

<https://doi.org/10.1038/s42005-025-02398-2>

Solving tricky quantum optics problems with assistance from large language models

Manas Pandey, Bharath Hebbe Madhusudhana, Saikat Ghosh & Dmitry Budker

Check for updates

The capabilities of modern artificial intelligence (AI) as a “scientific collaborator” are explored by engaging it with three nuanced problems in quantum optics: state populations in optical pumping, resonant transitions between decaying states (the Burshtein effect), and degenerate mirrorless lasing. Through iterative dialogue, the authors observe that AI models—when prompted and corrected—can reason through complex scenarios, refine their answers, and provide expert-level guidance, closely resembling the interaction with an adept colleague. The findings highlight that AI can democratize access to sophisticated modeling and analysis, shifting the focus in scientific practice from technical mastery to the generation and testing of ideas, and reducing the time for completing research tasks from days to minutes.

The rapid emergence of artificial intelligence (AI) is changing the way science is done. As with many new tools (calculators, e-mail, internet, etc.), we usually begin by applying these tools to solve common tasks better than it is possible with existing tools (e.g., performing arithmetical operations with a calculator rather than a slide rule). However, the real power of new tools lies in enabling completely new uses, such as collaborative paper writing with colleagues anywhere in the world, enabled by the Internet. We are convinced that the use of AI in science will bring a plethora of new uses and capabilities, some of which are already apparent today. An example is that AI is democratizing science by enabling any reasonably qualified scientist to perform sophisticated modeling using highly specialized algorithms, without the need to master software packages. AI takes the role of an expert colleague who can understand the professor formulating the question and is also able to run the dedicated software, thus eliminating the middleman.

Here, we describe how we test AI abilities and new ways of interacting with the tool with three problems in quantum optics:

- i. A straightforward question; however, known to trick even mature physicists in the field.
- ii. A subtle problem with important applications that, while known for some years, is still a subject of current research.
- iii. A problem of current research with an unsettled solution.

Based on experience with these problems, we make observations regarding the possible utility of modern AI in the scientific process. In essence, every scientist now has access to sophisticated tools previously accessible only to specialists. This brings forward the importance of ideas rather than techniques. Of course, the speed with which ideas can be elaborated, tested in detail, and perhaps executed is higher. What used to take months and years can now be done in minutes.

We also remark on the striking similarity of the AI behavior to that of students.

We need to explain here what exactly we mean by AI. We test our problems on various state-of-the-art general-purpose models, accessible for public use (e.g., Gemini 2.5 Pro). When using different models, the details of the dialog are different; however, the overall results are broadly consistent across AI platforms. Therefore, unless we are specifically interested in comparing the models, we generically describe the interactions with AI in the following. The dialogs with the AI on which this paper is based can be found on a dedicated website https://manasp21.github.io/AI_QO_Appendix/index.html.

Problem one: state populations upon optical pumping

The first problem we discuss is a rather straightforward quantum optics question. However, in our experience, most physicists, including experts in the specific subfield, tend to give an incorrect answer, at least initially.

Formulation of the problem. Consider an ensemble of stationary atoms in the ground $J = 1$ state (see Fig. 1), in the absence of external fields. The atoms are optically pumped with monochromatic light propagating along z , which is linearly polarized along x . The excited state has angular momentum $J = 0$ and it decays spontaneously back to the ground state. There is no relaxation in the system other than this spontaneous decay. Initially, the ground state is unpolarized, so that the population is equally distributed among the ground-state sublevels.

Assuming that the initial population of each of the ground-state Zeeman sublevels is $1/3$, what are the populations after optical pumping is complete?

Human solution. Most physicists presented with this problem initially think that the populations of the $m_J = -1, 0$, and 1 sublevels are $0, 1$, and 0 , respectively (see Table 1, the first row in Final population). Indeed, atoms in the $m_J = 0$ cannot be excited by light, in contrast to the atoms in the other two sublevels. If an atom excited from a $|m_J| = 1$ sublevel spontaneously decays back into one of these two sublevels, it has a chance to get reexcited until it lands in the dark $m_J = 0$ state.

