[2026 ICLR]

IR-Agent: Expert-Inspired LLM Agents for Structure Elucidation from Infrared Spectra

Heewoong Noh, Namkyeong Lee, Gyoung S. Na, Kibum Kim, Chanyoung Park*

Korean Advanced Institute of Science and Technology (KAIST)
Korea Research Institute of Chemical Technology (KRICT)

*Corresponding Author

DSAIL Data Science & Artificial Intelligence

Background Infrared Spectroscopy

Infrared Spectroscopy is an analytical technique that measures the interaction of infrared light with matter to identify and characterize molecular structures.

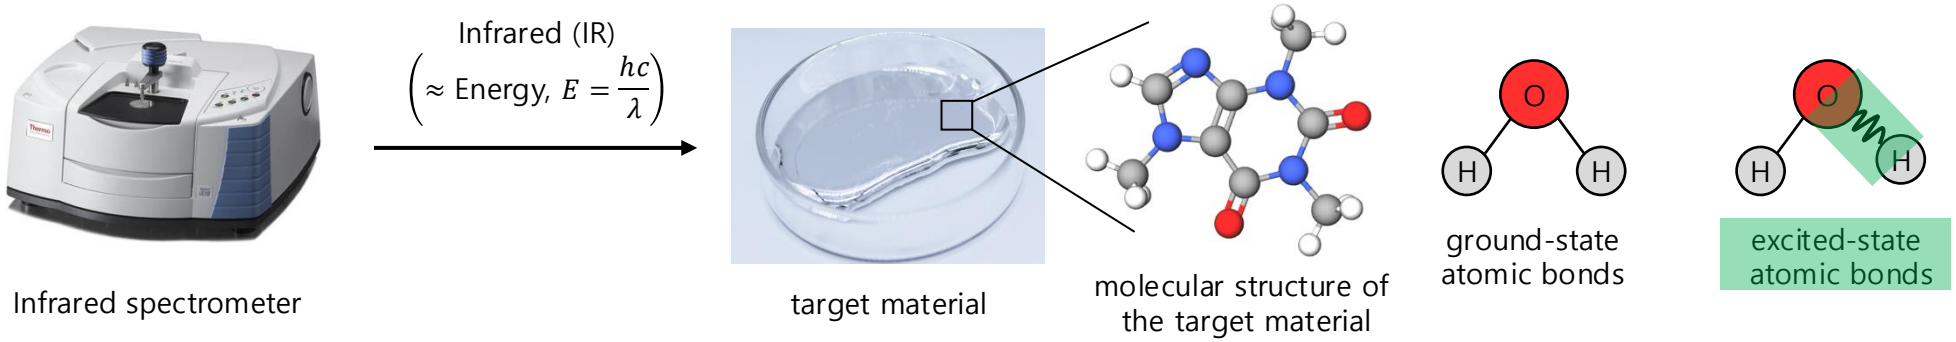

Infrared spectrometer

target material

molecular structure of the target material

ground-state atomic bonds

excited-state atomic bonds

- Specific atomic bonds vibrate only to specific energy levels
- Each bond absorbs infrared energy(light) and vibrates when exposed to it

  → Through IR spectra analysis, information about functional groups and atomic bonds can be obtained

