© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1722830

Beyond Water Quality Monitoring: Toward AI-Enabled Early Detection of Hidden Environmental Health Risks from Emerging Contaminants

ISHOLA ABDUL DIMEJI[1*], JUMOKE OLUWAFUNMITO OMOTOSHO[2], LINDA EGBUBINE[3], ADEJOKE OLOWOOKERE-ODUYEBO[4]

[1] *Chemical Sciences Department, Federal University of Agriculture and Development Studies, Iragbiji, Osun State, Nigeria*
ORCID iD: 0009-0005-6661-3234

[2] *Geology and Natural Resources, Ball State University, Muncie, Indiana, USA*
ORCID iD: 0009-0002-7159-0565

[3] *Department of Urban and Regional Planning, University of Colorado Denver, USA*
ORCID iD: 0009-0002-5534-7389

[4] *Department of Nanoengineering, North Carolina Agricultural and Technical State University, North Carolina, USA*
ORCID iD: 0009-0003-3409-2003

**Corresponding author: Ishola Abdul Dimeji*

Abstract- ***Water quality monitoring relies on a fixed, regulated list of things to test for, measured against limits set decades ago. But the chemicals actually present in water and in the human body are far more numerous, and increasingly dominated by substances routine testing does not cover, including PFAS ("forever chemicals"), pharmaceutical residues, microplastics, endocrine-disrupting chemicals, and pesticide breakdown products. This review brings together research from analytical chemistry, exposure science, and artificial intelligence (AI) to ask a question no existing review has directly answered: can AI move environmental health monitoring beyond routine compliance testing toward genuinely detecting hidden risks, meaning contaminants that are chemically real and biologically plausible but currently invisible to regulators? We examine four areas where AI is being used, finding unknown chemicals, real-time sensor monitoring, linking exposure to biological effects, and forecasting contamination, and identify where the field agrees and where claims remain unproven. The evidence is mostly early-stage: small datasets, testing on the same data used to build the model, limited interpretability, and little connection to real health outcomes. Building on these gaps, we propose a three-step AI-Enabled Sentinel Framework that keeps chemical detection separate from health-risk assessment and requires models to be interpretable and independently tested from the outset. We close with a research agenda and implications for regulators, utilities, and communities facing unequal exposure. Our central point: the next stage of water quality monitoring will come not from testing the same known list more precisely, but from AI systems that can responsibly flag the unknown, once the field closes the gaps described here.***

Indexed Terms- ***emerging contaminants; artificial intelligence; machine learning; early warning systems; exposome; non-target screening; environmental health surveillance; water quality***

I. INTRODUCTION

1.1 The problem: what routine testing cannot see

Every regulated water system is judged against a list. That list, a fixed set of pathogens, nutrients, metals, and a modest number of older chemicals, was built up over decades through a slow regulatory process, and it is the standard by which a water supply is publicly declared safe (National Research Council [US] Committee on Drinking Water Contaminants, 2001). But the chemicals that list was meant to represent have multiplied far beyond it. Tens of thousands of manufactured chemicals are in everyday commercial use, and a fast-growing group of them, PFAS, pharmaceutical and personal-care-product residues,

microplastics and the additives that leach out of them, endocrine-disrupting chemicals, and pesticide breakdown products, now turn up routinely in environmental and human samples wherever scientists look for them, usually at levels and in combinations the regulatory list never planned for (Mallek & Barceló, 2026). The result is a real blind spot: a water supply can pass every test it is legally required to run and still deliver a chemical mixture that has never been formally checked for human health risk.

1.2 Why this is timely now

Three changes have come together to make this blind spot something researchers can start to fix, rather than just describe. First, chemical analysis has shifted away from testing for a fixed, predetermined list of known compounds and toward "non-target screening," a technique using a highly sensitive lab instrument (high-resolution mass spectrometry) that can, in principle, detect almost any chemical present in a sample, known or unknown (Yuan et al., 2026). This has created a new problem: one test run can generate tens of thousands of unidentified chemical signals, far more than any lab can sort through by hand. Second, the science of measuring human exposure has moved toward the idea of the "exposome," meaning the entire mix of environmental exposures a person accumulates over a lifetime, rather than assessing risk one chemical at a time. This shift is driven by growing evidence that long-term, low-dose exposure to mixtures of chemicals can cause health effects that no single chemical's safety limit was designed to catch (Gago-Ferrero et al., 2026). Third, and most recently, machine learning (computer methods that learn patterns from data rather than following fixed rules) has become good enough to help across this entire process: sorting unknown chemical signals, combining sensor readings for real-time monitoring, and predicting how likely a chemical is to be biologically harmful based on its structure (Cui et al., 2023). This maturation is not unique to environmental chemistry either. Automation, continuous monitoring, and machine learning are advancing together across health-adjacent infrastructure more broadly, including the growing use of smart, sensor-integrated monitoring systems in hospital settings (Oluwadare, Adeoba, Akor, Yusuff, & Milimo, 2026), which suggests the underlying AI capabilities this review draws on are part of a wider shift rather than a narrow, field-specific trend. Each of these developments has its own separate body of research behind it. What has not been done yet is to look at all three together and ask what they mean, combined, for the field's actual goal: not finding more chemicals, but finding more real health risk.

