

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MSPECTMOL: A MULTI-MODAL SPECTROSCOPIC LEARNING FRAMEWORK FOR MOLECULAR STRUC- TURE IDENTIFICATION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Spectroscopic techniques are indispensable for the elucidation of molecular structures, particularly for novel molecules with unknown configurations. However, a fundamental limitation of any single spectroscopic modality is that it provides an inherently circumscribed and fragmented view, capturing only specific facets of the complete molecular structure, which is often insufficient for unequivocal and robust characterization. Consequently, the integration of data from multiple spectroscopic sources is imperative to overcome these intrinsic limitations and achieve a comprehensive and accurate structural characterization. In this work, we introduce **MSpecTmol**, a novel **Multi-modal Spectrum** information fusion learning framework for **Molecule** structure elucidation. By extending information bottleneck theory, our framework provides a principled and adaptive approach to fusing spectra. It designates a primary modality to extract core molecular features while leveraging auxiliary inputs to enrich the representation. To validate the end-to-end effectiveness of our framework, we design a two-fold evaluation: molecular substructure classification to probe its discriminative power in identifying substructures, and extends this knowledge to reconstruct plausible 3D structures. Our results not only demonstrate state-of-the-art performance in molecular substructure classification but also achieve near-experimental accuracy ( $\sim 0.68\text{\AA}$ ) in molecular conformation reconstruction. These findings underscore the model's capacity to learn interpretable features aligned with chemical intuition, thereby paving the way for future advances in automated and reliable spectroscopic analysis. Our code can be found at <https://anonymous.4open.science>.

## 1 INTRODUCTION

The rapid advancements in artificial intelligence (AI) have significantly propelled research in the chemical sciences Goh et al. (2017); Divya et al. (2024); Rial (2024); Ananikov (2024), enabling breakthroughs in molecular property prediction Feinberg et al. (2018); Walters & Barzilay (2020), drug design Blundell (1996); Riccardi et al. (2018), and drug-drug interaction studies Zhao et al. (2024); Wang et al. (2024). AI not only achieves high-precision predictions without compromising accuracy but also enhances trust in its applications through interpretable models Chander et al. (2024); Rane et al. (2024). These developments have increasingly integrated AI into chemistry as an indispensable tool. Notably, the vast majority of existing studies are *post-designed*, meaning that they operate on molecules with known structures, represented either as molecular graphs or SMILES strings Du et al. (2023); Xia et al. (2023).

However, for a novel, unknown molecule, chemists must first determine its fundamental structure before exploring its properties Hastings et al. (2021); Stanzione et al. (2021). In such cases, spectroscopic techniques serve as powerful tools for structural determination, fundamentally projecting high-dimensional chemical structures into lower-dimensional spectral representations as present in Figure 1 Barone et al. (2021); Meza Ramirez et al. (2021). Spectral techniques such as nuclear magnetic resonance (NMR), infrared (IR) spectroscopy, and mass spectrometry (MS) could provide critical insights into molecular structures Fontana & Widmalm (2023); Manogaran et al. (2024), including the presence or absence of functional groups Ge et al. (2021) which plays a crucial role in confirming structural assignments and ensuring the reliability of downstream chemical analysis.

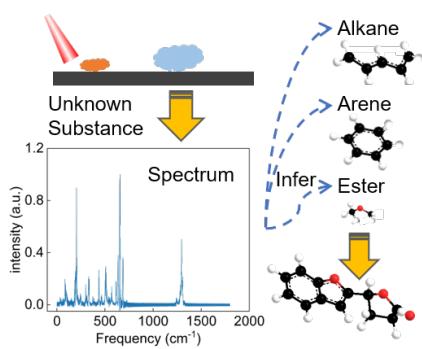


Figure 1: Schematic representation of inferring an unknown substance using spectral analysis techniques.

Yet, the inherent limitations of individual spectroscopic methods, due to their low-dimensional characteristics and the restricted information they contain Xue et al. (2024); Bose et al. (2021), necessitate the integration of multiple spectroscopic sources to achieve more precise molecular determination. For instance, IR spectroscopy focuses on functional group vibrations, ultraviolet-visible (UV-Vis) spectroscopy reflects overall molecular properties, and NMR provides information about the local atomic environment Chen et al. (2023); Manogaran et al. (2024). Each spectroscopic modality encapsulates distinct representational features and operates within different physically meaningful ranges Barone et al. (2021). Therefore, a key challenge lies in fully leveraging the available spectroscopic data to extract its physical significance and enable accurate molecular structural determination Meza Ramirez et al. (2021).

In this work, we propose a novel **Multi-modal SpecT**rum information fusion learning framework based on information bottleneck theory for **M**olecule confirmation, termed **MSpecT**mol, to integrate multi-modal spectroscopic data. Our framework adopts a primary-auxiliary synergistic modeling approach, where the roles of primary and auxiliary representations are clearly delineated. By extending the multi-objective information bottleneck theory to this setting, we enable the primary modality to capture core information by filtering out redundant or irrelevant features, while the auxiliary modalities supplement the primary representation to enhance and refine the results. To comprehensively validate our framework’s end-to-end effectiveness across the entire spectrum-to-molecule workflow, we rigorously applied it to two critically important tasks: molecular identification and spectrum-conditioned molecular conformation generation. In the molecular identification task, MSpecTmol significantly outperformed state-of-the-art baseline methods across both simulated and experimental spectra, achieving an F1-score of 0.959. Furthermore, the framework similarly demonstrated its capabilities for intricate structural elucidation in the challenging spectrum-conditioned conformation generation task, achieving an average RMSD of 0.682 Å. Moreover, MSpecTmol captures critical spectroscopic fragments that align well with chemical intuition, providing a degree of interpretability for its predictions. The synergy between primary and auxiliary modalities offers a flexible strategy for researchers to adapt to various chemical challenges, further improving performance outcomes. We envision that molecular identification through spectroscopic data will become a key research focus in automated laboratory workflows. MSpecTmol represents a promising solution to this challenge, offering a robust and interpretable framework for this.

## 2 METHODOLOGY

In this section, we introduce our proposed framework, called **MSpecT**mol, a novel *multi-modal information fusion learning* framework that refines representations based on the distinct roles of the underlying information. First, we formally define **MSpecT**mol (Section 2.1). Next, we present the overall model architecture (Section 2.2), followed by the final optimization process (Section 2.3).

### 2.1 PRIMARY-AUXILIARY INFORMATION BOTTLENECK

In this work, we focus on learning the core representations  $T_m$  and  $T_a$  from the input primary spectrum  $X_m$  and auxiliary spectra  $X_a$ .

*Primary-Auxiliary Information Bottleneck. (PA-IB)* Given the primary spectrum variables  $X_m$ , the auxiliary spectra variables  $X_a$ , and the target variable  $Y$ , the Primary-Auxiliary Information Bottleneck theory aims to compress  $X_m$  into a bottleneck variable  $T_m$  while preserving the information needed to predict  $Y$ , and to compress  $X_a$  into a bottleneck variable  $T_a$  while preserving the information needed to predict  $Y$  conditioned on  $X_m$ . Formally, we seek to solve:

$$\min -I(Y; T_m) - I(Y; T_a | T_m) + \alpha I(X_m; T_m) + \beta I(T_a; X_m, X_a), \quad (1)$$

where  $\alpha$  and  $\beta$  are Lagrange multipliers that balance the mutual information terms.

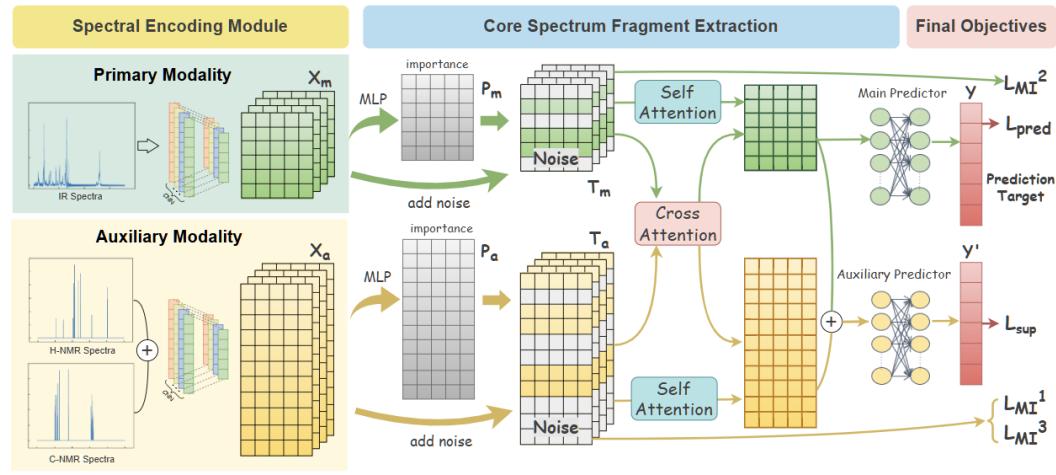


Figure 2: Illustration of the **MSpecTmol** framework. The model processes one primary and multiple auxiliary spectral modalities. Initially, vector representations are generated via linear interpolation and then fed into 1D-CNN layers to extract feature matrices. Subsequently, the core IB-Spectra module distills essential information from the primary modality and complementary features from the auxiliary inputs, producing a compact feature vector for downstream classification.

*IB-Spectra.* Given a set of spectra  $(X_m, X_a)$  and its corresponding label information  $\mathbf{Y}$ , we identify the optimal primary spectrum  $\mathcal{T}_{MIB}$  and auxiliary spectrum  $\mathcal{T}_{AIB}$  under the PA-IB principle:

$$\mathcal{T}_{MIB}, \mathcal{T}_{AIB} = \arg \min_{\mathcal{T}_{MIB}, \mathcal{T}_{AIB}} \left[ -I(Y; T_m) - I(Y; T_a | T_m) + \alpha I(X_m; T_m) + \beta I(T_a; X_m, X_a) \right], \quad (2)$$

This objective involves the following four components:

- $-I(Y; T_m)$ : Encourages the primary representation  $T_m$  to preserve the most predictive information about  $Y$ . This term corresponds to the classical IB objective and ensures that  $T_m$  serves as the main carrier of task-relevant information.
- $-I(Y; T_a | T_m)$ : Drives the auxiliary representation  $T_a$  to complement  $T_m$  by capturing additional information that is not contained in  $T_m$ , thus improving the overall predictive capacity.
- $\alpha I(X_m; T_m)$ : Regularizes the complexity of  $T_m$  by penalizing excessive mutual information with the input  $X_m$ , thereby promoting a compact and generalized encoding.
- $\beta I(T_a; X_m, X_a)$ : Limits the complexity of  $T_a$  by minimizing its mutual information with the combined inputs  $(X_m, X_a)$ , encouraging selective representation of auxiliary information.

Combining these four terms, the final objective is: The primary representation  $T_m$  captures the core information necessary for predicting  $Y$ . The auxiliary representation  $T_a$  complements  $T_m$  by providing any additional information needed for  $Y$ , while maintaining low complexity.