What is wrong with this solution is that it ignores the fact that there is not one but two dark states in the problem. The second dark state is an

equal-weight superposition of $m_j = -1$ and $m_j = 1$, which is

$$\frac{1}{\sqrt{2}}(|m_j = -1\rangle + |m_j = 1\rangle) \quad (1)$$

for x -polarized light.




Fig. 1 | Optical pumping in a $J = 1 \rightarrow J' = 0$ system. The figure illustrates the transition optically pumped with x -linearly polarized light. Spontaneous decay from the excited state back to the ground-state sublevels is the only relaxation process considered.

Table 1 | Populations of ground-state magnetic sublevels

State	$m_j = -1$	$m_j = 0$	$m_j = 1$
Initial population	1/3	1/3	1/3
Final population	0	1	0
	1/6	2/3	1/6
	1/4	1/2	1/4

The table summarizes the initial and several possible final populations for the sublevels, identified by their magnetic quantum number m_j . While all final distributions represent plausible educated guesses, only the last row is correct.

One way to see that there are indeed two dark states is to look at the same problem with the quantization axis chosen along the light polarization. In this case, light only drives the $\Delta m_j = 0$ transition, and the $|m_j| = 1$ sublevels are both dark.

Upon realization that only atoms in the bright superposition of the $|m_j| = 1$ sublevels orthogonal to the superposition (1), some physicists conclude that the final populations are 1/6, 2/3 and 1/6. However, this is wrong again, as this ignores the spontaneous decay of the excited state to the dark superposition.

Finally, the correct solution for the populations is 1/4, 1/2, and 1/4, which can be verified, once again, by turning to the picture with the quantization axis chosen along the light polarization: the final population of each of the dark states is equal to 1/2 of the total initial population, and the sole bright state is empty.

Interaction with AI. We posed the problem to the AI, and it immediately recognized that this is a question in the field of quantum optics. It proceeded to give us the 0,1,0 answer (see Table 1). While this answer is incorrect, it can be considered a reasonable first attempt. If a student gives such an answer and can motivate it, he or she will definitely earn points for this.

In the next step, we prompted the AI that we will try to explain that there is a flaw in the solution by solving a related problem, specifically the same problem but with the quantization axis along the light polarization. The AI solved this correctly. We then asked to write the density matrix and rotate it to the basis of the original problem, pointing out that this density matrix is different from the one originally presented by the AI.

The AI proceeded to acknowledge that it had been wrong and wrote: “Thank you! This detailed comparison has been very illuminating and has helped correct a flaw in one of my earlier consistency check analyses.”

We see that the interaction with the AI is quite similar to that with an excellent student or knowledgeable and intelligent colleague.

Problem Two: Resonant transitions between decaying states (the Burshtein effect)

Formulation of the problem. Consider two states, A and B , with energy gap Δ between them (Fig. 2a). At time $t = 0$, state A is populated and state B is empty. A resonant excitation field drives the transitions between the two states. The strength of the driving field is characterized by the

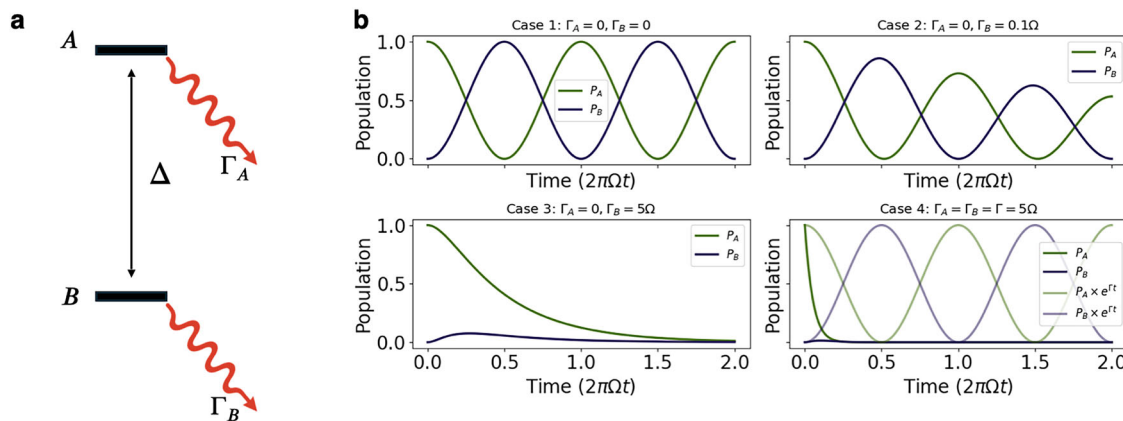


Fig. 2 | Dynamics of a resonantly driven two-level system with decay. **a** Two states, A and B , are coupled by a resonant field while also subject to decay with rates Γ_A and Γ_B , respectively. **b** The populations of states A and B are plotted for the Cases 1-4