1.3 Why this matters

The stakes here are practical, not abstract. Long-term, low-level exposure to endocrine-disrupting chemicals and PFAS mixtures has been linked, in studies of human populations, to developmental, reproductive, metabolic, and cancer-related health problems, and these links show up even at levels below current legal limits (Gaillard, Barouki, Blanc, Coumoul, & Andréau, 2025). Given that, telling the public that "everything tested met the standard" is not untrue, but it is an increasingly incomplete answer to the question a resident, doctor, or regulator is actually asking. Closing the gap between what is legally tested and what is actually biologically relevant is the reason for this review.

1.4 How this review differs from existing ones

Recent reviews have already covered parts of this ground well. Some summarize new AI-based detection technologies for emerging contaminants, others focus specifically on machine-learning methods for finding unknown chemicals, and others catalog the environmental and health risks of PFAS, microplastics, or endocrine-disrupting chemicals as a group (Qiu et al., 2026). These reviews are useful, but they share one limitation: they treat successful chemical detection, or a high accuracy score for a model, as the end of the story. None of them systematically asks whether that detection evidence is strong and well-tested enough, and connected enough to actual biology, to count as real evidence of hidden health risk, rather than just a longer list of chemically confirmed but medically unexamined compounds. This review's contribution is to make that distinction explicit, to critically judge how far the current evidence actually gets from "we found it" to "it matters for health," and to propose a framework that treats that gap as something to be built and tested, not assumed away.

1.5 Objectives and scope

This review has four goals: (i) bring together how AI and machine learning are being used across the whole process, from finding a chemical to assessing whether it matters for health, mainly in water; (ii) critically judge how strong, well-tested, and broadly applicable this evidence actually is; (iii) point out where the research agrees and where real disagreement remains, and explain why; and (iv) propose an original framework and a concrete research agenda for AI-based detection of hidden risk. The focus is on water, since that is where most of this research sits, while drawing on air, soil, and food studies where the same AI methods apply directly. We cover contaminants with established or likely relevance to human health: PFAS, pharmaceuticals, microplastics, endocrine-disrupting chemicals, and pesticide breakdown products, and we lean toward research from the last five years, since that is when most of these AI applications matured.

1.6 Roadmap

Section 2 pulls the evidence together by theme. Section 3 critically evaluates that evidence. Section 4 proposes a framework. Section 5 lays out knowledge gaps and future research priorities. Section 6 covers practical and policy implications, and Section 7 concludes.

II. BRINGING THE EVIDENCE TOGETHER

2.1 From a fixed testing list to open-ended discovery

The biggest methods shift behind this whole field is the move from targeted testing to non-target screening. Targeted methods ask, "how much of chemical X is here?" Non-target methods instead ask, "what is here at all?", and return thousands of chemical signals per sample, most of them unidentified at first (Yuan et al., 2026). Machine learning has become essential for handling this: models now help break down complex chemical signals, predict how a compound will behave during testing, and score how confident an identification is, turning an unmanageable pile of raw data into a shortlist a person can actually review (Sillé, Prasse, Luechtefeld, & Hartung, 2025). Here the research agrees clearly: machine-learning-assisted non-target screening reliably finds more chemicals than manual or rule-based methods, and does it faster and more consistently.

But this is exactly where the difference between detection and risk, raised in Section 1.4, becomes concrete. Finding more chemicals is not the same as finding more health-relevant chemicals. Most non-target screening studies stop at a tentative identification of a chemical's structure, sometimes adding a generic hazard label pulled from an existing chemical database, without checking how harmful that amount actually is, how it behaves in a mixture, or how people are actually exposed to it (Turkina et al., 2025). Put plainly, this body of work has solved a data-volume problem in chemistry labs. It has not yet solved a risk-assessment problem in environmental health, and treating those two things as the same is a recurring weakness across the literature covered here.