## 2.2 MODEL ARCHITECTURE

### 2.2.1 SPECTRAL ENCODING MODULE

Here,  $\{X_m, X_a\}$  denote a pair of input spectra, with  $X_m$  as the *primary* spectrum and  $X_a$  as the *auxiliary* spectra. To unify dimensions, each spectrum is interpolated to 600 uniformly spaced points by linear interpolation, where the point value  $x$  is updated to  $X'$ :

$$X'(x) = X(x_i) + \frac{(x - x_i)}{(x_{i+1} - x_i)} \cdot (X(x_{i+1}) - X(x_i)), \quad x \in [x_i, x_{i+1}] \quad (3)$$

where  $x_i$  and  $x_{i+1}$  are consecutive points in the original spectrum, and  $X(x)$  represents the updated value at  $x$ . This procedure is applied independently to two different spectra designated as auxiliary

162 inputs. After interpolation, both spectra are normalized and concatenated to create a single 1200-  
 163 dimensional auxiliary spectrum,  $X_a$ .  
 164

165 The input spectra  $X_m$  and  $X_a$  are encoded through two stages of 1D convolutional layers, each  
 166 followed by batch normalization, ReLU activation, and max pooling. The use of 1D convolutions  
 167 is motivated by the sequential nature of spectral data, where capturing local patterns is essential  
 168 for extracting meaningful features. Convolutional layers enable the model to automatically learn  
 169 hierarchical representations and spatially invariant characteristics from the spectra.  
 170

$$O^m = \text{MaxPool1D}(\text{ReLU}(\text{BatchNorm}(\text{Conv1D}(X_m)))), \quad (4)$$

$$O^a = \text{MaxPool1D}(\text{ReLU}(\text{BatchNorm}(\text{Conv1D}(x_a)))), \quad (5)$$

172 Here,  $O^m \in \mathbb{R}^{c \times d}$  and  $O^a \in \mathbb{R}^{c \times d}$ , where  $c$  denotes the number of convolutional channels and  $d$   
 173 represents the length after linear interpolation. The pool size is set to 2.  
 174

### 2.2.2 CORE SPECTRUM EXTRACTION

176 In this section, we extract core spectral segments by first transposing the frequency representations:  
 177  $H^m = O_m^{(t)}$ ,  $H^a = O_a^{(t)}$ , where  $(t)$  denotes the transpose operation. For the primary spectrum,  
 178 we compress  $X_m$  into  $T_m$  by injecting noise into its learned embedding, encouraging the model to  
 179 suppress less informative frequency bands. For the auxiliary spectrum, we similarly derive  $T_a$  from  
 180  $X_m$ ,  $X_a$ , and  $T_m$ , guided by equation 7. The key idea is to enable the model to inject noise into  
 181 insignificant frequency bands while injecting less noise into more informative ones Yu et al. (2022).  
 182 we could calculate the probability  $p_i^m$  and  $p_i^a$  using an MLP, i.e.,  
 183

$$p_i^m = \text{MLP}(\mathbf{H}_i^m) \quad p_i^a = \text{MLP}(\mathbf{H}_i^m \parallel \mathbf{H}_i^a \parallel \mathbf{T}_i^m). \quad (6)$$

184 With the  $p_i^m$  and  $p_i^a$ , we replace  $\mathbf{H}_i^m$  and  $\mathbf{H}_i^a$  of frequency band  $i$  with noise  $\epsilon$ , i.e.,  
 185

$$\mathbf{T}_i^m = \lambda_i^m \mathbf{H}_i^m + (1 - \lambda_i^m) \epsilon^m, \quad \mathbf{T}_i^a = \lambda_i^a \mathbf{H}_i^a + (1 - \lambda_i^a) \epsilon^a, \quad (7)$$

186 where  $\lambda_i^m \sim \text{Bernoulli}(p_i^m)$  and  $\epsilon^m \sim \mathcal{N}(\mu_m, \sigma_m^2)$ . Here,  $\mu_m$  and  $\sigma_m^2$  are mean and variance of  
 187  $\mathbf{H}^m$ , respectively. Thus, the information of  $X_m$  is compressed into  $T_m$  with the probability  $p_i^m$ , by  
 188 replacing unimportant frequency bands with noise. Similarly, for the core auxiliary spectrum, The  
 189 information from  $X_m$ ,  $T_m$ , and  $X_a$  is compressed into  $T_a$  with the same probability  $p_i^a$ .  
 190

191 Moreover, to make the sampling process differentiable, the Gumbel-Softmax is adopted Maddison  
 192 et al. (2016); Jang et al. (2016) for the discrete random variable  $\lambda_i$ , i.e.,  
 193

$$\lambda_i = \sigma \left( \frac{1}{t} \log \left( \frac{p_i}{1 - p_i} \right) + \log \left( \frac{u}{1 - u} \right) \right), \quad (8)$$

194 where  $u \sim \text{Uniform}(0, 1)$ , and  $t$  is the temperature hyperparameter, set to 1.0 in this work. A  
 195 detailed sensitivity analysis of  $t$  is provided in Appendix J.  
 196

### 2.3 MODEL OPTIMIZATION

200 To train the model while simultaneously detecting the core primary spectra and core auxiliary spec-  
 201 tra, we optimize the model with the objective function defined in equation 1 as follows:  
 202

$$\min -I(Y; T_m) - I(Y; T_a | T_m) + \alpha I(X_m; T_m) + \beta I(T_a; X_m, X_a), \quad (9)$$

204 where each term corresponds to prediction or compression, respectively. In the following sections,  
 205 we provide the upper bounds of each term, which should be minimized during training.  
 206

#### 2.3.1 MINIMIZING $-I(Y; T_m)$

208 **Proposition 3.1 (Upper bound of  $-I(Y; T_m)$ )** Given the primary spectra  $X_m$ , and its label in-  
 209 formation  $\mathbf{Y}$ , we have:  
 210

$$\begin{aligned} -I(\mathbf{Y}; T_m) &\leq \mathbb{E}_{T_m, \mathbf{Y}}[-\log p_\theta(\mathbf{Y} | T_m)] \\ &= \mathbb{E}_{(\mathbf{Y}, T_m)} \log [P_\theta(\mathbf{Y} | T_m)] + H(\mathbf{Y}) := \mathcal{L}_{\text{pred}}, \end{aligned} \quad (10)$$

213 where  $H(\mathbf{Y})$  is the entropy of the label  $\mathbf{Y}$ , which is constant across the dataset and can be omitted  
 214 in the optimization.  $p_\theta(\mathbf{Y} | T_m)$  is the variational approximation of the true posterior  $p(\mathbf{Y} | T_m)$ .  
 215 Minimizing this upper bound corresponds to minimizing the prediction loss  $\mathcal{L}_{\text{pred}}(\mathbf{Y}, T_m)$ , which  
 216 is modeled as the cross-entropy loss for classification. The proof can be found in Appendix F.1.1.  
 217

216 2.3.2 MINIMIZING  $-I(Y; T_a | T_m)$   
217218 **Proposition 3.2 (Upper bound of  $-I(Y; T_a | T_m)$ )** We decompose the term using the chain rule  
219 of mutual information:

220 
$$\begin{aligned} -I(Y; T_a | T_m) &= -I(Y; T_a, T_m) + I(T_a; T_m) \\ 221 &\leq \mathbb{E}_{(\mathbf{Y}, T_a, T_m)} \log [P_\theta(\mathbf{Y} | T_a, T_m)] + \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_a | t_m) \| q(t_a))] \quad (11) \\ 222 &:= \mathcal{L}_{\text{sup}} + \mathcal{L}_{\text{MI}^1}. \\ 223 \end{aligned}$$

224 Here,  $\mathcal{L}_{\text{sup}}$  represents the supervised prediction loss  $\mathcal{L}_{\text{pred}}(\mathbf{Y}, T_m, T_a)$ , which is implemented as  
225 cross-entropy for prediction. The second term,  $\mathcal{L}_{\text{MI}^1}$ , corresponds to the KL divergence between the  
226 posterior  $p(t_a | t_m)$  and a prior  $q(t_a)$ , regularizing the relationship between the auxiliary spectra  $T_a$   
227 and primary spectra  $T_m$ . This divergence is minimized using variational inference, and is estimated  
228 by averaging over samples of  $t_m$ . Detailed derivations can be found in Appendix F.1.2. Specifically,  
229 as shown in Appendix L, we investigate the impact of different prior distributions of  $q(t_m)$  and  $q(t_a)$   
230 on model performance, and select the best prior distribution as the distribution for MSpecTmol.231 2.3.3 MINIMIZING  $I(X_m; T_m)$   
232233 **Proposition 3.3 (Upper bound of  $I(X_m; T_m)$ )** We apply the variational approximation to bound  
234 the mutual information term:

235 
$$I(X_m; T_m) \leq \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_m | x_m) \| q(t_m))] := \mathcal{L}_{\text{MI}^2}. \quad (12)$$

236 Here,  $\mathcal{L}_{\text{MI}^2}$  corresponds to the KL divergence between the posterior  $p(t_m | x_m)$  and a prior  $q(t_m)$ .  
237 The KL divergence is computed using variational inference and is estimated by averaging over  
238 samples of  $x_m$ . The detailed derivation is provided in Appendix F.1.3.239 2.3.4 MINIMIZING  $I(T_a; X_m, X_a)$   
240241 **Proposition 3.4 (Upper bound of  $I(T_a; X_m, X_a)$ )** We minimize the mutual information between  
242 the auxiliary spectra  $T_a$  and both the primary spectra  $X_m$  as well as the auxiliary spectra  $X_a$ :

243 
$$I(T_a; X_m, X_a) \leq \mathbb{E}_{t_a, x_a \sim p(x_m, x_a)} [\text{KL}(p(t_a | x_m, x_a) \| q(t_a))] := \mathcal{L}_{\text{MI}^3}. \quad (13)$$

244 Here,  $\mathcal{L}_{\text{MI}^3}$  represents the KL divergence between the posterior  $p(t_a | x_m, x_a)$  and a prior  $q(t_a)$ . The  
245 KL divergence is estimated using variational inference, with derivations detailed in Appendix F.1.4.246 2.4 FINAL OBJECTIVES  
247248 The final objective function used for training is given by:  
249

250 
$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{sup}} + \mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{MI}^1} + \alpha \mathcal{L}_{\text{MI}^2} + \beta \mathcal{L}_{\text{MI}^3} \quad (14)$$

251 where  $\alpha$  and  $\beta$  control the trade-off between prediction accuracy and compression. The detailed  
252 derivations and proofs for  $\mathcal{L}_{\text{pred}}$ ,  $\mathcal{L}_{\text{sup}}$ ,  $\mathcal{L}_{\text{MI}^1}$ ,  $\mathcal{L}_{\text{MI}^2}$ , and  $\mathcal{L}_{\text{MI}^3}$  are provided in above.253 3 EXPERIMENT AND ANALYSES  
254255 We present experimental results to demonstrate the effectiveness of MSpecTmol under two tasks:  
256 molecular identification and spectrum-conditioned molecular conformation generation. In this sec-  
257 tion, we conduct extensive experiments to address the following research questions:258 

- 259 • **RQ1:** Can MSpecTmol accurately perform fine-grained classification of molecular substructures?
- 260 • **RQ2:** Can MSpecTmol accurately generate 3D molecular conformations by spectra?
- 261 • **RQ3:** Can MSpecTmol provide interpretable insights?

262 3.1 DATASETS AND SETUPS  
263264 **Datasets.** We utilize the large-scale dataset from Alberts et al. Alberts et al. (2024) for molecular  
265 structure elucidation, which contains 794K molecules with simulated IR,  $^1\text{H}$ -NMR,  $^{13}\text{C}$ -NMR, and

