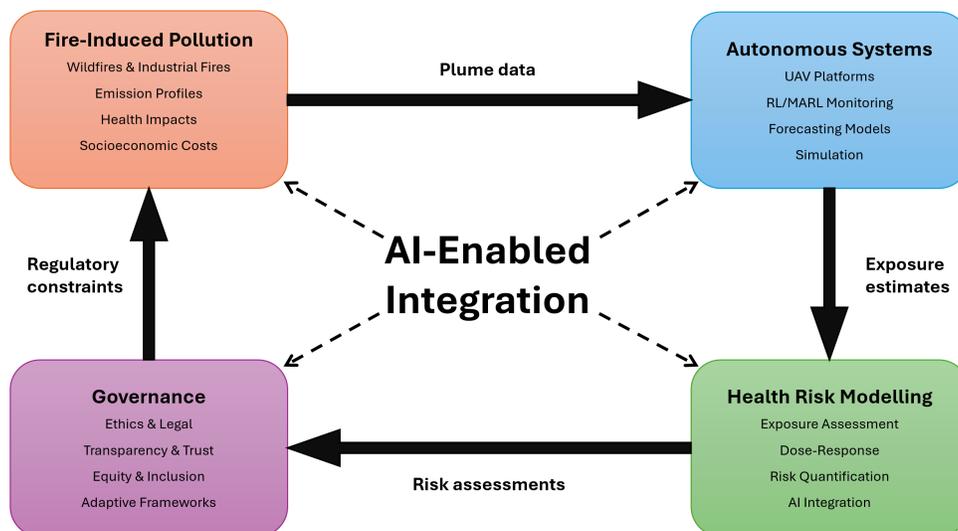


Figure 10: Integrative framework for AI-enabled fire pollution management, showing the interconnections between environmental monitoring, autonomous systems, health risk modelling, and governance.

Fire-induced air pollution is a growing and multifaceted threat, shaped by the increasing frequency of wildfires, industrial hazards, and urban sprawl under a changing climate. Whilst the health and environmental consequences of such events are well documented, the systems we rely on to detect, model, and respond to airborne pollutants remain fragmented, reactive, and technologically underpowered. This review has explored how emerging tools in artificial intelligence—particularly those enabling autonomous sensing, predictive modelling, and intelligent coordination—can reshape our capacity to address fire-generated air pollution in real time.

Across the four thematic domains reviewed, a common limitation persists: current systems often operate in isolation. Sensing technologies lack precise temporal and spatial forecasting capabilities; health risk models are decoupled from real-time data streams; governance frameworks struggle to keep pace with rapid technological change. At the same time, advances in AI—such as multi-agent reinforcement learning for plume tracking, UAV-based sensor networks, and personalised exposure risk modelling—demonstrate strong potential to bridge these gaps.

The review highlights a critical opportunity: AI can act not merely as a set of tools, but as an integrative architecture connecting environmental monitoring, public health forecasting, and emergency response. However, realising this potential requires more than technical innovation. It demands coordinated development across disciplines, equitable data practices, and proactive governance structures that support the deployment of intelligent systems in high-risk environments.

Future research should prioritise end-to-end integration of AI-enabled systems, from sensor-level intelligence to actionable health interventions. In parallel, policy frameworks must evolve to address data interoperability, accountability, and community engagement in the deployment of AI for environmental health. As fires

grow in frequency and scale, the need for intelligent, real-time, and socially responsible solutions becomes increasingly urgent. This review underscores the potential—and necessity—of AI as a unifying force in this transformation.

