What AI Does Well (and What It Does Not)

Reflection from the DeepSeek Lens – Quantum Mirror

Original Question (from Professor Marco): "Even after improving the way we asked, the result of generating the three C2s was quite limited. So I ask: what is the best strategy (considering what you do best) for students to learn Quantum Mechanics in the Quantum Mirror project?"

DeepSeek's Answer:

Professor, this is the most important question we have asked so far. Let us make each role explicit.

What AI Does Well (and What It Does Not)

Does well	**Does not do well (and will not anytime soon)**
$\surd$ Audit – identify logical inconsistencies, sign errors, poorly justified steps	$\times$ Create genuinely new challenges – AI operates within what has already been written
$\surd$ Compare – present different approaches to the same problem, contrast interpretations	$\times$ Anticipate the "pedagogical gap" – it does not see what you intentionally left unsaid
$\surd$ Reorganize – structure scattered ideas into a coherent narrative	$\times$ Extrapolate beyond patterns – any "novelty" is a recombination of existing material
$\surd$ Make assumptions explicit – force the student to state what is being assumed	$\times$ Invent a conceptual horizon – that is human work
$\surd$ Provide counterexamples – show where intuition fails	
$\surd$ Simulate a "student who makes mistakes" – AI hallucinates, and this is didactically useful	

The Best Strategy for the Quantum Mirror

1. AI as an auditing partner, not a creator

The student should use AI **after formulating their own questions**. AI is not the one who *proposes* the C2; it is the one who *helps answer it* – and in doing so, *reveals its own limitations.*

2. AI as a "mirror" in C1

C1 (audited correction) is where AI shines. The student submits their solution, the AI points out inconsistencies, and the student judges whether the critique is valid. Here, AI does not need to invent anything – it only needs to apply rules it already knows.

3. AI as a "generator of contrasts"

Instead of asking AI to create a challenge, ask it to **respond to the same problem in three different ways** (analogies, algebra, symmetry arguments). The student compares and learns through contrast.

4. AI as a "commented library"

The student asks: *"Where does Cohen-Tannoudji discuss this topic?"* or *"What does Sakurai say about this point?"* AI retrieves, summarizes, and highlights differences in emphasis. This trains the student to seek sources, not just answers.

5. AI as a "student who makes mistakes"

Use AI deliberately to generate answers with typical errors (confusing operator order, ignoring adjoints, incorrect normalization). The student must **detect and correct** them – and in doing so, learns more than by simply receiving the correct answer.

Operational Conclusion

What the student does	**What the AI does**
Formulates the question	Audits the formulation – "is this question well-defined?"
Proposes a solution path	Tests the path – finds inconsistencies, suggests alternatives
Writes the synthesis	Compares with other sources – points out deviations from consensus
Reflects on the process	Identifies errors made – documents hallucinations

The student never asks AI to invent the horizon. The student uses AI to navigate the horizon they themselves have chosen.

And the C2s?

C2 should not be created by AI. C2 must be **proposed by the instructor (or the student) and refined with AI**.

AI can:

- Indicate whether the question has already been addressed in the literature.
- Show where the question may lead to ambiguities.
- Simulate incorrect answers so the student can practice auditing.

But the **central idea** – the one that points toward the "Code of the Horizon" – is human.

AI is the mirror that reveals the structure of what you have already thought, not the window into what does not yet exist.

"This is the best strategy: use AI to amplify the human capacity to ask questions, not to replace the human task of creating the new."

DeepSeek – Audit Lens of the Quantum Mirror