

454 A MolPuzzle Benchmark Details

455 This section complements Section 3 with a fine-grained summary of the dataset collection, results
456 validation, and evaluation procedure, along with a fuller characterization of the task instances and the
457 corresponding prompts.

458 A.1 Data Collection

459 The initial molecules were selected by referencing the textbook *Organic Structures from Spectra, 4th*
460 *Edition*, available as an online PDF on ResearchGate. We chose 234 molecules based on spectrum
461 tasks involving IR, MS, $^1\text{H-NMR}$, and $^{13}\text{C-NMR}$ to reflect a difficulty level suitable for graduate
462 students[36].

463 To address copyright concerns, we excluded molecules with publicly available mass spectrometry
464 (MS) spectra in open-source databases from our study. The remaining spectra were sourced from
465 public resources, notably the PubChem database[37]. For additional spectra that were not available,
466 we used simulation methods[38][39] and provided a Jupyter notebook to generate these data, ensuring
467 high-quality spectra for analysis. Our final dataset comprised 200 molecules.

468 Given the challenges associated with NMR spectrum images, some spectra were obtained from
469 simulated data in text format for $^1\text{H-NMR}$ and $^{13}\text{C-NMR}$. This approach ensured clarity and accuracy
470 in the evaluation of molecular structures.

471 To assess the multiple-stage abilities of LLMs, we designed a unique question-and-answer evaluation.
472 This framework tested the LLMs' capabilities in interpreting and integrating data from different types
473 of spectra, simulating real-world challenges. Details of this evaluation framework are provided in the
474 next section.

475 A.2 Template design

476 Each template was crafted to target specific skills within molecular understanding. For instance,
477 saturation identification challenges the models' ability to discern the degree of saturation in a molecule,
478 which is crucial for understanding its chemical reactivity and stability. Aromatic ring identification
479 tests the models' ability to recognize benzene-like structures, which are fundamental in organic
480 chemistry due to their common occurrence and unique properties. Saturation degree calculation
481 pushes the models to apply quantitative analysis, requiring not just recognition but also computation
482 based on molecular structures.

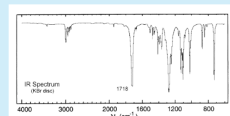
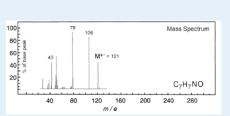
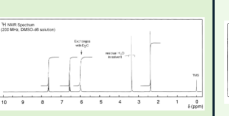
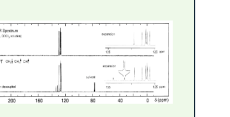
483 By diving deeper into the rationale behind each template and the kind of chemical knowledge they
484 are designed to test, we can better appreciate how these tasks simulate real-world applications in
485 chemistry. This approach not only tests the models' basic recognition abilities but also their capacity
486 to perform complex reasoning and apply theoretical knowledge practically. The template examples
487 are in A.3

488 **A.3 Stage1 QA Samples**

Table 3: QA samples for the molecule understanding task

| Task | Prompt |
|---------------------------------|---|
| Saturation Identification | Question: Could the molecule with the formula C ₈ H ₁₀ O potentially be Saturated? Answer: No Model response: No. |
| Aromatic Ring Identification | Question: Could the molecule with the formula C ₈ H ₁₀ O have aromatic rings? Answer: Yes Model response: Yes. |
| Functional Group Identification | Question: Could the molecule with the formula C ₆ H ₁₄ O ₂ potentially contain a Amine group, given the Degree of Unsaturation is 0.0? Answer: No Model response: No, the molecule doesn't contain Amine group |
| Saturation Degree Calculation | Question: Calculate the Degree of Unsaturation of the molecule with the formula C ₈ H ₁₀ O? Answer: 4.0 Model response: 2 |

489 **A.4 Stage2 QA Samples**

| IR Interpretation | MASS Interpretation | H-NMR Interpretation | C-NMR Interpretation |
|--|--|---|---|
|  <p>Question: Does the IR spectrum contains broad absorption peak of N-H stretching around 3200-3600 cm⁻¹?</p> |  <p>Question: Examine the MASS spectrum to determine if the molecule could potentially contain specific fragments: Ether.</p> |  <p>Question: Examine the H-NMR spectrum to determine if the molecule could potentially contain specific functional groups: Phenol?</p> |  <p>Question: Examine the C-NMR spectrum to determine if the molecule could potentially contain specific fragments: Ester.</p> |
| Answer: No Model response: No | Answer: No Model response: Yes | Answer: No Model response: No | Answer: No Model response: Yes |

