Innovative Research Journal of **Artificial Intelligence** (IRJAI)
2026; 4(1):1-16
DOI: https://doi.org/10.62497/irjai.238

Review Article

Can Artificial Intelligence Replace Traditional Public Health Disaster Management? A Critical Narrative Review

Abdussalam Khattak[1*], Safeer Gohar[2]

1 College of Environmental and Civil Engineering, Chengdu University of Technology, 610059 Chengdu, China
2 College of Computer Science and Cyber Security, Chengdu University of Technology, 610059 Chengdu, China

Citation: Khattak A, Gohar S. Can Artificial Intelligence Replace Traditional Public Health Disaster Management? A Critical Narrative Review. Innov. Res. J. AI. 2026;4(1):1-16. Available from: https://irjpl.org/irjai/article/view/238

Article Info
Received: March 11, 2026
Revised: May 20, 2026
Accepted: May 29, 2026

Keywords
artificial intelligence, AI, disaster management, disasters, public health, emergency, human health

Copyright © 2026 The Author(s).

Published by Innovative Research Journals.

This is an Open Access article under the CC BY NC 4.0 license. This license enables reusers to distribute, remix, adapt, and build upon the material in any medium or format for noncommercial purposes only, and only so long as attribution is given to the creator.

Abstract

Artificial intelligence (AI) has transitioned from experimental application to operational support in public health crises in a short amount of time. It is currently being applied, or considered, to outbreak detection, outlooks, emergency triage, damage assessment, planning supply chains, monitoring misinformation, and management of emergency operations centers. This growth has led to an intriguing assertion: that AI could someday become the replacement for conventional public health disaster management. In this important narrative review, we test this assertion by looking at the capacity of AI alongside the full range of duties of public health and disaster-management organizations. AI has been demonstrated to improve over traditional practice in certain tasks with a large amount of data. It can handle large and varied amounts of information, identify patterns that are hard to find manually, reduce reporting time and facilitate quick resource and manpower deployment. It is best when the amount of information is greater than what humans can process and when decisions need to be updated over and over. But good performance in a simple task does not constitute institutional replacement. Disasters result in incomplete and unstable data, in damage to digital systems, in behavioral changes by people, and in exacerbation of inequalities. Models can fail when distribution shifts, learn historical bias, and can give a very confident response to a question that has a mix of moral and political issues. Legal power, field verification, professional judgment, community knowledge, trust, negotiation and accountability are also key components of public health disaster management. These functions can't be meaningfully assigned to an algorithm. The experience of COVID-19, natural hazards, humanitarian crises and emergency care indicates that the future is not one consisting of traditional management without AI nor autonomous AI-based response. It is a human-in-the-loop and AI-assisted process, where algorithms handle analytical and administrative tasks with well-defined boundaries, and where accountable professionals have control, deal with uncertainty, uphold rights, and make decisions with high consequences. AI can't replace the public health disaster-management institution.

**Corresponding Author:*
Abdussalam Khattak
Email: abdussalamkhattak@stu.cdut.edu.cn

Introduction

Disasters seldom come in well-defined technical problems. They play out via fractured infrastructure, murky reporting, competing priorities, terrified communities, political pressure and extreme time limitations. Before the laboratory confirmation of an outbreak, public health officials might be called on to make decisions on protection, to protect hospitals where transport systems are breaking down, to communicate different advice, maintaining public trust, and to make decisions about the distribution of scarce medicines, shelters, staff and information. Public health disaster management was an established way of coordinating this challenging effort. It's not just a set of paper forms and tedious administrative tasks; it's an institutional system that links surveillance, preparedness, incident command, emergency care, risk communication, logistics, community engagement and recovery [1,2].

AI seems to be capable of filling in many of the gaps that have regularly plagued this system. Machine-learning models can analyze data in seconds faster than human teams; can integrate data that are typically siloed; can identify weak signals; can classify images; can estimate demand; can provide summaries or suggestions. These can help minimize information overload and time to action in a public health emergency operations centre [4,8]. AI devices that enable translation, remote triage, logistics, and disease surveillance could be particularly appealing in humanitarian contexts, where there are limited staffing resources and high demand [9].

Understandably, they'd be excited. A system that's able to process thousands of hospital records, feeds from weather and satellite data, and public messages in a single stroke of its pen sounds more powerful than a group of humans doing a spreadsheet and a phone call. But the problem of replacement is a more complex one than the problem of usefulness. An algorithm could be better and better, but fail to determine if the forecast calls for school closure. The fastest delivery route identified may not include a marginalized community. It can compose a comforting message but not as a representative of the person or organization issuing the message. The gap between prediction and public action is where law, ethics, politics, and trust come into play in the Response.

This review does not examine the possibility of AI adding to traditional public health disaster management; it is whether it can replace it. The difference is important because replacing suggests automation or semi-automation could be used to handle central functions of the current system. Therefore, a serious response to the disaster needs to address the entire disaster cycle – prevention, mitigation, preparedness, response, and recovery. It must also take into account the effects of the technologies in the presence of incomplete data, broken systems, and differential consequences of error.

This is a deliberate attempt to be balanced in the argument. AI isn't just a passing fancy and traditional practice is not simply regarded as "the way it's always been done. In some very limited areas, a continued reliance on manual processing would not be defensible. Meanwhile, statements of a full replacement overestimate the capabilities of the existing models and ignore what public health institutions are meant for. The more realistic scenario is that while some processes will be replaced by AI, many jobs will be transformed, and specific decisions will be reinforced with the power of AI, but there will remain a need for accountable human governance.

Review Aim and Narrative Method

This article has been written in the form of a critical narrative review. The strategy is selected because the question of the review is conceptual as well as empirical in nature. Technical, operational, ethical, public health, emergency-management and policy documents are included in the literature. These sources are not amenable to being meaningfully consolidated to a single combined estimate. They can be compared using a narrative method, focusing on the following question: what aspects of disaster management are technically replaceable, what aspects are better suited to be served by artificial intelligence, and what aspects still need human authority?

The targeted search included PubMed-indexed publications, database searches in publishers and official repositories of the World Health

Organization, United Nations Office for Disaster Risk Reduction, United States National Institute of Standards and Technology, and the European Union. The keywords used included artificial intelligence, machine learning, public health emergency, disaster management, emergency operations center, humanitarian health, risk communication, infodemic, emergency triage, bias, accountability, and human oversight. The peer-reviewed reviews, empirical evaluations, and governance frameworks were given priority. The historical context of documents and processes in disaster management was preserved because today's AI systems need to function in the confines of these institutional duties.