Background Molecular Structure Elucidation using Infrared Spectroscopy

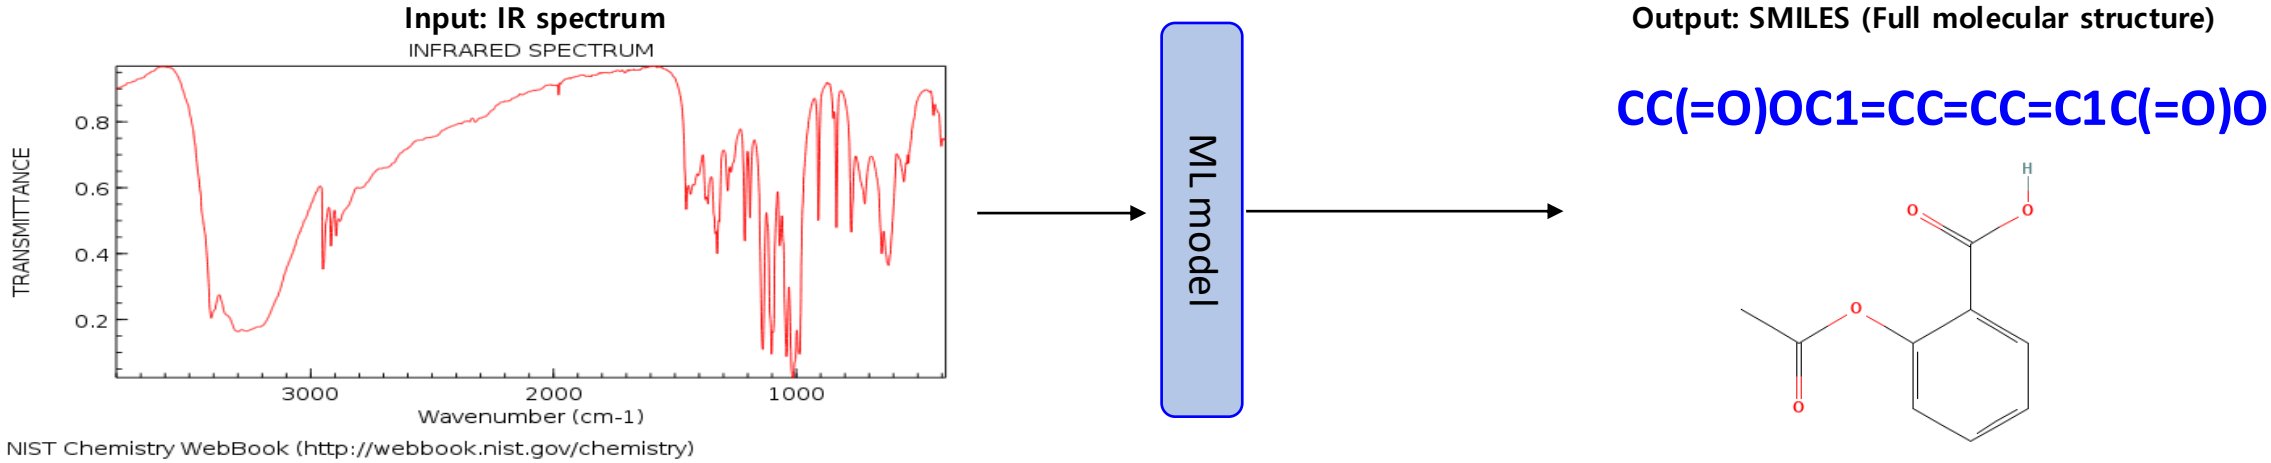

Task: Predicting the molecular structure of unknown materials using their IR spectra

Previous Approach

- Chemists manually find peaks of IR spectrum and interpreting the IR Absorption table (experience&intuition)
- Limited to the functional group detection task (the majority of ML approaches)
- Limited to Transformer based modeling (sequence to sequence)

Why is it challenging?

- Complexity and overlapping nature of spectra features → hard to interpret

Motivation Emulating Expert Analyses & Flexible Framework

- For IR spectra analysis, comprehensively integrating the analytical processes used by human experts is essential
  - Experts interpret IR absorption tables to infer **local substructures** and **bonding patterns** from peak positions.
  - They retrieve structurally similar molecules from spectral databases to provide **global contextual clues**

- IR spectra are often accompanied by **diverse chemical information**, such as atom types, scaffolds, and carbon counts.
  - Existing methods struggle to flexibly incorporate this information, requiring the redesign, retraining, and revalidation of new model architectures

  → **LLM Agent framework**

A new framework for IR spectra analysis is needed, one that
(1) emulates the integrated workflow of human experts and
(2) flexibly incorporates a wide range of chemical information

Preliminaries Problem Setup

Task Description

$\mathcal{X} \in \mathbb{R}^{1 \times L}$: IR spectrum of a molecule

: number of absorbance values corresponding to wavenumber positions

IR-Agent predicts the SMILES representation of the molecule given the input $\mathcal{X}$

Assumption

Predicting SMILES from **only IR spectra**, without ground-truth chemical formula of unknown material

While some studies assume that ground-truth chemical formula is always available along with the IR spectrum

→ In practice, obtaining the exact formula of an unknown material is often unrealistic

- Although mass spectrometry (MS) is commonly employed, (1) costly and time-consuming, (2) difficult to interpret, and (3) still leaves the derivation of an exact formula as a highly non-trivial challenge

Preliminaries Problem Setup

Tools

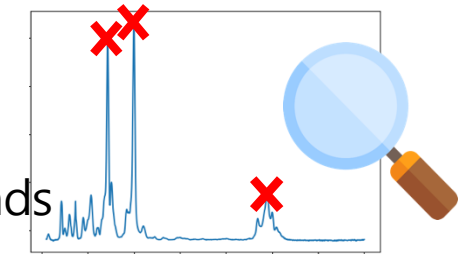

IR Peak Table Assigner: extracts the peaks from the spectrum and finds relevant substructures from the IR absorption table

IR Spectra Retriever: retrieves IR spectra from the IR Spectra Database that are similar to a given input spectrum

External Knowledge

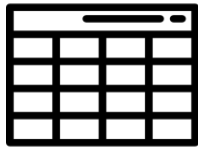

Wavenumber (cm^{-1})	Substructure
3700–3584	alcohol (O–H)
3550–3200	alcohol (O–H)
3500–3400	primary amine (N–H)
3400–3300	aliphatic primary amine (N–H)
3330–3250	aliphatic primary amine (N–H)
3350–3310	secondary amine (N–H)
3100–2900	carboxylic acid (O–H)
3200–2700	alcohol (O–H)
3000–2800	amine salt (N–H)

IR Absorption Table: summarizes the characteristic absorption frequencies associated with different molecular functional groups/substructures

IR Spectra Database contains a variety of IR spectra along with their corresponding molecules in SMILES format