490 **A.5 Stage3 QA Samples**

Table 4: QA samples for the molecule construction task

| Task | Prompt |
|-------------------|---|
| H-NMR Elucidation | Question: Calculate the number of different types of hydrogen atoms present in the molecule, based on the given H-NMR: 4.51-4.61 (4H, 4.56 (s), 4.56 (s)), 7.06-7.32 (10H, 7.13 (dddd, J = 7.9, 7.7, 1.8, 0.6 Hz), 7.13 (dddd, J = 7.9, 7.7, 1.8, 0.6 Hz), 7.25 (dddd, J = 7.9, 1.5, 1.3, 0.6 Hz), 7.25 (dddd, J = 7.9, 1.5, 1.3, 0.6 Hz), 7.26 (tt, J = 7.7, 1.5 Hz), 7.26 (tt, J = 7.7, 1.5 Hz)) Answer: 4 Model response: 3. |
| C-NMR Elucidation | Question: Analyze the given C-NMR data and determine the number of different types of carbon atoms present in the molecule, based on given C-NMR: 39.3 (1C, s), 63.4 (1C, s), 127.8 (1C, s), 128.4 (2C, s), 128.8 (2C, s), 134.2 (1C, s). Only output the number. Answer: 6 Model response: 8 |

491 **B Evaluation Experiments**492 **B.1 Experimental Setting**

493 During our testing phase, we selected 100 questions and employed two distinct prompting strategies
494 with the large language model (LLM). Initially, the LLM was tasked with directly answering the

495 questions. In a subsequent approach, the same queries were presented, but the model was prompted to
496 execute a chain-of-thought reasoning process before responding. Each question in our dataset begins
497 with a comprehensive description of the chemical context, along with specified answer formats and
498 detailed guiding rules. To ensure a balanced representation of each task category, for tasks in Stage 1,
499 the distribution ratio for Saturation Identification (SI), Functional Group Identification (FI), Aromatic
500 Ring Identification (AI), and Saturation Degree Calculation (SC) is set at 2:3:3:2. In Stage 2, we
501 have randomly selected 100 questions from each category of the spectrum. For Stage 3, we randomly
502 selected 100 questions focused on H-NMR and C-NMR analyses.

503 We carried out this evaluation over three rounds, analyzing responses using both accuracy and the
504 F1 score for tasks involving Saturation Identification (SI), Functional Group Identification (FI), and
505 Aromatic Ring Identification (AI). For Saturation Degree Calculation (SDC), which yields numerical
506 results, we assessed accuracy by comparing the count of correct matches to the ground truth data.
507 The detailed results are reported in Table [A.3](#). To ensure that all results are presented in a way that
508 facilitates direct comparison, only those using similar evaluation metrics (AI, FI, AI) are included
509 in the main table. For the SI, AI, and FI tasks, we use the F1 score and Accuracy to evaluate their
510 performance since they are classification tasks. For the SDC task, the answer is a numerical number,
511 so we only use the accuracy score to measure the performance of the LLMs. This approach helps to
512 keep the evaluation coherent and focused on comparable data points.

513 **B.2 Human Evaluation**

514 To evaluate the performance of large language models (LLMs) on specialized tasks against expert
515 humans, we recruited six graduate students from chemistry department to solve the MolPuzzle
516 benchmark. These students, having recently completed a graduate-level course in Molecular Structural
517 Elucidation, represented a highly skilled group of human participants.

518 For the experiment, we randomly selected six questions from the MolPuzzle dataset for each stage of
519 the study. These questions were presented to the students in different formats according to the stage:
520 In Stages 1 and 2, the questions were simple Yes/No or required short answers. In Stage 3, to align
521 with the conventional methods chemists use to express chemical structures, students were asked to
522 upload images of their hand-drawn structures instead of using SMILES strings. These images were
523 manually compared to the ground truth to calculate scores.