The review is not comprehensive for all technologies or disaster types. It is instead an interpretive synthesis that is informed by accepted principles for the quality of a narrative review, such as having a clearly formulated question, logical reasoning, appropriate referencing, and presenting evidence that supports and challenges the central question [3]. The analysis is centered on public health impacts, with examples from emergency medicine, remote sensing, logistics, and humanitarian response where applicable, where they impact the health of populations.

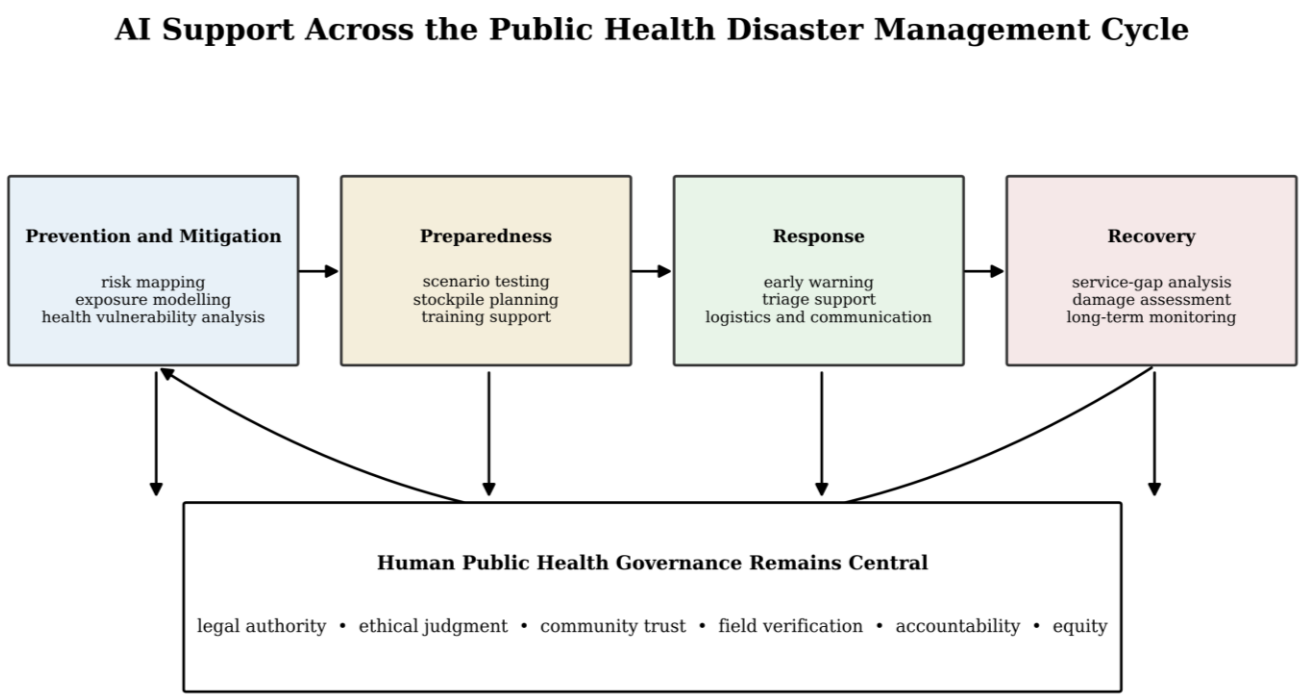

Figure 1: Although AI can support the disaster cycle, it was argued that public health governance is the key centerpiece of the cycle.

The traditional public health disaster management is more than a legacy system

When considering technological replacement, the discussions start with an incomplete system being replaced. The traditional public health disaster management can be referred as a chain of reporting disasters manually. It is, in fact, a network of agencies, hospitals, laboratories, emergency services, local authorities, non-governmental organizations, community leaders and international organizations. The core issue is coordination. There are various organizations with varying information, power, and responsibilities, and they need to work together even when they are facing increased pressure from normal rules and communication. The Health Emergency and Disaster Risk Management framework is conceptualized in a way which integrates prevention, preparedness, response, and recovery into a single process [1]. The Sendai Framework also highlights the need to know the risk, better governance, investment in resilience and preparedness for effective response and recovery [2]. The importance of these frameworks lies in the fact that the impacts of disasters are determined long before the disaster itself occurs. Establishing a well-trained workforce, a trusted surveillance system, a resilient primary care network, water supply security, and clear legal authority can't happen when a crisis is already underway, and can't be replaced by an algorithm.

These general duties are implemented at the work level through emergency operations centers and incident-management systems. They create roles and escalation procedures, communication pathways and reporting cycles, and accountability for decision-making. A situation report is worth only as much as those who made it know who made it, how it has been confirmed, and what action will be taken next. A public warning is not effective just because the message is understandable; it must be trusted by the source, and it must be easily feasible in the population to which it is given.

Disagreements are also handled by public health disaster management. Disease control, hospital capacity, economic disruption by the disease, and community concerns about safety, access and fairness are all priorities for epidemiologists, hospital leaders, political leaders and communities. Which is the best choice is usually a compromise of legitimate, but often conflicting, interests. AI can offer evidence to this process, but it can't make the conflict go away. When a model chooses one of its objectives, e.g. minimize deaths, travel time, or cost, it has already been selected by people.

This is why the replacement argument needs to start by making a distinction between tasks and institutions. A task is clearly specified with input, process, and output: classify an image, make a prediction about admissions, translate a message, optimize a route. An institution has authority, obligations, memory and accountability, and relationships. While AI can be more efficient than current systems for certain tasks, it is no simple matter to replace an institution.

Where AI Can Genuinely Improve Disaster Public Health

Prevention, mitigation, and risk mapping

Disaster impacts on public health outcomes are in part shaped by the geographic locations of people, the state of infrastructure, chronic disease rates, and the availability of essential services. AI can merge geospatial and demographic, environmental and health-system information to find intersectionality of hazards and vulnerability. The models can be used for heat-health planning, mapping flood risks, monitoring vectors, assessing exposure to wildfire, and determining communities at risk of service interruptions [5-7].

The usefulness of the practical aspect isn't prediction for prediction's sake. Pre-emergency planning can be informed by better risk maps, including inspection, vaccination, stockpiling, educating the public, and investment. They can also provide insights into relationships that would be hard to identify from a single data source. For instance, the health effects of a flood event can be influenced by the depth of water, as well as the accessibility of roads, age distribution, dependence on medications, housing quality, and distance to functional clinics.

AI still suffers from the limitations of the data on which the maps are based. Administrative datasets traditionally lack data on informal settlements, migrants, displaced populations, and rural communities. A tidy-looking map may then be misleading as to completeness. Local validation and community knowledge is still critical when model outputs dictate investment or areas are determined to be of low priority.

