



GLOBAL JOURNAL OF MEDICAL RESEARCH: K  
INTERDISCIPLINARY  
Volume 25 Issue 4 Version 1.0 Year 2025  
Type: Double Blind Peer Reviewed International Research Journal  
Publisher: Global Journals  
Online ISSN: 2249-4618 & Print ISSN: 0975-5888

# From the Instinctive Drowning Response to Intelligent Water Safety: A Research Agenda for AI-Augmented Aquatic Risk Science

By Francesco Pia

*Abstract-* This article proposes a research program that integrates recent advances in artificial intelligence with established theory and practice in drowning prevention, lifeguard operations, and aquatic risk management. Building on prior contributions that characterized the Instinctive Drowning Response and clarified the cognitive and perceptual demands placed on rescuers, I outline how intelligent systems can support earlier recognition, better prioritization, and more reliable intervention in aquatic environments. The paper introduces a modular architecture for AI-enhanced water safety that spans perception, causal risk modeling, decision support, education and training, and system-level prevention. It presents practical evaluation criteria, interdisciplinary collaboration paths, and governance principles tailored to the realities of beaches, pools, waterparks, and open water. The agenda emphasizes human-centered design, ethical deployment, and translational research that connects laboratory methods with lifeguard stands, facility control rooms, and public health practice.

*GJMR-K Classification: NLMC Code: WA 292*



*Strictly as per the compliance and regulations of:*



© 2025. Francesco Pia. This research/review article is distributed under the terms of the Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0). You must give appropriate credit to authors and reference this article if parts of the article are reproduced in any manner. Applicable licensing terms are at <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

# From the Instinctive Drowning Response to Intelligent Water Safety: A Research Agenda for AI-Augmented Aquatic Risk Science

Francesco Pia

**Abstract-** This article proposes a research program that integrates recent advances in artificial intelligence with established theory and practice in drowning prevention, lifeguard operations, and aquatic risk management. Building on prior contributions that characterized the Instinctive Drowning Response and clarified the cognitive and perceptual demands placed on rescuers, I outline how intelligent systems can support earlier recognition, better prioritization, and more reliable intervention in aquatic environments. The paper introduces a modular architecture for AI-enhanced water safety that spans perception, causal risk modeling, decision support, education and training, and system-level prevention. It presents practical evaluation criteria, interdisciplinary collaboration paths, and governance principles tailored to the realities of beaches, pools, waterparks, and open water. The agenda emphasizes human-centered design, ethical deployment, and translational research that connects laboratory methods with lifeguard stands, facility control rooms, and public health practice.

## I. INTRODUCTION

Drowning remains a leading cause of unintentional injury death worldwide. In aquatic settings, seconds matter, and the window between subtle onset and catastrophic outcome is narrow. Earlier work on the Instinctive Drowning Response helped practitioners recognize that drowning rarely looks like dramatic surface struggle, but instead presents as quiet, time-limited, and constrained by human physiology and motor control. This understanding informed training, scanning techniques, and supervisory protocols that shift attention toward pattern recognition of the atypical rather than search for cinematic cues.

Despite progress, three structural barriers persist. First, recognition is difficult because signal quality is low: glare, occlusion, and overlapping swimmers degrade perceptual clarity. Second, data are sparse and heterogeneous. Many events are near-misses or unreported, and incident narratives vary across venues and jurisdictions. Third, the cognitive demands placed on lifeguards and supervisors are high, with sustained vigilance, divided attention, and dynamic risk that can change with weather, crowd composition, and program type.

Author: e-mail: piafranc@hotmail.com

This paper explores how modern intelligent systems can augment the research and practice that grew from the recognition of the Instinctive Drowning Response. Rather than pursue abstract AI development, the work focuses on integration and application that respect the constraints and values of aquatic safety. The central question is not whether machines replace lifeguards, but how the field can leverage sensing, pattern discovery, causal inference, simulation, and decision support to make human action more timely, precise, and equitable.

## II. RECENT AI DEVELOPMENTS RELEVANT TO AQUATIC SAFETY

Over the past five years, three streams of AI progress have practical relevance for water safety.

First, perception has improved through advances in computer vision and multimodal learning. Models trained on video can represent movement patterns across time, not only static images. They can perform unsupervised anomaly detection and learn from small labeled datasets combined with large unlabeled corpora. Edge computing makes it feasible to run such models on-site, reducing reliance on cloud processing.