524 We also imposed self-regulated time constraints to mirror the challenging nature of molecular
525 structural elucidation. Beyond individual stage evaluations, we presented each participant with a
526 complete molecule puzzle, consisting of a formula and four spectral images. The students were tasked
527 with solving these puzzles within a 20-minute time frame. Impressively, all participants successfully
528 submitted their solutions within the allotted period.

529 Our study included a component where human evaluators were involved to assess the performance
530 of the AI models. To ensure the protection and ethical treatment of all participants, we conducted a
531 thorough risk assessment. Potential risks identified included privacy concerns and the mental strain
532 of repetitive tasks. Mitigation strategies, such as ensuring anonymity and providing breaks, were
533 implemented to protect our evaluators.

534 The study was submitted for review and received approval from our Institutional Review Board (IRB).
535 The IRB approval number is [insert approval number], which verifies that our protocols met all ethical
536 guidelines for research involving human subjects. Throughout the project, we adhered strictly to
537 these protocols to ensure ongoing compliance with ethical standards.

538 **B.3 Stage1**

539 Molecule understanding requires comprehensive analysis and interpretation of molecular structures,
540 with a focus on chemical properties and spectroscopic data. In our study, we created a dataset of
541 234 molecules and developed eight distinct question templates across four categories: **Saturation**
542 **Identification(SI), Functional Group Identification(FI), Aromatic Ring Identification(AI), and**

543 **Saturation Degree Calculation(SC)**. These templates assess the ability to identify substructures,
 544 compute saturation levels, and infer structural presence, incorporating concepts in the chemistry
 545 reasoning process. Each question also necessitates a deep understanding of molecular bonding,
 546 stereochemistry, and functional group identification. Responses were generated using the RDKit
 547 library, ensuring precise and reliable answers grounded in established chemical informatics.

Table 3: The accuracy(\uparrow), F1 score(\uparrow) in 4 different molecule understanding categories, the best LLMs are in bold font.

| Model | CoT | SI | | AI | | FI | | SC |
|---------------|-----|--------------|--------------|--------------------|--------------------|--------------------|-------------|--------------------|
| | | F1 | Acc | F1 | Acc | F1 | Acc | Acc |
| GPT-4o | - | 1±0.0 | 1±0.0 | 0.943±0.016 | 0.944±0.015 | 0.934±0.005 | 0.966±0.0 | 0.667±0.003 |
| GPT-4o | ✓ | 1±0.0 | 1±0.0 | 0.911±0.031 | 0.911±0.031 | 0.689±0.025 | 0.766±0.027 | 0.816±0.062 |
| GPT-3.5 | - | 0.451±0.025 | 0.825±0.075 | 0.816±0.017 | 0.816±0.075 | 0.826±0.075 | 0.683±0.016 | 0.5±0.099 |
| GPT-3.5 | ✓ | 0.448±0.026 | 0.816±0.008 | 0.798±0.025 | 0.800±0.027 | 0.526±0.053 | 0.622±0.031 | 0.533±0.131 |
| Claude-3-opus | - | 0.361±0.009 | 0.556±0.023 | 0.988±0.015 | 0.988±0.015 | 0.934±0.001 | 0.966±0.001 | 0.856±0.016 |
| Claude-3 | ✓ | 0.760±0.189 | 0.903±0.046 | 0.878±0.025 | 0.867±0.001 | 0.547±0.112 | 0.843±0.081 | 0.900±0.025 |
| Gemini-pro | - | 0.285±0.020 | 0.399±0.040 | 0.775±0.093 | 0.788±0.083 | 0.646±0.052 | 0.748±0.051 | 0.200±0.004 |
| Gemini-pro | ✓ | 0.391±0.045 | 0.651±0.108 | 0.685±0.088 | 0.688±0.087 | 0.562±0.018 | 0.629±0.023 | 0.283±0.062 |
| LLama3 | - | 0.367±0.018 | 0.583±0.047 | 0.490±0.030 | 0.533±0.027 | 0.472±0.133 | 0.588±0.0 | 0.0±0.0 |
| LLama3 | ✓ | 0.473±0.011 | 0.899±0.040 | 0.384±0.026 | 0.533±0.0 | 0.570±0.035 | 0.799±0.047 | 0.017±0.001 |
| Vicuna-13b | - | 0.031±0.022 | 0.033±0.025 | 0.500±0.087 | 0.522±0.083 | 0.308±0.038 | 0.311±0.041 | 0.0±0.0 |
| Vicuna-13b | ✓ | 0.380±0.023 | 0.616±0.062 | 0.342±0.006 | 0.522±0.157 | 0.516±0.080 | 0.855±0.016 | 0.0±0.0 |
| Mistral-7b | - | 0.221±0.014 | 0.283±0.025 | 0.384±0.005 | 0.500±0.0 | 0.319±0.014 | 0.322±0.157 | 0.0±0.0 |
| Mistral-7b | ✓ | 0.433±0.007 | 0.766±0.023 | 0.342±0.006 | 0.522±0.016 | 0.601±0.102 | 0.877±0.031 | 0.0±0.0 |