2.2 AI for sensor-based, real-time monitoring

A separate but related line of research applies AI to real-time or near-real-time monitoring, using low-cost sensor networks, biosensors (devices that use biological material, such as enzymes or cells, to detect a substance), and satellite or remote observation feeding models that flag unusual readings. Some of this work comes from outside the emerging-contaminants field entirely. AI-linked biosensors, first built for tracking antibiotic-resistant bacteria and pathogens, are increasingly being adapted for chemical detection, a useful example of methods moving across related monitoring fields (Lawal et al., 2025). A concrete recent example of this same convergence is calibration-free, on-site detection of microcystin-LR (a toxin produced by harmful algal blooms) using integrated biosensing, multi-parameter water-quality monitoring, and machine learning together in a single field-deployable system (Kim et al., 2026), which illustrates exactly the kind of low-cost, real-time, AI-assisted sensing this section describes. This kind of finding matters because harmful algal toxins are themselves a clear example of an emerging contaminant whose occurrence depends on complex, correlated water-quality conditions rather than a single fixed threshold, a pattern documented directly in freshwater-estuarine systems (Egon, 2025). "AutoML," a method that automatically builds and fine-tunes a machine-learning model with little manual setup, has also been used to predict water-

quality readings from a smaller set of sensor inputs, aimed at making continuous, low-cost monitoring workable in places that cannot support full laboratory equipment (Campos et al., 2026).

The pattern across this research is consistent and worth stating plainly: these systems work well in controlled or pilot settings, a single watershed, a single set of sensors, a fixed testing period, but the evidence gets much thinner at the scale that would actually matter in practice: whole city water systems, multiple watersheds, or year-round monitoring through seasonal changes (Fan, Li, Ren, Liu, & Zhang, 2026). This is not a small caveat. A monitoring system is only as useful as its reliability under exactly the conditions pilot studies tend not to test.

2.3 AI linking chemical exposure to health effects

The third theme, and the one that matters most for this review's central question, is the newer work that starts to connect AI-based chemical detection to actual biological or health effects. This is the point where "hidden risk" becomes more than a phrase. AI-assisted chemical analysis methods, including models that predict a compound's identity and behavior in support of human exposure research, are starting to connect environmental chemical detection with measuring how much of a chemical is actually inside the human body, within the same computational process (Sillé et al., 2025). Separately, predictive toxicology, meaning methods that estimate how harmful a chemical is likely to be without testing it directly in animals or cells, is advancing through models such as QSAR (quantitative structure-activity relationship models, which estimate a chemical's likely biological effects from its molecular shape). These are being used to estimate whether newly found or entirely novel compounds are likely to be biologically active or disrupt hormones (Turkina et al., 2025).

This is the area where existing reviews say the least, and where the real novelty in the current research lies. Put together, machine-learning-assisted chemical screening, AI-linked exposure analysis, and predictive toxicology sketch out, for the first time, a computer-assisted path from an unknown chemical signal in a water sample, to a plausible chemical identity, to an estimate of its likely health effect, without that path needing years of conventional lab-based toxicology first. That is a genuinely new capability. It is also, as Section 3 explains, a path where every step still carries real uncertainty, and those uncertainties build up rather than cancel out as you move along it.

2.4 Predicting and forecasting contamination

A further body of research applies AI more narrowly to forecasting: predicting future contaminant levels, concentration trends, or water-quality changes from historical and real-time data, often using AutoML or "ensemble" methods (approaches that combine several simpler models into one) built to work with fewer inputs (Wang et al., 2025). One clear real-world example is machine-learning-based forecasting of cholera outbreaks, a disease directly driven by contaminated water, using data-driven methods in a West African setting (Onyijen, Olaitan, Olayinka, & Oyelola, 2023). This is exactly the kind of health-outcome-linked forecasting this review argues the field needs more of, and it stands out precisely because studies connecting an AI-based water-related signal to a named disease outcome remain uncommon (see Section 3.2). Ensemble feature-selection methods, which combine and rank multiple candidate predictors to improve a model's accuracy, have shown similar promise in adjacent health-prediction problems, such as forecasting health insurance coverage in Sierra Leone (Olawade, Osborne, Soladoye, Oluwadare, Awogbindin, & Wada, 2026), suggesting the same feature-selection techniques could usefully be adapted to contaminant forecasting in comparably resource-limited settings. The consistent finding across this broader body of forecasting research is a genuine trade-off, not a settled question: models that need fewer inputs are easier to deploy in resource-limited settings, but the research has not established whether that simplicity costs them the sensitivity needed to catch exactly the kind of low-concentration, mixed-chemical signal that defines "hidden" risk in the first place. Studies so far mostly report how accurately a model predicts the specific measurements it was built to predict, rather than testing whether a simpler model would have caught a genuinely new or unexpected contamination event.