Preliminaries IR Spectra Translator

IR Spectra Translator

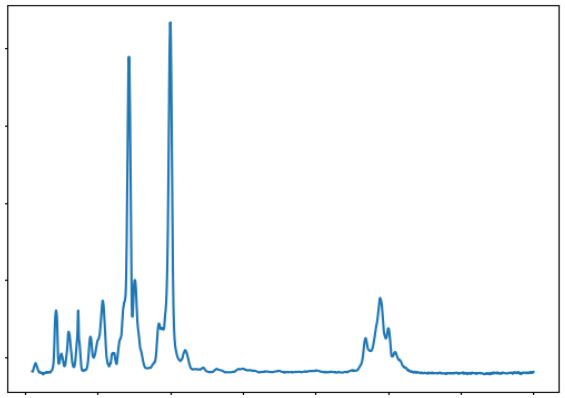

IR Spectrum $\mathcal{X}$

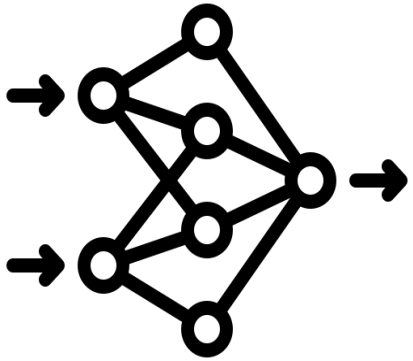

IR Spectra Translator

$\mathcal{C}$

: a set of SMILES candidates

- Since deriving reliable SMILES directly from thousands of real-values IR abosrbance is challenging for LLMs

→ Using the IR spectra translator to obtain plausible starting structures, which are then expanded and refined

$$\mathcal{C} = \{\mathbf{s}_1, \ldots, \mathbf{s_K}\} = \text{Transformer}(\mathcal{X})$$

Methodology Table Interpretation (TI) Expert (1/2)

Table Interpreation (TI) Expert

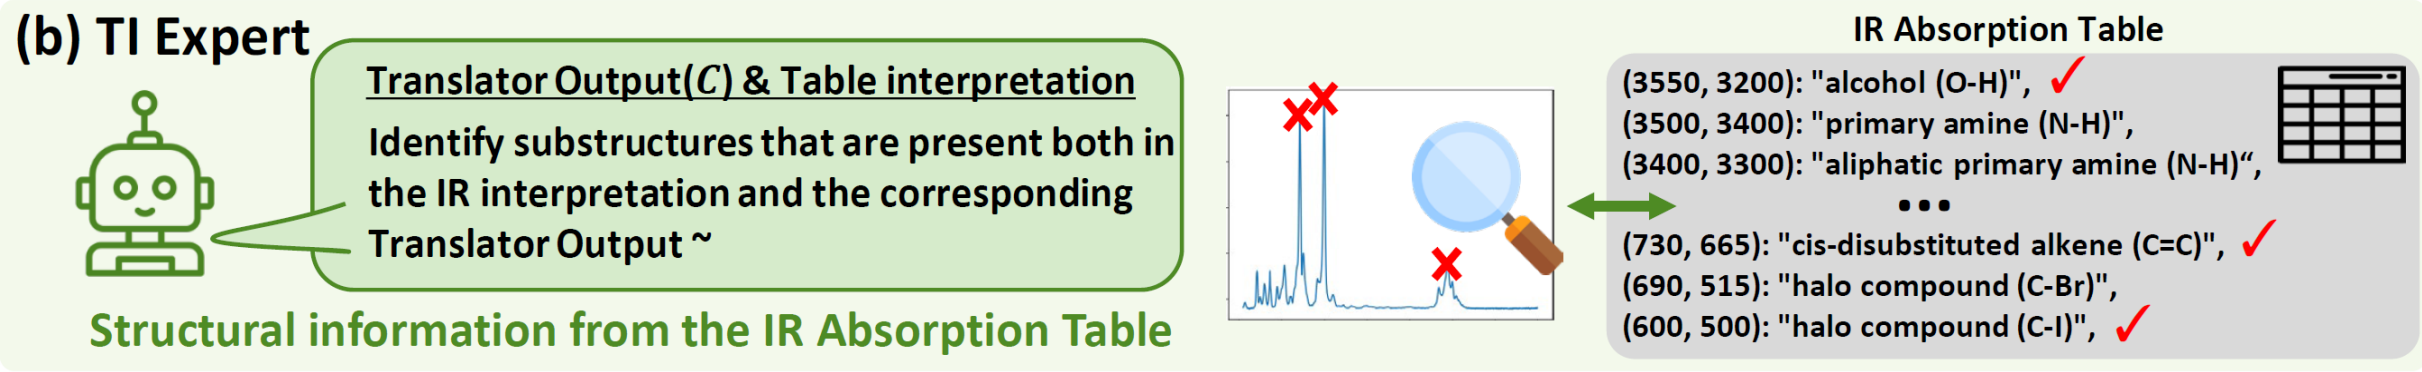

- IR absoprtion tables offer fine-grained localized structural features, such as substitution patterns, cis/trans isomerism, and conjugation, which makes it a crucial component in structure elucidation
- Effective utilization of the table requires accurate idenfication of spectral peaks → challenging to LLM
- Employing the curated IR Peak Table Assigner tool to detect peaks and assign substructures from their wavenumber ranges using the IR absoprtion table

a set of SMILES candidates

$$\mathcal{A}_{\text{TI Expert}} = \text{TI Expert}\left(\mathbf{P}_{\text{TI Expert}}, \text{IR Peak Table Assigner}(\mathcal{X}, \mathbf{T}), \mathcal{C}\right)$$