548 B.4 Stage2

549 The Spectrum interpretation tasks mainly measure the capability of LLMs in analyzing images
 550 related to identifying key substructures indicated by the spectrum plot. In this study, we utilize
 551 four distinct types of spectral images: nuclear magnetic resonance (NMR), infrared spectroscopy
 552 (IR), mass spectrometry, and others. Each type of data offers insights into various aspects of the
 553 molecular structure. We’ve created specific question templates for each spectrum, targeting peak
 554 and substructure identification factors. These templates are designed manually and emphasize the
 555 intricate connection between the spikes or troughs in the figures and the structures of the molecules.
 556 Responses were generated using the RDKit library to ensure correctness.

557 The findings from Stage 2 are presented in Table 4. We exclusively focus on the zero-shot learning
 558 outcomes, as our observations indicate that implementing chain-of-thought prompting leads to a
 559 deterioration in model performance. To address this limitation, we offer qualitative insights in B.6.

Table 4: The accuracy(\uparrow), F1 score(\uparrow) for IR, MASS spectrum, H-NMR, and C-NMR interpretation tasks. "-" means the results are not interoperable

| Model | Stage-2 Tasks | | | | | | | |
|---------------------|--------------------|-------------------|---------------------|--------------------|----------------------|--------------------|----------------------|--------------------|
| | IR Interpretation | | MASS Interpretation | | H-NMR Interpretation | | C-NMR Interpretation | |
| | F1 | Acc | F1 | Acc | F1 | Acc | F1 | Acc |
| GPT-4o | 0.656±0.052 | 0.713±0.06 | 0.609±0.042 | 0.767±0.042 | 0.618±0.026 | 0.864±0.007 | 0.639±0.107 | 0.892±0.049 |
| Claude-3-opus | 0.440±0.006 | 0.476±0.055 | 0.398±0.032 | 0.466±0.019 | 0.572±0.190 | 0.842±0.017 | 0.554±0.075 | 0.716±0.042 |
| Gemini-3-pro-vision | 0.194±0.002 | 0.119±0.016 | 0.116±0.036 | 0.124±0.038 | 0.545±0.048 | 0.851±0.062 | 0.492±0.016 | 0.619±0.044 |
| LLava1.5-8b | 0.256±0.026 | 0.414±0.044 | 0.101±0.021 | 0.104±0.26 | 0.118±0.008 | 0.186±0.011 | 0.254±0.015 | 0.472±0.023 |
| Qwen-VL-Chat | 0.243±0.027 | 0.392±0.043 | 0.125±0.006 | 0.116±0.021 | 0.255±0.007 | 0.611±0.031 | - | - |
| InstructBLIP-7b | 0.239±0.020 | 0.263±0.014 | 0.101±0.021 | 0.104±0.26 | - | - | 0.044±0.006 | 0.064±0.023 |
| InstructBLIP-13b | 0.239±0.020 | 0.263±0.014 | 0.101±0.021 | 0.104±0.26 | - | - | 0.047±0.014 | 0.067±0.025 |

560 B.5 Stage-3

561 Constructing a molecule involves a detailed analysis of NMR data, which is critical for understanding
 562 its structure. H-NMR data are essential as they provide information about the hydrogen environments
 563 within the molecule, including the number and types of hydrogen atoms (such as aliphatic or

564 aromatic), as well as their connectivity. Conversely, C-NMR data offer in-depth insights into the
 565 carbon framework, illustrating the distribution and linkage of carbon atoms within the molecule.
 566 In our study, to evaluate the ability of large language models (LLMs) to interpret NMR data, we
 567 generated 1,171 question-and-answer (QA) pairs. These pairs focus on key NMR interpretation tasks,
 568 such as counting hydrogen atom types and identifying substructures, which are critical for accurate
 569 analysis.

