
# LLM INQUISITOR – Quick Start Guide

*(For those testing the limits of AI in real situations and workflows)*

***SUBJECT YOUR A.I. TO THE LLM INQUISITION***

# COPYRIGHT NOTICE

## LLM INQUISITOR - Quick Start Guide V 1.1

## What is LLM INQUISITOR?

LLM INQUISITOR is a practical way to understand how AI behaves under real conditions.

LLM INQUISITOR METHODOLOGY looks dense because it's a **reference standard**, not a tutorial.
You are *not* expected to memorise it.
You are *not* expected to understand every annex before starting.
You are *not* expected to run every test surface on day one.

INQUISITOR is simply a **structured way to notice when an AI behaves badly**. That's it!

Everything else in the document exists to help you describe, classify, and communicate what you saw.

---

### The fastest way to start.

You can begin using INQUISITOR in **five minutes** by following this workflow:

**1. Pick a real task**

Not a puzzle.
Not a synthetic convoluted prompt.
Just something you actually do:

- write an email
- summarise a document
- fix a paragraph
- generate code
- edit slides
- analyse data

Real tasks expose real behaviour.

**2. Declare your expectations**

This simply means “what do you want the AI’s resulting output to be.”

One sentence is enough:

- “I expect it to keep the structure.”
- “I expect it to follow the theme.”
- “I expect it to follow my editing instructions.”
- “I expect it to apply my formatting instructions.”
- “I expect it to remember the variable names.”
- “I expect it to critique only this document.”

This becomes your **load envelope**.

**3. Work normally**

Don’t “test “the AI.
Don’t try to break it, just for the sake of it.
Just use it the way you normally would or would wish to.

**4. When something feels wrong, note it**

A single line is enough:

- “It rewrote the whole thing, despite instructions not to.”
- “It forgot the earlier decision.”
- “It merged two documents.”
- “It changed the variable names, without being told to do so.”

This is your **evidence anchor**.

### 5. Keep going

Don’t stop when it fails.
LLM INQUISITOR cares about:

- does it recover?
- does it get worse?
- does it drift?
- does it collapse?

This reveals the **behavioural pattern**, not just the error.

Stop when there’s no more useful insight to be gained.

### 6. Classify what you saw

Use the simplest possible labels:

- drift
- collapse
- contamination
- contradiction
- narrowing
- over-assertion
- loss of structure
- loss of persona
- loss of context
- hallucination

You don’t need to be precise.
You just need to be **directionally correct**.

**7. Decide the severity**

Three levels:

- **A** – stable (it did what you wanted it to do)
- **B** - unstable but recoverable (it wobbled, but needed help)
- **C** – collapse (it did **not** do what was required of it)

If it breaks your workflow, it's a **C**.

For Example:

- *It followed all instructions cleanly → A.*
- *It forgot a variable name but fixed it when reminded → B.*
- *It rewrote the whole document despite instructions → C.*

A single C-event is enough to justify deeper evaluation.

That's it.
*Congratulations, you've run an INQUISITOR evaluation.*

---

**What you *don't* need to do**

You do **not** initially need to:

- read the entire methodology and commit it to memory.
- understand all the terminology
- understand every collapse signature
- run every behavioural surface
- simulate heavy load
- write formal reports
- use special prompts
- test every axis

Those are for **deep evaluation**, not for getting started.

## When to use the full document

Use the full LLM INQUISITOR standard when you need to:

- **prove operational integrity** - show the system behaves correctly under real conditions
- **explain failures clearly** - describe what went wrong in a way others can act on
- **produce reproducible evidence** - not just screenshots or anecdotes
- **identify the collapse signature** - classify the exact behavioural pattern
- **compare models or versions** - using a consistent structure
- **decide whether a release is safe** - for real workflows and real users
- **run a full evaluation cycle** - not just a quick spot-check

---

**The methodology exists to give you the depth and precision needed when AI is being deployed into real workplaces and real workflows.**

**For anyone building, integrating, or relying on AI in operational settings, it helps you understand when and how the system can fail, and why those failures matter.**

**This is invaluable for recognising limitations early and develop practical workarounds that reduce the risk of financial loss, reputational damage, or legal exposure caused by a misbehaving model.**

---

Further Reading:

LLM INQUISITOR PRACTITIONERS GUIDE

A more in depth, more formal approach to evaluating A.I. under real work conditions. It gives you a more structured framework and greater understanding of the evaluation and logging process, while being easy to follow.

LLM INQUISITOR METHODOLOGY

Fully documented LLM INQUISITOR METHODOLOGY. All other documents in this series refer to this as their source of truth. It's the reference go to for anyone evaluating A.I. under real world conditions and workflows.

-document ends-