Task-specific prompt

IR Absoprtion Table

Methodology Table Interpretation (TI) Expert (2/2)

Table Interpreation (TI) Expert

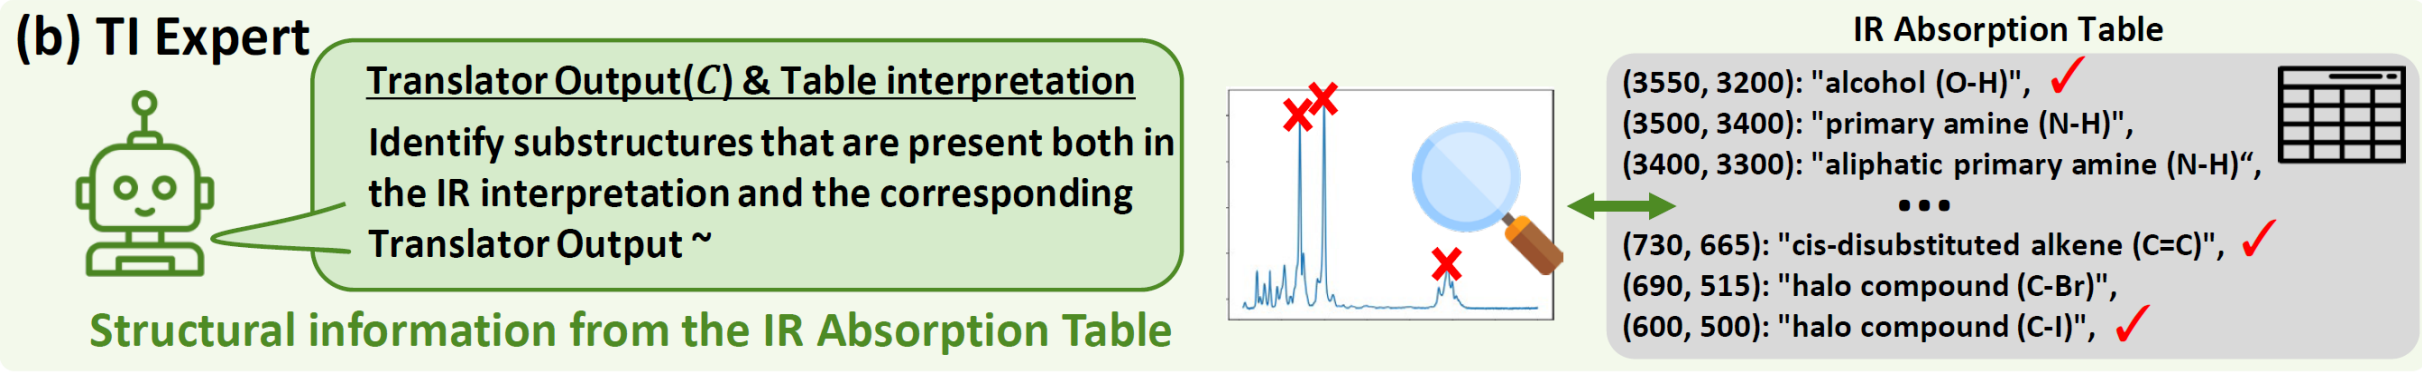

- Despite its utility, table-based interpretation has inherent limitations: IR spectra contain noise, and multiple substructures can share similar wavenumber regions, creating ambiguity in peak assignments

- To mitigate the possible misinterpretaion, we design a TI Expert prompt that guides the agent to compare the IR Peak Table Assigner output with the SMILES candidates, identify shared substructures, and produce a confidence score with a brief rationale for each match

  → **enhance the reliability of table-based interpretation**

e.g., substructure → confidence → brief rationale

Methodology Retriever (Ret) Expert (1/2)