570 Despite observing moderate accuracy from the LLMs in Stage 2 of our testing, we enhanced the
 571 quality of the QA pairs in Stage 3 by providing the LLMs with verified NMR data, generated by using
 572 nmrdB [40]. This approach ensures that the data used is reliable and helps maintain the integrity of
 573 our results. The findings from Stage 2 are presented in Table. We exclusively focus on the zero-shot
 574 learning outcomes, as our observations indicate that implementing chain-of-thought prompting leads
 575 to a deterioration in model performance. To address this limitation, we offer qualitative insights in

Table 5: The F1 score(\uparrow) for H-NMR, and C-NMR Structure Elucidation

| Method | H-NMR Elucidation | C-NMR Elucidation |
|---------------|--------------------|--------------------|
| GPT-4o | 0.433±0.013 | 0.411±0.034 |
| Claude-3-opus | 0.395±0.008 | 0.313±0.029 |
| Gemini-pro | 0.333±0.012 | 0.308±0.031 |
| Llama3 | 0.211±0.012 | 0.342±0.007 |
| Vicuna-13b | 0.181±0.013 | 0.244±0.001 |
| Mistral-7b | 0.131±0.032 | 0.122±0.027 |

576 B.6 Qualitative Results

577 In this section, we present several examples using GPT-4’s chain-of-thought (CoT) reasoning to
 578 facilitate a clearer understanding of the results. We have enlisted two Ph.D. candidates from the
 579 chemistry department to evaluate these CoT outcomes. The analysis uses color coding to indicate the
 580 accuracy of the generated text: green signifies correct responses, red indicates incorrect ones, and
 581 yellow denotes responses that are partially correct.

