

MWA LEARNING EVENT: TOOLS, TRADE-OFFS, AND TAKEAWAYS: EXPLORING PRACTICAL APPLICATIONS OF AI FOR WATER SECURITY – Q&A SUMMARY

Prepared by: Audrey Caprio, Millennium Water Alliance | Contact: audrey.caprio@mwawater.org

Audience: Event Participants

Question: For organizations working in real-time water monitoring: How do you decide if and when artificial intelligence (AI) makes sense? What questions do you ask, and what holds you back?

Answer: The starting point is asking whether AI is actually needed in the given context. AI is only as effective as the data used to train it, and in many regions, particularly in sub-Saharan Africa, data systems are still developing. As a result, many organizations continue to rely on rule-based logic for predictive and filtering analysis. There is hesitation to adopt AI until there are more deployments and more relevant datasets available, rather than relying on models built for other regions. Cost and sustainability are also major constraints, as organizations must consider whether AI systems can be realistically maintained and afforded by users over time.

Question: Where does AI realistically help in water systems, and where is it overkill or unnecessary?

Answer: There is a clear distinction between using AI as a practical tool and applying it unnecessarily. AI can be especially useful for making complex or dense information more accessible, such as legislation or long reports, and for connecting qualitative and quantitative data sources. It also helps unlock insights from existing but underutilized data. However, not all problems require AI, and simpler solutions are often sufficient. A key shift is needed in how organizations prepare their data, ensuring it is well-documented and structured so both humans and AI systems can interpret and use it effectively.

Question: How do you ensure that AI-generated insights actually translate into better decisions in practice?

Answer: Contextualization is critical to ensuring AI outputs are useful, particularly for non-technical users who need insights translated into practical terms. Co-creation and co-design with stakeholders help ensure that outputs are relevant and actionable. It is also important to integrate AI outputs directly into decision-making processes rather than treating them as standalone tools. Building trust through transparency is essential, including clearly showing what data is used and how results are generated. At the same time, maintaining a balance between automation and human oversight helps ensure accuracy and credibility.

Question: How can systems increase support to ground data collection? How do we address this challenge?

Answer: Improving field data collection has historically been a challenge, but progress is being made through increased digitalization, including the use of tools like drones. There is also a growing recognition that organizations often duplicate efforts by collecting data that already exists, and platforms like Project W aim to address this by surfacing existing datasets and making them easier to find and use. The integration of AI further enhances this process by improving search and discovery and helping users identify relevant data more efficiently. AI can also help clean and

integrate data across many different in-situ sources without manual coding and cleaning. In addition, colleagues have pointed to initiatives like iNaturalist to show how [AI can enable more people to monitor, understand, and engage with nature](#), raising the possibility of using similar approaches to crowdsource water data collection.

Question: Many people are struggling with the ethical aspects of using AI such as transparency, honesty, and stigma in scientific work. Where is the case for these concerns, and are we moving toward universal adoption?

Answer: These concerns are valid and reflect real limitations of AI. AI systems can produce incorrect outputs, reflect bias, and are only as good as the data they are trained on, which makes transparency and human oversight essential. This is why there is a strong emphasis on keeping a “human in the loop” and being clear about what AI is doing, what data is being used, and where human judgment remains.

There is also growing recognition that AI is already being used informally in many organizations, which increases the need for clear guidance and responsible use.

Rather than universal adoption in all areas, the discussion pointed to more targeted and practical use of AI, where it solves a clear problem and where its outputs can be validated. The focus is less on adopting AI everywhere and more on using it where it adds value, while continuing to address ethical considerations through transparency, human oversight, and responsible design.

Question: What specific infrastructure upgrades were needed for Project W to enable AI? What about data cleaning?

Answer: One key lesson was that the platform needed supportive infrastructure before AI could be effectively integrated. This included ensuring that the database and backend systems were structured in a way that AI could read and interact with them and allow for “semantic” search which is the natural language search referenced in the case study presentation. In terms of data cleaning, Project W is not directly editing the datasets as the platform hosts in the “library” role. But, as we move forward with developing AI capabilities for basic analysis, there would likely be elements of standardization across similar datasets for standalone analysis tasks without changing the original datasets. Human oversight remains essential to review flagged issues and maintain the integrity of the platform.

Question: There are many AI models, with different costs and privacy concerns. What is a good process for choosing the right tool for the right problem?

Answer: A good starting point is to determine whether AI is actually needed for the problem and to clearly define the domain and type of problem being solved. Reviewing benchmarks for different models can help identify which tools are best suited for specific use cases. A practical approach is to identify someone within the organization who has a strong interest in AI, provide them with access to premium tools, and allow them to experiment to build internal understanding of what works best. This is also an important issue because there is often “shadow use” of AI by staff who are using tools without telling anyone, which can exceed official use, especially where no policy is in place. Putting a policy in place, involving different parts of the organization in decision making, and providing training and clear guidance on what tools can be used and what should never happen, such as recording or transcribing meetings without permission is important but not sufficient on its own. It is also important to make it easier to use data that is appropriate to share by

making it available online with proper documentation and metadata to support responsible use. Additional guidance on the best open source LLMs in 2026 can be found in [this article](#) and within [this online resource](#).

Question: Is the Well Chlorination Shock Calculator designed for one-time disinfection or routine chlorine dosing?

Answer: The initial use case for the tool was to replace one-time shock chlorination dosing calculations rather than support routine monitoring of residual chlorine.

Question: Can the Shock Calculator be used across different regions, or does it require adaptation for local contexts?

Answer: The tool was designed to be adaptable across regions by allowing users to adjust presets such as well depth, well type, and chlorine concentrations. If different areas require varying shock levels, these can be configured within the app to suit specific use cases. There is also potential to support multiple languages, including the ability to self-translate fields through a configuration file, making the tool highly customizable for different contexts.

Question: What does it mean to “build in configuration” in the Shock Calculator?

Answer: Rather than hard-coding elements like well dimensions or chlorine concentrations, these inputs are stored in a plain-language configuration file that can be updated by non-coders, making the tool easier to adapt and maintain.

Question: How does the Shock Calculator account for conditions like biofilm or iron in boreholes?

Answer: The tool itself does not determine how to account for conditions like biofilm or iron; instead, it relies on established research to define appropriate chlorine concentrations. Users can adjust the intended concentration within the app based on these factors. Additional features, such as options to account for biofilm, could be added in the future using fixed scientific parameters rather than AI-generated judgments.