Retriever (Ret) Expert

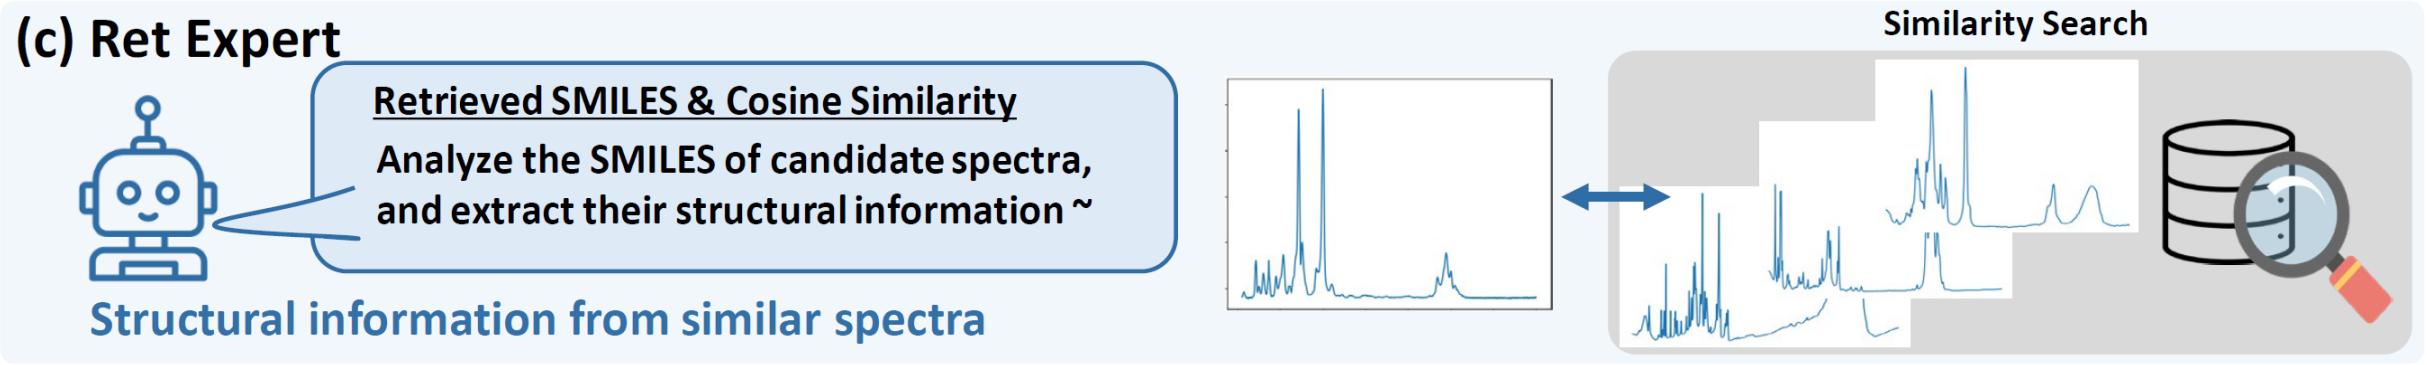

- Local structural information from TI expert is often insufficient to elucidate complete molecular structure
- Like human experts, the Ret Expert leverages molecular structures with similar IR spectra from the database to provide global structural context, helping link local substructures to a more complete molecular structure
- (1) computes the cosine similarity between target spectrum and all spectra in the database, (2) then retrieves the top-N most similar spectra with its SMILES structure

$$\{\text{candi}_1 : \text{sim}_1, \ldots, \text{candi}_N : \text{sim}_N\} = \text{IR Spectra Retriever}(\mathcal{X})$$

Methodology Retriever (Ret) Expert (2/2)

Retriever (Ret) Expert

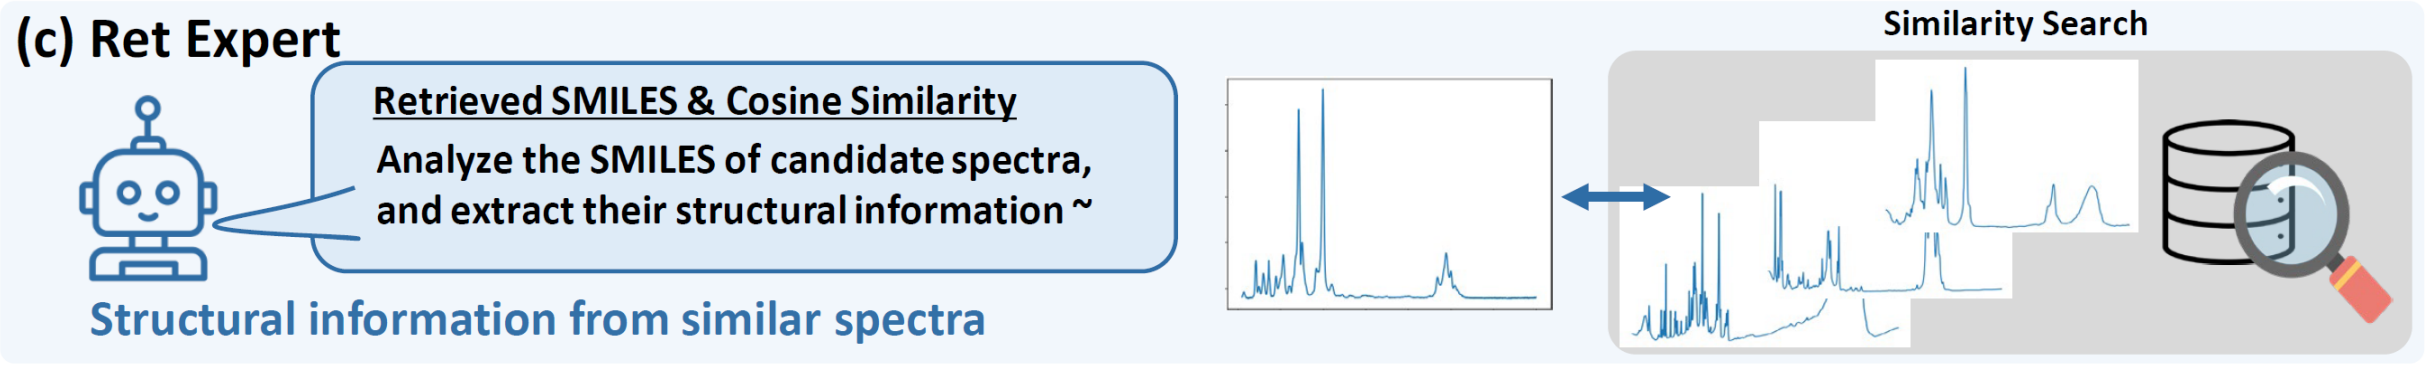

$$\mathcal{A}_{\text{Ret Expert}} = \text{Ret Expert}\left(\mathbf{P}_{\text{Ret Expert}}, \text{IR Spectra Retriever}(\mathcal{X})\right)$$