582 B.6.1 Stage 2 examples

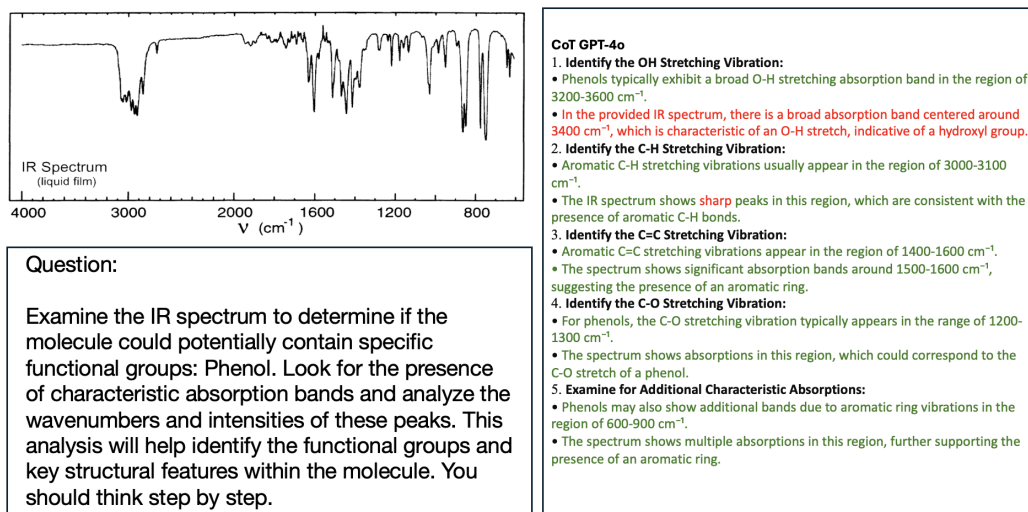


Figure 6: Human annotated IR spectrum interpretation

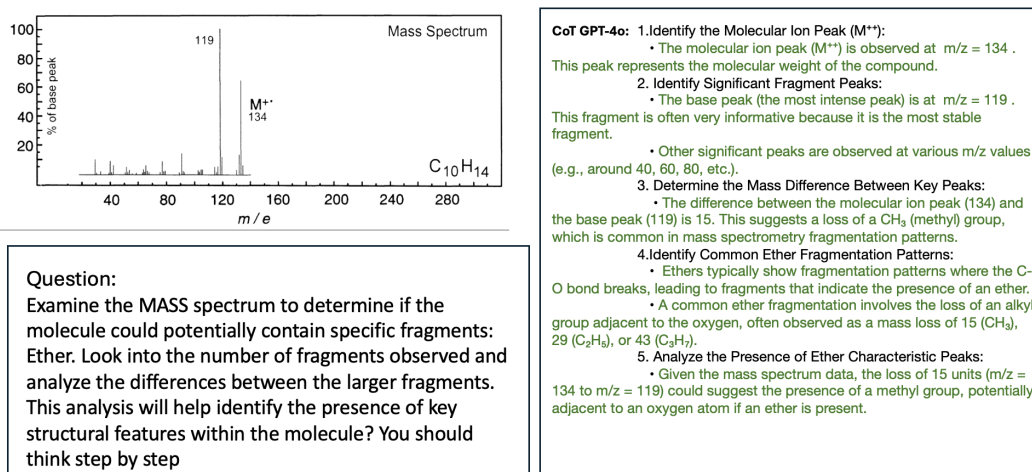


Figure 7: Human annotated MASS spectrum interpretation

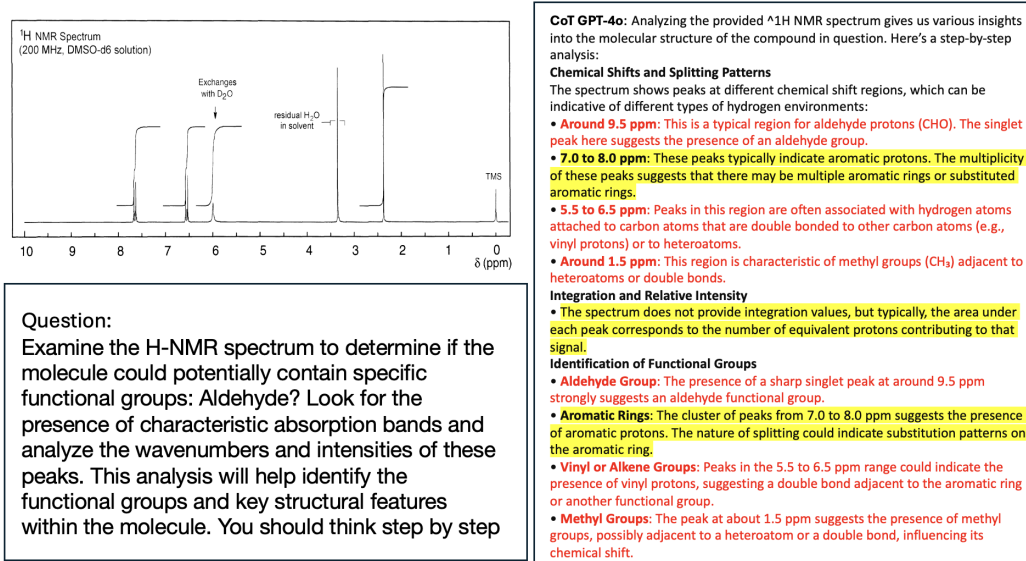


Figure 8: Human annotated H-NMR spectrum interpretation

583 B.6.2 Stage 3 examples

584 Based on the observation in spectrum interpretation, it is safe to say that the GPT-4o model has a
 585 basic grasp of the concept and the task at hand, but their deficiency mainly resides in obtaining fully
 586 correct information from the spectrum images and also they lack in-depth understanding. As for the
 587 problems in Stage 3, although the GPT-4o outputs are mostly correct, they are too general which
 588 shows the poor understanding of GPT-4o models in solving this intricate task. The stage 3 task can
 589 be roughly broken down into 3 subtasks: obtaining the correct information from the spectrum image,
 590 deducing the correct structural information from the spectral information, and finally translating this
 591 structural information into a correct molecular structure. GPT-4o models seem to perform well in the
 592 second subtask, and moderately for simple structures in the third subtask but seem to be especially
 593 struggling with the first subtask in the case of NMR spectra. This indicates the gap in current LLMs
 594 in fully interpreting data therefore more advanced models and approaches should be developed to
 595 tackle the problem.