Second, models that reason about cause and effect have matured. Structural causal modeling and counterfactual inference enable systems that do more than correlate risk with conditions. They can estimate the effect of feasible interventions such as adjusting staffing, changing program mix, or reallocating attention to hotspots.

Third, simulation and digital twins now allow realistic rehearsal and evaluation. Agent-based models and physics-informed environments can recreate pool and surf dynamics, pedestrian flow, and supervision lines-of-sight. Synthetic scenarios can stress-test procedures, train personnel, and generate diverse data that would be impractical or unsafe to collect otherwise.

Together these capabilities can be tuned to the distinctive features of aquatic environments. The research challenge is to translate them into methods, tools, and governance suitable for public spaces and public health outcomes.



### III. IMPACT ON AQUATIC SAFETY: A MODULAR ARCHITECTURE

I propose a modular architecture for AI-enhanced aquatic risk science that aligns with operational realities. The architecture links five layers: perception, causal risk modeling, decision support, education and training, and system-level prevention. The modules are designed to be independently evaluable and interoperable.

#### a) Perception: Multimodal Sensing and Behavior Recognition

*Goal:* detect early signs of distress and unsafe conditions with high sensitivity and acceptable false alarm rates.

##### *Components:*

- Camera-based video analysis with spatiotemporal models that track motion patterns, postural stability, and head position relative to the waterline. Models focus on the micro-patterns consistent with the Instinctive Drowning Response such as limited voluntary movement, inability to call out, and vertical posture with minimal leg action.
- Auxiliary sensing such as water turbidity, acoustic cues in indoor pools, and crowd density mapping to contextualize detection confidence.
- Unsupervised anomaly detection tuned to venue baselines, so the system learns what is typical for a specific lane swim, open swim, or wave cycle and flags departures rather than relying on universal templates.
- On-device processing to respect privacy and latency constraints, with privacy-preserving transforms such as pose abstraction that discard personally identifiable detail while preserving dynamics necessary for safety.

##### *Research Tasks:*

- Create a benchmark that spans pools, waterparks, and surf zones with varying light and crowd conditions, labeled for near-miss and confirmed incidents.
- Quantify performance with metrics that include sensitivity, false positives per operational hour, time to recognition, robustness to glare and occlusion, and subgroup fairness across skin tones, body types, and swimwear.
- Establish protocols for external validation at new sites without extensive retraining, with performance guarantees under distribution shift.

#### b) Causal Risk Modeling and Early Warning

*Goal:* move from detection to understanding by estimating how conditions and interventions change risk.

##### *Components:*

- A knowledge graph that encodes hazards, protective factors, and facility attributes such as pool depth profile, staff positions, program types, and environmental states.
- Structural causal models that can evaluate “what if” questions, for example the effect of adding a roving lifeguard during peak density periods or reducing wave intensity during youth sessions.
- Near miss capture and learning loops that treat minor events as informative signals, not noise. The models ingest lifeguard notes, supervisor logs, and patron reports to update risk estimates continuously.

##### *Research Tasks:*

- Develop standard data schemas that harmonize incident reports across sites while preserving local specificity.
- Compare causal estimators on retrospective incident data and prospective quality improvement cycles. Report effect sizes in units meaningful to practice such as predicted reduction in seconds to recognition or number of prevented rescues per season.

#### c) Decision Support at the Edge

*Goal:* deliver alerts and recommendations that improve human action without overwhelming staff.

##### *Design Principles:*

- Human-in-the-loop workflow. Alerts escalate from soft prompts to hard alarms based on confidence, proximity, and staff workload. Operators acknowledge and annotate alerts, creating feedback for continual improvement.
- Spatial and temporal triage. The system prioritizes alerts by time-to-catastrophe and by resource availability, for example prompting a nearby roving guard rather than a distant tower.
- Fail-safes and explainability. When possible, alerts include a concise rationale such as “vertical posture with submergence attempts for 3.2 seconds” supported by a short video clip or pose trace, retained under strict retention policies.

##### *Evaluation:*

- Measure reduction in time-to-recognition and time-to-contact during drills and real operations.
- Track alert acceptance rates, response times, secondary task interference, and operator trust across shifts and seasons.