The planning, training, and readiness components of the program

Organizations need to envision situations that haven't occurred yet, which is what preparedness entails. Traditional activities are scenarios, tabletop discussions and drills to test plans. AI can support this by creating a range of potential scenarios, modelling demand based on various assumptions and pinpointing where plans are likely to fail. Using digital twins and simulation can help planners evaluate various evacuation routes, hospital surge plans, stockpile locations or staffing arrangements without incurring resource commitments.

AI also plays a significant role in training.AI can also assist with training. Conversational systems are able to present changing scenarios, tailor questions to the learner, and offer immediate feedback. With proper use and in a responsible way, these can be used to increase staff training opportunities for those unable to attend frequent in-person drills. They can be particularly helpful for practicing standard coordination activities, record-keeping and information management.

This is not just a modelling exercise, though. Develops connections between people who will need to cooperate in the future. If there is trust that's created in meetings, drills and joint planning it is put to the test in crisis operations. AI-generated exercises can't take the place of institutional learning that happens when agencies identify conflicting processes and work out a process that they're both working with.

Prevention and protection measures

One of the most obvious use cases where AI exceeds traditional processing is surveillance. There's no doubt about the use case here, and that's surveillance. Laboratory reports, clinical notifications, event-based alerts, news reports, environmental signals, and community information are all things that public health teams need to monitor. The volume becomes too big to be continuously monitored. Weak signals that are disconnected can be amalgamated by machine learning; anomaly detection models can detect unusual patterns, and multilingual text can be screened by NLP [4,5].

This is most clearly seen as attention support. AI can narrow down a vast amount of information down to a smaller number of signals worth investigating by humans. It could allow for the detection of an unusual cluster of calls earlier, recognition of a shift in descriptions of symptoms, or the identification of an increase in calls before the formal diagnosis increases. This acceleration can be important if it avoids the spread of transmission or allows services to get ready.

An Alert is NOT an actual event. The digital attention may not be in line with the incidence of the disease, but may rather reflect fear of the disease, media coverage of the disease, or unequal access to the internet. Urban areas can be flagged repeatedly in the models while communities that lack a strong digital representation remain unflagged. Laboratory testing, field epidemiology and clinical expertise are therefore important and still required. AI can expedite the investigation process, but it should NOT automate the investigation.

Formulating forecasts and operational situational awareness

Forecasts are important because public health decisions must be made in the absence of complete information. AI models can predict the number of cases, admission numbers, demand in the emergency department, mortality risk, consumption of supplies, and spread geographically. These outputs can be aggregated with dashboards in an emergency operations center that can provide a summary of capacity and identify where capacity is shifting (8,10).

The best operational use is sometimes short-term, not far-out prediction. A prediction of the number of beds required tomorrow or the number of ambulatory requests or oxygen consumption over the next few days may be more useful than a prediction month in advance that is very uncertain. AI can make these predictions swiftly as new data come in. This is helpful in shift planning, supply redistribution, and activation of surge arrangements.

The risk is of an incorrect degree of precision. However, despite the large uncertainty, software needs to return a value, and the model can do this. The relationship between past and future can change due to behavioral or policy change, new variants, and failure of infrastructure. Ranges, data-quality indicators, and statements of model non-validated regions are necessary for decision-makers.

Emergency triage and clinical decision support

A disaster can create more patients than there are clinicians available to assess "on the spot" or "the same day". AI-powered triage systems can aid in the identification of patients at risk of deterioration, improve the accuracy of triage decisions, and help standardize initial assessment. There have been recent systematic reviews reporting encouraging results for emergency and disaster triage, especially when vital signs and other clinical data are used in combination, using machine learning [10,11].

The added value of these tools is that they complement and do not substitute for a professional evaluation. A model can also warn a nurse when there is a subtle risk pattern or can help to prioritize reassessment. Can help to prevent cognitive fatigue over extended shifts and offer a second perspective. A prehospital environment may support decision-making in a setting where there is limited access to specialists.

Clinical triage is also a moral practice. In times of scarcity, choices will be made about who gets treated right away and who waits. Common data shortcomings include missing data, differences between patients and population, and varying definitions of risk as resources become available. The ethical policy by which the model should be used cannot be decided by a high-performing model.

Logistics, routing, and resource allocation

Emergency logistics provides a series of choices to make on where to place ambulances, route supplies, where to take patients, and when we have to move limited supplies. These decisions can be updated by using AI and optimization techniques as the roads close or demand changes, or facilities become unavailable. A study of emergency-management models indicates significant improvements in incident prediction and emergency dispatching and resource allocation systems [12].

These are the tools that will help us become efficient, but efficiency is not our only public goal. The route that covers the widest area can often pass through isolated areas more than once. A stock-allocation model might choose to invest in facilities that have good reporting mechanisms rather than facilities with the highest unmet need. Equity needs to be incorporated into the objective function and considered by decision-makers, rather than being put in as a peripheral statement of good intention.

When it comes to using AI, the ideal scenario is to make the compromises transparent. A model can illustrate what impact varying priorities on an allocation have, whether minimizing travel time, maximizing population cover, maintaining a reserve, or protecting high-risk groups. Humans can then make choices based on these options and understand the consequences.

Risk communication and infodemic management

An information emergency accompanies public health emergencies. People are in need of practical information as well as misinformation, rumors, confused media, and politicized claims. AI can be used to track public sentiment, pinpoint common questions, translate messages, and generate content for various audiences. Classification tools can be used to monitor misinformation, and generative systems can save time on the preparation of routine updates [13].

While these are helpful skills, they're easily exaggerated. Communication is not community engagement. A chatbot can give the same answers, but not speak for public authority, engage in negotiations with local leaders, and understand the reasons behind people's distrust of a recommendation. It can convey a respectful tone without being aware of the distrustful history.

When it comes to advice that affects freedom, the means of a living and one's identity, human review is particularly relevant. The message relating to quarantine, evacuation, burial, vaccination or shelter must be both medically accurate and culturally appropriate and feasible. While AI can assist with drafting and listening at scale, it is the accountable individuals who need to speak and explain uncertainty and react to disagreement.

Damage assessment and recovery

Satellite, drone, and ground imagery can be analyzed using computer vision to find flooded areas, damaged structures, obstructed roads, and land use changes. This could help to speed up the initial situational awareness when field access is not safe. The analysis of social media can be used as an additional source of information from affected populations, but verification must be done [27,28].