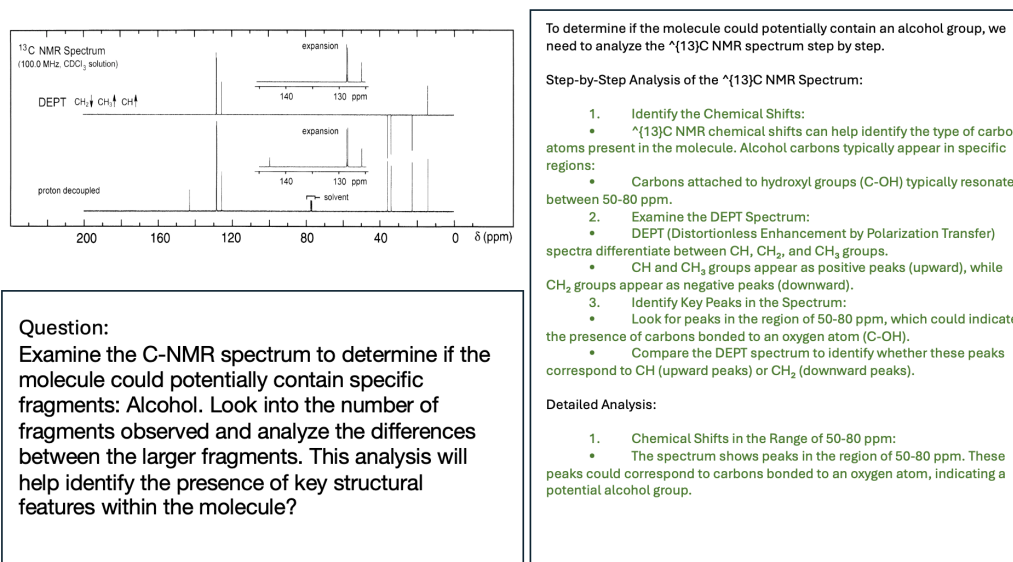


Figure 9: Human annotated C-NMR spectrum interpretation

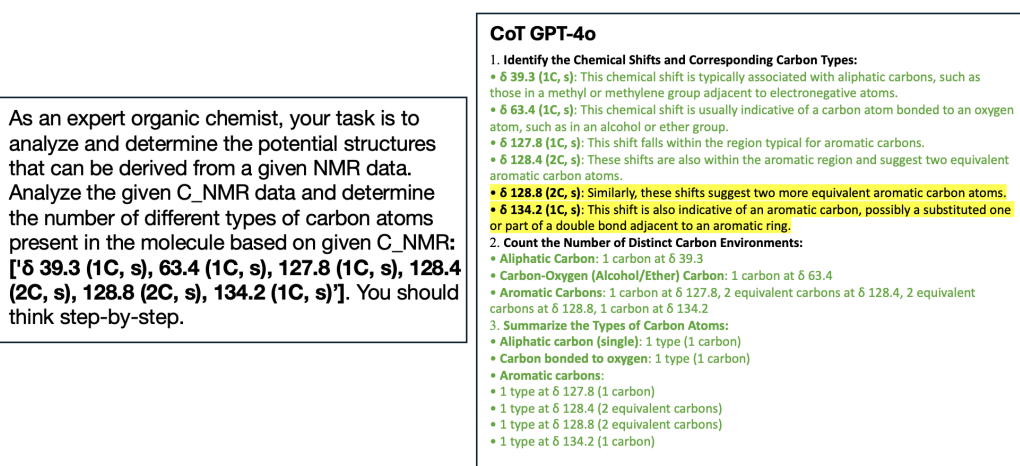


Figure 10: Human annotated H-NMR Elucidation

596 B.6.3 Complex Molecules

597 In addition to presenting molecules extracted from textbooks, we also demonstrate how the large
 598 language model (LLM) handles complex molecular structures. As illustrated in Figure 12, complex
 599 molecules typically have a larger pool of fragments. This expansion results in a greater number
 600 of valid elucidation paths, complicating the selection process for an appropriate starting point.
 601 Successfully navigating this enlarged pool necessitates an in-depth understanding of each fragment's
 602 properties and the associated, more intricate NMR data. In this context, LLMs may struggle because
 603 they often lack the nuanced chemical intuition and detailed analytical capabilities that human experts
 604 possess. Such limitations can lead to inaccuracies in interpreting complex interactions within NMR
 605 spectra, making LLMs less reliable for this task.