2.5 Where the field agrees, and where it does not

Pulling these four areas together, this review finds one clear point of agreement and several real, unresolved

disagreements. The agreement: AI and machine-learning methods measurably improve the speed, sensitivity, and consistency of chemical detection compared with manual or fixed-list methods, across chemical screening, sensor monitoring, and forecasting alike. This finding holds up consistently across the research and across different types of contaminants.

The disagreements matter more for this review's purpose. First, the field is split, more by implication than open debate, on whether a compound or unusual reading flagged by a model is actually meaningful for health or just statistically unusual. Papers rarely spell out which one they mean, leaving readers to guess. Second, there is real, unresolved tension over how well these models work outside the exact conditions they were built for: models trained and tested on one type of water, one region, or one instrument show, when tested at all, inconsistent results elsewhere, and a large share of the research does not test this at all (Fan et al., 2026). Third, and most fundamentally, there is no agreement, because there is essentially no established practice, on what regulatory or public-health weight an AI-flagged, previously unknown finding should carry. Right now it sits entirely outside formal risk-assessment and regulatory systems. We think these disagreements come from three underlying features of the field, rather than from any one study's mistakes: widely different training data from one research group to the next, no shared benchmark datasets that would let different groups' results be compared fairly, and a publishing culture that, as in most applied machine-learning fields, favors reporting successful pilot results over weak or negative ones. Naming these underlying causes, instead of just noting that studies disagree, is the synthesis this review aims to provide.

III. HOW MUCH SHOULD WE TRUST THIS EVIDENCE?

Having summarized what the research claims, this section asks how much confidence that research actually earns. That is a separate task from describing its findings, and it is the one most often skipped in reviews of fast-moving technical fields.

3.1 Data quality and representativeness

Most of the AI/ML environmental-health studies covered here train and test their models on small, single-site, or lab-prepared datasets (Zhu, Yang, & Ren, 2023). What is consistently missing is data spanning multiple seasons, multiple types of water (surface water, groundwater, finished drinking water), and multiple instruments, the kind of variety a model would need to work reliably once deployed in the real world. This is less a criticism of any one study's care than a structural feature of a young field: the large, shared datasets that more established machine-learning fields rely on for solid training simply do not exist here yet.

3.2 Testing gaps

A related but separate problem is that this field almost always relies on internal testing (checking a model's performance on data drawn from the same pool used to build it) rather than testing it on new sites, populations, or time periods, which remains rare (Zhu et al., 2023). More important still: essentially no studies in this review's scope test AI-flagged chemical or biological signals against independent, real-world health outcome data. This is arguably the single biggest gap between the field's detection claims and genuine health-risk claims, and it is the gap our proposed framework in Section 4 is built to address.

3.3 Understanding why a model made a decision

Many of the best-performing models here, especially deep learning models (a complex type of machine learning good at finding subtle patterns) applied to chemical or sensor data, act as "black boxes": they produce accurate results without a clear, checkable reason for why a given sample was flagged (Zhu et al., 2023). This matters more here than in many other uses of machine learning, because regulatory and public-health decisions need reasoning that can be reviewed and defended. A model that cannot explain its own output is hard to act on with legal or clinical confidence, however accurate it is on average.

3.4 Publication bias

As in most applied machine-learning fields, weak or negative results are rarely published, which likely makes this field look more ready for real-world use

than it actually is (Kapoor & Narayanan, 2023). This review's own synthesis is necessarily built on published research shaped by that same bias, a limitation worth stating plainly rather than pretending the picture assembled here is unfiltered.

3.5 No shared benchmarks

Finally, unlike more established machine-learning fields such as medical imaging or language processing, which settled early on shared benchmark datasets that let different models be compared directly, environmental AI for emerging contaminants has no equivalent (Chang et al., 2026). This makes it hard for anyone, including this review, to rigorously compare studies or combine their results in a systematic way.

Taken together, these five limitations point to one overall conclusion: the current evidence is genuinely promising early-stage work, not yet solid enough for direct regulatory or clinical use. That is not a dismissal of the field. It is the precise, useful diagnosis the framework in the next section is built to respond to.

IV. A PROPOSED FRAMEWORK: THE AI-ENABLED SENTINEL FRAMEWORK

The evaluation above points toward a specific fix, not just a general call for "more research." We propose a three-step framework for detecting hidden risk, called the AI-Enabled Sentinel Framework.