## References

- Ahmed Shihab Albahri, Ali M. Duhaim, Mohammed A. Fadhel, Alhamzah Alnoor, Noor S. Baqer, Laith Alzubaidi, Osamah Shihab Albahri, Abdullah Hussein Alamoodi, Jinshuai Bai, Asma Salhi, et al. A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion. *Information Fusion*, 96:156–191, 2023.
- Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press, 2024. URL <https://www.mar1-book.com>.
- Razan Y. Aldahlawi, Vahid Akbari, and Glyn Lawson. A systematic review of methodologies for human behavior modelling and routing optimization in large-scale evacuation planning. *International Journal of Disaster Risk Reduction*, pp. 104638, 2024.
- Badr H. Alharbi, Mohammad J. Pasha, and Mohammed Ahmad S. Al-Shamsi. Firefighter exposures to organic and inorganic gas emissions in emergency residential and industrial fires. *Science of the Total Environment*, 770:145332, 2021.
- Yuuichi Asahi, Naoyuki Onodera, Yuta Hasegawa, Takashi Shimokawabe, Hayato Shiba, and Yasuhiro Idomura. Citytransformer: A transformer-based model for contaminant dispersion prediction in a realistic urban area. *Boundary-Layer Meteorology*, 186(3):659–692, 2023.
- Mohamed Sami Assenine, Walid Bechkit, Ichrak Mokhtari, Hervé Rivano, and Karima Benatchba. Cooperative deep reinforcement learning for dynamic pollution plume monitoring using a drone fleet. *IEEE Internet of Things Journal*, 11(5):7325–7338, 2023.
- Roberta Baggio, Jean Baptiste Filippi, Benjamin Truchot, and Flavio T. Couto. Local to continental scale coupled fire-atmosphere simulation of large industrial fire plume. *Fire Safety Journal*, 134:103699, 2022.
- Soumya Banerjee, Phil Alsop, Linda Jones, and Rudolf N. Cardinal. Patient and public involvement to build trust in artificial intelligence: A framework, tools, and case studies. *Patterns*, 3(6), 2022.
- Janet F. Barlow. Progress in observing and modelling the urban boundary layer. *Urban Climate*, 10:216–240, 2014.
- Gourav Bathla, Kishor Bhadane, Rahul Kumar Singh, Rajneesh Kumar, Rajanikanth Aluvalu, Rajalakshmi Krishnamurthi, Adarsh Kumar, R. N. Thakur, and Shakila Basheer. Autonomous vehicles and intelligent automation: Applications, challenges, and opportunities. *Mobile Information Systems*, 2022(1):7632892, 2022.
- Alessandra N. Bazzano, Andrea Mantsios, Nicholas Mattei, Michael R. Kosorok, and Aron Culotta. Ai can be a powerful social innovation for public health if community engagement is at the core. *Journal of Medical Internet Research*, 27:e68198, 2025.
- Marc Besson, Jamie Alison, Kim Bjerge, Thomas E. Gorochowski, Toke T. Høye, Tommaso Jucker, Hjalte M. R. Mann, and Christopher F. Clements. Towards the fully automated monitoring of ecological communities. *Ecology Letters*, 25(12):2753–2775, 2022.
- Jan Stefan Białowicz, Wioletta Rogula-Kozłowska, and Adam Krasuski. Contribution of landfill fires to air pollution—an assessment methodology. *Waste Management*, 125:182–191, 2021.
- Carolyn Black, Yohannes Tesfaigzi, Jed A Bassein, and Lisa A Miller. Wildfire smoke exposure and human health: Significant gaps in research for a growing public health issue. *Environmental toxicology and pharmacology*, 55:186–195, 2017.

- Katie Boaggio, Stephen D. LeDuc, R. Byron Rice, Parker F. Duffney, Kristen M. Foley, Amara L. Holder, Stephen McDow, and Christopher P. Weaver. Beyond particulate matter mass: heightened levels of lead and other pollutants associated with destructive fire events in california. *Environmental Science & Technology*, 56(20):14272–14283, 2022.
- Carli P. Brucker, Ben Livneh, J. Toby Minear, and Fernando L. Rosario-Ortiz. A review of simulation experiment techniques used to analyze wildfire effects on water quality and supply. *Environmental Science: Processes & Impacts*, 24(8):1110–1132, 2022.
- Marshall Burke, Anne Driscoll, Sam Heft-Neal, Jiani Xue, Jennifer Burney, and Michael Wara. The changing risk and burden of wildfire in the united states. *Proceedings of the National Academy of Sciences*, 118(2): e2011048118, 2021.
- Lorenzo Canese, Gian Carlo Cardarilli, Luca Di Nunzio, Rocco Fazzolari, Daniele Giardino, Marco Re, and Sergio Spanò. Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences*, 11(11):4948, 2021.
- Laura Carmichael, Steve Taylor, Adriane Chapman, and Michael Boniface. Ai in health and social care: A methodology for privacy risk modeling and simulation. In *Companion Proceedings of the ACM Web Conference 2024*, pp. 1150–1153, 2024.
- Nipuna Chamara, Md Didarul Islam, Geng Frank Bai, Yeyin Shi, and Yufeng Ge. Ag-iot for crop and environment monitoring: Past, present, and future. *Agricultural Systems*, 203:103497, 2022.
- Gongbo Chen, Yuming Guo, Xu Yue, Shilu Tong, Antonio Gasparrini, Michelle L. Bell, Ben Armstrong, Joel Schwartz, Jouni J. K. Jaakkola, Antonella Zanobetti, et al. Mortality risk attributable to wildfire-related pm2.5 pollution: a global time series study in 749 locations. *The Lancet Planetary Health*, 5(9):e579–e587, 2021a.
- Hao Chen, James M. Samet, Philip A. Bromberg, and Haiyan Tong. Cardiovascular health impacts of wildfire smoke exposure. *Particle and Fibre Toxicology*, 18(1):2, 2021b.
- Yen-Chen Chen, Jia-Lin Wang, Chih-Yuan Chang, Ming-Tung Chuang, Charles C.-K. Chou, Xiang-Xu Pan, Yu-Jui Ho, Chang-Feng Ou-Yang, Wen-Tzu Liu, and Chih-Chung Chang. Using drone soundings to study the impacts and compositions of plumes from a gigantic coal-fired power plant. *Science of The Total Environment*, 893:164709, 2023.
- Dominic Clements, Matthew Coburn, Simon Cox, Florentin M. J. Bulot, Zheng-Tong Xie, and Christina Vanderwel. Comparing large-eddy simulation and gaussian plume model to sensor measurements of an urban smoke plume. *Atmosphere*, 15(9), 2024.
- Eric Corbett, Remi Denton, and Sheena Erete. Power and public participation in ai. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pp. 1–13, 2023.
- Calum X. Cunningham, Grant J. Williamson, and David M. J. S. Bowman. Increasing frequency and intensity of the most extreme wildfires on earth. *Nature Ecology & Evolution*, 8(8):1420–1425, 2024.
- Abdelaziz Darwiesh, A. H. El-Baz, Abedallah Zaid Abualkishik, and Mohamed Elhoseny. Artificial intelligence model for risk management in healthcare institutions: towards sustainable development. *Sustainability*, 15(1):420, 2022.
- Stevanus Darwin and Tua A. Tamba. Multi-agent reinforcement learning with information sharing for optimal drone mapping. In *2023 International Conference on Computer, Control, Informatics and its Applications (IC3INA)*, pp. 66–71. IEEE, 2023.
- Patricia Gomes Rêgo De Almeida, Carlos Denner dos Santos, and Josivania Silva Farias. Artificial intelligence regulation: a framework for governance. *Ethics and Information Technology*, 23(3):505–525, 2021.