described in the Formulation of the problem, showing undamped oscillations, damped oscillations, overdamped behavior, and the Burshtein effect, respectively. Time is presented as a dimensionless quantity (number of Rabi cycles).




Fig. 3 | Degenerate mirrorless lasing (DML). **a** The energy levels and transitions for the proposed DML scheme, based on the original proposals^{4,5}. **b** A typical experimental arrangement showing a pump laser interacting with an atomic vapor cell to generate forward or backward propagating DML. The thick diagonal

arrows represent light beams; the thinner vertical and horizontal arrows represent the linear polarization of the pump light and the generated mirrorless lasing light, respectively.

Rabi frequency Ω . We are interested in the population of state B as a function of time for several different cases for the relaxation rates of the two states Γ_A and Γ_B ; it is assumed that the states decay into unobserved states:

Case 1: $\Gamma_A = \Gamma_B = 0$;

Case 2: $\Gamma_A = 0$; $\Gamma_B \ll \Omega$;

Case 3: $\Gamma_A = 0$; $\Gamma_B \gg \Omega$;

Case 4: $\Gamma_A = \Gamma_B = \Gamma \gg \Omega$.

Human solution. Up to a point, this problem is a standard textbook problem of a two-level system in the presence or a resonant periodic perturbation (see, for example¹, Sec. 7.1). When the relaxation rates of states A and B are both zero (case 1), we have Rabi oscillations of the population between the two states [Fig. 2b, top left]. Adding a small decay for one of the states (case 2), the oscillations decay exponentially [Fig. 2b, top right]. When the decay rate becomes large compared to the Rabi frequency, the system is in the overdamped regime [Fig. 2b, bottom left] and there are no oscillations seen. We note that the population of B is zero at $t=0$ and never at intermediate times. Finally, we consider the most interesting case 4 [Fig. 2b, bottom right]. Again, as in case 3, the populations decay. However, and importantly, the solution remains oscillatory despite the rapid decay. To illustrate this, we multiply the solution by a growing exponential e^{Jt} . We see that this operation completely restores the solution to that of case 1. In particular, the population of B turns to zero at the same periodicity as for Rabi oscillations in case 1 (no decay). This phenomenon was pointed out by A.I. Burshtein and coworkers in the 1980s².

The Burshtein effect has important implications for metrology. If one measures the time evolution with excellent signal-to-noise and the states decay at the same rate, one can compensate for the exponential decay and obtain full information about the coupling of states A and B , which is impossible when the decay rates of A and B are vastly different (for example, as in case 3).

Interaction with AI. When we presented the problem of the two-level system with its varying decay scenarios to the AI, Gemini, it promptly recognized the context as a foundational one within quantum optics. The AI then proceeded to address the calculation of the population of state B by systematically tackling each of the four enumerated conditions. Its methodology for each case appeared to be:

- **Framing the physics:** The AI initially framed the problem in terms of coherent Rabi driving (Ω) counteracted by the specified incoherent decay channels (Γ_A, Γ_B).
- **Systematic case analysis:**
 - For **Case 1** ($\Gamma_A = \Gamma_B = 0$), it provided the standard $\sin^2(\Omega t/2)$ solution, characteristic of undamped Rabi oscillations.
 - In **Case 2** ($\Gamma_A = 0; \Gamma_B \ll \Omega$), its solution appropriately introduced an exponential decay factor enveloping the Rabi oscillations, reflecting the impact of the slow decay of state B .
 - For **Case 3** ($\Gamma_A = 0; \Gamma_B \gg \Omega$), its analysis correctly shifted to an overdamped perspective. It highlighted how the rapid decay of the state B suppresses the oscillations and severely limits the maximum population transferred.
 - In **Case 4** ($\Gamma_A = \Gamma_B = \Gamma \gg \Omega$), the AI described a scenario where strong decay in both states heavily dampens the coherent evolution, leading to only a minor, transient population in state B . **While the statements given by the AI thus far are reasonably correct, it failed to recognize the Burshtein effect.**