Step 1: Broad screening. Low-cost sensor and biosensor networks, of the kind covered in Section 2.2, continuously watch for unusual readings against an established baseline, producing low-cost but low-specificity flags (meaning: cheap to run, but not yet telling us exactly what caused the flag) across a wide monitoring area.

Step 2: AI-assisted chemical identification. A flag from Step 1 triggers focused non-target chemical screening at that location, using machine-learning-assisted analysis (Section 2.1) to generate a tentative chemical identity along with a clearly reported confidence score, rather than a simple yes-or-no answer.

Step 3: Health-relevance assessment. Only compounds that clear Step 2 move on to predictive-toxicology and exposure-linked models (Section 2.3), which estimate whether the compound is likely to be biologically active or health-relevant. These estimates are labeled clearly as a starting point for further investigation, meant to guide confirmatory lab testing and, where warranted, follow-up health studies, never presented as a final answer on risk by itself.

Two design principles run through all three steps and are what set this framework apart from the largely single-step detection pipelines common in current research. First, being interpretable and independently tested are treated as requirements for a model to be used at any step, not nice-to-haves added later. A model that cannot report its own confidence, and cannot show it performs well on data it was not trained on, does not move a flag to the next step. Second, the framework keeps chemical detection (Steps 1 and 2) permanently separate from health-relevance assessment (Step 3), with each one needing its own independent testing. This directly addresses the tendency, seen in Sections 2.1 and 3.2, to treat "we detected it" and "it matters for health" as the same claim. Figure 1 shows this pipeline with a traffic-light color code (green, amber, red) at each step, reflecting how strong the current evidence actually is, as described in Section 3, so the figure itself communicates the review's critical view and not just the technology flow. This makes clear to readers and regulators exactly which parts of the pipeline are ready to be relied on today, and which are not.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1722830

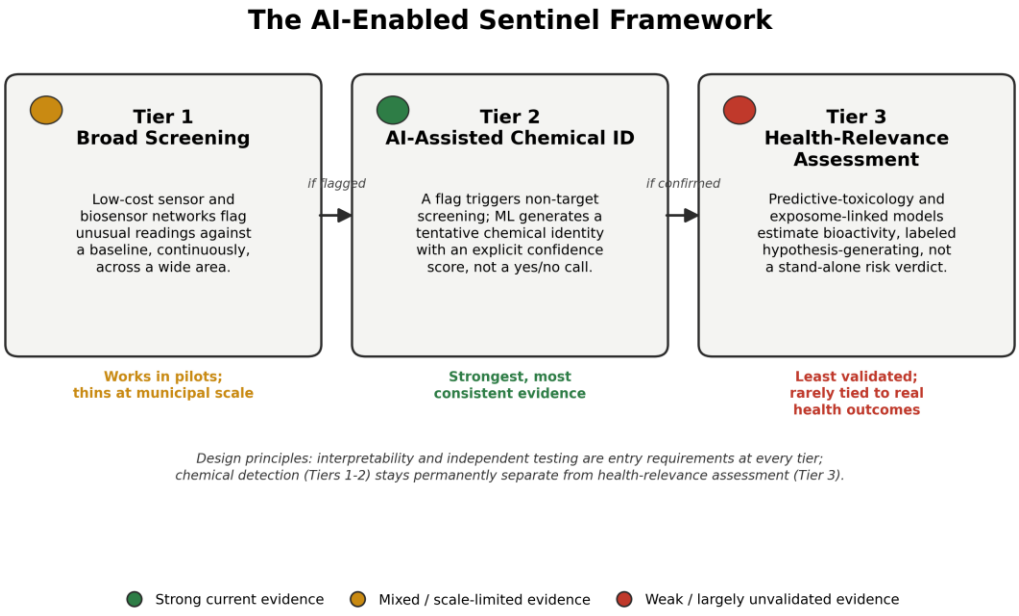

Figure 1. The AI-Enabled Sentinel Framework. Chemical detection (Tiers 1-2) is kept permanently separate from health-relevance assessment (Tier 3); each tier requires its own independent testing before a flag advances. Traffic-light colors reflect the strength of current evidence at each tier as assessed in Section 3 of this review (green = strong current evidence; amber = mixed or scale-limited evidence; red = weak or largely unvalidated evidence), not a claim about the underlying science itself.

V. WHAT WE STILL DON'T KNOW, AND WHERE RESEARCH SHOULD GO NEXT

Four types of gap follow directly from the evaluation above, each with a matching research priority.

