LINKS GLOBAL
Communicating with Purpose. Powered by Insight.

THOUGHT LEADERSHIP SERIES

AI Literacy Is Not What You Think It Is

A field guide for public health professionals, emergency managers, municipal leaders, and businesswomen

A Links Global thought leadership article by M. Sanchez, founder of Links Global, a communications firm working at the intersection of trust science, narrative design, and civic engagement. Part of the Links Global article series: AI and Technology. Drawn from 25 years of field practice and evidence-based research in strategic communication, public engagement, and institutional trust.

In Brief

Knowing how to use an AI tool is not AI literacy. This article describes the four levels the term should actually cover, what today's systems do well and what that hides, and what the literacy gap means in practice for the people whose decisions affect public wellbeing.

Every professional field that has ever encountered a powerful new technology has made the same mistake in roughly the same sequence. First, enthusiasm: the technology will solve problems that have resisted solution for decades. Then, adoption without understanding: deploying the tool before developing the conceptual framework needed to use it well. Then, a reckoning: discovering that the tool has amplified existing inequities, introduced new failure modes, and created dependencies that are difficult to exit. Then, finally, the development of genuine literacy: the nuanced, critical, contextual understanding of what the technology can and cannot do, where it serves and where it harms, and what governance is required to keep it accountable.

Artificial intelligence is currently somewhere between the second and third phase for most of the professional communities this article is addressed to. The enthusiasm is enormous. The adoption is accelerating. The reckoning is beginning: in lawsuits, in documented bias findings, in emergency response failures, in public health interventions that performed brilliantly in well-resourced populations and catastrophically in underserved ones. The clearest example is a widely used health care algorithm that predicted health costs as a proxy for health need. Because less money was spent on Black patients with the same level of illness, the algorithm systematically rated them as healthier than they were. Nobody chose that outcome. The data did.

The genuine literacy has not yet arrived. And that gap (between the pace of AI adoption and the development of the conceptual tools needed to use it responsibly) is one of the most consequential professional challenges of this decade for anyone whose work affects public wellbeing.

This article is not a guide to AI tools. It is a guide to AI literacy: the capacity to engage with artificial intelligence technologies critically, contextually, and with enough understanding of their actual mechanisms to make sound judgments about when to use them, how to evaluate their outputs, and when to refuse them.

What AI Literacy Actually Requires

AI literacy is commonly reduced to tool familiarity: knowing how to write a prompt, how to evaluate an output, how to use a specific platform for a specific task. This is useful. It is also the shallowest layer of what the concept requires.

Genuine AI literacy for professionals in public-interest fields (public health, emergency management, municipal governance, business leadership) requires understanding at four levels.

Conceptual literacy: Understanding what AI systems actually are and what they actually do, not the marketing version, not the science fiction version, but the technical reality. Current AI systems, including large language models and machine learning classifiers, are sophisticated pattern-matching systems. They find and reproduce patterns in the data they were trained on. On most technical accounts, they do not reason or understand in the way humans do. They do not have beliefs or intentions. They produce outputs that look like reasoning, that look like understanding, that look like belief, because they were trained on enormous quantities of human text that is the product of reasoning, understanding, and belief. The resemblance is genuine and sometimes stunning. The mechanism is categorically different from what it resembles.

This matters because the limitations of AI systems are the limitations of pattern matching, not the limitations of reasoning. An AI system produces confident outputs in domains where its training data is sparse, biased, or simply wrong, because it has no mechanism for recognizing the difference between a pattern well-supported by its training data and a pattern that appears in a narrow, biased, or incorrect subset of it. Understanding this is the foundation of every other component of AI literacy.

Domain literacy: Understanding how AI systems perform in your specific professional domain, what the training data looked like, where it was collected, whose experiences and outcomes it represents, and what the documented failure modes are in your domain of practice. An AI diagnostic tool trained primarily on data from large academic medical centers will reflect the patterns of populations served by large academic medical centers. Deployed in a community health setting serving populations substantially different from those represented in the training data, its performance will degrade in ways that are not always visible without careful evaluation. An emergency response AI trained on historical disaster data will reflect the resource allocation patterns of historical disaster responses, including their documented inequities. The tool does not know this about itself. The professional using it needs to.