270  
 271 Table 1: F1-scores for predicting functional groups. For multi-modal settings, the primary modality  
 272 is indicated in **bold**. Baseline models are invariant to the choice of primary modality, whereas  
 273 MSpecTmol leverages this information to achieve superior performance. The best results are high-  
 274 275  
 275 **276**  
 276 **277**  
 277 **278**  
 278 **279**  
 279 **280**  
 280 **281**  
 281 **282**  
 282 **283**  
 283 **284**  
 284 **285**  
 285 **286**  
 286 **287**  
 287 **288**  
 288 **289**  
 289 **290**  
 290 **291**  
 291 **292**  
 292 **293**  
 293 **294**  
 294 **295**  
 295 **296**  
 296 **297**  
 297 **298**  
 298 **299**  
 299 **300**  
 300 **301**  
 301 **302**  
 302 **303**  
 303 **304**  
 304 **305**  
 305 **306**  
 306 **307**  
 307 **308**  
 308 **309**  
 309 **310**  
 310 **311**  
 311 **312**  
 312 **313**  
 313 **314**  
 314 **315**  
 315 **316**  
 316 **317**  
 317 **318**  
 318 **319**  
 319 **320**  
 320 **321**  
 321 **322**  
 322 **323**  
 323 **324**  
 324 **325**  
 325 **326**  
 326 **327**  
 327 **328**  
 328 **329**  
 329 **330**  
 330 **331**  
 331 **332**  
 332 **333**  
 333 **334**  
 334 **335**  
 335 **336**  
 336 **337**  
 337 **338**  
 338 **339**  
 339 **340**  
 340 **341**  
 341 **342**  
 342 **343**  
 343 **344**  
 344 **345**  
 345 **346**  
 346 **347**  
 347 **348**  
 348 **349**  
 349 **350**  
 350 **351**  
 351 **352**  
 352 **353**  
 353 **354**  
 354 **355**  
 355 **356**  
 356 **357**  
 357 **358**  
 358 **359**  
 359 **360**  
 360 **361**  
 361 **362**  
 362 **363**  
 363 **364**  
 364 **365**  
 365 **366**  
 366 **367**  
 367 **368**  
 368 **369**  
 369 **370**  
 370 **371**  
 371 **372**  
 372 **373**  
 373 **374**  
 374 **375**  
 375 **376**  
 376 **377**  
 377 **378**  
 378 **379**  
 379 **380**  
 380 **381**  
 381 **382**  
 382 **383**  
 383 **384**  
 384 **385**  
 385 **386**  
 386 **387**  
 387 **388**  
 388 **389**  
 389 **390**  
 390 **391**  
 391 **392**  
 392 **393**  
 393 **394**  
 394 **395**  
 395 **396**  
 396 **397**  
 397 **398**  
 398 **399**  
 399 **400**  
 400 **401**  
 401 **402**  
 402 **403**  
 403 **404**  
 404 **405**  
 405 **406**  
 406 **407**  
 407 **408**  
 408 **409**  
 409 **410**  
 410 **411**  
 411 **412**  
 412 **413**  
 413 **414**  
 414 **415**  
 415 **416**  
 416 **417**  
 417 **418**  
 418 **419**  
 419 **420**  
 420 **421**  
 421 **422**  
 422 **423**  
 423 **424**  
 424 **425**  
 425 **426**  
 426 **427**  
 427 **428**  
 428 **429**  
 429 **430**  
 430 **431**  
 431 **432**  
 432 **433**  
 433 **434**  
 434 **435**  
 435 **436**  
 436 **437**  
 437 **438**  
 438 **439**  
 439 **440**  
 440 **441**  
 441 **442**  
 442 **443**  
 443 **444**  
 444 **445**  
 445 **446**  
 446 **447**  
 447 **448**  
 448 **449**  
 449 **450**  
 450 **451**  
 451 **452**  
 452 **453**  
 453 **454**  
 454 **455**  
 455 **456**  
 456 **457**  
 457 **458**  
 458 **459**  
 459 **460**  
 460 **461**  
 461 **462**  
 462 **463**  
 463 **464**  
 464 **465**  
 465 **466**  
 466 **467**  
 467 **468**  
 468 **469**  
 469 **470**  
 470 **471**  
 471 **472**  
 472 **473**  
 473 **474**  
 474 **475**  
 475 **476**  
 476 **477**  
 477 **478**  
 478 **479**  
 479 **480**  
 480 **481**  
 481 **482**  
 482 **483**  
 483 **484**  
 484 **485**  
 485 **486**  
 486 **487**  
 487 **488**  
 488 **489**  
 489 **490**  
 490 **491**  
 491 **492**  
 492 **493**  
 493 **494**  
 494 **495**  
 495 **496**  
 496 **497**  
 497 **498**  
 498 **499**  
 499 **500**  
 500 **501**  
 501 **502**  
 502 **503**  
 503 **504**  
 504 **505**  
 505 **506**  
 506 **507**  
 507 **508**  
 508 **509**  
 509 **510**  
 510 **511**  
 511 **512**  
 512 **513**  
 513 **514**  
 514 **515**  
 515 **516**  
 516 **517**  
 517 **518**  
 518 **519**  
 519 **520**  
 520 **521**  
 521 **522**  
 522 **523**  
 523 **524**  
 524 **525**  
 525 **526**  
 526 **527**  
 527 **528**  
 528 **529**  
 529 **530**  
 530 **531**  
 531 **532**  
 532 **533**  
 533 **534**  
 534 **535**  
 535 **536**  
 536 **537**  
 537 **538**  
 538 **539**  
 539 **540**  
 540 **541**  
 541 **542**  
 542 **543**  
 543 **544**  
 544 **545**  
 545 **546**  
 546 **547**  
 547 **548**  
 548 **549**  
 549 **550**  
 550 **551**  
 551 **552**  
 552 **553**  
 553 **554**  
 554 **555**  
 555 **556**  
 556 **557**  
 557 **558**  
 558 **559**  
 559 **560**  
 560 **561**  
 561 **562**  
 562 **563**  
 563 **564**  
 564 **565**  
 565 **566**  
 566 **567**  
 567 **568**  
 568 **569**  
 569 **570**  
 570 **571**  
 571 **572**  
 572 **573**  
 573 **574**  
 574 **575**  
 575 **576**  
 576 **577**  
 577 **578**  
 578 **579**  
 579 **580**  
 580 **581**  
 581 **582**  
 582 **583**  
 583 **584**  
 584 **585**  
 585 **586**  
 586 **587**  
 587 **588**  
 588 **589**  
 589 **590**  
 590 **591**  
 591 **592**  
 592 **593**  
 593 **594**  
 594 **595**  
 595 **596**  
 596 **597**  
 597 **598**  
 598 **599**  
 599 **600**  
 600 **601**  
 601 **602**  
 602 **603**  
 603 **604**  
 604 **605**  
 605 **606**  
 606 **607**  
 607 **608**  
 608 **609**  
 609 **610**  
 610 **611**  
 611 **612**  
 612 **613**  
 613 **614**  
 614 **615**  
 615 **616**  
 616 **617**  
 617 **618**  
 618 **619**  
 619 **620**  
 620 **621**  
 621 **622**  
 622 **623**  
 623 **624**  
 624 **625**  
 625 **626**  
 626 **627**  
 627 **628**  
 628 **629**  
 629 **630**  
 630 **631**  
 631 **632**  
 632 **633**  
 633 **634**  
 634 **635**  
 635 **636**  
 636 **637**  
 637 **638**  
 638 **639**  
 639 **640**  
 640 **641**  
 641 **642**  
 642 **643**  
 643 **644**  
 644 **645**  
 645 **646**  
 646 **647**  
 647 **648**  
 648 **649**  
 649 **650**  
 650 **651**  
 651 **652**  
 652 **653**  
 653 **654**  
 654 **655**  
 655 **656**  
 656 **657**  
 657 **658**  
 658 **659**  
 659 **660**  
 660 **661**  
 661 **662**  
 662 **663**  
 663 **664**  
 664 **665**  
 665 **666**  
 666 **667**  
 667 **668**  
 668 **669**  
 669 **670**  
 670 **671**  
 671 **672**  
 672 **673**  
 673 **674**  
 674 **675**  
 675 **676**  
 676 **677**  
 677 **678**  
 678 **679**  
 679 **680**  
 680 **681**  
 681 **682**  
 682 **683**  
 683 **684**  
 684 **685**  
 685 **686**  
 686 **687**  
 687 **688**  
 688 **689**  
 689 **690**  
 690 **691**  
 691 **692**  
 692 **693**  
 693 **694**  
 694 **695**  
 695 **696**  
 696 **697**  
 697 **698**  
 698 **699**  
 699 **700**  
 700 **701**  
 701 **702**  
 702 **703**  
 703 **704**  
 704 **705**  
 705 **706**  
 706 **707**  
 707 **708**  
 708 **709**  
 709 **710**  
 710 **711**  
 711 **712**  
 712 **713**  
 713 **714**  
 714 **715**  
 715 **716**  
 716 **717**  
 717 **718**  
 718 **719**  
 719 **720**  
 720 **721**  
 721 **722**  
 722 **723**  
 723 **724**  
 724 **725**  
 725 **726**  
 726 **727**  
 727 **728**  
 728 **729**  
 729 **730**  
 730 **731**  
 731 **732**  
 732 **733**  
 733 **734**  
 734 **735**  
 735 **736**  
 736 **737**  
 737 **738**  
 738 **739**  
 739 **740**  
 740 **741**  
 741 **742**  
 742 **743**  
 743 **744**  
 744 **745**  
 745 **746**  
 746 **747**  
 747 **748**  
 748 **749**  
 749 **750**  
 750 **751**  
 751 **752**  
 752 **753**  
 753 **754**  
 754 **755**  
 755 **756**  
 756 **757**  
 757 **758**  
 758 **759**  
 759 **760**  
 760 **761**  
 761 **762**  
 762 **763**  
 763 **764**  
 764 **765**  
 765 **766**  
 766 **767**  
 767 **768**  
 768 **769**  
 769 **770**  
 770 **771**  
 771 **772**  
 772 **773**  
 773 **774**  
 774 **775**  
 775 **776**  
 776 **777**  
 777 **778**  
 778 **779**  
 779 **780**  
 780 **781**  
 781 **782**  
 782 **783**  
 783 **784**  
 784 **785**  
 785 **786**  
 786 **787**  
 787 **788**  
 788 **789**  
 789 **790**  
 790 **791**  
 791 **792**  
 792 **793**  
 793 **794**  
 794 **795**  
 795 **796**  
 796 **797**  
 797 **798**  
 798 **799**  
 799 **800**  
 800 **801**  
 801 **802**  
 802 **803**  
 803 **804**  
 804 **805**  
 805 **806**  
 806 **807**  
 807 **808**  
 808 **809**  
 809 **810**  
 810 **811**  
 811 **812**  
 812 **813**  
 813 **814**  
 814 **815**  
 815 **816**  
 816 **817**  
 817 **818**  
 818 **819**  
 819 **820**  
 820 **821**  
 821 **822**  
 822 **823**  
 823 **824**  
 824 **825**  
 825 **826**  
 826 **827**  
 827 **828**  
 828 **829**  
 829 **830**  
 830 **831**  
 831 **832**  
 832 **833**  
 833 **834**  
 834 **835**  
 835 **836**  
 836 **837**  
 837 **838**  
 838 **839**  
 839 **840**  
 840 **841**  
 841 **842**  
 842 **843**  
 843 **844**  
 844 **845**  
 845 **846**  
 846 **847**  
 847 **848**  
 848 **849**  
 849 **850**  
 850 **851**  
 851 **852**  
 852 **853**  
 853 **854**  
 854 **855**  
 855 **856**  
 856 **857**  
 857 **858**  
 858 **859**  
 859 **860**  
 860 **861**  
 861 **862**  
 862 **863**  
 863 **864**  
 864 **865**  
 865 **866**  
 866 **867**  
 867 **868**  
 868 **869**  
 869 **870**  
 870 **871**  
 871 **872**  
 872 **873**  
 873 **874**  
 874 **875**  
 875 **876**  
 876 **877**  
 877 **878**  
 878 **879**  
 879 **880**  
 880 **881**  
 881 **882**  
 882 **883**  
 883 **884**  
 884 **885**  
 885 **886**  
 886 **887**  
 887 **888**  
 888 **889**  
 889 **890**  
 890 **891**  
 891 **892**  
 892 **893**  
 893 **894**  
 894 **895**  
 895 **896**  
 896 **897**  
 897 **898**  
 898 **899**  
 899 **900**  
 900 **901**  
 901 **902**  
 902 **903**  
 903 **904**  
 904 **905**  
 905 **906**  
 906 **907**  
 907 **908**  
 908 **909**  
 909 **910**  
 910 **911**  
 911 **912**  
 912 **913**  
 913 **914**  
 914 **915**  
 915 **916**  
 916 **917**  
 917 **918**  
 918 **919**  
 919 **920**  
 920 **921**  
 921 **922**  
 922 **923**  
 923 **924**  
 924 **925**  
 925 **926**  
 926 **927**  
 927 **928**  
 928 **929**  
 929 **930**  
 930 **931**  
 931 **932**  
 932 **933**  
 933 **934**  
 934 **935**  
 935 **936**  
 936 **937**  
 937 **938**  
 938 **939**  
 939 **940**  
 940 **941**  
 941 **942**  
 942 **943**  
 943 **944**  
 944 **945**  
 945 **946**  
 946 **947**  
 947 **948**  
 948 **949**  
 949 **950**  
 950 **951**  
 951 **952**  
 952 **953**  
 953 **954**  
 954 **955**  
 955 **956**  
 956 **957**  
 957 **958**  
 958 **959**  
 959 **960**  
 960 **961**  
 961 **962**  
 962 **963**  
 963 **964**  
 964 **965**  
 965 **966**  
 966 **967**  
 967 **968**  
 968 **969**  
 969 **970**  
 970 **971**  
 971 **972**  
 972 **973**  
 973 **974**  
 974 **975**  
 975 **976</b**