#### d) Education, Training, and Competency Maintenance

*Goal:* use intelligent systems to teach and sustain the mental models that underpin effective scanning and intervention.

*Approach:*

- Digital twins of specific venues recreate lines-of-sight, glare patterns, and patron flow. Scenario libraries expose staff to rare but critical events.
- Adaptive training that personalizes difficulty and pace based on performance, with spaced repetition for retention of pattern recognition cues derived from the Instinctive Drowning Response.
- Case-based learning that curates de-identified local incidents, allowing teams to reflect on near misses and practice improved responses.

*Research Questions:*

- What combinations of simulation fidelity and practice spacing yield durable improvements in recognition latencies under real conditions?
- How can on-shift micro-drills maintain vigilance without adding cognitive overload?

#### e) System-Level Prevention and Policy

*Goal:* Translate local insights into population-level drowning prevention.

*Methods:*

- Aggregate, de-identified analysis of incident and near-miss data to identify inequities, design flaws, and high-risk programs.
- Agent-based modeling of patron flow and supervision zones to support facility redesign, staffing plans, and policy adjustments such as entry screening, flotation device rules, and class composition.
- Integration with weather and surf forecasts to support dynamic operational decisions at beaches and open water.

*Outcomes:*

- Evidence-informed policies that reduce risk without unduly restricting access or enjoyment.
- Continuous quality improvement cycles that connect measurement, change, and re-measurement.

### IV. INTERDISCIPLINARY POTENTIAL

Aquatic safety naturally crosses boundaries. Intelligent systems create new joints of work among:

- Public health for surveillance, burden estimation, and prevention strategies.
- Human factors engineering for interface design, workload, and vigilance management.

- Computer vision and signal processing for robust detection under challenging optics.
- Causal inference and biostatistics for rigorous evaluation.
- Urban design and mechanical engineering for facility layout and surf-zone management.
- Law and policy for privacy, liability, consent, and procurement standards.

An effective research network would include aquatic operators, unions, manufacturers, standards bodies, and insurers. Shared datasets and open protocols would accelerate progress. Multi-site trials can evaluate external validity and foster common safety baselines, while allowing local tailoring.

### V. ETHICAL, POLICY, AND SOCIAL IMPLICATIONS

Deployment must be responsible and human-centered.

Privacy and dignity. Many aquatic patrons are minors, and attire reveals more skin than typical public settings. Systems should prefer on-device processing, pose-level analysis rather than identifiable imagery, strict retention limits, and segregated data pathways for safety functions. Consent practices must be transparent and understandable, with accommodations for those who opt out where feasible without compromising safety.

Fairness and bias. Vision systems can struggle with varied skin tones, body types, and cultural swim practices. Datasets and evaluation protocols must include diverse populations and environments. Performance reporting should include subgroup metrics and corrective actions when disparities appear.

Accountability and liability. Clear allocation of responsibility is necessary if a system fails to alert or generates too many false alarms. Contracts should enshrine operator control, delineate roles of vendors, and require safety cases that document hazard analyses and mitigations.

Workforce impact. Intelligent systems should elevate professional practice, not deskill it. Training and staffing plans must anticipate new competencies such as interpreting alerts, managing escalations, and contributing to incident learning loops. Unions and professional associations should help shape curricula and evaluation.

Standards and oversight. Adoption should align with established risk management frameworks and emerging AI governance standards. Independent assessment, post-deployment monitoring, and pathways to decommission unsafe systems are essential.

## VI. A PRACTICAL RESEARCH PROGRAM

The field can make measurable progress through staged research that couples laboratory development with operational trials.

*Phase 1: Benchmarking and method validation.*

- Build a curated, de-identified video benchmark with emphasis on near-miss and early distress. Include metadata on lighting, density, program type, and vantage.
- Evaluate perception models against baseline heuristics used by staff. Report sensitivity, false alarms per hour, latency, and fairness metrics. Publish pre-registered analysis plans.

*Phase 2: Human-in-the-Loop Prototypes.*

- Deploy edge-based prototypes in controlled trials at indoor pools and a guarded beach. Use silent mode first to collect alerts without staff exposure, then graduated exposure with simulated drills.
- Measure time-to-recognition during staged events, vigilance indicators, task interference, and staff sentiment. Iterate interface and escalation logic.