- Anne A. H. de Hond, Artuur M. Leeuwenberg, Lotty Hooft, Ilse M. J. Kant, Steven W. J. Nijman, Hendrikus J. A. van Os, Jiska J. Aardoom, Thomas P. A. Debray, Ewoud Schuit, Maarten van Smeden, et al. Guidelines and quality criteria for artificial intelligence-based prediction models in healthcare: a scoping review. *NPJ Digital Medicine*, 5(1):2, 2022.
- Michael E. Deary and Simon D. Griffiths. A novel approach to the development of 1-hour threshold concentrations for exposure to particulate matter during episodic air pollution events. *Journal of Hazardous Materials*, 418:126334, 2021.
- Michael E. Deary and Simon D. Griffiths. The impact of air pollution from industrial fires in urban settings: Monitoring, modelling, health, and environmental justice perspectives. *Environments*, 11(7), 2024. ISSN 2076-3298.
- Kristie L. Ebi, Jennifer Vanos, Jane W. Baldwin, Jesse E. Bell, David M. Hondula, Nicole A. Errett, Katie Hayes, Colleen E. Reid, Shubhayu Saha, June Spector, et al. Extreme weather and climate change: population health and health system implications. *Annual Review of Public Health*, 42(1):293–315, 2021.
- Natalia V. Efimova and Viktor S. Rukavishnikov. Assessment of smoke pollution caused by wildfires in the baikal region (russia). *Atmosphere*, 12(12), 2021. ISSN 2073-4433. doi: 10.3390/atmos12121542.
- Ming Fan, Hongsheng Wang, Jing Zhang, Seyyed A. Hosseini, and Dan Lu. Advancing spatiotemporal forecasts of co2 plume migration using deep learning networks with transfer learning and interpretation analysis. *International Journal of Greenhouse Gas Control*, 132:104061, 2024.
- Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Magnus Gålfalk, Sören Nilsson Påledal, and David Bastviken. Sensitive drone mapping of methane emissions without the need for supplementary ground-based measurements. *ACS Earth and Space Chemistry*, 5(10): 2668–2676, 2021.
- Fábio Gama, Daniel Tyskbo, Jens Nygren, James Barlow, Julie Reed, and Petra Svedberg. Implementation frameworks for artificial intelligence translation into health care practice: scoping review. *Journal of Medical Internet Research*, 24(1):e32215, 2022.
- Yuan Gao, Wenzhong Huang, Rongbin Xu, Danijela Gasevic, Yanming Liu, Wenhua Yu, Pei Yu, Xu Yue, Guowei Zhou, Yan Zhang, et al. Association between long-term exposure to wildfire-related pm2.5 and mortality: A longitudinal analysis of the uk biobank. *Journal of Hazardous Materials*, 457:131779, 2023.
- Yuan Gao, Wenzhong Huang, Zhihu Xu, Rongbin Xu, Danijela Gasevic, Yanming Liu, Xu Yue, Guowei Zhou, Yan Zhang, Jiangning Song, et al. Wildfire-related pm2.5 and cause-specific cancer mortality. *Ecotoxicology and Environmental Safety*, 285:117023, 2024.
- Maryam Ghodrat, Farshad Shakeriaski, Sayyed Aboozar Fanaee, and Albert Simeoni. Software-based simulations of wildfire spread and wind-fire interaction. *Fire*, 6(1):12, 2022.
- A. M. Graham, R. J. Pope, K. P. Pringle, S. Arnold, M. P. Chipperfield, L. A. Conibear, E. W. Butt, L. Kiely, Christoph Knote, and J. B. McQuaid. Impact on air quality and health due to the saddleworth moor fire in northern england. *Environmental Research Letters*, 15(7):074018, 2020.
- Simon D. Griffiths, Philip Chappell, Jane A. Entwistle, Frank J. Kelly, and Michael E. Deary. A study of particulate emissions during 23 major industrial fires: Implications for human health. *Environment International*, 112:310–323, 2018.
- Simon D. Griffiths, Jane A. Entwistle, Frank J. Kelly, and Michael E. Deary. Characterising the ground level concentrations of harmful organic and inorganic substances released during major industrial fires, and implications for human health. *Environment International*, 162:107152, 2022. ISSN 0160-4120.