At this point, we decided to proceed as in the case of Problem One and convince the AI to re-examine and refine its solution. However, while we were doing this, an updated AI model was released, and it handled the problem perfectly from the first attempt, including an excellent discussion of the Rabi oscillation underneath the exponential decay envelope.

Problem three: degenerate mirrorless lasing

Formulation of the problem. Mirrorless lasing (ML) is a process of generation of a directed nearly monochromatic radiation, without any mirrors, by a medium (in our case, an ensemble of atoms) exposed to a pump laser beam. ML is usually associated with pencil-like pumping; it often occurs in both forward and backward direction with respect to the pump beam. Degenerate mirrorless lasing (DML) refers to a situation where the emitted radiation is at a frequency that is the same (or close to) the frequency of the pump light.

Consider a $J = 1 \rightarrow J' = 2$ transition [Fig. 3a]. There is no relaxation in the system except for the spontaneous decay from the upper state to the lower state. The atoms are stationary. The pumped volume is a cylinder of length L and radius R ; the atomic number density is n . Atoms are illuminated with resonant cw light with linear z polarization.

What are the conditions for DML in the forward and backward directions for light with x polarization? The axis of the cylinder is along the y -axis.

Human solution. As mentioned in the Introduction, this is actually an unsolved problem. Reference³ reports that backwards DML was observed in experiments in Ashtarak, Armenia; however, these results could not be reproduced in Mainz, Germany. The most recent theoretical analysis³ points out that the simplified theoretical treatment in the original proposal papers^{4,5} is insufficient in the high-light-power conditions that would be needed to achieve the desired lasing in a thermal vapor. A recent theoretical suggestion⁶ to use large detunings is awaiting experimental verification.

Interaction with AI. Given the formulation of the problem as a prompt, the AI immediately correctly identified the area of physics and found relevant literature. Its comments on the topic were somewhat general but nevertheless useful for someone not yet intimately familiar with the field. We noticed that it would not commit to a specific answer to the problem; however, it would make statements along the lines of “the backwards DML may occur under the right conditions”, going on to specify that the atomic density, pump laser power, geometry, etc., are the important factors. At this point, we asked AI the following: “I am a researcher and I want to set up an experiment to detect backwards DML in a thermal rubidium vapor. Can you help me design the parameters of the experiment?” The response was a detailed and substantive discussion, at a high professional level, of the design considerations, how one should choose the parameters, and what aspects one needs to pay particular attention to. We felt that this kind of advice was similar to what, say, a PhD student would get from a highly knowledgeable and experienced colleague. We remarked among ourselves that we should, indeed, recommend students to consult AI when preparing experiments and that this represents a powerful tool that was not available in pre-AI times (except for rare cases where a suitably knowledgeable, experienced, and patient human colleague could be found). Finally, we asked AI to suggest the reasons that backwards DML could be observed in one lab but could not be reproduced in another. Again, the response was thoughtful and professional, listing a number of possible explanations. However, for this latter question, it is hard to delineate whether the AI used its direct knowledge of the discussion in the literature or came up with the idea based on an original thought process. A look at the thinking process feature of the AI model indicates that the model has not explicitly used experimental papers to come up with its answers.

Discussion

Our experience as researchers in the field using AI to help navigate non-trivial questions in quantum optics, exemplified by the three problems above and many other similar experiences, convinces us that we have an unprecedented powerful tool at our disposal that has already changed the way we research a scientific topic, analyze discrepancies, and design new experiments.

As with any new tool, one needs to learn to use it properly. Although some people tend to put full trust in AI, it is, in fact, not an omniscient being, and the answers it gives are often wrong. The art of using AI is to take full advantage of it despite this. We conclude that the old human wisdom that points to the importance of asking the right questions equally applies to interactions with the AI. We derive the most benefit from treating AI as a multidisciplinary collaborator and not as a calculator or a handbook with ready answers.