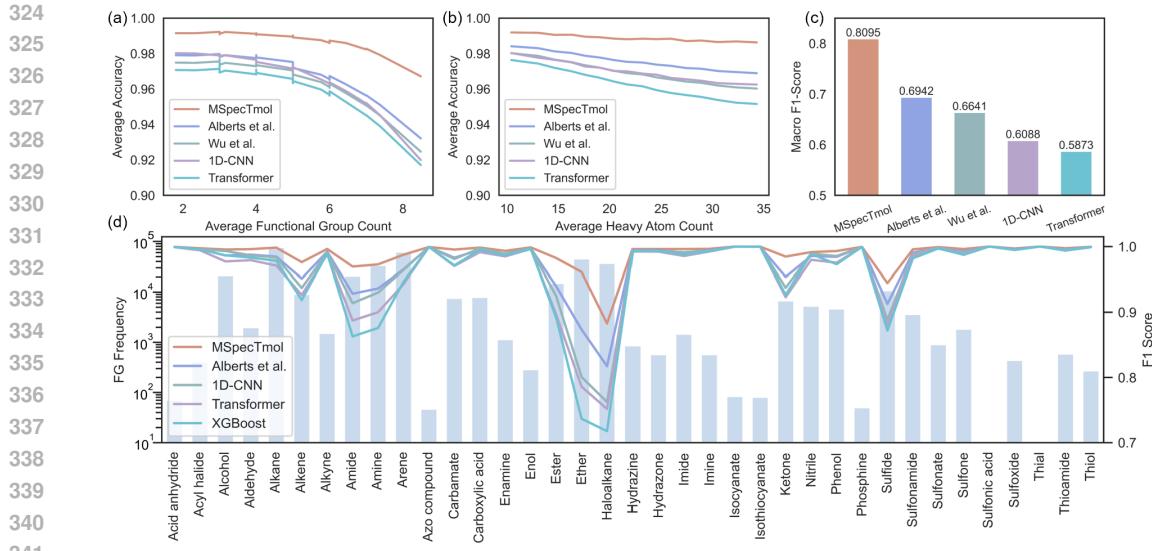


Figure 3: Comparison of our model against baselines in terms of accuracy and macro-F1 score. (a) Performance across different numbers of functional groups; (b) Performance under varying heavy atom counts; (c) Macro-F1 score results; (d) Correlation between the number of functional groups and prediction performance. Samples are binned in chunks of 5000.

lecture in capturing both local and global spectral features. However, when auxiliary spectra are introduced in the multi-modal setting, performance improves consistently across all models. Notably, MSpecTmol benefits the most, indicating its ability to effectively leverage additional information from auxiliary spectra to enhance substructure recognition and prediction accuracy. Furthermore, to demonstrate the broader applicability of MSpecTmol, we extended our evaluation to other state-of-the-art baselines with different objectives. Specifically, SpectraLLM leverages LLMs to treat spectral analysis as a sequence generation task, while DiffSpectra focuses on end-to-end molecular conformation generation. As shown in Table 2, MSpecTmol consistently outperforms both baselines in terms of Functional Group Similarity scores on the QM9S dataset, underscoring its robustness in identifying chemically meaningful substructures. In Appendix H, we analyze the computational overhead (time and memory) of MSpecTmol across varying numbers of input modalities and outline a decision framework for selecting the primary modality. To balance efficiency and performance, we selected an ideal combination of three modalities: IR,  $^1\text{H}$ -NMR, and  $^{13}\text{C}$ -NMR, which delivers strong predictive power with acceptable complexity.

**Table 2: Functional Group Similarity comparison on the QM9S dataset.** We compare MSpecTmol against generative baselines under their respective evaluation standards (17 functional groups for SpectraLLM, 13 for DiffSpectra).

Modality	17 Functional Groups		13 Functional Groups	
	SpectraLLM	MSpecTmol	DiffSpectra	MSpecTmol
IR	0.6599	<b>0.9328</b>	0.9322	<b>0.9501</b>
Raman	0.7317	<b>0.9334</b>	0.9279	<b>0.9417</b>
UV-Vis	0.3713	<b>0.5449</b>	0.4354	<b>0.5621</b>
All (IR+Ram.+UV)	0.7934	<b>0.9781</b>	0.9495	<b>0.9830</b>

**Obs.2: MSpecTmol exhibits superior robustness under increasing structural complexity.** In the process of functional group classification, as shown in Figure 3(a) and (b), we observe that the prediction accuracy tends to decrease as the number of functional groups and heavy atoms in a molecule increases. This is likely because greater structural complexity leads to more intricate and overlapping spectral signals, making it challenging to disentangle the features corresponding to individual substructures. Despite these challenges, our model consistently achieves superior performance rel-

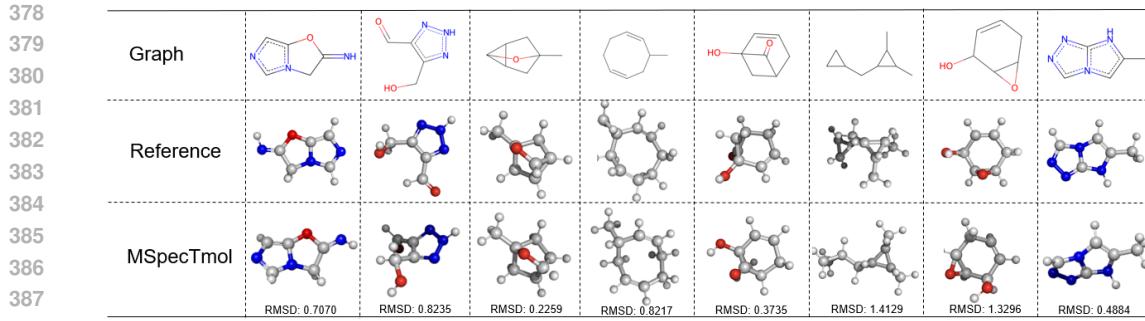


Figure 4: Example conformations generated from the QM9S dataset. The top row displays the 2D graphs, the middle row shows the reference structures, while the bottom shows the conformations generated by MSpecTmol with their corresponding RMSD values.

ative to baselines, particularly for molecules with complex structures. This advantage is primarily attributed to the PA-IB framework, which selectively filters out redundant or non-informative signals and retains only the most relevant structural information, thereby enhancing predictive reliability. **To further validate the model’s generalization capability on unseen complex structures, we performed a stress test on the top 10% largest molecules in the dataset. Results provided in Appendix K demonstrate that MSpecTmol maintains robust performance (F1-score of 0.925) even under significant distribution shifts.**

**Obs.3: MSpecTmol exhibits superior performance across diverse functional groups.** As shown in Figure 3(c), MSpecTmol achieves the highest macro-F1 score, indicating balanced performance across both common and rare substructures. In classifying 37 functional groups, prediction accuracy varies due to intrinsic spectral differences. Groups with weak or overlapping signals, such as alkyl (-CH<sub>3</sub>) and ether (-O-), are more challenging than those with distinct peaks like carbonyl (C=O) and hydroxyl (-OH). To provide more interpretable insights into these variations, we include a detailed analysis based on confusion matrices in Appendix P. Nevertheless, MSpecTmol consistently outperforms baselines, benefiting from the PA-IB framework that filters redundant signals and preserves the most informative structural cues.

**Obs.4: MSpecTmol exhibits superior performance on experimental spectra.** To evaluate the effectiveness of MSpecTmol on real-world data, we constructed a dataset of approximately 12K samples collected from the SDBS Web. As shown in Table 1, all models suffer from a noticeable performance drop under single-modality settings, which can be attributed to the limited amount of data and the inherent complexity of experimental spectra. However, in multi-modal settings, MSpecTmol achieves a substantial improvement over the baselines, reaching a F1 score of 0.913. This result highlights the practical value of MSpecTmol. Furthermore, given that CNNs-based model are inherently data-hungry and the scarcity of multi-modal datasets containing experimental spectra, we investigated various data augmentation strategies in Appendix Q, which yielded additional performance gains. Our analysis in Appendix R confirms that MSpecTmol is also highly robust against noisy and missing spectra, highlighting its reliability for practical deployment.

### 3.3 MODEL PERFORMANCE ON CONFORMATION GENERATION (RQ2)

In this section, we evaluate the effectiveness of MSpecTmol for spectrum-conditioned molecular conformation generation. Specifically, we integrated MSpecTmol as a spectral encoder into a diffusion model, which is trained to generate corresponding atom coordinates by input spectrum and SMILES representations. The detailed algorithmic procedure is presented in the Appendix S.

**Obs.5: MSpecTmol facilitates high-fidelity conformation generation.** **To thoroughly evaluate the benefit of incorporating spectral data, we compared MSpecTmol against both spectrum-free graph-based methods (RDKit, OpenBabel, ConfGF) and spectral-conditioned baselines (GeoDiff, Attention-based model).** As shown in Table 3, MSpecTmol consistently outperforms both baselines across all spectral inputs. This performance advantage becomes more pronounced when multi-

432  
 433 **Table 3: Comparison of mean RMSD for molecular conformation generation. We evaluate base-**  
 434 **lines including graph-based methods (RDKit, OpenBabel, ConfGF) and spectral-conditioned mod-**  
 435 **els. Lower RMSD indicates better agreement with reference structures.**

436 437 438 439 440 441 442 443	444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485	476 <b>Graph-based Baselines</b>			477 <b>Spectral-Conditioned Models</b>		
		478 <b>RDKit</b>	479 <b>OpenBabel</b>	480 <b>ConfGF</b>	481 <b>GeoDiff</b>	482 <b>Attention</b>	483 <b>MSpecTmol</b>
UV	1.350	1.279	1.143	1.125 <sub>(0.022)</sub>	0.718 <sub>(0.009)</sub>	0.697 <sub>(0.007)</sub>	
IR	1.350	1.279	1.143	1.233 <sub>(0.020)</sub>	0.726 <sub>(0.010)</sub>	0.706 <sub>(0.009)</sub>	
Raman	1.350	1.279	1.143	1.042 <sub>(0.032)</sub>	0.735 <sub>(0.011)</sub>	0.701 <sub>(0.008)</sub>	
UV + (IR, Raman)	1.350	1.279	1.143	0.882 <sub>(0.019)</sub>	0.714 <sub>(0.008)</sub>	0.682 <sub>(0.006)</sub>	
IR + (UV, Raman)	1.350	1.279	1.143	0.882 <sub>(0.019)</sub>	0.714 <sub>(0.007)</sub>	0.689 <sub>(0.007)</sub>	

modal spectra are utilized, indicating that our model generates conformations that are highly consistent with the input spectral data. This advantage stems from MSpecTmol’s superior ability to fuse information from multiple spectra into a rich, unified representation for the generative task. Furthermore, the qualitative results in Figure 4 corroborate these findings, illustrating our model’s capacity to translate complex spectral patterns into high-fidelity molecular structures.

**Obs.6: MSpecTmol achieves superior stability through effective multi-modal fusion.** Figure 6 demonstrates the superior performance of MSpecTmol, which consistently generates conformations with a lower and more tightly concentrated RMSD distribution than all baselines. This result highlights the necessity of multi-modal fusion, as single-modality models yield higher errors and greater variance. Moreover, the stark performance degradation observed in the ablation variant, which is evidenced by its high and widely dispersed RMSD—confirms that our proposed fusion architecture is critical for ensuring both the accuracy and stability of the generative process.

### 3.4 INTERPRETALITY AND ABLATION STUDY AND SENSITIVITY ANALYSIS (RQ3)

In this section, we further investigate the intrinsic relationships between different spectral segments and molecular substructures. Additionally, we analyze the contributions of individual model components and examine the model’s sensitivity to hyperparameter variations, detailed in Figure 11.