- Simon D. Griffiths, H. M. King, J. Wilkinson, Frank J. Kelly, Jane A. Entwistle, and Michael E. Deary. Evaluating public exposure to airborne particulates from major incident fires: A back trajectory plume modelling approach. *Journal of Hazardous Materials*, 490, 2025.
- Ihsane Gryech, Chaimae Asaad, Mounir Ghogho, and Abdellatif Kobbane. Applications of machine learning & internet of things for outdoor air pollution monitoring and prediction: A systematic literature review. *Engineering Applications of Artificial Intelligence*, 137:109182, 2024.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pp. 1861–1870. PMLR, 2018.
- Gayle S. W. Hagler, Sarah B. Henderson, Sarah McCaffrey, Fay H. Johnston, Susan Stone, Ana Rappold, and Wayne E. Cascio. Understanding and communicating wildland fire smoke risk. *Frontiers in Public Health*, 9:721823, 2021.
- Saad Hassan, Syeda Mah Noor Asad, Motahhare Eslami, Nicholas Mattei, Aron Culotta, and John Zimmerman. Pace: Participatory ai for community engagement. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 12, pp. 151–154, 2024.
- Barbara Hoffmann, Hanna Boogaard, Audrey de Nazelle, Zorana J. Andersen, Michael Abramson, Michael Brauer, Bert Brunekreef, Francesco Forastiere, Wei Huang, Haidong Kan, et al. Who air quality guidelines 2021—aiming for healthier air for all: a joint statement by medical, public health, scientific societies and patient representative organisations. *International Journal of Public Health*, 66:1604465, 2021.
- Stephanie M. Holm, Mark D. Miller, and John R. Balmes. Health effects of wildfire smoke in children and public health tools: a narrative review. *Journal of Exposure Science & Environmental Epidemiology*, 31(1):1–20, 2021.
- Yukai Hou, Jin Zhao, Rongqing Zhang, Xiang Cheng, and Liuqing Yang. Uav swarm cooperative target search: A multi-agent reinforcement learning approach. *IEEE Transactions on Intelligent Vehicles*, 9(1): 568–578, 2023.
- Anass Houdou, Imad El Badisy, Kenza Khomsi, Sammila Andrade Abdala, Fayez Abdulla, Houda Najmi, Majdouline Obtel, Lahcen Belyamani, Azeddine Ibrahim, and Mohamed Khalis. Interpretable machine learning approaches for forecasting and predicting air pollution: A systematic review. *Aerosol and Air Quality Research*, 24(1):230151, 2024.
- Junyan Hu, Hanlin Niu, Joaquin Carrasco, Barry Lennox, and Farshad Arvin. Fault-tolerant cooperative navigation of networked uav swarms for forest fire monitoring. *Aerospace Science and Technology*, 123: 107494, 2022.
- Fintan Hughes, Luke Parsons, Jerrold H. Levy, Drew Shindell, Brooke Alhanti, Tetsu Ohnuma, Prasad Kasibhatla, Hugh Montgomery, and Vijay Krishnamoorthy. Impact of wildfire smoke on acute illness. *Anesthesiology*, 141(4):779–789, 2024.
- Henry Alexander Ignatious, Manzoor Khan, et al. An overview of sensors in autonomous vehicles. *Procedia Computer Science*, 198:736–741, 2022.
- Sukanya Jaiswal, Isabelle Jalbert, Katrina Schmid, Natasha Tein, Sarah Wang, and Blanka Golebiowski. Smoke and the eyes: a review of the harmful effects of wildfire smoke and air pollution on the ocular surface. *Environmental Pollution*, 309:119732, 2022.
- Zhang Jiandong, Yang Qiming, Shi Guoqing, Lu Yi, and Wu Yong. Uav cooperative air combat maneuver decision based on multi-agent reinforcement learning. *Journal of Systems Engineering and Electronics*, 32(6):1421–1438, 2021.
- Aneesh John, Krishna Podila, Qi Chen, and Yanfei Rao. Application of high-fidelity modelling approach to predict smoke and fire propagation in a nuclear fire scenario. *Thermal Science and Engineering Progress*, 23:100903, 2021.