Where the research has and hasn't happened. Validated studies in this field come overwhelmingly from well-resourced, high-income places, leaving low-resource and high-exposure communities, often the ones most affected by unregulated contaminants, largely missing from the evidence (Khan, Umer, & Faruqe, 2024). This is not because the underlying AI capacity is absent in those regions: West African research groups have already applied machine-learning forecasting to a water-related disease outbreak (Onyijen et al., 2023) and ensemble feature-selection methods to a health-prediction problem in Sierra Leone (Olawade et al., 2026). Rather, that capacity currently sits in adjacent domains, disease outbreak forecasting and health insurance analytics, and has not yet been turned toward environmental contaminant surveillance itself. Priority: collaborative studies across multiple sites and regions that deliberately include under-monitored communities, building on the local AI expertise these examples demonstrate, and using the simpler, low-input models discussed in Section 2.4 as a starting point.

Gaps in method and comparability. The lack of shared benchmark datasets and standard testing methods (Section 3.5) is not just inconvenient for reviewers like us. It actively slows the field's progress by making it hard to compare models fairly. Priority: coordinated, community-built shared datasets covering multiple types of water, similar to the ones that sped up progress in other applied machine-learning fields.

Gaps connecting detection to real health outcomes. Almost no studies in this review's scope connect an AI-flagged contaminant to actual health outcome data (Section 3.2); the path from detection to health-risk assessment proposed in Section 4 remains, for now, almost entirely untested in the real world. Priority: forward-looking studies specifically designed to pair AI-detected exposure signals with follow-up health outcomes, instead of treating detection research and health-outcome research as two separate literatures that happen to share a subject.

Gaps in regulation and oversight. No established path currently exists for how an AI-flagged, previously unknown finding would enter formal regulatory decision-making (de Paula Souza et al., 2026). Priority: building interpretability standards and reporting rules specifically for environmental regulatory use, developed jointly by chemists, toxicologists, and regulators rather than by any one group alone.

VI. WHAT THIS MEANS IN PRACTICE: PUBLIC HEALTH AND POLICY

For water utilities and regulators, this review supports moving from a fixed, static testing list toward a more adaptive approach, where AI-based screening, used carefully within the three-step logic in Section 4, helps decide where and when to run additional targeted tests, rather than replacing existing regulatory testing outright. For communities facing unequal exposure to contaminants, the implications cut both ways. Low-cost, AI-enabled sensing could genuinely extend real monitoring to communities that have historically had little of it, but only if the models used there are tested on their specific water types and populations rather than borrowed uncritically from studies done somewhere else. Used without that step, the same technology risks giving false reassurance exactly where communities can least afford it. For industry, including chemical manufacturers and water utilities, the growing ability of AI-assisted detection to find previously unlisted compounds means those compounds are increasingly likely to be found and reported. That is a reason to prepare for scrutiny on chemicals not yet subject to any regulation, rather than waiting for formal rules to catch up. For public-health communication, AI-flagged but unconfirmed findings create a genuinely new challenge: explaining a lead worth investigating (Step 3 of the proposed framework) without either dismissing it too quickly or overstating it as a confirmed health risk, a distinction that current communication practice is not yet built to make.

VII. CONCLUSION

The next stage of water quality monitoring will not come from measuring the same regulated list of contaminants with ever more precision. It will come from whether the field can responsibly build AI systems that surface chemically real and biologically plausible unknowns: contaminants, and combinations of contaminants, that are actually present and potentially significant, yet currently invisible to routine testing. This review has shown that the individual pieces of that capability, non-target chemical screening, real-time sensor monitoring, exposure-linked predictive toxicology, and forecasting, already exist and are advancing quickly. What does not yet exist is the well-tested, interpretable, and health-outcome-connected pipeline that would let any of these tools be trusted with a real health-risk judgment, rather than just an interesting chemical flag. Closing that gap, not simply building more detection capability on its own, is the field's real task for the next decade. It is a task that will need chemists, toxicologists, data scientists, and regulators to build testing standards together from the start, rather than bolting them on afterward to technology that was never designed with that goal in mind.