*Phase 3: Causal Risk and Decision Support.*

- Implement a knowledge graph and structural causal model that integrates facility attributes, schedules, and weather. Run prospective quality improvement cycles that test concrete changes such as adding a roving guard in the shallow end during peak periods.
- Estimate causal effects on leading indicators such as observed submergence attempts and response times.

*Phase 4: Training Integration and Competency Maintenance.*

- Develop venue-specific digital twins and adaptive training modules. Randomize teams to different training regimens. Track performance over a season.

*Phase 5: System-Level Prevention.*

- Pool de-identified data across sites to analyze patterns by age, program, and design features. Inform policy recommendations and equipment standards.
- Throughout, governance should include community advisory input, ethical review, and clear success and stop criteria. Data stewardship must be rigorous, with strong access controls, audit trails, and privacy-by-design.

## VII. EVALUATION METRICS AND REPORTING

The following measures support comprehensive evaluation and comparability across studies.

*Core Safety Metrics:*

- Sensitivity and specificity for early distress detection.
- False alarms per operational hour, stratified by program type.
- Median and distribution of time-to-recognition and time-to-contact during drills and real events.
- Interruption cost measured as impact on other supervisory tasks.

*Fairness and Generalization:*

- Performance by skin tone, age group, body size, attire, and activity.
- Robustness to lighting, glare, water conditions, and crowd density.
- External validity across sites without retraining.

*Human Factors:*

- Trust and acceptance trajectories over time.
- Alert acknowledgement and override patterns.
- Team coordination and communication effects.

*System Metrics:*

- On-device latency and bandwidth use.
- Energy consumption and thermal profile for edge devices.
- Uptime, failure modes, and safe degradation.
- Reporting should include transparent limitations, failure cases, and corrective actions.

## VIII. FUTURE OUTLOOK: TOWARD GENERAL INTELLIGENT SUPPORT

As intelligent systems advance, the aquatic safety community can aim for integrated assistants that unify perception, causal reasoning, and policy analysis across facilities and regions. A general assistant could:

- Coordinate multi-venue operations during heat waves or holiday surges, adjusting staffing, signage, and programming in near real time.
- Synthesize evidence from international datasets to recommend targeted prevention campaigns for communities at elevated risk.
- Guide facility design using learned models that predict risk hotspots given geometry, usage, and staffing patterns.

Autonomous rescue platforms such as drones or robotic buoys may complement human responders where conditions permit. Even in these futures, human judgment remains central. The instincts and experience that underpinned the Instinctive Drowning Response are not replaced, but rather amplified by earlier warning, better context, and more effective team coordination.

## IX. CONCLUSION

Aquatic safety advances when theory, observation, and practice reinforce each other. The recognition that drowning often follows a constrained, instinctive pattern reshaped training and scanning. Today, intelligent systems can extend that progress by improving early detection, quantifying causal risk, supporting rapid decisions, personalizing training, and guiding prevention policy. The agenda presented here aims to channel AI into forms that respect privacy, fairness, and professional judgment, while delivering measurable gains in safety.

A collaborative program that links aquatic operators, researchers, and technologists can build shared benchmarks, test human-in-the-loop tools, and institutionalize learning from near misses. By aligning intelligent systems with domain-specific needs and human-centered goals, we can reduce time-to-recognition, prevent harm, and widen access to safe enjoyment of the water.

## ACKNOWLEDGMENTS

I thank colleagues in aquatic safety, lifeguard education, and public health who have advanced the science and practice of drowning prevention.

## REFERENCES RÉFÉRENCES REFERENCIAS

1. F. A. Pia, *The Instinctive Drowning Response*, On Scene, United States Coast Guard, 2006.
2. World Health Organization, *Global Report on Drowning: Preventing a Leading Killer*, WHO Press, Geneva, 2014.
3. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework 1.0*, NIST, Gaithersburg, 2023.
4. International Organization for Standardization, *ISO 31000:2018 Risk Management — Guidelines*, ISO, Geneva, 2018.
5. International Organization for Standardization, *ISO/IEC 23894:2023 Information Technology — Artificial Intelligence — Guidance on Risk Management*, ISO, Geneva, 2023.
6. International Life Saving Federation, *Drowning Prevention Strategies: A Framework to Reduce Drowning Deaths*, ILS, Leuven, 2015.