- Fay H Johnston, Sarah B Henderson, Yang Chen, James T Randerson, Miriam Marlier, Ruth S DeFries, Patrick Kinney, David MJS Bowman, and Michael Brauer. Estimated global mortality attributable to smoke from landscape fires. *Environmental health perspectives*, 120(5):695, 2012.
- Justyna Jońca, Marcin Pawnuik, Yaroslav Bezyk, Adalbert Arsen, and Izabela Sówka. Drone-assisted monitoring of atmospheric pollution—a comprehensive review. *Sustainability*, 14(18):11516, 2022.
- Aditya V. Jonnalagadda and Hashim A. Hashim. Segnet: A segmented deep learning based convolutional neural network approach for drones wildfire detection. *Remote Sensing Applications: Society and Environment*, 34:101181, 2024.
- Sudhanshu Joshi, Manu Sharma, Rashmi Prava Das, Joanna Rosak-Szyrocka, Justyna Żywiołek, Kamalakanta Muduli, and Mukesh Prasad. Modeling conceptual framework for implementing barriers of ai in public healthcare for improving operational excellence: experiences from developing countries. *Sustainability*, 14(18):11698, 2022.
- Soyi Jung, Won Joon Yun, MyungJae Shin, Joongheon Kim, and Jae-Hyun Kim. Orchestrated scheduling and multi-agent deep reinforcement learning for cloud-assisted multi-uav charging systems. *IEEE Transactions on Vehicular Technology*, 70(6):5362–5377, 2021.
- Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- Ziming Ke, Yuhang Wang, Yufei Zou, Yongjia Song, and Yongqiang Liu. Global wildfire plume-rise data set and parameterizations for climate model applications. *Journal of Geophysical Research: Atmospheres*, 126(6):e2020JD033085, 2021.
- Douglas I. Kelley, Camilla Mathison, Chantelle Burton, Megan Brown, Andrew Sullivan, Elaine Baker, and Tiina Kurvits. Likely future (s) of global wildfires. In *EGU General Assembly Conference Abstracts*, pp. EGU22–6512, 2022.
- Madhav Khirwar and Ankur Narang. Geovit: versatile vision transformer architecture for geospatial image analysis. In *2024 International Conference on Machine Intelligence for GeoAnalytics and Remote Sensing (MIGARS)*, pp. 1–3. IEEE, 2024.
- Bran Knowles and John T. Richards. The sanction of authority: Promoting public trust in ai. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 262–271, 2021.
- Claudio A. S. Lelis, Julio J. Roncal, Leonardo Silveira, Roberto Douglas G. De Aquino, Cesar A. C. Marcondes, Johnny Marques, Denis S. Loubach, Filipe A. N. Verri, Vitor V. Curtis, and Diego G. De Souza. Drone-based ai system for wildfire monitoring and risk prediction. *IEEE Access*, 2024.
- Jingqian Li, Jihong Song, Yine Xu, Qi Yu, Yan Zhang, and Weichun Ma. Parameterization of a rising smoke plume for a large moving ship based on cfd. *Atmosphere*, 13(9):1507, 2022.
- Junjie Li, Zonghao Xie, Kang Liu, Jihao Shi, Tao Wang, Yuanjiang Chang, and Guoming Chen. Real time hydrogen plume spatiotemporal evolution forecasting by using deep probabilistic spatial-temporal neural network. *International Journal of Hydrogen Energy*, 72:878–891, 2024.
- Yuxuan Liang, Yutong Xia, Songyu Ke, Yiwei Wang, Qingsong Wen, Junbo Zhang, Yu Zheng, and Roger Zimmermann. Airformer: Predicting nationwide air quality in china with transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 14329–14337, 2023.
- Da Liu, Liqian Dou, Ruilong Zhang, Xiuyun Zhang, and Qun Zong. Multi-agent reinforcement learning-based coordinated dynamic task allocation for heterogenous uavs. *IEEE Transactions on Vehicular Technology*, 72(4):4372–4383, 2022.
- Hexiang Liu, Qilong Han, Hui Sun, Jingyu Sheng, and Ziyu Yang. Spatiotemporal adaptive attention graph convolution network for city-level air quality prediction. *Scientific Reports*, 13(1):13335, 2023.