**Obs.7: Different spectral modalities emphasize distinct molecular features.** To probe the relationship between spectral segments and molecular substructures, we designed an experiment where a model predicts a single functional group from a concatenated input of three spectra. By averaging the importance scores in the dataset, we obtained attention distribution images for the three modalities, as shown in Figure 5. In the overall image, we can observe that different functional groups tend to focus on different spectra, which may indicate that each spectrum contains information with a distinct emphasis. The varying roles of spectral information highlight the necessity of utilizing information bottleneck theory to extract supplementary information from auxiliary modalities that enhance the main modality. More results could be found in Appendix O. **Additionally, to demonstrate the superiority of our PA-IB framework over conventional integration paradigms, we compared MSpecTmol against various standard fusion strategies. Detailed comparisons and analysis are provided in Appendix N.**

**Obs.8:  $\alpha$  and  $\beta$  regulate the trade-off between prediction accuracy and information compression.** We analyze the joint effect of  $\alpha$  and  $\beta$  in balancing prediction accuracy and information compression, as defined in equation 1. As shown in Figure 11(a), setting  $\alpha = \beta = 1 \times 10^{-6}$  consistently yields the best performance. Larger values of  $\alpha$  and  $\beta$  impose excessive compression, hindering the model’s ability to retain crucial essential spectral features. Conversely, smaller values preserve more input information but fail to effectively suppress redundancy, ultimately compromising generalization. These results highlight the critical importance of properly tuning  $\alpha$  and  $\beta$  to maintain an optimal trade-off.

**Obs.9: Regulating Information Processing at Different Levels via the Primary–Auxiliary Information Bottleneck.** We conducted ablation studies on the multi-modal spectroscopic dataset. As shown in Figure 11(b), with all loss terms, the model achieved an F1-score of 0.9589. Removing the KL divergence loss for primary spectra ( $\mathcal{L}_{MI^2}$ ), which regulates compression in the primary modal-

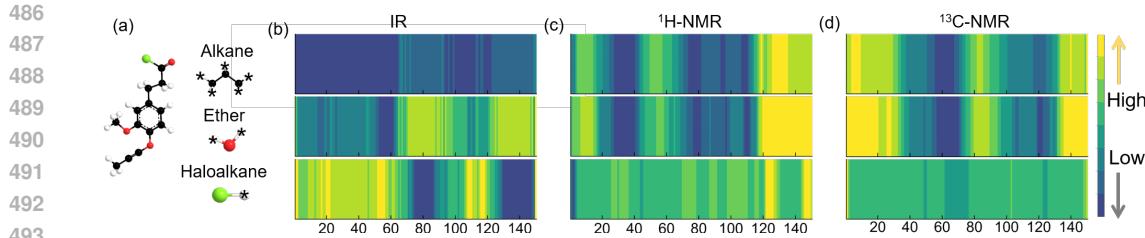


Figure 5: Illustration of spectral information importance. (a) Molecular structure of CC#COc1ccc(CCC(=O)Cl)cc1OC, which contains Alkane, Ether, and Haloalkane functional groups. (b) Importance map of IR spectra, (c)  $^{13}\text{C}$ -NMR spectra, (d) and  $^1\text{H}$ -NMR spectra. The x-axis of all three spectral plots is normalized to the range [0, 150].

Table 4: Performance comparison (F1-score) between our asymmetric PA-IB and the symmetric Uniform IB baseline across different modality configurations.

Fusion Strategy	IR	$^{13}\text{C}$ -NMR	$^1\text{H}$ -NMR	Multi-Modal (All)
Symmetric (Uniform IB)	0.906	0.904	0.905	0.934
<b>Asymmetric (Ours)</b>	<b>0.923</b>	<b>0.920</b>	<b>0.927</b>	<b>0.959</b>

ity, reduced the F1-score to 0.9543. Removing the KL divergence for auxiliary spectra ( $\mathcal{L}_{MI^1}$  and  $\mathcal{L}_{MI^3}$ ) led to a more substantial decrease to 0.9521, indicating that uncompressed auxiliary information introduces noise that undermines prediction quality. These results underscore the importance of regulating information at both levels to suppress redundancy and preserve relevant features.

**Obs.10: Superiority of Asymmetric Design over Symmetric Fusion.** To demonstrate the effectiveness of our asymmetric fusion strategy, we compared MSpecTmol against a symmetric Uniform IB baseline, which applies uniform information bottleneck constraints across all modalities independently without the primary-auxiliary distinction. As presented in Table 4, MSpecTmol consistently outperforms the symmetric approach. Notably, in the full multi-modal setting, our method achieves a 2.5% absolute gain (0.959 vs. 0.934), confirming that the asymmetric design effectively suppresses cross-modal redundancy that the symmetric strategy fails to address.

## 4 CONCLUSION AND FUTURE OUTLOOK

In this work, we introduce **MSpecTmol**, a multi-modal spectrum information fusion framework based on the information bottleneck principle, designed for molecular structure determination. Our framework adopts a primary-auxiliary synergistic modeling approach, which distills core information from a primary modality while leveraging auxiliary spectra to supplement and refine the final representation. Rigorous experimental evaluations validate MSpecTmol’s end-to-end effectiveness, achieving a SOTA F1-score of 0.959 in molecular identification and a low average RMSD of 0.682Å in 3D conformation generation. Meanwhile, our model provides chemically interpretable spectroscopic fragment importance, bridging the gap between ML predictions and domain knowledge.

Looking forward, this framework not only assists chemists in unraveling complex molecular systems but also accelerates the analysis of novel compounds. MSpecTmol holds potential to benefit diverse scientific domains—such as drug discovery, materials science, and chemical forensics—where accurate and reliable molecular identification is critical. MSpecTmol paves the way toward democratized, efficient, and interpretable molecular analysis for broad scientific and industrial applications.

## 5 REPRODUCIBILITY

We provide the complete implementation in the repository along with guidance on how to reproduce our results. Our code is available at <https://anonymous.4open.science/r/MspecTmol-6B4D>.

540 6 ETHICS STATEMENT  
541542 Our study does not involve human participants, personal data, or sensitive information. The datasets  
543 and resources used are either publicly available or released under appropriate licenses. We confirm  
544 that our research does not raise any ethical concerns related to privacy, safety, fairness, or potential  
545 misuse. The contributions of this work are intended solely for advancing scientific research and are  
546 not designed or evaluated for harmful applications.  
547548 REFERENCES  
549550 Marvin Alberts, Oliver Schilter, Federico Zipoli, Nina Hartrampf, and Teodoro Laino. Unraveling  
551 molecular structure: A multimodal spectroscopic dataset for chemistry. In *Annual Conference on*  
552 *Neural Information Processing Systems*, 2024.553 Marvin Alberts, Federico Zipoli, and Teodoro Laino. Setting new benchmarks in ai-driven infrared  
554 structure elucidation. *Digital Discovery*, 2025.  
555556 Valentine P Ananikov. Artificial intelligence chemistry. *Artificial Intelligence*, 2:100075, 2024.  
557558 Vincenzo Barone, Silvia Alessandri, Małgorzata Biczysko, James R Cheeseman, David C Clary,  
559 Anne B McCoy, Ryan J DiRisio, Frank Neese, Mattia Melosso, and Cristina Puzzarini. Compu-  
560 tational molecular spectroscopy. *Nature Reviews Methods Primers*, 1(1):38, 2021.561 Tom L Blundell. Structure-based drug design. *Nature*, 384(6604):23, 1996.  
562563 Riya Bose, Xiaohe Zhou, Tianle Guo, Haoze Yang, Jun Yin, Aditya Mishra, Jason D Slinker, Os-  
564 man M Bakr, Omar F Mohammed, and Anton V Malko. Single-particle spectroscopy as a versatile  
565 tool to explore lower-dimensional structures of inorganic perovskites. *ACS Energy Letters*, 6(10):  
566 3695–3708, 2021.567 Bhanu Chander, Chinju John, Lekha Warrier, and Kumaravelan Gopalakrishnan. Toward trustwor-  
568 thy artificial intelligence (tai) in the context of explainability and robustness. *ACM Computing*  
569 *Surveys*, 2024.  
570571 Jian Chen, Jie-qing Li, Tao Li, Hong-gao Liu, and Yuan-zhong Wang. Application of uv-vis and  
572 infrared spectroscopy on wild edible bolete mushrooms discrimination and evaluation: a review.  
573 *Critical Reviews in Analytical Chemistry*, 53(4):852–868, 2023.574 V Divya, V Aparna, Sunitha Ashna Rasheed, Lakshmy M V R, and K Anu. The transformative  
575 role of artificial intelligence in advancing chemistry: A comprehensive review. *Library Progress*  
576 *International*, 44(3):23917–23927, 2024.  
577578 Wenjie Du, Xiaoting Yang, Di Wu, FenFen Ma, Baicheng Zhang, Chaochao Bao, Yaoyuan Huo,  
579 Jun Jiang, Xin Chen, and Yang Wang. Fusing 2d and 3d molecular graphs as unambiguous  
580 molecular descriptors for conformational and chiral stereoisomers. *Briefings in Bioinformatics*,  
581 24(1):bbac560, 2023.582 Evan N Feinberg, Debnil Sur, Zhenqin Wu, Brooke E Husic, Huanghao Mai, Yang Li, Saisai Sun,  
583 Jianyi Yang, Bharath Ramsundar, and Vijay S Pande. Potentialnet for molecular property predic-  
584 tion. *ACS central science*, 4(11):1520–1530, 2018.  
585586 Carolina Fontana and Goran Widmalm. Primary structure of glycans by nmr spectroscopy. *Chemical*  
587 *Reviews*, 123(3):1040–1102, 2023.588 Meihong Ge, Pan Li, Guoliang Zhou, Siyu Chen, Wei Han, Feng Qin, Yuman Nie, Yaoxiong Wang,  
589 Miao Qin, Guangyao Huang, et al. General surface-enhanced raman spectroscopy method for  
590 actively capturing target molecules in small gaps. *Journal of the American Chemical Society*, 143  
591 (20):7769–7776, 2021.592 Garrett B Goh, Nathan O Hudas, and Abhinav Vishnu. Deep learning for computational chemistry.  
593 *Journal of computational chemistry*, 38(16):1291–1307, 2017.

594 Janna Hastings, Martin Glauer, Adel Memariani, Fabian Neuhaus, and Till Mossakowski. Learning  
 595 chemistry: exploring the suitability of machine learning for the task of structure-based chemical  
 596 ontology classification. *Journal of Cheminformatics*, 13:1–20, 2021.

597

598 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv  
 599 preprint arXiv:1611.01144*, 2016.

600 Hongchao Ji, Hanzi Deng, Hongmei Lu, and Zhimin Zhang. Predicting a molecular fingerprint from  
 601 an electron ionization mass spectrum with deep neural networks. *Analytical chemistry*, 92(13):  
 602 8649–8653, 2020.

603

604 Guwon Jung, Son Gyo Jung, and Jacqueline M Cole. Automatic materials characterization from in-  
 605 frared spectra using convolutional neural networks. *Chemical Science*, 14(13):3600–3609, 2023.

606 Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

607

608 Guillaume Klein, Yoon Kim, Yuntian Deng, Vincent Nguyen, Jean Senellart, and Alexander M  
 609 Rush. Opennmt: Neural machine translation toolkit. *arXiv preprint arXiv:1805.11462*, 2018.

610

611 Namkyeong Lee, Dongmin Hyun, Gyoung S Na, Sungwon Kim, Junseok Lee, and Chanyoung  
 612 Park. Conditional graph information bottleneck for molecular relational learning. *arXiv preprint  
 613 arXiv:2305.01520*, 2023.

614

615 Yanmin Liu, Xuan Zhang, Wei Zhao, Daming Zhu, and Xuefeng Cui. De novo molecular structure  
 616 generation from mass spectra. In *2023 IEEE International Conference on Bioinformatics and  
 617 Biomedicine (BIBM)*, pp. 373–378. IEEE, 2023.

618

Marcus Ludwig, Louis-Félix Nothias, Kai Dührkop, Irina Koester, Markus Fleischauer, Martin A  
 619 Hoffmann, Daniel Petras, Fernando Vargas, Mustafa Morsy, Lihini Aluwihare, et al. Database-  
 620 independent molecular formula annotation using gibbs sampling through zodiac. *Nature Machine  
 621 Intelligence*, 2(10):629–641, 2020.