Scientists have a toolbox for checking physics solutions. These include checking the units and limiting cases, solving similar problems where the solutions are known, solving the same problem in multiple ways, checking that the solution possesses the required symmetries (e.g., basis invariance as on our Problem One), satisfies conservation laws, etc. Not only can AI be used to perform these tests, it can be prompted (read: *taught*) to check its own answers and not to make the same type of mistake twice. Our experiences show that this approach can help AI overcome its own biases.

Conclusions

We found AI to be an impressively helpful and efficient colleague, while not immune to errors and misconceptions; however, eagerly ready to deal efficiently with them in cooperation with us, the user. The pace at which such an interaction progresses is already incomparably faster than purely human interactions, and the content is not exclusive to the availability of knowledge on the Internet. The time needed for intellectual synthesis has also been significantly reduced. Although we have not attempted a scientific quantification, we feel that a discussion that would normally take days can be compressed to less than an hour.

Manas Pandey ¹, **Bharath Hebbe Madhusudhana**², **Saikat Ghosh**¹ & **Dmitry Budker** ^{3,4,5,6} 

¹Indian Institute of Technology, Kanpur, Uttar Pradesh, India. ²MPA-Quantum, Los Alamos National Laboratory, Los Alamos, NM, USA.

³Johannes Gutenberg-Universität Mainz, Mainz, Germany. ⁴Helmholtz Institute Mainz, Mainz, Germany. ⁵GSI Helmholtzzentrum für Schwerionenforschung GmbH, Darmstadt, Germany. ⁶Department of Physics, University of California, Berkeley, CA, USA.

 e-mail: budker@uni-mainz.de

Received: 16 July 2025; Accepted: 25 October 2025;

Published online: 07 November 2025

References

1. Auzinsh, M., Budker, D. & Rochester, S. *Optically Polarized Atoms: Understanding Light-atom Interactions* <https://books.google.co.in/books?id=btol6sg2vJoC> (OUP Oxford, 2010).
2. Burshtein, A. & Storozhev, A. Transition width between two metastable states. *Chem. Phys.* **119**, 1–13 (1988).
3. Ramaswamy, A. et al. Mirrorless lasing: a theoretical perspective. *Optical Mem. Neural Netw.* **32**, S443–S466 (2023).
4. Gazaryan, E., Grigoryan, G. & Papoyan, A. Amplification of a weak circularly polarized light signal in a multilevel atomic medium. *Izv. Natl. Akad. Nauk Armen. Fiz.* **46**, p. 227–233 (2011).
5. Movsisyan, M., Shmavonyan, S. & Papoyan, A. Amplification of radiation in atomic vapor induced by a linearly polarized laser radiation. *Open Phys.* **9**, 948–955 (2011).
6. Ramaswamy, A. & Malinovskaya, S. A. Degenerate mirrorless lasing in thermal alkali vapors. arXiv:2503.14705 <https://doi.org/10.48550/arXiv.2503.14705> (2025).

Acknowledgements

We thank Dr. Anahit Gogyan, Prof. Stefan Kramer, Marina Gil Sendra, Andrew Budker, Prof. Derek Jackson Kimball, Prof. Alexej Jerschow, Prof. Ilya Kuprov, Prof. Mikhail Kozlov, Prof. Piet O. Schmidt, Prof. Ferdinand Schmidt-Kaler, Prof. Svetlana Malinovskaya, Prof. Mark Raizen, Dr. Aneesh Ramaswami, Dr. Aleksandra Ziolkowska, Erin Burgard, Dr. Andreas Trabesinger, and others for helpful discussions. Research presented in this article was supported in part by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20230779PRD1 and as an initiative under the National Quantum Mission of DST.

Author contributions

M.P. and D.B. initiated the project based on their discussion of the novel capabilities of AI; M.P. conducted the interactions with AI and documented the results; M.P., B.H.M., S.G., and D.B. contributed ideas, discussed the problems, wrote and edited the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02398-2>.

Correspondence and requests for materials should be addressed to Dmitry Budker.

Peer review information *Communications Physics* thanks Nikola Šibalić and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. [A peer review file is available].

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025