Task-specific prompt

- The output of the Ret Expert includes shared structural features among the top-N retrieved SMILES
- Given the SMILES and the cosine similarity between the target and retrieved spectra, the Ret Expert identifies common substructures and assigns greater weight to more similar spectra

→ **global contextual clues for guiding molecular structure reasoning**

Methodology Structure Elucdiation (SE) Expert & Overall Framework

Structure Elucidation (SE) Expert

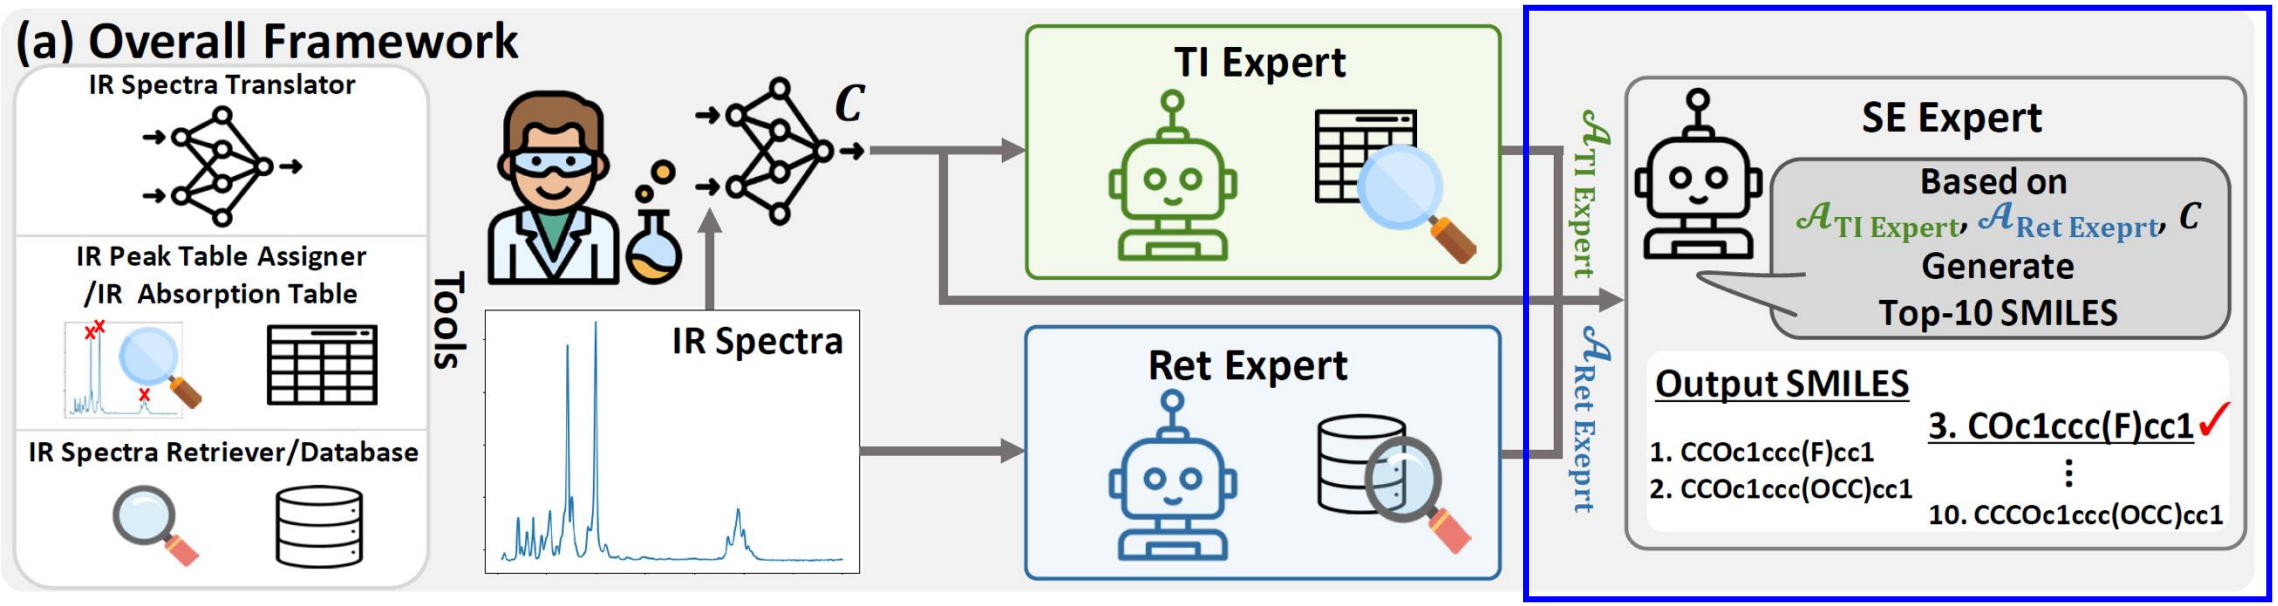

$$\mathcal{A}_{\text{SE Expert}} = \text{SE Expert}\left(\mathbf{P}_{\text{SE Expert}}, \mathcal{A}_{\text{TI Expert}}, \mathcal{A}_{\text{Ret Expert}}, \mathcal{C}\right)$$