622

Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous  
 623 relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

624

Kumaran Manogaran, T Sivarajani, Periandy Sengeny, VSK Venkatachalapathy, M Mahadevan,  
 625 K Elangovan, Stevan Armaković, Sanja J Armaković, and Biljana F Abramović. Investigation  
 626 on molecular and biomolecular spectroscopy of the novel 2bca molecule to analyse its biological  
 627 activities and binding interaction with nipah viral protein. *Spectrochimica Acta Part A: Molecular  
 628 and Biomolecular Spectroscopy*, 321:124737, 2024.

630

Carlos A Meza Ramirez, Michael Greenop, Lorna Ashton, and Ihtesham Ur Rehman. Applications  
 631 of machine learning in spectroscopy. *Applied Spectroscopy Reviews*, 56(8-10):733–763, 2021.

632

Nitin Rane, Mallikarjuna Paramesha, and Jayesh Rane. Trustworthy artificial intelligence: Enhanc-  
 633 ing trustworthiness through explainable ai (xai). Available at SSRN 4873183, 2024.

635

Rafael Cardoso Rial. Ai in analytical chemistry: Advancements, challenges, and future directions.  
 636 *Talanta*, pp. 125949, 2024.

637

Laura Riccardi, Vito Genna, and Marco De Vivo. Metal–ligand interactions in drug design. *Nature  
 639 Reviews Chemistry*, 2(7):100–112, 2018.

640

Shuaike Shen, Jiaqing Xie, Zhuo Yang, Antong Zhang, Shuzhou Sun, Ben Gao, Tianfan Fu,  
 641 Biqing Qi, and Yuqiang Li. Molspectllm: A molecular foundation model bridging spectroscopy,  
 642 molecule elucidation, and 3d structure generation. *arXiv preprint arXiv:2509.21861*, 2025.

643

Francesca Stanzione, Ilenia Giangreco, and Jason C Cole. Use of molecular docking computational  
 645 tools in drug discovery. *Progress in medicinal chemistry*, 60:273–343, 2021.

646

Michael A Stravs, Kai Dührkop, Sebastian Böcker, and Nicola Zamboni. Msnovelist: de novo  
 647 structure generation from mass spectra. *Nature Methods*, 19(7):865–870, 2022.

648 Yunyue Su, Jiahui Chen, Zao Jiang, Zhenyi Zhong, Liang Wang, and Qiang Liu. Language models  
 649 can understand spectra: A multimodal model for molecular structure elucidation. *arXiv preprint*  
 650 *arXiv:2508.08441*, 2025.

651  
 652 W Patrick Walters and Regina Barzilay. Applications of deep learning in molecule generation and  
 653 molecular property prediction. *Accounts of chemical research*, 54(2):263–270, 2020.

654 GuiShen Wang, Hui Feng, and Chen Cao. Birnn-ddi: A drug-drug interaction event type prediction  
 655 model based on bidirectional recurrent neural network and graph2seq representation. *Journal of*  
 656 *Computational Biology*, 2024.

657  
 658 Liang Wang, Yu Rong, Tingyang Xu, Zhenyi Zhong, Zhiyuan Liu, Pengju Wang, Deli Zhao, Qiang  
 659 Liu, Shu Wu, and Yang Zhang. Diffspectra: Molecular structure elucidation from spectra using  
 660 diffusion models. *arXiv preprint arXiv:2507.06853*, 2025.

661 Jennifer N Wei, David Belanger, Ryan P Adams, and D Sculley. Rapid prediction of electron-  
 662 ionization mass spectrometry using neural networks. *ACS central science*, 5(4):700–708, 2019.

663  
 664 Wenjin Wu, Ales Leonardis, Jianbo Jiao, Jun Jiang, and Linjiang Chen. Transformer-based models  
 665 for predicting molecular structures from infrared spectra using patch-based self-attention. *The*  
 666 *Journal of Physical Chemistry A*, 129(8):2077–2085, 2025.

667  
 668 Jun Xia, Chengshuai Zhao, Bozhen Hu, Zhangyang Gao, Cheng Tan, Yue Liu, Siyuan Li, and Stan Z  
 669 Li. Mole-bert: Rethinking pre-training graph neural networks for molecules. In *The Eleventh*  
 670 *International Conference on Learning Representations*, 2023.

671 Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geo-  
 672 metric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*,  
 673 2022.

674 Qian Xue, Yang Yang, Wenkai Ma, Hanqiu Zhang, Daoli Zhang, Xinzhen Lan, Liang Gao, Jian-  
 675 bing Zhang, and Jiang Tang. Advances in miniaturized computational spectrometers. *Advanced*  
 676 *Science*, 11(47):2404448, 2024.

677 Zhuo Yang, Jiaqing Xie, Shuaike Shen, Daolang Wang, Yeyun Chen, Ben Gao, Shuzhou Sun, Biqing  
 678 Qi, Dongzhan Zhou, Lei Bai, et al. Spectrumworld: Artificial intelligence foundation for spec-  
 679 troscopy. *arXiv preprint arXiv:2508.01188*, 2025.

680 Junchi Yu, Tingyang Xu, Yu Rong, Yatao Bian, Junzhou Huang, and Ran He. Graph information  
 681 bottleneck for subgraph recognition. *arXiv preprint arXiv:2010.05563*, 2020.

682  
 683 Junchi Yu, Jie Cao, and Ran He. Improving subgraph recognition with variational graph informa-  
 684 tion bottleneck. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 685 *Recognition*, pp. 19396–19405, 2022.

686  
 687 Yan Zhao, Jun Yin, Li Zhang, Yong Zhang, and Xing Chen. Drug–drug interaction prediction:  
 688 databases, web servers and computational models. *Briefings in Bioinformatics*, 25(1):bbad445,  
 689 2024.

690 Zihan Zou, Yujin Zhang, Lijun Liang, Mingzhi Wei, Jiancai Leng, Jun Jiang, Yi Luo, and Wei Hu.  
 691 A deep learning model for predicting selected organic molecular spectra. *Nature Computational*  
 692 *Science*, 3(11):957–964, 2023.

693  
 694  
 695  
 696  
 697  
 698  
 699  
 700  
 701

702 A USE OF LARGE LANGUAGE MODELS (LLMs)  
703704 In preparing this manuscript, we used a large language model (LLM) solely for writing assistance  
705 and text refinement (e.g., grammar correction, stylistic polishing, and conciseness). The LLM was  
706 not involved in research design, data analysis or model implementation. All technical content, ex-  
707 periments, and conclusions were conceived, executed, and validated by the authors.  
708709 B TRAINING SETTINGS  
710711 For the functional group classification task, the model was trained for 100 epochs with a batch size  
712 of 128, using the Adam optimizer with an initial learning rate of  $4 \times 10^{-4}$  and a cosine annealing  
713 scheduler. The loss coefficients  $\alpha$  and  $\beta$  were set to  $1 \times 10^{-6}$ , while the weight for the auxiliary  
714 prediction loss was 0.7. The temperature for the information bottleneck's stochastic gating was  
715 maintained at 1.0. For the spectrum-conditioned molecular conformation generation task, the model  
716 was trained for 10,000 iterations with a batch size of 64. We employed the Adam optimizer with a  
717 learning rate of  $1 \times 10^{-3}$ , which was adjusted by a plateau scheduler based on validation loss. For  
718 this task, the loss coefficients  $\alpha$  and  $\beta$  were both set to  $1 \times 10^{-6}$ . Both models were trained on two  
719 NVIDIA A100 GPUs (80 GB each). The classification model required approximately 6 hours of  
720 training, while the conformation generation model took around 30 hours.  
721722 C RELATED WORK  
723724 D RELATED WORK  
725726 D.1 SPECTROSCOPY-BASED MOLECULAR MODELING  
727728 Machine learning has advanced spectroscopy-based molecular structure prediction significantly. Liu  
729 et al. introduced MS2SMILES Liu et al. (2023), treating hydrogen atoms as implicitly linked to  
730 heavy atoms, improving molecular generation accuracy. Ji et al. Ji et al. (2020) presented DeepEI,  
731 a deep learning framework for elucidating structures from EI-MS spectra. Wei et al Wei et al. (2019).  
732 developed NEIMS, a neural network model that captures fragmentation patterns from electron ion-  
733 ization for rapid small molecule mass spectrum prediction. Marcus proposed ZODIAC Ludwig et al.  
734 (2020), leveraging tandem mass spectrometry (MS/MS) for molecular formula generation. Michael  
735 A. et al. Stravs et al. (2022) further explored de novo molecular structure generation using RNN  
736 models.  
737738 More recently, the rapid development of generative AI has expanded spectral analysis into the realm  
739 of Large Language Models and Diffusion Models. For instance, DiffSpectra Wang et al. (2025)  
740 introduces a generative framework that formulates structure elucidation as a conditional diffusion  
741 process, enabling the end-to-end generation of 3D molecular conformations from multi-modal spec-  
742 tra. In parallel, Large Language Models have been adapted for this domain: SpectraLLM Su et al.  
743 (2025) and MolSpectLLM Shen et al. (2025) leverage the reasoning capabilities of heavy trans-  
744 former backbones to treat spectrum-to-structure translation as a sequence generation task, bridg-  
745 ing spectroscopy with textual molecular representations. Furthermore, SpectrumWorld Yang et al.  
746 (2025) expands this frontier by introducing a multi-modal agent framework and benchmark suite to  
747 systematize deep learning research in spectroscopy.748 While these methods have achieved notable success, previous approaches often rely heavily on sin-  
749 gle modalities like mass spectrometry or require massive computational resources. Additionally,  
750 mass spectrometry is costly, sensitive to interference, and challenging to standardize in automated  
751 workflows. Recently, Marvin et al. Alberts et al. (2024) released a 790k Multimodal Spectroscopic  
752 Dataset, providing a foundation for integrating multi-spectroscopic data. Their work introduced  
753 baseline models for tasks like molecular structure prediction and functional group identification  
754 from spectral data, forming a key resource for our research. These studies highlight the potential  
755 and limitations of current methods, motivating our approach to integrate multi-spectroscopic modal-  
756 ities for enhanced molecular structure determination.

756 D.2 INFORMATION BOTTLENECK (IB) THEORY  
757

758 The IB theory provides a principled framework for extracting compact and informative substructures  
759 from complex data, playing a key role in challenges like denoising and compression. PGIB Yu et al.  
760 (2020) extends IB by introducing a framework with a mutual information estimator for irregular  
761 graph data, and a connectivity loss to stabilize information extraction. VGIB Yu et al. (2022) further  
762 improves stability by injecting Gaussian noise into node representations, regulating information flow  
763 between original and perturbed graphs. Lee et al. Lee et al. (2023) expanded IB to paired graphs  
764 with the Conditional GIB, optimizing compressed information extraction by retaining only the most  
765 relevant information. While effective, these approaches focus on single-target tasks and lack strategies  
766 for redundancy reduction and complementary integration under multi-modal conditions. This  
767 growing body of work underscores the versatility of IB theory while highlighting opportunities for  
768 further refinement, particularly in handling multi-modal scenarios, where redundancy removal and  
769 cross-modal synergy are essential.

770 E BROADER IMPACTS AND LIMITATION DISCUSSION  
771772 E.1 BROADER IMPACTS  
773

774 This work promotes automated and interpretable molecular structure elucidation via multi-modal  
775 spectroscopic learning. MSpecTmol may assist domains such as drug discovery, materials science,  
776 and chemical forensics by providing chemically intuitive insights and reducing reliance on manual  
777 spectral interpretation. Its interpretable design supports broader and more accessible molecular  
778 analysis. The proposed framework can reduce reliance on extensive manual spectral interpretation,  
779 democratizing molecular analysis for broader scientific and industrial use.

780 E.2 LIMITATIONS  
781

782 While MSpecTmol demonstrates strong performance, several limitations remain. First, the model's  
783 effectiveness depends on the availability of complete, multi-modal spectra, which are often scarce  
784 in practice and may hinder its deployment on incomplete datasets. Additionally, its training on  
785 a fixed vocabulary of functional groups restricts its ability to identify rare or novel substructures,  
786 particularly when analyzing new chemical entities. Future work will focus on addressing these  
787 challenges to improve the model's robustness and expand its chemical scope.