Methods-Transparency Note

This narrative review is based on research identified through structured literature searches conducted between August and September 2026, using PubMed, Google Scholar, ScienceDirect, Nature-portfolio journals, ACS Publications, Frontiers, MDPI, and bibliographic metadata services (OpenAlex and Crossref) to confirm full citation details. Searches used combinations of terms including "emerging contaminants," "artificial intelligence," "machine learning," "non-target screening," "exposome," and "water quality," covering mainly publications from 2015 to 2026 and weighted toward the most recent three years. Studies were included where they addressed AI/ML applications to contaminant detection, environmental health monitoring, or predictive toxicology in water or closely related environmental settings. Studies with enough methodological detail to support the evaluation in Section 3 were given deeper discussion. As a narrative rather than a systematic review, study selection reflects relevance and strength of evidence rather than an exhaustive, fully reproducible search process. This is acknowledged as standard for the narrative review format.

Declaration of Interest

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

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contributions

Conceptualization, I.A.D.; Methodology, I.A.D. and A.O.O.; Investigation (literature search and synthesis), J.O.O. and L.E.; Writing (original draft), I.A.D.; Writing (review and editing), J.O.O., L.E., and A.O.O.; Visualization, L.E. and A.O.O.; Supervision, I.A.D. All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

No new data were generated or analyzed in this narrative review.

Acknowledgments

The authors thank their respective institutions for their support. No external assistance beyond the named authors contributed to this work.

REFERENCES

[1] Campos, D., Galvão, V., de Rezende, M. L., Braga, A., Bodini, M., Aires, U. R. V., Yonaba, R., & Goliatt, L. (2026). Automated machine learning achieves accurate water quality prediction with reduced parameter requirements. *Scientific Reports, 16*(1), 4431. https://doi.org/10.1038/s41598-025-34448-8

[2] Chang, S. Y., Aboelyazeed, D., Sawadekar, K., Chavda, D., Shen, C., & Van Meter, K. J. (2026). A community dataset for large-scale river nitrogen modeling in the United States. *Scientific Data, 13*(1), 836. https://doi.org/10.1038/s41597-026-06873-5

[3] Cui, S., Gao, Y., Huang, Y., Shen, L., Zhao, Q., Pan, Y., & Zhuang, S. (2023). Advances and applications of machine learning and deep learning in environmental ecology and health. *Environmental Pollution, 335*, 122358. https://doi.org/10.1016/j.envpol.2023.122358

[4] de Paula Souza, J., Blum, J., Maran, U., Sild, S., Dawson, L., Čavoški, A., Holden, L., Lee, R., Karnel, V., Meusburger, L., Fraize-Frontier, S., Walsh, A., Rivière, G., Raitano, G., Roncaglioni, A., Di Consiglio, E., Tcheremenskaia, O., Bossa, C., Wendt-Rasch, L., Puzyn, T., & Fritsche, E. (2026). Advancing the implementation of artificial intelligence in regulatory frameworks for chemical safety assessment by defining robust readiness criteria. *Frontiers in Artificial Intelligence, 8*, 1738770. https://doi.org/10.3389/frai.2025.1738770

[5] Egon, S. S. (2025). *Investigating algal toxin occurrence and water quality correlations in a freshwater-estuarine system* [Master's thesis, University of Central Florida]. STARS. https://stars.library.ucf.edu/etd2024/444

[6] Fan, Y., Li, J., Ren, Y., Liu, C., & Zhang, W. (2026). Where machine learning fails in predicting emerging contaminant adsorption: A decision-oriented framework for model credibility and transferability. *Environmental Science & Technology, 60*(27), 18954-18974. https://doi.org/10.1021/acs.est.6c00882

[7] Gago-Ferrero, P., Boehm, A. B., Hsu-Kim, H., Li, X., Gibson, M., Vrijheid, M., Wang, B., & Zimmerman, J. (2026). The human exposome: Integrating the environment, human health, and society for the next 60 years. *Environmental Science & Technology, 60*(19), 13817-13824. https://doi.org/10.1021/acs.est.6c03080

[8] Gaillard, L., Barouki, R., Blanc, E., Coumoul, X., & Andréau, K. (2025). Per- and polyfluoroalkyl substances as persistent pollutants with metabolic and endocrine-disrupting impacts. *Trends in Endocrinology and Metabolism, 36*(3), 249-261. https://doi.org/10.1016/j.tem.2024.07.021

[9] Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-

learning-based science. *Patterns, 4*(9), 100804. https://doi.org/10.1016/j.patter.2023.100804

[10] Khan, M. S., Umer, H., & Faruqe, F. (2024). Artificial intelligence for low income countries. *Humanities and Social Sciences Communications, 11*(1), 1422. https://doi.org/10.1057/s41599-024-03947-w