SMILES candidates from Translator

Task-specific prompt | Output of TI Expert | Output of Ret Expert

- Structure Elucidation (SE) expert conducts integrative structure reasoning based on both outputs of TI and Ret expert
- Comprehensive reasoning process that integrates both local and global molecular structures

  → final ranked list of top-K predicted molecular structures

Experiments Dataset & Evaluation Protocol

Datasets & Evaluation Protocol

- **9,052 experimental IR spectra from NIST**
  - To ensure diversity, include spectra from all phases (solid, liquid, and gas), do not exclude compounds with stereochemistry or ionic features, and impose no restrictions on heavy-atom count or the presence of mixtures

- **Knolwedge Base:** Training set
- **Backbone LLM**: GPT-4o-mini, GPT-4o, o3-mini
- **Random Split**: train/valid/test of 80/10/10%
- After converting SMILES into **InChI representation**, calculate the **Top-K Exact Match** (K=1,3,5,10)

Experiments Effectivenss of IR-Agent in Structure Elucidation

Method	Agent	Top-K Accuracy			
		Top-1	Top-3	Top-5	Top-10
Transformer	-	0.098 (0.007)	0.169 (0.000)	0.176 (0.003)	0.176 (0.003)
IR-Agent (GPT-4o-mini)	single	0.072 (0.008)	0.118 (0.002)	0.133 (0.002)	0.157 (0.003)
	multi	0.093 (0.005)	0.152 (0.003)	0.167 (0.005)	0.176 (0.005)
IR-Agent (GPT-4o)	single	0.083 (0.004)	0.135 (0.002)	0.165 (0.007)	0.194 (0.008)
	multi	0.093 (0.007)	0.153 (0.005)	0.177 (0.005)	0.204 (0.005)
IR-Agent (o3-mini)	single	0.087 (0.006)	0.153 (0.005)	0.179 (0.002)	0.197 (0.004)
	multi	**0.103** (0.005)	**0.178** (0.007)	**0.199** (0.004)	**0.216** (0.001)

- Performance improvement over Transformer
→ Comprehensive reasoning that integrates the analyses of the TI and Ret Experts enables the SE Expert to refine candidates and generate more accurate molecular structures

- The multi-agent framework consistently outperforms the single-agent approach

Experiments Structure Elucidation with Chemical Information

Chemical Information	o3-mini				IR-Agent (single) (o3-mini)				IR-Agent (multi) (o3-mini)			
	Top-1	Top-3	Top-5	Top-10	Top-1	Top-3	Top-5	Top-10	Top-1	Top-3	Top-5	Top-10
No Knowledge	0.073 (0.010)	0.131 (0.011)	0.157 (0.011)	0.185 (0.005)	0.087 (0.010)	0.153 (0.011)	0.179 (0.011)	0.197 (0.005)	0.103 (0.005)	0.178 (0.007)	0.199 (0.004)	0.216 (0.001)
Scaffold	0.096 (0.002)	0.160 (0.006)	0.177 (0.003)	0.198 (0.003)	0.112 (0.003)	0.195 (0.009)	0.208 (0.008)	0.228 (0.010)	0.118 (0.003)	0.208 (0.009)	0.232 (0.008)	0.258 (0.010)
Carbon Count	**0.105** (0.009)	0.158 (0.014)	0.186 (0.014)	0.214 (0.013)	0.121 (0.003)	0.177 (0.008)	0.194 (0.010)	0.219 (0.009)	0.123 (0.003)	0.190 (0.0005)	0.215 (0.009)	0.252 (0.007)
Atom Types	0.104 (0.011	**0.182** (0.007)	**0.209** (0.005)	**0.237** (0.003)	**0.123** (0.006)	**0.208** (0.003)	**0.235** (0.011)	**0.266** (0.009)	**0.127** (0.006)	**0.213** (0.003)	**0.250** (0.011)	**0.278** (0.009)

- Chemical information is appended as a textual sentence to the prompt of the expert agents
- Even a brief textual prompt containing chemical information can enhance the model's ability to predict accurate molecular structures (without any architecutre modifications or retraining)

→ A flexible framework capable of integrating various types of available chemical information

Experiments In-Depth Analysis

Expert	Top-K Accuracy			
	Top-1	Top-3	Top-5	Top-10
No Expert	0.073 (0.010)	0.131 (0.011)	0.157 (0.011)	0.185 (0.005)
TI Expert only	0.089 (0.011)	0.154 (0.004)	0.171 (0.002)	0.190 (0.002)
Ret Expert only	0.098 (0.003)	0.169 (0.006)	0.188 (0.001)	0.211 (0.003)
IR-Agent (TI + Ret)	**0.103** (0.005)	**0.178** (0.007)	**0.199** (0.004)	**0.216** (0.001)