788 F PROOF  
789790 F.1 PROOF OF PROPOSITION  
791

792  $p_\theta(\mathbf{Y}|T_m)$  is variational approximation of  $p(\mathbf{Y}|T_m)$ . We model  $p_\theta(\mathbf{Y}|T_m)$  as a predictor  
793 parametrized by  $\theta$ , which outputs the model prediction  $\mathbf{Y}$  based on the core primary spectra  $T_m$ .

$$\begin{aligned} 794 I(\mathbf{Y}; T_m) &= \mathbb{E}_{\mathbf{Y}, T_m} [\log \frac{p(\mathbf{Y}|T_m)}{p(\mathbf{Y})}] \\ 795 &= \mathbb{E}_{\mathbf{Y}, T_m} [\log \frac{p_\theta(\mathbf{Y}|T_m)}{p(\mathbf{Y})}] \\ 796 &\quad + \mathbb{E}_{T_m} [KL(p(\mathbf{Y}|T_m) || p_\theta(\mathbf{Y}|T_m))] \end{aligned} \tag{15}$$

802 According to the non-negativity of the KL divergence, we have:

$$\begin{aligned} 803 I(\mathbf{Y}; T_m) &\geq \mathbb{E}_{\mathbf{Y}, T_m} [\log \frac{p_\theta(\mathbf{Y}|T_m)}{p(\mathbf{Y})}] \\ 804 &= \mathbb{E}_{\mathbf{Y}, T_m} [\log p_\theta(\mathbf{Y}|T_m)] + H(\mathbf{Y}) \end{aligned} \tag{16}$$

805 Thus, we can minimize the upper bound of  $-I(\mathbf{Y}; T_m)$  by minimizing the model prediction loss  
806  $\mathcal{L}_{\text{pred}}(\mathbf{Y}, T_m)$ , which can be modeled as the cross entropy loss for classification and the mean  
807 square loss for regression.

810 F.1.1 MINIMIZING  $-I(Y; T_m)$   
811812 (Upper bound of  $-I(Y; T_m)$ ) Given the primary spectra  $X_m$ , and its label information  $\mathbf{Y}$ , we have  
813

814 
$$\begin{aligned} -I(\mathbf{Y}; T_m) &\leq \mathbb{E}_{T_m, \mathbf{Y}}[-\log p_\theta(\mathbf{Y}|T_m)] \\ &= \mathbb{E}_{(\mathbf{Y}; T_m)} \log [P_\theta(\mathbf{Y} | T_m)] + H(\mathbf{Y}) := \mathcal{L}_{pred}, \end{aligned} \quad (17)$$
  
815

816 where  $H(\mathbf{Y})$  is constant across all data, it will be omitted in the model optimization process.  
817  $p_\theta(\mathbf{Y}|T_m)$  is variational approximation of  $p(\mathbf{Y}|T_m)$ . A detailed proof for proposition is given in  
818 Appendix F.1.2.819 F.1.2 MINIMIZING  $-I(Y; T_a | T_m)$   
820821 For the second term of Equation 9, i.e.,  $-I(Y; T_a | T_m)$ , we decompose the term into the sum of  
822 two terms based on the chain rule of mutual information as follows:  
823

824 
$$I(Y; T_a | T_m) = I(Y; T_a, T_m) - I(T_a; T_m). \quad (18)$$

825 For the upper bound of  $-I(Y; T_a, T_m)$ , Given the core primary spectra  $T_m$  and core auxiliary spectra  
826  $T_a$ , and its label information  $\mathbf{Y}$ , we have  
827

828 
$$\begin{aligned} -I(\mathbf{Y}; T_a, T_m) &\leq \mathbb{E}_{(T_a, T_m, \mathbf{Y})}[-\log p_\theta(\mathbf{Y}|T_a, T_m)] \\ &= \mathbb{E}_{(\mathbf{Y}, T_a, T_m)} \log [P_\theta(\mathbf{Y} | T_a, T_m)] + H(\mathbf{Y}) := \mathcal{L}_{sup}, \end{aligned} \quad (19)$$
  
829

830 where  $p_\theta(\mathbf{Y}|T_a, T_m)$  is variational approximation of  $p(\mathbf{Y}|T_a, T_m)$ . We model  $p_\theta(\mathbf{Y}|T_a, T_m)$  as a  
831 predictor parametrized by  $\theta$ , which outputs the model prediction  $\mathbf{Y}$  based on the core spectra  $T_a$  and  
832  $T_m$ . Thus, we can minimize the upper bound of  $-I(\mathbf{Y}; T_a, T_m)$  by minimizing the supplementary  
833 prediction loss  $\mathcal{L}_{sup}$ ,834 For the upper bound of  $I(T_a; T_m)$ , drawing inspiration from the experiences derived in Variational  
835 Autoencoders (VAE) Kingma (2013), we attempt to replace  $p(t_a)$  with  $q(t_a)$  and consolidate the  
836 additional components to form a Kullback-Leibler (KL) divergence:  
837

838 
$$\begin{aligned} I(T_a; T_m) &= \mathbb{E}_{(t_m, t_a) \sim p(t_m, t_a)} \left[ \log \frac{p(t_a | t_m)}{p(t_a)} \right] \\ &= \mathbb{E}_{(t_m, t_a) \sim p(t_m, t_a)} \left[ \log \frac{p(t_a | t_m)}{q(t_a)} \cdot \frac{q(t_a)}{p(t_a)} \right] \\ &= \mathbb{E}_{(t_m, t_a) \sim p(t_m, t_a)} \left[ \log \frac{p(t_a | t_m)}{q(t_a)} \right] \\ &\quad + \mathbb{E}_{(t_m, t_a) \sim p(t_m, t_a)} \left[ \log \frac{q(t_a)}{p(t_a)} \right] \end{aligned} \quad (20)$$
  
839

840 For the first term, both  $p(t_a | t_m)$  and  $q(t_a)$  have analytical forms, allowing the function within the  
841 brackets to be computed analytically. By utilizing the relationship  $p(t_m, t_a) = p(t_m)p(t_a | t_m)$ , we  
842 can rewrite the first term in a more elegant manner:  
843

844 
$$\begin{aligned} &\mathbb{E}_{(t_m, t_a) \sim p(t_m, t_a)} \left[ \log \frac{p(t_a | t_m)}{q(t_a)} \right] \\ &= \iint p(t_m)p(t_a | t_m) \log \frac{p(t_a | t_m)}{q(t_a)} dt_a dt_m \\ &= \int p(t_m) \left( \int p(t_a | t_m) \log \frac{p(t_a | t_m)}{q(t_a)} dt_a \right) dt_m \\ &= \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_a | t_m) \| q(t_a))] \\ &:= L_{MI^1} \approx \frac{1}{N} \sum_{i=1}^N \text{KL}[p(t_a | t_{mi}) \| q(t_a)], \quad t_{mi} \sim p(t_m) \end{aligned} \quad (21)$$
  
845

846 The term  $L_{MI^1} := \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_a | t_m) \| q(t_a))]$  is often referred to as the *rate* in rate-  
847 distortion theory. This rate component can be optimized using mini-batch gradient descent. Specifi-  
848 cally, by sampling a batch of training samples  $t_{m1}, \dots, t_{mN}$  from the training set, we can minimize  
849 the KL divergence  $\text{KL}[p(t_a | t_{mi}) \| q(t_a)]$  for each  $t_{mi}$ .  
850

864 Since both distributions  $p(t_a | t_m)$  and  $q(t_a)$  are Gaussian, the KL divergence between them has an  
 865 analytical solution:

$$\begin{aligned}
 866 \quad & \text{KL}[p(t_a | t_m) \| q(t_a)] \\
 867 \quad & = \text{KL}[\mathcal{N}(\mu(t_m), \sigma^2(t_m) \mathbf{I}) \| \mathcal{N}(\mathbf{0}, \mathbf{I})] \\
 868 \quad & = \sum_{j=1}^J \text{KL}[\mathcal{N}(\mu_j, \sigma_j^2) \| \mathcal{N}(0, 1)] \\
 869 \quad & = \sum_{j=1}^J \frac{1}{2} (-\log \sigma_j^2 - 1 + \mu_j^2 + \sigma_j^2)
 \end{aligned} \tag{22}$$

870 Here,  $\mu(t_m), \sigma^2(t_m)$  are mean and variance of  $\mathbf{H}^m$ , respectively.

### 871 F.1.3 MINIMIZING $-I(X_m; T_m)$

872 For the upper bound of  $-I(X_m; T_m)$ , Similarly, we attempt to replace  $p(t_m)$  with  $q(t_m)$  and  
 873 consolidate the additional components to form a Kullback-Leibler (KL) divergence:

$$\begin{aligned}
 874 \quad & I(X_m; T_m) = \mathbb{E}_{(t_m, x_m) \sim p(t_m, x_m)} \left[ \log \frac{p(t_m | x_m)}{q(t_m)} \right] \\
 875 \quad & + \mathbb{E}_{(t_m, x_m) \sim p(t_m, x_m)} \left[ \log \frac{q(t_m)}{p(t_m)} \right]
 \end{aligned} \tag{23}$$

876 By utilizing the relationship  $p(t_m, x_m) = p(x_m)p(t_m | x_m)$ , we can rewrite the first term in a more  
 877 elegant manner:

$$\begin{aligned}
 878 \quad & \mathbb{E}_{(t_m, x_m) \sim p(t_m, x_m)} \left[ \log \frac{p(t_m | x_m)}{q(t_m)} \right] \\
 879 \quad & = \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_m | x_m) \| q(t_m))] := L_{MI^2}
 \end{aligned} \tag{24}$$

880 The term  $L_{MI^2} := \mathbb{E}_{t_m \sim p(t_m)} [\text{KL}(p(t_m | x_m) \| q(t_m))]$  is often referred to as the *rate* in rate-  
 881 distortion theory. This rate component can be optimized using mini-batch gradient descent. Specifically,  
 882 by sampling a batch of training samples  $t_{m1}, \dots, t_{mN}$  from the training set, we can minimize  
 883 the KL divergence  $\text{KL}[p(t_m | x_{mi}) \| q(t_m)]$  for each  $x_{mi}$ .

884 Since both distributions  $p(t_m | x_m)$  and  $q(t_m)$  are Gaussian, the KL divergence between them has  
 885 an analytical solution:

$$\text{KL}[p(t_m | x_m) \| q(t_m)] = \sum_{j=1}^J \frac{1}{2} (-\log \sigma_j^2 - 1 + \mu_j^2 + \sigma_j^2) \tag{25}$$

### 886 F.1.4 MINIMIZING $-I(T_a; X_m, X_a)$

887 For the upper bound of  $-I(T_a; X_m, X_a)$ , Similarly, we attempt to replace  $p(t_a)$  with  $q(t_a)$  and  
 888 consolidate the additional components to form a Kullback-Leibler (KL) divergence:

$$\begin{aligned}
 889 \quad & I(T_a; X_m, X_a) = \mathbb{E}_{(t_a, x_a, x_m) \sim p(t_a, x_a, x_m)} \left[ \log \frac{p(t_a | x_m, x_a)}{q(t_a)} \right] \\
 890 \quad & + \mathbb{E}_{(t_a, x_a, x_m) \sim p(t_a, x_a, x_m)} \left[ \log \frac{q(t_a)}{p(t_a)} \right]
 \end{aligned} \tag{26}$$

891 By utilizing the relationship  $p(t_a, x_a, x_m) = p(t_a, x_a)p(t_a | x_m, x_a)$ , we can rewrite the first term  
 892 in a more elegant way:

$$\begin{aligned}
 893 \quad & \mathbb{E}_{(t_a, x_a, x_m) \sim p(t_a, x_a, x_m)} \left[ \log \frac{p(t_a | x_m, x_a)}{q(t_a)} \right] \\
 894 \quad & = \mathbb{E}_{t_a, x_a \sim p(x_m, x_a)} [\text{KL}(p(t_a | x_m, x_a) \| q(t_a))] := L_{MI^3}
 \end{aligned} \tag{27}$$

895 The term  $L_{MI^3} = \mathbb{E}_{x_m, x_a \sim p(x_m, x_a)} [\text{KL}(p(t_a | x_m, x_a) \| q(t_a))]$  is often referred to as the *rate* in  
 896 rate-distortion theory. This rate component can be optimized using mini-batch gradient descent.