As an expert organic chemist, your task is to analyze and determine the potential structures that can be derived from a given NMR data. Generate the number of different types of hydrogen atoms present in the molecule based on given H_NMR: [δ 2.23 (3H, s), 3.79 (3H, s), 7.04 (2H, ddd, J = 8.3, 1.3, 0.4 Hz), 7.97 (2H, ddd, J = 8.3, 1.6, 0.4 Hz)]. You should think step by step.

CoT GPT-4o

1. Identify the Chemical Shifts and Corresponding Hydrogen Types:

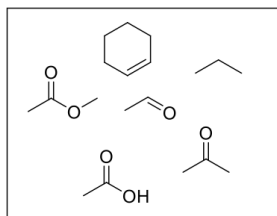
- δ 2.23 (3H, s): This singlet indicates 3 equivalent hydrogen atoms. These are likely part of a methyl group (CH₃) that is not split by neighboring hydrogens.
- δ 3.79 (3H, s): This singlet also indicates 3 equivalent hydrogen atoms, suggesting another methyl group (CH₃), possibly adjacent to an electronegative atom like oxygen (e.g., methoxy group -OCH₃).
- δ 7.04 (2H, ddd, J = 8.3, 1.3, 0.4 Hz): This complex splitting pattern (doublet of doublet of doublets) indicates 2 equivalent hydrogen atoms in an aromatic environment, with coupling constants suggesting interactions with neighboring hydrogens.
- δ 7.97 (2H, ddd, J = 8.3, 1.6, 0.4 Hz): Another complex splitting pattern, also indicating 2 equivalent hydrogen atoms in an aromatic environment, with similar coupling constants indicating interactions with neighboring hydrogens.

2. Determine the Number of Different Types of Hydrogen Environments:

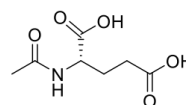
- **Aliphatic Hydrogens:**
 - 3 equivalent hydrogens at δ 2.23 (likely CH₃ group)
 - 3 equivalent hydrogens at δ 3.79 (likely CH₃ group, possibly part of a methoxy group)
- **Aromatic Hydrogens:**
 - 2 equivalent hydrogens at δ 7.04
 - 2 equivalent hydrogens at δ 7.97
 - The molecule has a total of 4 distinct types of hydrogen atoms:
- **Methyl group hydrogens:**
 - 3 equivalent hydrogens at δ 2.23 (CH₃)
 - 3 equivalent hydrogens at δ 3.79 (likely part of a methoxy group, -OCH₃)
- **Aromatic hydrogens:**
 - 2 equivalent hydrogens at δ 7.04
 - 2 equivalent hydrogens at δ 7.97

Figure 11: Human annotated C-NMR Elucidation

Fragment pool:



Target Molecule



H_NMR: 1.85-1.97 (5H, 1.91 (q, J = 7.3 Hz), 1.91 (q, J = 7.3 Hz), 1.91 (s)), 2.39-2.51 (2H, 2.45 (t, J = 7.4 Hz), 2.45 (t, J = 7.4 Hz)), 4.29 (1H, t, J = 7.3 Hz)

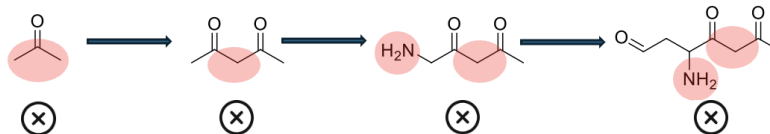


Figure 12: Complex molecule Structure Elucidation

606 C Compute Resources

607 For the execution of various models in our experiments, distinct compute resources were utilized
 608 based on the model's accessibility and computational requirements. Specifically, for models like
 609 Claude 3, GPT, and Gemini, we employed API calls to facilitate their operation, leveraging the
 610 existing infrastructure provided by their respective platforms. This approach allowed us to access
 611 these models without the need for local computational resources, thereby streamlining the process.
 612 Conversely, for all other open-sourced models employed in our study, we conducted the experiments
 613 locally using an NVIDIA A100 GPU. This high-performance computing unit was chosen due to its
 614 advanced capabilities in handling extensive computations and large model requirements efficiently.