[11] Kim, M., Adjei-Nimoh, S., Park, J., Egon, S. S., Choi, E., Pak, G., You, K., Sadmani, A. H. M. A., Kim, K. T., & Lee, W. H. (2026). Calibration-free on-site detection of microcystin-LR using integrated biosensing, multi-parameter water quality monitoring, and machine learning. *Water Research, 298*, 125832. https://doi.org/10.1016/j.watres.2026.125832

[12] Lawal, O. P., Opara, I. J., Ayo-ige, A., Eboh, N. A., Cos-Ibe, U., Forson, K. A. M., Mensah, E. K., Olaitan, O. F., Nii-Okai, E., Yeboah, A., Gabriels, N., & Olaniyi, A. O. (2025). Artificial intelligence-integrated biosensors for antimicrobial resistance detection and surveillance: A review and future perspectives for global biosecurity. *Cureus, 17*(11), e98098. https://doi.org/10.7759/cureus.98098

[13] Mallek, M., & Barceló, D. (2026). Analysis and risks of emerging contaminants and microplastics in natural and treated waters and human health: A critical review. *Journal of Xenobiotics, 16*(3), 93. https://doi.org/10.3390/jox16030093

[14] National Research Council (US) Committee on Drinking Water Contaminants. (2001). Drinking water contaminant candidate list: Past, present, and future. In *Classifying drinking water contaminants for regulatory consideration*. National Academies Press. https://www.ncbi.nlm.nih.gov/books/NBK207426/

[15] Olawade, D. B., Osborne, A., Soladoye, A. A., Oluwadare, O. E., Awogbindin, E. O., & Wada, O. Z. (2026). Smart insurance analytics: A novel ensemble feature selection approach to unlock health insurance coverage predictions in Sierra Leone. *International Journal of Medical Informatics*. https://www.sciencedirect.com/science/article/pii/S1386505626000535

[16] Oluwadare, O. E., Adeoba, M. I., Akor, J. T., Yusuff, T. A., & Milimo, P. (2026). The rise of smart hospitals: Biomedical engineering innovations in automation, monitoring, and digital healthcare. *Journal of Medicine and Health Research, 11*(2), 121-138. https://doi.org/10.56557/jomahr/2026/v11i210794

[17] Onyijen, O. H., Olaitan, E. O., Olayinka, T. C., & Oyelola, S. (2023). Data-driven machine learning techniques for the prediction of cholera outbreak in West Africa. *International Journal of Applied and Natural Sciences*. https://bluemarkpublishers.com/index.php/IJANS

[18] Qiu, W., Gago-Ferrero, P., Hollender, J., Leusch, F. D. L., Richardson, S. D., Samanipour, S., Shi, H., & Wang, Z. (2026). Advancing the discovery of emerging contaminants: A leap in technology and data. *Environmental Science & Technology, 60*(22), 15530-15543. https://doi.org/10.1021/acs.est.6c03027

[19] Scheringer, M., & Schulz, R. (2025). The state of the world's chemical pollution. *Annual Review of Environment and Resources, 50*, 381-408. https://doi.org/10.1146/annurev-environ-111523-102318

[20] Sillé, F. C. M., Prasse, C., Luechtefeld, T., & Hartung, T. (2025). AI redefines mass spectrometry chemicals identification: Retention time prediction in metabolomics and for a Human Exposome Project. *Frontiers in Public Health, 13*, 1687056. https://doi.org/10.3389/fpubh.2025.1687056

[21] Turkina, V., Gringhuis, J. T., Boot, S., Petrignani, A., Corthals, G. L., Praetorius, A., O'Brien, J., & Samanipour, S. (2025). Prioritization of unknown LC-HRMS features based on predicted toxicity categories. *Environmental Science & Technology, 59*(16), 8004-8015. https://doi.org/10.1021/acs.est.4c13026

[22] Wang, Q., Zhang, Y., Wang, W., Wu, X., Zhou, H., Chen, L., & Wu, B. (2025). A review of AI-driven monitoring, forecasting, and source attribution of aquatic biocontaminants.

Biocontaminant, 1, e025. https://doi.org/10.48130/biocontam-0025-0025

[23] Yuan, R., Zhang, H., Liu, J., Yu, M., Hou, X., & Jiang, G. (2026). From target-nontarget to nontarget screening: A review on screening methods for organic pollutants based on high-resolution mass spectrometry. *Environment & Health, 4*(5), 843-861. https://doi.org/10.1021/envhealth.5c00520

[24] Zhu, J.-J., Yang, M., & Ren, Z. J. (2023). Machine learning in environmental research: Common pitfalls and best practices. *Environmental Science & Technology, 57*(46), 17671-17689. https://doi.org/10.1021/acs.est.3c00026