In recovery, AI can be used to facilitate recovery service comparisons across areas, identify health gaps that remain, analyze changes in mental-health demand, and help evaluate the priorities for reconstruction. It can also be used to facilitate the large amount of paperwork created during an event to make after-action review more systematic.

Health need is not necessarily health damage. Buildings may be intact, but staff, water, medicines, and/or referral pathways have failed. Equity, restitution, social healing, and the allocation of long-term equity investments are also issues of recovery. These are decisions regarding politics and ethics, not image class.

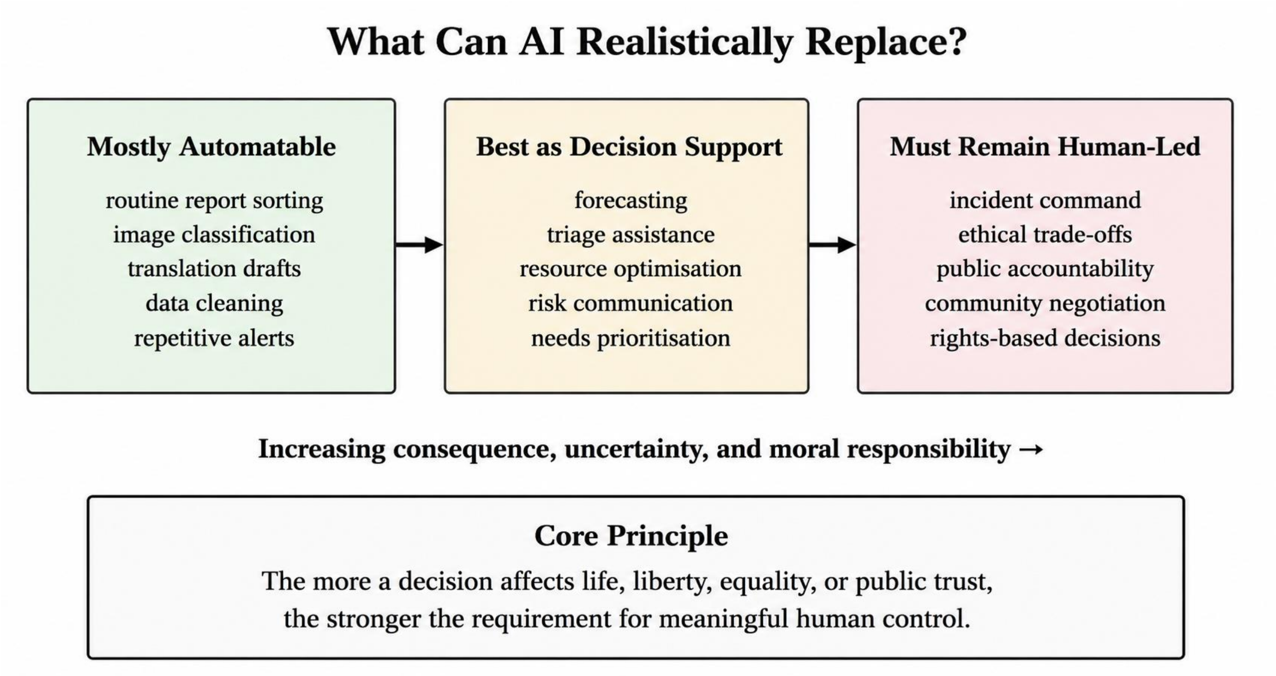

Figure 2: The replacement continuum: automation is most defensible for routine tasks and least defensible for decisions involving rights, equity, and public accountability.

Replacement or Augmentation? A Function-by-Function Assessment

The comparison exemplifies the fact that replacement statements are misleading in general. Some of the procedures are fairly automated, as they have a stable goal and it is possible to verify the result. Sorting incoming reports, converting speech to text, identifying duplicate records, or flagging an image for review are examples. Other functions are appropriate for decision support, as AI can enhance the speed and the interpretation of the output varies according to context. This includes forecasting, triage assistance, and logistics. A third group should be clearly human-designed and have legal, moral, or democratic responsibility.

This boundary isn't set in stone. Individual tasks could be increasingly automated with technological advances. The boundary should, however, be drawn around the consequence and not on the novelty. The greater the impact on life, liberty, access to essential services and/or the allocation of risk, the greater the need for meaningful human control.

Table 1: Comparison of traditional practice, AI contribution, and the limit of replacement

Function	Traditional public health role	Realistic AI contribution	Why complete replacement is inappropriate
Surveillance	Case reporting, laboratories, field investigation, event verification	Anomaly detection, text screening, rapid integration of multiple data streams	Signals may reflect digital access or media attention; confirmation requires field and laboratory work
Forecasting	Expert modelling, scenario interpretation, policy discussion	Rapid prediction, dynamic updating, comparison of multiple scenarios	Forecasts depend on assumptions and can fail when behaviour, policy, or hazard conditions change
Incident command	Named authority, role assignment, interagency coordination, documentation	Dashboards, summaries, task tracking, decision support	Algorithms have no legal mandate and cannot negotiate institutional conflict

Triage	Clinical assessment, ethical protocols, reassessment, professional accountability	Risk scoring, deterioration alerts, decision support	High-consequence clinical decisions require context, consent, professional judgment, and appeal
Resource allocation	Balancing efficiency, need, equity, and political feasibility	Demand prediction, routing, optimization, stock monitoring	Optimisation objectives are value choices and may disadvantage poorly represented groups
Risk communication	Trusted spokespersons, dialogue, community engagement, explanation of uncertainty	Drafting, translation, rumour monitoring, audience segmentation	Trust and legitimacy are relational; communities need accountable people who can listen and respond
Damage assessment	Field teams, facility reports, local knowledge, needs assessment	Image classification, remote sensing, geospatial analysis	Visible damage can miss interrupted services, vulnerable people, and hidden health needs
Recovery	Restoring services, community consultation, compensation, institutional learning	Trend analysis, gap detection, documentation support	Recovery involves justice, political choice, memory, and long-term responsibility

Why AI Cannot Replace the System as a Whole

Disaster information is incomplete, delayed and politicized

The ability to respond to an overload of disaster data is often cited as one of the top applications of AI. The reality of the operation is usually quite the opposite. Electrical power goes out, communication lines become clogged, facilities go dark, staff resort to making up their own records, and population denominators shift as people move. The data produced may be the least for the most affected areas. A model can only process what it receives in, and “no” is much easier to mistake for “no need”.

Data are also reflective of institutional priorities. It is the amount measured that is taken care of, and the amount that is not measured may disappear from the model. Some areas may be less well captured, such as informal care, unpaid care, disability related barriers, fear of authorities, and interrupted medication use. These are no random gaps. They tend to target groups that are already disadvantaged.

