DGXprt

FREE GUIDE

Getting AI answers you can verify

A one-page guide for dangerous goods professionals

Independent research tested 13 AI models against the IMDG Code. Asked to identify the specific provision behind an answer, they managed it **14.2%** of the time. An answer you cannot trace is an answer you have to verify from scratch, so the way you ask matters as much as what you ask. Four things make the difference.

Context, Task, Sources, Constraints

C Context

What situation are we dealing with?

Jurisdiction, site, quantities, packaging, who is asking and why.

T Task

What exactly do you want done?

A comparison, a list, a draft, a check against a document. Be specific.

S Sources

What should it rely on?

Name the documents. Without this it uses its training data, which may be a superseded edition or another country’s rules, and it will not warn you.

C Constraints

What should it not assume or do?

Do not infer. Do not fill gaps. Say plainly when the information is insufficient.

What that looks like as a prompt

> "Using only the attached SDS and the Victorian dangerous goods regulations at this link, for a Victorian workplace storing 800 litres of this product in packaged form, identify the requirements relevant to its storage. Do not infer requirements that are not contained in the sources. State the clause, table or section supporting each requirement. Tell me where the available information is insufficient."

Reading it back, clause by clause:

"Using only the attached SDS and the Victorian dangerous goods regulations at this link,"	Sources. This is the clause doing most of the work.
"for a Victorian workplace storing 800 litres of this product in packaged form,"	Context. Jurisdiction, quantity and packaging, stated up front.
"identify the requirements relevant to its storage."	Task, stated narrowly enough to be checkable.
"Do not infer requirements that are not contained in the sources."	Constraint. Left alone, it will fill a gap rather than leave one.
"State the clause, table or section supporting each requirement."	Makes the answer verifiable in minutes instead of hours.
"Tell me where the available information is insufficient."	Changes the failure mode. Instead of a confident gap-fill, you get the gap flagged.

Before you rely on the answer

☐ Trace it. Can every requirement be traced to a named clause, table or section?

☐ Check the source. Is it the current edition, and the right jurisdiction?

☐ Look for the gaps. Did it tell you what it could not determine? Silence is not the same as completeness.

☐ Repeat the question. A threshold or segregation outcome should be calculated, not generated. Ask twice and compare.

☐ Keep the decision with a person. The higher the consequence, the more oversight the work needs.

Read the research behind this guide: Can you trust AI with dangerous goods? What new research says.
