AI FORUM

Year: 2026 | Edition: 32 |
Cambodia | 25 November 2026 |

INSIGHTS

AI Literacy in Practice: Five Questions for Using Generative AI Responsibly

LIM Sengkhoun[1], MD, MHPE, International FAIMER Fellow

[1] ***LIM Sengkhoun*** *is a Health Advisor at GIZ Cambodia and Research Fellow at AI Forum*

Generative artificial intelligence (GenAI) is rapidly becoming part of everyday learning and work. People increasingly use it to search for information, summarize documents, generate ideas, draft text, analyze data, and automate tasks.

But being able to use GenAI does not necessarily mean being AI literate.

AI literacy does not mean becoming an AI expert or mastering every new GenAI tool. So, how much do we need to know - and be able to do - to use GenAI well?

There is no universal threshold. AI literacy is increasingly understood as a multidimensional capability involving knowledge, practical skills, critical judgement, ethical responsibility, contextual awareness, and continued learning. A recently developed scale for healthcare workers, for example, identified seven domains: AI Concept, Data Fluency, AI Evaluation, AI in Practice, Ethics and Regulation, AI in System, and Continuous Learning (Ang et al., 2026).

For everyday practice, this article translates these domains into five practical questions.

1. Do I Understand Enough About the GenAI I Am Using?

We do not need to know exactly how GenAI is built. But we should understand enough about how it works to make sense of both its capabilities and limitations.

At its simplest, GenAI learns patterns from large amounts of data and uses those patterns to generate new outputs. For example, a language model generates text by predicting likely next tokens based on input, available context, and patterns learned during training.

At a practical level: Data → Training → Patterns learned → New input and context → Generated output

This helps explain an important limitation: a convincing answer is not necessarily a correct answer.

GenAI can generate, summarize, compare, analyze, and increasingly perform multi-step tasks. But it can also produce inaccurate or fabricated information, miss context, make unsupported assumptions, and reflect biases.

Different GenAI systems also have different capabilities. A tool suited for drafting may not be best for retrieving current evidence, analyzing data, or handling sensitive information. No single GenAI system is equally suitable for every task.

So, how much do we need to understand? Not technical details, but enough to know what GenAI can do, what it cannot reliably do, where it may fail, and when greater verification is needed.

2. Do I Understand Enough About the Data Behind It?

GenAI does not learn from nowhere. What it learns - and how well it performs - is shaped by data.

We do not need to become data scientists, but four basic questions matter:

- Source - Where might information or data come from?
- Quality - Are they likely to be accurate, complete, and reliable?
- Representativeness - Do they adequately represent the people, languages, and contexts where the output will be used?
- Bias - Who or what might be missing, underrepresented, or misrepresented?

For example, GenAI may provide less complete or contextually appropriate answers about Cambodia when relevant Khmer-language information is limited, outdated, or poorly represented in information available to the system.

The familiar principle "*garbage in, garbage out*" still matters. but the issue is broader than poor-quality data alone. Even high-quality data may not adequately represent a particular population, language, culture, or local context.

We should also distinguish training data from information available to a system when we use it. Some GenAI rely mainly on what was learned during training, while others can search the web, retrieve documents, or use information we provide. Knowing where an answer appears to come from - and whether that source can be checked - is increasingly important.

So, how much is enough? Not knowing everything about a system's training data, which is often impossible, but understanding that GenAI outputs are shaped by information available to the system and may not be neutral, complete, current, or equally applicable across contexts.

3. Can I Critically Judge What GenAI Gives Me?

Understanding GenAI is not enough. AI literacy also requires evaluative judgement - the ability to critically assess what GenAI produces.

GenAI outputs can be accurate, partly accurate, incomplete, misleading, biased, or fabricated - yet presented with similarly confident and convincing language. Fluency should therefore never be mistaken for reliability.

When an output matters, ask:

- Is it accurate and supported by evidence?
- Is anything important missing?
- Is it relevant and appropriate to my context?

- Could it contain bias or unsupported assumptions?
- What should I verify independently before using it?

Not every GenAI output requires the same level of verification. An error while brainstorming presentation titles may matter little; an error used to inform a high-stakes decision could have serious consequences.

A practical rule is: The greater the consequence of being wrong, the stronger the verification should be.

So, how much is enough? Enough judgement to know what to accept provisionally, what to question, what to verify, and what to reject - and enough humility to recognize when our own knowledge is insufficient to judge answers confidently.

4. Can I Use GenAI Appropriately in Practice?

Knowing what GenAI can do is different from knowing when, where, and how to use it appropriately.

Before using GenAI, we should be able to ask: *What am I trying to accomplish? Is GenAI appropriate for this task*? *What can it help me do? What should I still do myself? How will I check the result?*

The same GenAI tool may be useful for brainstorming ideas, summarizing a non-sensitive document, or producing a first draft, but inappropriate as the sole basis for an important or high-stakes decision.

Using GenAI can also change how work gets done. If it helps draft a report, analyze information, or recommend an action, we should be clear about what GenAI does, what people still need to do, who checks the output, and who remains responsible for the result (Dakan & Feller, 2025).

So, how much is enough? Enough to know when to use GenAI, what to use it for, and what should remain under human judgement and responsibility.

5. Do I Know the Boundaries of GenAI Use?

AI literacy is not only about knowing what GenAI can do, but also what we should - and should not - do with it.

Several boundaries matter: privacy and confidentiality, fairness and harm, intellectual property, transparency, institutional policies, professional standards, and accountability.

GenAI makes it easy to upload documents and share information, including personal, organizational, or confidential information. But being able to share something does not mean we are authorized to share it. Likewise, being able to generate content does not mean we can use it without considering authorship, intellectual property, disclosure, or appropriate review.

GenAI makes it easy to share information, including sensitive or confidential information. But being able to share something does not mean we should. Likewise, GenAI may suggest what to do, but responsibility for the final decision remains with us.

A simple principle applies: Just because GenAI can do something does not mean we should use it that way. Ultimately, GenAI does not bear the consequences of our decisions. Responsibility remains with us.

So, how much is enough? Enough to recognize the boundaries of GenAI use, know when additional guidance or permission is needed, and remain accountable for what we use, share, or act upon.

So, How Much Is Enough?

For everyday use, these five areas provide a practical foundation for AI literacy: understanding GenAI we use, recognizing the role of data and context, critically evaluating its outputs, using it appropriately, and knowing the boundaries of responsible use.

But “enough” does not mean finished. As GenAI continues to evolve, so must our understanding. Continuous learning is part of being AI literate - not chasing every new tool, but keeping our knowledge, skills, and judgement relevant to how we use GenAI.

But “enough” does not mean finished. As GenAI continues to evolve, so must our understanding. Continuous learning is therefore part of being AI literate - not chasing every new tool, but keeping our knowledge, skills and judgement relevant to how we use GenAI.

Learn enough to use GenAI well. Keep learning as GenAI evolves. And never stop thinking for ourselves.

References

Ang, C. S., Ito, S., Bajaj, S., Chow, M., Cleland, J., & Heaversedge, J. (2026). The Scale for AI Literacy in Health Care Workers: Development and Validation. *Journal of Medical Internet Research*, *28*, e92373.

Dakan, R., & Feller, J. (2025). *The 4D framework for AI fluency.*

* Edited with ChatGPT for clarity; all views expressed are solely those of the author.

* The opinions expressed are the authors' own and do not reflect the views of AI Forum.