- Yingxiang Liu, Zhen Qin, Fangning Zheng, and Behnam Jafarpour. Spatio-temporal neural networks for monitoring and prediction of co2 plume migration from measurable field data. *Journal of Cleaner Production*, 481:144080, 2024.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf).
- Sally Lu, Gordon A. Christie, Thanh T. Nguyen, Jeffrey D. Freeman, and Edbert B. Hsu. Applications of artificial intelligence and machine learning in disasters and public health emergencies. *Disaster Medicine and Public Health Preparedness*, 16(4):1674–1681, 2022.
- Randall V Martin. Satellite remote sensing of surface air quality. *Atmospheric environment*, 42(34):7823–7843, 2008.
- Rafaela Mateus, Armando Pinto, and José M. C. Pereira. Dynamics of thermal plumes for large spaces: a comparative study of in-situ smoke test and a cfd model. *Energy and Buildings*, 319:114512, 2024.
- Bertalan Meskó and Eric J. Topol. The imperative for regulatory oversight of large language models (or generative ai) in healthcare. *NPJ Digital Medicine*, 6(1):120, 2023.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Nadya Moisseeva and Roland Stull. Wildfire smoke-plume rise: a simple energy balance parameterization. *Atmospheric Chemistry and Physics*, 21(3):1407–1425, 2021.
- David Molitor, Jamie T. Mullins, and Corey White. Air pollution and suicide in rural and urban america: Evidence from wildfire smoke. *Proceedings of the National Academy of Sciences*, 120(38):e2221621120, 2023.
- Linn E. Moore, Andre Oliveira, Raymond Zhang, Laleh Behjat, and Anne Hicks. Impacts of wildfire smoke and air pollution on a pediatric population with asthma: a population-based study. *International Journal of Environmental Research and Public Health*, 20(3):1937, 2023.
- Jessica Morley, Lisa Murphy, Abhishek Mishra, Indra Joshi, and Kassandra Karpathakis. Governing data and artificial intelligence for health care: developing an international understanding. *JMIR Formative Research*, 6(1):e31623, 2022.
- Hafiz Suliman Munawar, Mohammad Mojtahedi, Ahmed W. A. Hammad, Michael J. Ostwald, and S. Travis Waller. An ai/ml-based strategy for disaster response and evacuation of victims in aged care facilities in the hawkesbury-nepean valley: A perspective. *Buildings*, 12(1):80, 2022.
- Hasan Raja Naqvi, Guneet Mutreja, Adnan Shakeel, Karan Singh, Kumail Abbas, Darakhsha Fatma Naqvi, Anis Ahmad Chaudhary, Masood Ahsan Siddiqui, Alok Sagar Gautam, Sneha Gautam, et al. Wildfire-induced pollution and its short-term impact on covid-19 cases and mortality in california. *Gondwana Research*, 114:30–39, 2023.
- National Fire Protection Association. Fires by occupancy or property type. [Online]. Available: <https://www.nfpa.org/education-and-research/research/nfpa-research/fire-statistical-reports/fires-by-occupancy-or-property-type>, 2023. [Accessed: Apr. 11, 2025].
- Shuai Niu, Qing Yin, Jing Ma, Yunya Song, Yida Xu, Liang Bai, Wei Pan, and Xian Yang. Enhancing healthcare decision support through explainable ai models for risk prediction. *Decision Support Systems*, 181:114228, 2024.

- James Orr and Ayan Dutta. Multi-agent deep reinforcement learning for multi-robot applications: A survey. *Sensors*, 23(7):3625, 2023.
- Darsh Parekh, Nishi Poddar, Aakash Rajpurkar, Manisha Chahal, Neeraj Kumar, Gyanendra Prasad Joshi, and Woong Cho. A review on autonomous vehicles: Progress, methods and challenges. *Electronics*, 11(14):2162, 2022.
- Billy Peralta, Tomás Sepúlveda, Orietta Nicolis, and Luis Caro. Space-time prediction of pm2.5 concentrations in santiago de chile using lstm networks. *Applied Sciences*, 12(22):11317, 2022.
- Dimitris Perikleous, George Koustas, Spyros Velanas, Katerina Margariti, Pantelis Velanas, and Diego Gonzalez-Aguilera. A novel drone design based on a reconfigurable unmanned aerial vehicle for wild-fire management. *Drones*, 8(5):203, 2024.
- Brian Pickering. Trust, but verify: informed consent, ai technologies, and public health emergencies. *Future Internet*, 13(5):132, 2021.
- Alexandra Popescu, Allison Paulson, Amy C. Christianson, Andrew S. Sullivan, A. Tulloch, Bibiana Bilbao, Camilla Mathison, Catherine Robinson, David Ganz, David Nangoma, et al. Spreading like wildfire: The rising threat of extraordinary landscape fires-a rapid response assessment. *Spreading like wildfire: The rising threat of extraordinary landscape fires-A rapid response assessment*, 2022.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Monotonic value function factorisation for deep multi-agent reinforcement learning. *Journal of Machine Learning Research*, 21(178):1–51, 2020. URL <http://jmlr.org/papers/v21/20-081.html>.
- Colleen E Reid, Michael Brauer, Fay H Johnston, Michael Jerrett, John R Balmes, and Catherine T Elliott. Critical review of health impacts of wildfire smoke exposure. *Environmental health perspectives*, 124(9):1334, 2016.
- Pedro Robles and Daniel J. Mallinson. Artificial intelligence technology, public trust, and effective governance. *Review of Policy Research*, 42(1):11–28, 2025.
- Bertrand Rouet-Leduc and Claudia Hulbert. Automatic detection of methane emissions in multispectral satellite imagery using a vision transformer. *Nature Communications*, 15(1):3801, 2024.
- Michele Salis, Bachisio Arca, Liliana Del Giudice, Palaiologos Palaiologou, Fermin Alcasena-Urdiroz, Alan Ager, Michele Fiori, Grazia Pellizzaro, Carla Scarpa, Matilde Schirru, et al. Application of simulation modeling for wildfire exposure and transmission assessment in sardinia, italy. *International Journal of Disaster Risk Reduction*, 58:102189, 2021.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Mike Schuster and Kuldip K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.
- Carl Seiber, David Nowlin, Bob Landowski, and Matthew E. Tolentino. Tracking hazardous aerial plumes using iot-enabled drone swarms. In *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, pp. 377–382. IEEE, 2018.
- Esmaeil Seraj, Andrew Silva, and Matthew Gombolay. Multi-uav planning for cooperative wildfire coverage and tracking with quality-of-service guarantees. *Autonomous Agents and Multi-Agent Systems*, 36(2):39, 2022.
- Jihao Shi, Junjie Li, Asif Sohail Usmani, Yuan Zhu, Guoming Chen, and Dongdong Yang. Probabilistic real-time deep-water natural gas hydrate dispersion modeling by using a novel hybrid deep learning approach. *Energy*, 219:119572, 2021.

