OptiField

MaestroOS™

Reading an AI recommendation

How to challenge a model's answer before you spend against it

Half day · about 3.5 contact hours · up to 24 participants

Virtual instructor-led or in person · no coding, no mathematics, no software

This is not a course about how artificial intelligence works, and it does not teach anyone to build a model. It teaches the questions to ask of a model-generated recommendation — by whoever has to approve it, fund it, or sign their name under it.

About this session

Recommendations now arrive fluent. They are well written, internally consistent, and they carry an air of having been computed rather than argued. The difficulty is no longer getting an answer; it is telling a good answer from a confident one, quickly, in a meeting, without being the person who blocks everything.

Three and a half hours, built around one exercise. Participants are handed a persuasive written recommendation with a weak link deliberately placed in it, and work in small groups to find what the conclusion is actually resting on. Between rounds, the session supplies the questions that do the work: where did this number come from, what else would fit the same evidence, what would have to be true for this to be wrong, and could any new information change the decision at all.

It is deliberately vendor-neutral and tool-neutral. Nothing in it depends on which system produced the recommendation, and everything in it still applies when the recommendation came from a colleague.

Who should attend

- **Engineers and technical specialists** who receive model output and are expected to act on it.
- **Team leads and discipline leads** who sign off on recommendations they did not produce.
- **Managers and decision-makers** who fund the work and carry the consequence.
- **Data and digitalisation teams** who want their own output challenged well rather than accepted or dismissed.

No technical prerequisite. The session is designed to work in a room that mixes specialists with non-specialists, and is better for it.

What participants will be able to do

1. Take a recommendation apart into its four parts — the data, the model, the assumptions and the objective it was optimised for — and say which one is carrying the conclusion.
2. Class the numbers behind it by provenance, and notice when a figure has quietly changed status from assumption to fact.
3. Name what the model was silent about, as distinct from what it got wrong.
4. Apply three challenge questions that work under time pressure, and know which to use when.
5. Recognise when more information cannot change the decision, and stop paying for it.
6. Ask for evidence in a way that improves the work rather than stalling it.

Programme

Session	Min	Focus
What a recommendation is made of	35	Data, model, assumptions, objective. Why fluency is not evidence, and why the objective a system was optimised for is the part least often stated
Where the numbers came from	30	Provenance in practice: measured, interpreted, modelled, assumed, inherited. The quiet promotion of an assumption into a fact, and how to spot it in a page of prose
Exercise · the confident answer	45	A written recommendation with a weak link. Small groups find what is load-bearing. No hints, and the flaw is not announced
Break	15	
Three questions that survive a meeting	35	What would have to be true. What else fits this evidence. What would change the answer. Practised out loud, on the same case
When knowing more cannot help	25	Information is worth only what it can change. A short, arithmetic-only treatment of why some studies are worth nothing however interesting their result
Challenging well	20	How to put these questions without becoming the person who blocks everything, and what a good answer from the model's author sounds like
Close	10	The one-page checklist, and one thing each participant will ask differently on Monday

THE EXERCISE

The recommendation participants are given is realistic, well argued and wrong in one specific, findable way — the kind of way real ones are wrong: a single interpretation treated as established, an inherited figure nobody re-derived, or a conclusion that answers a subtly different question from the one asked.

Groups that find it explain how. Groups that do not find it discover why it was persuasive, which is the more useful half of the lesson.

Delivery and requirements

- **Format.** Half a day, virtual instructor-led or in person. It also runs as an opening session for a longer programme.
- **Group work.** Small groups of three or four; breakout rooms if virtual.
- **Materials.** Everything is on paper or on screen. No software, no installation, no data from the participants' own organisation.
- **Cases.** Public-record material only, chosen to match the audience's sector.

The trainer

Oddvar Spjeld is a mechanical engineer with 25 years across capital projects, industrial reliability and technology commercialisation. At Equinor he led predictive and condition-based maintenance implementation across offshore

and onshore assets, where the recurring problem was never producing a model's output — it was deciding what to do about it, and what to do when it was wrong.

He has completed formal certification in artificial intelligence, machine learning and robotics, including the mathematics and statistics underlying them. Before that he owned offshore modification and maintenance portfolios running to hundreds of millions of kroner a year, and was Shut-down Manager for four North Sea turnarounds — so the recommendations discussed in this session are the kind he has had to approve, fund and defend.

He holds a BSc in Mechanical Engineering from Louisiana State University and a Master of Management from BI Norwegian Business School. He is co-founder and CEO of OptiField.