A safe system should therefore include explicit data quality tests, as well as the capability of identifying when inputs have exceeded acceptable levels. It also requires field teams and communities that can tell the story of what the digital system is lacking.

Models do not work if the emergency is different from the training scenario

The relationship between the training data and the situation in which the machine-learning model is used is a critical factor affecting machine-learning performance. Distribution shift is an inherent aspect of disaster. A new disease emerges, a heat wave comes, a power outage occurs, or a conflict affects the flow of people, or a hospital changes its admission policy. It may not be new variables, but the relationships between the variables may change.

There have been multiple studies on clinical AI that signal that accuracy is not necessarily the same in the real world. The success of a model can vary between institutions depending on differences in coding, equipment, population characteristics or workflows. With the pressure to deploy, in case of disaster, external and local validation can be more challenging, and in fact more important than ever.

There should be clear operating range and suspension rules for models. If the data quality decreases, the performance drifts or the event falls outside of the validated conditions, the model should be downgraded or removed from the decision pathway. Doing more with a familiar system than admitting uncertainty can be more dangerous.

The COVID-19 experience brought to light the price of premature confidence

The COVID-19 crisis generated an unprecedented number of AI models to diagnose, predict, image and forecast epidemics. It was also a powerful message to be innovative in emergencies. Critical reviews identified high risk of bias, poor reporting, small or unrepresentative samples, weak external validation, and overstated clinical readiness as issues that frequently occurred [19,20].

These problems were not due to lack of effort by the researchers. They emerged due to urgency, lack of data, and publication pressure that resulted in fast technical development in relation to careful evaluation. When presented as a typical predictive model, it is difficult to see the added value, but this is what sometimes happens with the label AI.

Even the pandemic showed the way AI can be useful. Several teams were assisted by natural-language processing, operational forecasting, research support, and information management. It's not about dismissing the use of AI. It's because public health protections can never be paused because the tool is complex and/or the crisis is immediate.

Bias can go viral at machine speed

All groups are not affected by disaster in the same way. Survival and recovery are influenced by poverty, housing, disability, immigration status, discrimination, location, and health care provision. These inequalities are present in historical data. A model that learns from past service usage, spending, digital behavior or allocation can interpret past exclusions as evidence of less need.

One of the common health-management algorithms was underestimating the needs of blacks, because spending on health services was employed as an indicator of illness [21]. This logic can also manifest itself in a disaster system. A less well-reported service use level in a community can be taken as both a low level of risk if it is indeed so, or a low access level if the actual risk is higher. A place that seems stable may be because the residents just don't have the connectivity.

Fairness calls for more than just the elimination of a sensitive variable. The same pattern can be obtained by proxy variables. Subgroup evaluation, community engagement, clear intentions, and a means of challenge are all essential to safe deployment. These are duties that must be performed by humans at institutions and are not something that can be mechanized.

Explainability does not imply Accountability

An explanation method can provide an idea of how variables might have affected a prediction. This can make it easier for users to check a model, but it does not provide an answer to the question of who's to blame. Even when a system suggests not providing a scarce intervention, a restrictive order, or a recommendation for deprioritizing a district, it is still a human and institutional decision.

The responsibility is divided among developers, vendors, agencies, managers and professionals. The choices of data and objectives are left up to the developers, the deployments are made by organizations, the policy is set by the leaders, and the outputs are interpreted by users. If roles are not clear, then automation can lead to an accountability deficit as everyone blames someone else.

Governance, transparency, safety, accountability and human oversight are therefore emphasized by WHO and NIST guidance [14,15,22]. During a public health emergency, it is important to see who is on duty at each high-consequence decision point.

Infrastructure dependence creates a new layer of fragility

The electricity, connectivity, data pipelines, cloud services, software maintenance, and cybersecurity are all required for AI systems. These dependencies are likely to be solid when things are going well, but brittle when disaster strikes. Hospitals and public agencies are also potential targets for cyberattacks, and manipulated or poisoned data can affect model output.

Even traditional systems are susceptible, and they may have low-tech alternatives like radio, paper forms, manual triage, local maps, and in-person coordination. The system should gracefully degrade when powered by an AI system. The quality of the test in optimal conditions is not the most important, but whether the organization can function safely once the tool is gone.

There's also a risk of relying too heavily on vendors. Proprietary systems can be hard to audit, modify, and/or sustain. During a long emergency, a public authority must not lose operational control due to a change in a commercial interface, an expiration of a public authority's license or a withdrawal of a vendor's support.

Generating trust and legitimacy is not something that can be manufactured

Disasters often involve people having to accept some risk, restriction, or unequal priority. The level of public cooperation depends on whether the authorities are honest, competent, and fair. Relationships establish trust before a crisis, and responsibility and respectful communication during a crisis.

While AI can produce language that is calm and empathetic, it doesn't lend institutional legitimacy. It cannot do a "morally meaningful" apology, accept responsibility, or negotiate a policy with communities affected. Nor can it know that proper counsel is unattainable unless someone tells it that the person is travelling, or taking care of the child, or earning money, or a member of a culture, or disabled, or afraid.

Local knowledge does not replace data; it is a piece of evidence that informs on feasibility and lived conditions. Taking it out of the decision-making process would not render the response objective. It would place a limited concept of evidence in the model.

Full replacement could widen global inequality

Many advanced AI systems make assumptions about digitized records, the reliability of the sensors, special staff, and connectivity. These conditions are unevenly spread. In higher-resource health systems, AI tools could enhance existing surveillance and logistics systems, whereas in lower-resource settings, they are either lacking or reliant on imported tools that are not in the local language or reflect local customs and institutions.

The outcome could be a two-tiered disaster system: high-income places have locally developed, continuously monitored models; vulnerable places have generic models, receiving limited validation. Responsible innovation should thus focus on light, multilingual, offline and locally controlled systems. Capacity building is a more important process than transferring a finished product.

The issues of ownership also come into play with data extraction. Communities may be able to submit data in the context where meaningful consent may be challenging. Their information should not become an ongoing resource for commercial model development, unless well-governed, with limits on reuse and benefit sharing.

Lessons from Different Disaster Contexts

Infectious disease emergencies

The value of information integration is evident in infectious-disease emergencies. AI can be used for syndromic surveillance, laboratory prioritization, genome analysis, forecasting and monitoring of public concerns. The tools are most beneficial when supported by robust testing, reporting and response systems. Without the means to investigate, isolate, treat and communicate, a prediction is of little public health value.