Governance literacy: Understanding the regulatory, legal, and accountability frameworks that apply to AI use in your professional domain, what exists, what is missing, and what the implications are for professional and institutional liability. Most public health, emergency management, and municipal governance contexts are operating with AI governance frameworks that are significantly behind the pace of AI adoption. A usable starting point exists: the NIST AI Risk Management Framework, released in January 2023, is voluntary, free, and written for exactly this kind of organization. This is not unique to these fields, but it has specific implications for professionals whose decisions affect vulnerable populations: in our view, the absence of adequate governance does not reduce professional accountability. It increases the burden on individual professional judgment.

Critical literacy: The capacity to evaluate specific AI tools and outputs with genuine skepticism, asking not "does this work?" but "for whom does this work, under what conditions, and at whose expense?" This is the hardest component of AI literacy to develop and the most important. It requires the ability to hold two things simultaneously: genuine appreciation for what AI systems can do, and genuine alertness to what they obscure, distort, and systematically miss.

What AI Systems Do Well, and What That Obscures

AI systems have genuine capabilities that are relevant and valuable for the professionals this article addresses.

Large language models can synthesize large volumes of text quickly, identify patterns across extensive literature, generate first drafts of communications materials, and assist with language accessibility: translating complex technical content into plain language, adapting materials for different reading levels, and producing translations that require professional review but are substantially faster than starting from scratch. For public health communicators managing large volumes of materials, these are real efficiencies.

Machine learning classifiers can identify patterns in large datasets that human analysts would miss: detecting early outbreak signals, identifying populations at elevated risk, flagging anomalies in infrastructure performance data before they become failures. For emergency managers and public health officials working with large administrative datasets, these capabilities are genuinely valuable when the underlying data is adequate and the validation is rigorous.

Predictive modeling tools can improve resource allocation decisions: identifying where to position emergency assets, which populations are most likely to need specific interventions, how different policy choices will likely affect health outcomes. When well-designed and properly validated, these tools can contribute meaningfully to decisions that affect many lives.

But here is what the enthusiasm around these capabilities consistently obscures: every one of these capabilities depends entirely on the quality, representativeness, and appropriateness of the underlying data, and the data almost never meets the standard the applications require.

The populations most underserved by public health systems, emergency management, and municipal governance are systematically underrepresented in the data that AI systems are trained on. Not because anyone chose to exclude them, but because the administrative systems that generate data collect more data from populations that interact more frequently with administrative systems, which are typically the populations with greater access, greater resources, and longer histories of institutional engagement. The populations with the greatest need are, in many domains, the populations least visible in the training data. AI systems trained on this data will systematically underperform for the populations most in need. And the systems will do this invisibly: producing confident outputs without any indication of the data sparsity underlying them.

This is not a future risk. It is a documented pattern across healthcare AI, criminal justice AI, and social services AI. The professional who treats AI system outputs as reliable for underserved populations without domain-specific validation is not practicing AI literacy. They are practicing AI credulity.

Specific Implications by Professional Domain

Public health professionals: The most urgent literacy need in public health AI is the ability to evaluate claims about AI diagnostic and risk-stratification tools. Every major technology company and many well-funded startups are producing AI tools that claim to predict disease, identify high-risk individuals, and optimize intervention targeting. The evidence base for many of these tools is significantly weaker than the marketing suggests. Independent validation in populations different from the training population is frequently absent. Equity analyses (examining whether the tool performs differently across racial, economic, and geographic subgroups) are often not conducted or not reported. Before adopting any AI tool that affects clinical or public health decisions, public health professionals need to be asking: what was the training data? Where is the independent validation? What are the documented performance differences across population subgroups? These are not hostile questions. They are the minimum standard of professional due diligence.