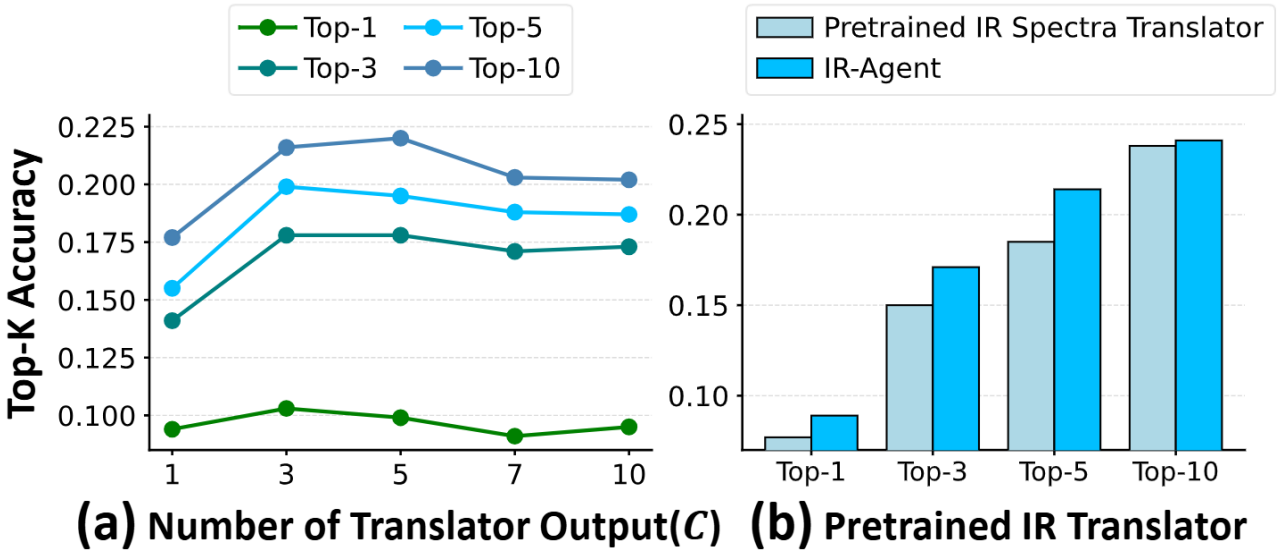

- Both experts are essential for providing complementary structural insights

- (a) An appropriate number of candidates is crucial

- (b) Boost performance when using a translator pretrained on simulated spectra and finetuned on our experimental dataset
→ Robust to different spectra translator choices

Experiments How IR-Agent performs the structure elucidation process

Translator Output(C)

'Nc1ccc(C(F)(F)F)cc1Br', 'Nc1ccc(Br)c(C(F)(F)F)c1', 'Nc1ccc(N)c(C(F)(F)F)c1'

TI Expert

• C-F group → High confidence → All three SMILES include a trifluoromethyl (CF_3) substituent, which corresponds to the **IR peaks between 1200–1000 cm^{-1} typically seen for C–F bonds.**
• Halogen substitution (Br) → Low confidence → SMILES 1 and 2 show bromine on the aromatic ring, but **the IR interpretation specifies C–Cl**; while both are halogens, the match is imprecise.

Ret Expert

All candidate SMILES share an aromatic core bearing **a trifluoromethyl group (–C(F)(F)F)**, suggesting that the target compound likely features **a benzene ring substituted with a CF3 group.** ... This pattern indicates that the target spectrum's molecule is an aromatic system with electron-donating and/or electron-withdrawing substituents that can influence IR absorptions (e.g., CF3 stretches near 1200 cm^{-1}, possible N–H stretches for amines, and C≡N stretches). The similarities in the candidates' SMILES imply that the target likely shares these structural motifs with variations in the nature and position of these substituents.

SE Expert

1. Nc1ccc(C(F)(F)F)cc1Br, 2. Nc1ccc(Br)c(C(F)(F)F)c1 ..., 4. **Nc1ccc(Cl)c(C(F)(F)F)c1** ✓, ..., 10. Nc1ccc(O)c(C(F)(F)F)c1

Complementary Analyses

- The TI Expert infers a C–F group with high confidence because both the Translator output and table interpretation consistently indicate C–F bonds, whereas halogen substitutions like Br receive low confidence due to conflicting evidence (C- Cl bonds from table interpretation)
- The Ret Expert provides global context by identifying a benzene ring with a CF_3 substituent as the dominant structural motif among the retrieved candidates

Conclusion

- We introduce **IR-Agent**, a novel multi-agent framework for molecular structure elucidation from infrared (IR) spectra. This framework models expert-driven IR spectrum analysis processes and is designed to be highly extensible.

- While each agent specializes in a specific aspect of IR spectrum analysis, their complementary roles enable an integrative analysis process, ultimately improving molecular structure elucidation from IR spectrum.

- Through extensive experimentation, **IR-Agent** not only improves baseline performance on experimental IR spectra but also demonstrates strong adaptability to diverse types of chemical information.

Thank you!

[Full Paper] https://openreview.net/forum?id=6bthH14pD8
[Source Code] https://github.com/HeewoongNoh/IR-Agent

Full Paper

Source Code