Disease outbreaks are also testament to the rapidity at which models can become obsolete. What happens is that a clinical presentation changes, variants arise, test behavior changes, public policy changes, and the data-generating process changes. These are important for continuous monitoring and re-validation.

Floods, storms, wildfires, and earthquakes

Natural hazards are an ideal application of remote sensing and geospatial analysis. AI can identify and classify damage, estimate exposure, and aid in determining obstructed routes. These outputs can be used to inform initial deployment and to minimize the risk to field teams. They are especially useful in the initial hours, when data is limited and can't wait.

But the problems in health go beyond apparent destruction. Medications: Withdrawal, or exposure to CO, may develop over days or weeks. Exposure to CO, use of medications that can lead to withdrawal, unsafe water, displacement, mental-health distress, and loss of routine care can take days or weeks to develop. Post-production of the imagery-based map, field assessment and local health surveillance are still required.

Conflict and complex humanitarian emergencies

Claims of autonomous AI-led management are most difficult in situations involving conflict. Information can be intentionally misrepresented, access to the information is challenged, populations are rapidly moving, and information can pose security threats. Humanitarian organizations are also obligated to remain neutral, impartial and protect affected populations [9,23].

While AI can help with translation, logistics, distance assessments, and pattern recognition, it is not able to address access, consent, political negotiation, and protection issues. Recommendations may be technically efficient but not practical or safe in the local security environment. Judgement is vital in human affairs.

Heatwaves and slow onset emergencies

Heatwaves, drought, and air-pollution episodes illustrate another useful role for AI: linking environmental AI can also serve as a powerful tool for connecting environmental forecasts with health vulnerability and service needs, as seen in instances like heatwaves, drought, and air-pollution episodes. Targeted early warnings can be issued to neighborhoods that feature older residents, substandard housing and fewer green spaces. Health services can be ready for more cardiovascular, respiratory and renal disease.

These emergencies also prove that "warning is not protection. No words of encouragement are going to help a person if they don't have a safe place to go, power to heat or cool, water, or the means to eliminate work. Predicting the need for public health action must be linked with social support. AI can be used to detect risk but institutions have to work to mitigate risk.

Human-Led, AI-Augmented Disaster Management: The Stronger Alternative

A practical way to go around replacement is not to use it at all. It is a conscious division of labor in accordance with comparative strength. AI can be used to perform tasks that demand fast processing, repeated calculations, pattern recognition, and information organization. Humans should have jobs that involve authority, ethical judgment, context, negotiation, empathy, and responsibility.

This model may be called bounded augmentation. Each tool comes with clear goals, target groups, geographical scope, time horizon and degree of autonomy. The output is fed into a decision rather than it is the decision itself. Users can view uncertainty and data restrictions. Overrides can be made, of course; they can be documented and they can be protected from the assumption that the algorithm is more objective.

Human control is more than just the fact that a person is at the end of an automated process. The person should be knowledgeable enough, have time available, authority, and competence to disagree. A nominal reviewer who is supposed to review a hundred or more model outputs can't really give oversight. There has to be a match between the number and impact of decisions and workflow design.

A system of community participation should also be embedded. Those impacted by a model can expose variables not accounted for, negative assumptions, and practical obstacles that developers might not be aware of. The participation of women and marginalized groups is especially crucial as regards resource allocation, surveillance, and communication aimed at the marginalized groups.

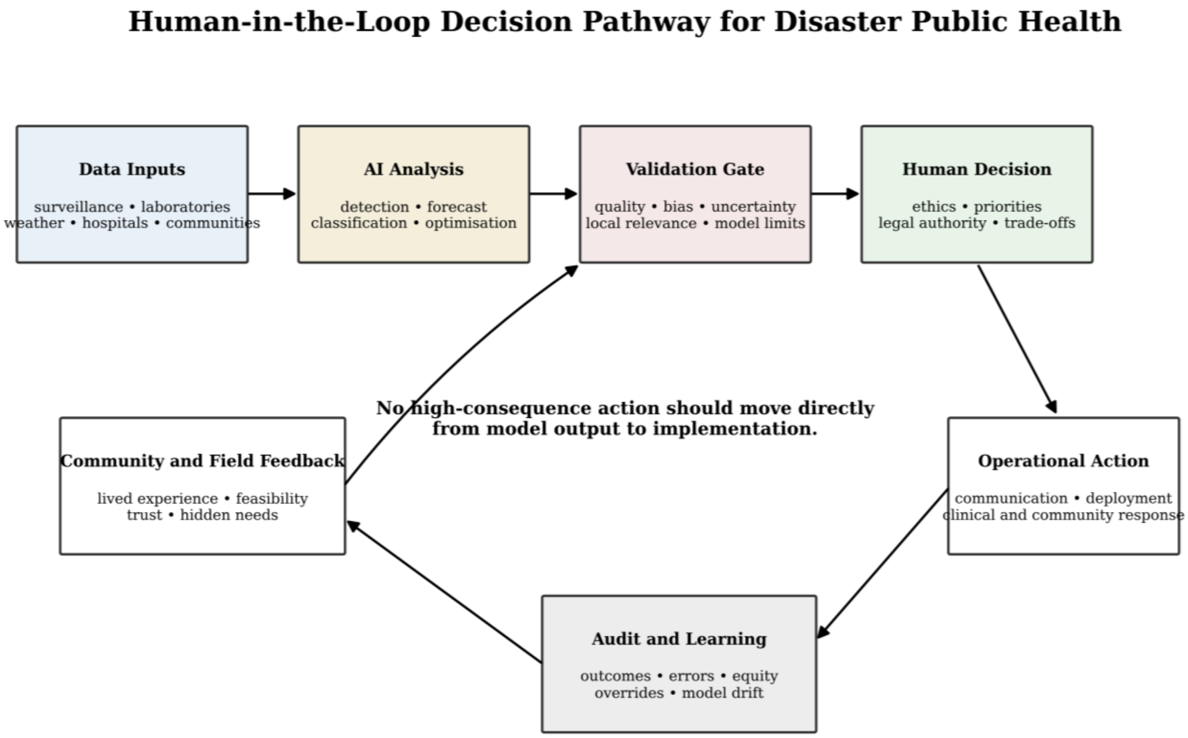

Figure 3: A human-in-the-loop pathway in which model outputs are validated, interpreted, acted upon, and audited rather than implemented automatically.

A Practical Governance Framework for Safe Use

Public health agencies need more than general ethical principles. They need operational requirements that can be checked before, during, and after deployment. The following framework translates responsible-AI guidance into disaster-management practice [14,15,22,24].