Emergency managers: The specific AI literacy need for emergency management is understanding the difference between AI tools that support decision-making and AI tools that replace it, and maintaining the human judgment capacity to function when AI systems fail, which they will, at the moments of highest stress. Emergency AI tools tend to fail in novel situations: the combination of conditions that has not appeared in the training data, the disaster type that is uncommon in the jurisdiction's historical record, the cascading failure mode that no model anticipated. These are precisely the situations emergency managers most need good decision-making tools. The professional who has allowed AI systems to atrophy their own analytical capacity is least equipped at the moment of greatest need.

Municipal leaders: The critical AI literacy need for municipal governance is procurement. The capacity to evaluate vendor claims about AI tools with enough technical understanding to identify the gap between what is promised and what is delivered, and enough governance awareness to build accountability into contracts before the tools are deployed. Most AI procurement in municipal government is currently conducted without adequate technical expertise, without meaningful validation requirements, and without accountability mechanisms that give the municipality leverage when tools perform below specification. The consequences (biased algorithmic decision-making in benefits determination, policing, and social services) are already well-documented across multiple jurisdictions. The Gender Shades study found commercial gender-classification systems with error rates of up to 34.7 percent for darker-skinned women and a maximum of 0.8 percent for lighter-skinned men, and those systems were already on the market.

Businesswomen: The specific AI literacy need in business leadership is the ability to evaluate AI adoption decisions with an honest cost-benefit framework that includes the costs that vendor presentations consistently omit: the cost of data governance and compliance, the cost of staff displacement and the human capital losses that accompany it, the cost of AI system errors in customer-facing contexts, and the reputational cost of AI failures in domains where trust is the primary asset. Many AI adoption decisions are being made on the basis of efficiency gains that are real but narrower than projected and risk exposures that are real but not modeled.

The Literacy the Field Actually Needs

The AI literacy conversation in most professional development contexts is currently a tool-familiarity conversation: here is how to use this platform, here is how to write a better prompt, here is how to integrate this tool into your workflow. This is useful at the level of individual productivity. It is insufficient at the level of professional responsibility.

The professionals this article addresses (the ones whose decisions affect vulnerable populations, community safety, public health, and civic governance) need AI literacy at a different level. Not the level that makes them more efficient users of AI tools. The level that makes them capable of evaluating AI tools critically, governing AI adoption responsibly, and maintaining the professional judgment capacity to recognize when AI systems are failing in ways the systems cannot announce for themselves.

That literacy is harder to develop than tool familiarity. It requires engaging with the actual mechanisms of AI systems, with the documented evidence of AI failure modes, with the governance frameworks that are adequate and those that are not, and with the honest assessment of which AI applications have a genuine evidence base and which are being adopted on the strength of marketing rather than performance.

Developing it is not optional for professionals who take their responsibility to vulnerable populations seriously. Because the populations most likely to be harmed by poorly implemented AI systems are the same populations these professionals exist to serve. And the professionals who lack the literacy to recognize that harm when it is occurring are not innocent bystanders. They are the last line of defense that failed.

Where This Fits

AI literacy and trust literacy turn out to be close cousins: both ask you to read the conditions before you act. The Trust Literacy article makes that case for communications, and the Links Global Trust Audit™ is where the reading starts.

Schedule a conversation with Links Global: www.linksglobal.org

Sources

- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science, 366*(6464), 447–453. doi.org/10.1126/science.aax2342
- Buolamwini, J., & Gebru, T. (2018). Gender Shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 77–91. proceedings.mlr.press/v81/buolamwini18a.html
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). nist.gov/itl/ai-risk-management-framework

Ownership and use. © 2026 Links Global. All rights reserved. The Links Global Trust Audit™, Narrative Stewardship Infrastructure™, NSI™, Epistemic Re-Regulation Architecture™, and ERA™ are trademarks of Links Global. You are welcome to read this article, link to it, share it, and quote brief passages with attribution to Links Global and a link to this page. Reproducing it in whole or in substantial part, adapting or translating it, or using the four-level literacy model in training, consulting, publications, or derivative works requires prior written permission from Links Global. Please cite as: Sanchez, M. (2026). *AI Literacy Is Not What You Think It Is: A Field Guide for Public Health Professionals, Emergency Managers, Municipal Leaders, and Businesswomen*. Links Global. www.linksglobal.org.