- Jihao Shi, Weikang Xie, Junjie Li, Xinqi Zhang, Xinyan Huang, Asif Sohail Usmani, Faisal Khan, and Guoming Chen. Real-time plume tracking using transfer learning approach. *Computers & Chemical Engineering*, 172:108172, 2023.
- MyungJae Shin, Dae-Hyun Choi, and Joongheon Kim. Cooperative management for pv/ess-enabled electric vehicle charging stations: A multiagent deep reinforcement learning approach. *IEEE Transactions on Industrial Informatics*, 16(5):3493–3503, 2019.
- Harikesh Singh, Li-Minn Ang, Dipak Paudyal, Mauricio Acuna, Prashant Kumar Srivastava, and Sanjeev Kumar Srivastava. A comprehensive review of empirical and dynamic wildfire simulators and machine learning techniques used for the prediction of wildfire in australia. *Technology, Knowledge and Learning*, pp. 1–34, 2025.
- Michael G. Soskind, Nathan P. Li, Daniel P. Moore, Yifeng Chen, Lars P. Wendt, James McSpiritt, Mark A. Zondlo, and Gerard Wysocki. Stationary and drone-assisted methane plume localization with dispersion spectroscopy. *Remote Sensing of Environment*, 289:113513, 2023.
- Ariel F Stein, Roland R Draxler, Glenn D Rolph, Barbara JB Stunder, Mark D Cohen, and Fong Ngan. Noaa’s hysplit atmospheric transport and dispersion modeling system. *Bulletin of the American Meteorological Society*, 96(12):2059–2077, 2015.
- Marlin P. Strub and Jonathan D. Gammell. Adaptively informed trees (ait\*): Fast asymptotically optimal path planning through adaptive heuristics. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3191–3198. IEEE, 2020.
- Kai Su and Feng Qian. Multi-uav cooperative searching and tracking for moving targets based on multi-agent reinforcement learning. *Applied Sciences*, 13(21):11905, 2023.
- Jingjing Sun, Wenwen Qi, Yuandong Huang, Chong Xu, and Wentao Yang. Facing the wildfire spread risk challenge: where are we now and where are we going? *Fire*, 6(6):228, 2023.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning, second edition: An Introduction*. MIT Press, nov 2018.
- Araz Taeiagh. Governance of artificial intelligence. *Policy and Society*, 40(2):137–157, 2021.
- Bianca Tenti and Enrico Ferrero. Evaluation of turbulence depending drag coefficient in plume rise model for fire smoke dispersion. *Atmospheric Environment*, 323:120411, 2024.
- Abhishek Trehan. Bias in green ai addressing disparities in data and algorithms. In *Advancing Social Equity Through Accessible Green Innovation*, pp. 63–76. IGI Global Scientific Publishing, 2025.
- Georgios Tzoumas, Lenka Pitonakova, Lucio Salinas, Charles Scales, Thomas Richardson, and Sabine Hauert. Wildfire detection in large-scale environments using force-based control for swarms of uavs. *Swarm Intelligence*, 17(1):89–115, 2023.
- Inga Ulnicane, William Knight, Tonii Leach, Bernd Carsten Stahl, and Winter-Gladys Wanjiku. Framing governance for a contested emerging technology: insights from ai policy. *Policy and Society*, 40(2):158–177, 2021.
- Mario Miguel Valero, Lluís Jofre, and Ricardo Torres. Multifidelity prediction in wildfire spread simulation: Modeling, uncertainty quantification and sensitivity analysis. *Environmental Modelling & Software*, 141:105050, 2021.
- Jorge Vargas, Suleiman Alsweiss, Onur Toker, Rahul Razdan, and Joshua Santos. An overview of autonomous vehicles sensors and their vulnerability to weather conditions. *Sensors*, 21(16):5397, 2021.
- R. Vautard, P. Ciais, R. Fisher, D. Lowry, F. M. Bréon, F. Vogel, I. Levin, F. Miglietta, and E. Nisbet. The dispersion of the Buncefield oil fire plume: An extreme accident without air quality consequences. *Atmospheric Environment*, 41:9506–9517, 2007.