The starting point should always be the decision or problem, not the technology itself. Before selecting a tool, it is important to understand the current workflow, the expected benefit, and the level of error that can be accepted. Any system should then be validated locally and under realistic conditions, using the actual target population, data systems, language, staffing, and time pressures. Research performance alone is not enough to justify emergency deployment.

High-consequence decisions must remain clearly human-led, with named authorities and processes for reviewing decisions such as evacuation, quarantine enforcement, treatment denial, and allocation of scarce resources. Equity should also be built into the system by considering vulnerability, access, and geographic fairness alongside efficiency, while regularly examining errors across different subgroups. At the same time, systems should communicate uncertainty through data-quality indicators, confidence ranges, and warnings when operating outside validated conditions.

Emergency systems should also be designed to fail safely. Manual procedures, communication channels, paper forms, offline access, and clear suspension criteria should remain available so that essential services can continue without the model. Disagreements and human overrides should be recorded and systematically reviewed, not treated simply as user failures, because they can reveal limitations in the model and important local context.

Finally, emergency data must remain protected after the crisis, with clear rules governing retention, secondary use, commercial access, and future model training. Safe implementation also requires mixed teams: public health professionals need enough AI literacy to critically assess models, while data scientists need expertise in epidemiology, ethics, and field realities. Neither technical nor public health expertise alone is sufficient to build and operate a safe system.

Table 2: Minimum requirements for AI-supported public health disaster management

Requirement	Operational question	Minimum evidence before use
Problem definition	What decision will the tool improve, and what happens without it?	Clear use case, responsible owner, expected benefit, and foreseeable harm
Data fitness	Do the data represent the affected population and current event?	Coverage assessment, missingness review, subgroup analysis, and data-quality thresholds
Local validation	Does performance hold in the intended setting?	Independent or local testing using realistic workflows and relevant subgroups
Human authority	Who can approve, reject, or suspend the output?	Named decision-maker, override procedure, escalation route, and documented accountability
Equity protection	Who may be systematically missed or harmed?	Fairness testing, community input, alternative access routes, and equity constraints
Transparency	Can users understand the purpose, limits, and uncertainty?	Model documentation, confidence information, data provenance, and plain-language explanation
Resilience	What happens if the model, network, or vendor fails?	Offline or manual fallback, backups, cybersecurity controls, and suspension triggers
Monitoring	How will drift, error, and unintended effects be detected?	Real-time performance review, incident logging, outcome audit, and post-event evaluation
Data governance	Is emergency data use proportionate and time-limited?	Legal basis, privacy controls, retention limits, access logs, and rules for secondary use

10. Implications for the Public Health Workforce

By automating the process, AI will transform the work of professionals, even if it does not eliminate the institution. Analysts will end up spending less time on routine data cleaning and more time on examining model behavior. Emergency managers are at a disadvantage in terms of getting the summaries quickly, but will be at a greater advantage in uncertainty management and algorithmic risk assessment. Clinicians can consult with scores from the triage, which need explanation and override. Communication teams can draft using generative tools and are responsible for more verification and tone.

Training should therefore go beyond "digital literacy". Staff should be aware of the model's purpose, the data it relies on, how it performs, and when it might fail. They also require hands-on expertise in debating outputs. If you look at it from a no-big-deal-about-AI perspective, you're likely to develop automation bias.

Technical teams need to be adapted too. Graduates in the field of disaster public health should have experience in field operations, ethics, public administration, and community outreach. Not a successful model if it is statistically elegant but not easy to use. Performance encompasses human factors, workflow, and organizational incentives.

Displacement must also be acknowledged as a part of workforce policy. Some administrative duties will be reduced in size and some administrative jobs will undergo significant transformation. The aim is not to maintain all the routine procedures. It should be to enhance the capability of under-resourced functions to use the productivity gains: field epidemiology, community engagement, preparedness, evaluation and support of vulnerable people.

Research Gaps and Future Directions

There are many prototypes in the literature and only a few relatively few evaluations of sustained operational benefit. Future studies need to focus not just on accuracy, but also on the outcomes of AI, such as decreased mortality, shorter response times, more equitable outcomes, reduced workloads, and increased public trust. Further investigation is required as the retrospective performance may mask workflow and implementation issues, and prospective or multi-site and real-world studies are required.

More attention is needed for human-AI team performance. Even if the model is correct, it can be dangerous if it is misinterpreted by the user, if it is automatic and users ignore it after false alarms, or if it is ignored by the user. Research should focus on the following aspects of interface design, workload, disagreement, automation bias, trust calibration, and when and why employees disregard or contradict recommendations.

Low resource areas must not be considered a second market for systems which are developed elsewhere. The research agenda focuses on offline-enabled tools, low bandwidth data exchange, multilingual models, local governance and strategies that are still applicable when digital records are incomplete. Describe and measure community ownership and capacity building, rather than just state that this is a strategy for implementation.

There is a need to do more work on compound and cascading disaster. A heatwave can happen when there is wildfire smoke and power outages, an outbreak can happen during conflict/displacement. Isolating hazards to train models may not capture these interactions. Rather than trying to predict the one optimal scenario, it might be more appropriate to do scenario testing and robust decision making.

A huge omission is failure reporting. Success stories are published but false alarms, abandoned systems, biased allocations, cyber security incidents and near misses are still hard to come by. To achieve these in a mature field, incident reporting must be transparent, auditing must be independent, and there must be learning from failure.

Lastly, cost of the environment should be studied. Energy and infrastructure are used by big models and continuous data processing. Any AI-based disaster-management tool needs to be judged against simpler options, and not taken for granted as being useful. There are three dimensions of proportionality: technical, social/financial and environmental costs.

Conclusion

AI is moving from private to public health disaster management, and in some respects it's good. Traditional manual tools often fail to meet the challenges of today's fast-paced, large-volume, and varied emergency data. AI can identify patterns, categorize information, provide assistance to forecasts, categorize damage and support triage, and optimize logistics. If these capabilities are denied, then unfilled gaps within the response system would be left.

The argument for total replacement is a lot less convincing. Public health disaster management is not a single task of prediction. It is an open institution of authority, with a protection of rights, a coordination of organizations, an interpretation of uncertainty, a listening ear for communities and consequences to be held accountable for. The absence of data, failure of infrastructure, distribution shift, political pressure, and lack of representation are all the conditions in which AI can be most fragile in times of disaster.

The answer to the question in the review is thus conditional and clear. Selected repetitive and analytical procedures can be replaced by AI. It can support various operational decisions. It should not be used instead of incident command; ethical judgment; community involvement; nor public accountability. Those functions are social and institutional in nature, so they need to be carried out with meaning by humans, rather than through technology that is currently inadequate.

The future need not be a dichotomy between the old and the new: public health versus advanced AI. It should be created as a more robust public health system in which machines do what they do well; in which people do what only people and institutions can properly decide; in which each technological advancement is evaluated based on its ability to generate safer, fairer, and more trusted disaster response.

Declarations

Conflict of interest: The author should insert the appropriate declaration before submission.

Ethical approval: Not applicable because the article reviews published literature and policy documents.

References

1. World Health Organization. Health Emergency and Disaster Risk Management Framework. Geneva: World Health Organization; 2019.
2. United Nations Office for Disaster Risk Reduction. Sendai Framework for Disaster Risk Reduction 2015-2030. Geneva: UNDRR; 2015.
3. Baethge C, Goldbeck-Wood S, Mertens S. SANRA-a scale for the quality assessment of narrative review articles. Res Integr Peer Rev. 2019;4:5. doi:10.1186/s41073-019-0064-8.
4. Lu S, Christie GA, Nguyen TT, Freeman JD, Hsu EB. Applications of artificial intelligence and machine learning in disasters and public health emergencies. Disaster Med Public Health Prep. 2022;16(4):1674-1681. doi:10.1017/dmp.2021.125.
5. Bari LF, Ahmed I, Ahamed R, Zihan TA, Sharmin S, Pranto AH, et al. Potential use of artificial intelligence in disaster risk and emergency health management: a critical appraisal on environmental health. Environ Health Insights. 2023;17:11786302231217808. doi:10.1177/11786302231217808.
6. Kolivand P, et al. AI applications in disaster governance with health approach: a scoping review. Arch Public Health. 2025;83:218. doi:10.1186/s13690-025-01712-2.
7. Albahri AS, et al. A systematic review of trustworthy artificial intelligence applications in natural disasters. Comput Electr Eng. 2024;117:109409. doi:10.1016/j.compeleceng.2024.109409.
8. Ali Maher O, Cegolon L, Bellizzi S. Artificial intelligence as a tool for enhancing the performance of public health emergency operation centres. BMJ Glob Health. 2025;10:e016097. doi:10.1136/bmjgh-2024-016097.
9. Haykal D, Goldust M, Cartier H, Treacy P. AI in humanitarian healthcare: a game changer for crisis response. Front Artif Intell. 2025;8:1627773. doi:10.3389/frai.2025.1627773.
10. Almulihi QA, et al. Applications of artificial intelligence and machine learning in emergency department triage: a systematic review. BMC Emerg Med. 2025.
11. Tahernejad A, et al. Application of artificial intelligence in triage in emergencies and disasters: a systematic review. BMC Public Health. 2024.
12. Mukhopadhyay A, Pettet G, Vazirizade SM, Lu D, Jaimes A, El Said S, et al. A review of incident prediction, resource allocation, and dispatch models for emergency management. Accid Anal Prev. 2022;165:106501. doi:10.1016/j.aap.2021.106501.
13. Mahl D, et al. Responsible artificial intelligence in public health: a Delphi study on risk communication, community engagement and infodemic management. BMJ Glob Health. 2025;10:e018545. doi:10.1136/bmjgh-2024-018545.
14. World Health Organization. Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva: World Health Organization; 2021.
15. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Gaithersburg, MD: NIST; 2023. doi:10.6028/NIST.AI.100-1.
16. Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. BMC Med. 2019;17:195. doi:10.1186/s12916-019-1426-2.
17. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019;25:1337-1340. doi:10.1038/s41591-019-0548-6.
18. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25:44-56. doi:10.1038/s41591-018-0300-7.
19. Wynants L, Van Calster B, Collins GS, Riley RD, Heinze G, Schuit E, et al. Prediction models for

diagnosis and prognosis of COVID-19: systematic review and critical appraisal. BMJ. 2020;369:m1328. doi:10.1136/bmj.m1328.
20. Roberts M, Driggs D, Thorpe M, Gilbey J, Yeung M, Ursprung S, et al. Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. Nat Mach Intell. 2021;3:199-217. doi:10.1038/s42256-021-00307-0.
21. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447-453. doi:10.1126/science.aax2342.
22. World Health Organization. Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models. Geneva: World Health Organization; 2025.
23. Burkle FM Jr, Khorram-Manesh A, Goniewicz K. Artificial intelligence assisted decision-making in current and future complex humanitarian emergencies. Disaster Med Public Health Prep. 2025;19:e63. doi:10.1017/dmp.2025.63.
24. United Nations Educational, Scientific and Cultural Organization. Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO; 2021.
25. European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Off J Eur Union. 2024.
26. Vollmer S, Mateen BA, Bohner G, Kiraly FJ, Ghani R, Jonsson P, et al. Machine learning and artificial intelligence research for patient benefit: 20 critical questions on transparency, replicability, ethics, and effectiveness. BMJ. 2020;368:l6927. doi:10.1136/bmj.l6927.
27. Zhang C, Fan C, Yao W, Hu X, Mostafavi A. Social media for intelligent public information and warning in disasters: an interdisciplinary review. Int J Inf Manag. 2019;49:139-159. doi:10.1016/j.ijinfomgt.2019.04.004.
28. Imran M, Castillo C, Diaz F, Vieweg S. Processing social media messages in mass emergency: a survey. ACM Comput Surv. 2015;47(4):67. doi:10.1145/2771588.
29. Yuan X, Guo Q, Dadson YA, Goodarzi M, Jung J, Dong Y, et al. A review of ethical challenges in AI for emergency management. Knowledge. 2025;5(3):21. doi:10.3390/knowledge5030021.
30. Floridi L, Cowls J. A unified framework of five principles for AI in society. Harv Data Sci Rev. 2019;1(1). doi:10.1162/99608f92.8cd550d1.
31. Rajkomar A, Dean J, Kohane I. Machine learning in medicine. N Engl J Med. 2019;380:1347-1358. doi:10.1056/NEJMra1814259.
32. World Health Organization. International Health Regulations (2005). 3rd ed. Geneva: World Health Organization; 2016.
33. Centers for Disease Control and Prevention. Public Health Emergency Preparedness and Response Capabilities: National Standards for State, Local, Tribal, and Territorial Public Health. Atlanta, GA: CDC; 2018.
34. Wibowo A, et al. Leveraging artificial intelligence in disaster management: a bibliometric and systematic review. 2025.
1. 35. Hattab G, Irrgang C, Korber N, Kuhnert D, Ladewig K. The way forward to embrace artificial intelligence in public health. Am J Public Health. 2024. doi:10.2105/AJPH.2024.307888.

Disclaimer: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article or claim that may be made by its manufacturer is not guaranteed or endorsed by the publisher