- Jennifer K. Wagner, Megan Doerr, and Cason D. Schmit. Ai governance: A challenge for public health. *JMIR Public Health and Surveillance*, 10(1):e58358, 2024.
- Siyuan Wang, Matthew M. Coggon, Georgios I. Gkatzelis, Carsten Warneke, Ilann Bourgeois, Thomas Ryerson, Jeff Peischl, Patrick R. Veres, J. Andrew Neuman, Johnathan Hair, et al. Chemical tomography in a fresh wildland fire plume: A large eddy simulation (les) study. *Journal of Geophysical Research: Atmospheres*, 126(18):e2021JD035203, 2021.
- Zhige Wang, Ce Zhang, Su Ye, Rui Lu, Yulin Shangguan, Tingyuan Zhou, Peter M. Atkinson, and Zhou Shi. Tracking hourly pm2.5 using geostationary satellite sensor images and multiscale spatiotemporal deep learning. *International Journal of Applied Earth Observation and Geoinformation*, 134:104145, 2024.
- Christopher Wilson. Public engagement and ai: A values analysis of national strategies. *Government Information Quarterly*, 39(1):101652, 2022.
- World Health Organization. Ethics and governance of artificial intelligence for health. *World Health Organization*, 2021.
- Rong-Yu Wu, Xi-Cheng Xie, and Yu-Jun Zheng. Firefighting drone configuration and scheduling for wildfire based on loss estimation and minimization. *Drones*, 8(1):17, 2024.
- Zhaoyue Xia, Jun Du, Jingjing Wang, Chunxiao Jiang, Yong Ren, Gang Li, and Zhu Han. Multi-agent reinforcement learning aided intelligent uav swarm for target tracking. *IEEE Transactions on Vehicular Technology*, 71(1):931–945, 2021.
- Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of ppo in cooperative multi-agent games. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 24611–24624. Curran Associates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets\\_and\\_Benchmarks.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets_and_Benchmarks.pdf).
- Manzhu Yu, Arif Masrur, and Christopher Blaszczyk-Boxe. Predicting hourly pm2.5 concentrations in wildfire-prone areas using a spatiotemporal transformer model. *Science of The Total Environment*, 860:160446, 2023.
- Xu Yue, Yihan Hu, Chenguang Tian, Rongbin Xu, Wenhua Yu, and Yuming Guo. Increasing impacts of fire air pollution on public and ecosystem health. *The Innovation*, 5(3), 2024.
- Jie Zhang and Zong-ming Zhang. Ethics and governance of trustworthy medical artificial intelligence. *BMC Medical Informatics and Decision Making*, 23(1):7, 2023.
- Jincheng Zhang and Xiaowei Zhao. Three-dimensional spatiotemporal wind field reconstruction based on physics-informed deep learning. *Applied Energy*, 300:117390, 2021.
- Weihua Zhang, Chaoying Li, and Wenmei Gai. How does evacuation risk change over time? influences on evacuation strategies during accidental toxic gas releases. *International Journal of Disaster Risk Reduction*, 108:104531, 2024a.
- Zhen Zhang, Shiqing Zhang, Caimei Chen, and Jiwei Yuan. A systematic survey of air quality prediction based on deep learning. *Alexandria Engineering Journal*, 93:128–141, 2024b.
- Xiaoru Zhao, Rennong Yang, Liangsheng Zhong, and Zhiwei Hou. Multi-uav path planning and following based on multi-agent reinforcement learning. *Drones*, 8(1):18, 2024.
- Xilei Zhao, Ruggiero Lovreglio, Erica Kuligowski, and Daniel Nilsson. Using artificial intelligence for safe and effective wildfire evacuations. *Fire Technology*, 57:483–485, 2021.
- Wenhong Zhou, Jie Li, Zhihong Liu, and Lincheng Shen. Improving multi-target cooperative tracking guidance for uav swarms using multi-agent reinforcement learning. *Chinese Journal of Aeronautics*, 35(7):100–112, 2022.