918 Specifically, by sampling a batch of training samples  $t_{m1}, \dots, t_{mN}$  from the training set, we can  
 919 minimize the KL divergence  $\text{KL}[p(t_a | t_{mi}) \| q(t_a)]$  for each  $t_{mi}$ .  
 920

921 Since both distributions  $p(t_a | x_m, x_a)$  and  $q(t_a)$  are Gaussian, the KL divergence between them  
 922 has an analytical solution:

$$923 \quad \text{KL}[p(t_a | x_m, x_a) \| q(t_a)] = \sum_{j=1}^J \frac{1}{2} (-\log \sigma_j^2 - 1 + \mu_j^2 + \sigma_j^2) \quad (28)$$

## 927 G DEFINITION OF FUNCTIONAL GROUPS

930 Functional groups play a crucial role in determining the chemical reactivity and properties of  
 931 molecules. To systematically analyze molecular structures, we employ a set of predefined patterns  
 932 to identify key functional groups within a given molecular dataset.

933 Table 5 lists the functional groups considered in this study, along with their corresponding SMARTS  
 934 representations. These functional groups were selected based on their relevance to organic and  
 935 medicinal chemistry, including common moieties such as hydroxyl (-OH), carbonyl (C=O), and  
 936 amine (-NH<sub>2</sub>) groups. The identification process involves scanning molecular structures using sub-  
 937 graph matching algorithms, ensuring accurate detection of these structural motifs.

938  
 939 Table 5: Predefined Functional Groups and Their SMARTS Patterns

940 <b>Functional Group</b>	941 <b>SMARTS Pattern</b>
942 Acid anhydride	<chem>[CX3] (=[OX1]) [OX2] [CX3] (=[OX1])</chem>
943 Acyl halide	<chem>[CX3] (=[OX1]) [F,Cl,Br,I]</chem>
944 Alcohol	<chem>[#6] [OX2H]</chem>
945 Aldehyde	<chem>[CX3H1] (=[O]) [#6,H]</chem>
946 Alkane	<chem>[CX4;H3,H2]</chem>
947 Alkene	<chem>[CX3]=[CX3]</chem>
948 Alkyne	<chem>[CX2]#[CX2]</chem>
949 Amide	<chem>[NX3] [CX3] (=[OX1]) [#6]</chem>
950 Amine	<chem>[NX3;H2,H1,HO;!\$(NC=O)]</chem>
951 Arene	<chem>[cX3]1[cX3] [cX3] [cX3] [cX3]1[cX3]</chem>
952 Azo compound	<chem>[#6] [NX2]=[NX2] [#6]</chem>
953 Carbamate	<chem>[NX3] [CX3] (=[OX1]) [OX2H0]</chem>
954 Carboxylic acid	<chem>[CX3] (=[O]) [OX2H]</chem>
955 Enamine	<chem>[NX3] [CX3]=[CX3]</chem>
956 Enol	<chem>[OX2H]#[#6X3]=[#6]</chem>
957 Ester	<chem>[#6] [CX3] (=[O]) [OX2H0] [#6]</chem>
958 Ether	<chem>[OD2] ([#6]) [#6]</chem>
959 Haloalkane	<chem>[#6] [F,Cl,Br,I]</chem>
960 Hydrazine	<chem>[NX3] [NX3]</chem>
961 Hydrazone	<chem>[NX3] [NX2]=[#6]</chem>
962 Imide	<chem>[CX3] (=[OX1]) [NX3] [CX3] (=[OX1])</chem>
963 Imine	<chem>[\$([CX3] ([#6]) [#6]), \$([CX3H] [#6])]=[\$([NX2] [#6]), \$([NX2H])]</chem>
964 Isocyanate	<chem>[NX2]=[C]=[O]</chem>
965 Isothiocyanate	<chem>[NX2]=[C]=[S]</chem>
966 Ketone	<chem>[#6] [CX3] (=[O]) [#6]</chem>
967 Nitrile	<chem>[NX1]#[CX2]</chem>
968 Phenol	<chem>[OX2H] [cX3]:[c]</chem>
969 Phosphine	<chem>[PX3]</chem>
970 Sulfide	<chem>[#16X2H0]</chem>
971 Sulfonamide	<chem>[#16X4] ([NX3]) (=[OX1]) (=[OX1]) [#6]</chem>
Sulfonate	<chem>[#16X4] (=[OX1]) (=[OX1]) ([#6]) [OX2H0]</chem>
Sulfone	<chem>[#16X4] (=[OX1]) (=[OX1]) ([#6]) [#6]</chem>
Sulfonic acid	<chem>[#16X4] (=[OX1]) (=[OX1]) ([#6]) [OX2H]</chem>
Sulfoxide	<chem>[#16X3]=[OX1]</chem>
Thial	<chem>[CX3H1] (=[S]) [#6,H]</chem>
Thioamide	<chem>[NX3] [CX3]=[SX1]</chem>
Thiol	<chem>[#16X2H]</chem>

972 The functional group identification is performed using cheminformatics libraries such as RDKit,  
 973 which allows for efficient substructure searches within molecular datasets. This approach enables  
 974 us to extract chemically meaningful information and facilitate downstream tasks such as molecular  
 975 property prediction, reactivity analysis, and structure-based clustering.  
 976

## 977 H COMPLEXITY ANALYSIS AND CHOICE OF PRIMARY SPECTRA

978 We conducted a comprehensive analysis to determine the optimal multi-modal configuration for  
 979 MSpecTmol, balancing predictive performance with computational efficiency. The time and space  
 980 complexity of our model and several baselines are presented in Table 6. This analysis reveals a  
 981 clear trade-off between the number of input modalities and the required resources. **As shown in**  
 982 **table 7**, while expanding from three to five spectral inputs nearly doubled the resource consumption,  
 983 it yielded only marginal performance improvements. This finding led us to select a three-modality  
 984 fusion as the most balanced and efficient configuration.  
 985

986 A critical aspect of our framework is the selection of the primary spectrum, as MSpecTmol is de-  
 987 signed to prioritize its features while using auxiliary spectra for supplementary information. Our  
 988 recommended procedure is to first identify the single best-performing modality in standalone ex-  
 989 periments and assign it the primary role, thereby ensuring that the most informative stream is preserved.  
 990

991 To implement this strategy for the functional group classification task, we first assessed the pre-  
 992 dictive power of each individual spectrum (Table 1 and Figure 9(a)). The results revealed that IR  
 993 spectroscopy delivered relatively high and stable accuracy, making it the ideal candidate for the pri-  
 994 mary spectrum. Conversely, MS/MS spectra exhibited the lowest performance. In our multi-modal  
 995 evaluations, we observed that fusing multiple spectra consistently improved performance. Notably,  
 996 the combination of IR,  $^1\text{H}$ -NMR, and  $^{13}\text{C}$ -NMR not only outperformed other fusion strategies—  
 997 achieving the best results for 35 out of 37 functional groups—but was also more effective than using  
 998 all available spectra, all while maintaining lower computational complexity (Figure 9(b)).  
 999

1000 Consequently, we established the optimal configuration for this task as using IR spectroscopy as the  
 1001 primary input, with the  $^1\text{H}$ -NMR and  $^{13}\text{C}$ -NMR modalities serving as powerful auxiliary inputs.  
 1002

1003 Table 6: Comparison of resource usage  
 1004 and performance.

Model	Mem.(GB)	Time(h)	F1-score
1D-CNN	5.7	2	0.900
Transformer	1.7	35	0.911
MSpecTmol	6	3	0.959

1005 Table 7: Comparison across different  
 1006 numbers of modalities.

#Mod.	Mem.(GB)	Time(h)	F1-score
1	3.4	2	0.923
3	6.6	2.5	0.959
5	11.3	5	0.963

1007 Table 8: Inference time for processing 10,000  
 1008 samples.

Model	Inference Time (s)
1D-CNN	10.4
Transformer	45.1
Wu et al.	55.4
Alberts et al.	60.1
<b>MSpecTmol</b>	<b>14.0</b>

1009 Table 9: Training time vs. Molecular Size  
 1010 (Heavy Atom Count).

Heavy Atom Count	Training Time (s)
5 - 15	564
16 - 25	556
25 - 35	558

## 1011 I INFERENCE TIME AND SCALABILITY

1012 **Inference Time** To assess the model’s suitability for high-throughput screening, we measured the  
 1013 total inference time for processing 10,000 samples on a single NVIDIA A100 GPU. Table 8 shows  
 1014 that MSpecTmol completes the task in just 14.0 seconds. This speed is comparable to the simple  
 1015 1D-CNN (10.4 s) and drastically faster than the Transformer (45.1s), confirming its efficiency for  
 1016 real-time applications.  
 1017

1026  
 1027 **Scalability with Molecular Size** To verify whether the model’s computational cost is sensitive  
 1028 to molecular complexity, we measured the training time on subsets of data sorted by Heavy Atom  
 1029 Count. As presented in Table 9, the training time remains remarkably consistent ( $\approx 560$ s) across  
 1030 different molecular sizes. This is because our model takes fixed-dimension interpolated spectra as  
 1031 input and outputs functional group probabilities; consequently, the physical size or complexity of  
 1032 the molecule does not alter the input tensor dimensions or the model architecture.

## 1033 J SENSITIVITY ANALYSIS OF GUMBEL-SOFTMAX TEMPERATURE

1035 In the Core Spectrum Extraction module, we employ Gumbel-Softmax to enable differentiable sam-  
 1036 pling of the discrete importance masks. The temperature parameter  $t$  plays a pivotal role in control-  
 1037 ling the sharpness of this distribution. To evaluate its impact on model performance, we conducted  
 1038 a sensitivity analysis across various temperatures. As shown in Table 10, MSpecTmol achieves the  
 1039 optimal F1-score at  $t = 1.0$ . When the temperature is set to a lower value ( $t = 0.5$ ), the performance  
 1040 experiences a slight decline. Conversely, increasing the temperature to higher values ( $t = 1.5, 2.0$ )  
 1041 leads to a more noticeable degradation in prediction quality.

1042  
 1043 **Table 10: Sensitivity analysis of the Gumbel-Softmax temperature parameter  $t$ .**

1044	Temperature ( $t$ )	F1-score
1045	0.5	0.952
1046	<b>1.0</b>	<b>0.959</b>
1047	1.5	0.954
1048	2.0	0.945

1050 Intuitively, the temperature controls the sparsity and sharpness of the gating over spectral frequency  
 1051 bands. In our framework, each gate determines whether a local region of the spectrum is preserved  
 1052 or replaced by noise. A very low temperature makes these decisions almost binary. While this  
 1053 promotes sparsity, it risks discarding weak but chemically informative peaks (e.g., small shoulders  
 1054 or minor bands) that are critical for distinguishing fine-grained functional groups and isomers. On  
 1055 the other hand, a high temperature yields overly soft gates, causing most bands to be partially re-  
 1056 tained. This weakens the model’s ability to suppress redundancy and blurs the importance patterns  
 1057 across modalities. The superior performance observed at  $t = 1.0$  confirms that this setting achieves  
 1058 an optimal balance, allowing the PA-IB framework to learn selective yet stable masks that retain  
 1059 structurally informative spectral regions while effectively filtering out redundant or noisy segments.

## 1060 K GENERALIZATION ANALYSIS ON COMPLEX MOLECULAR STRUCTURES

1063 To evaluate the model’s generalization capability on samples with more complex molecular struc-  
 1064 tures and denser, overlapping spectral peaks, we performed a supplementary stress test on the dataset  
 1065 from Alberts et al. Specifically, instead of a standard random split, we sorted the entire dataset by  
 1066 heavy atom count. We utilized the bottom 90% for training and reserved the top 10% strictly for test-  
 1067 ing. This setup introduces a significant distribution shift, requiring the model to infer the structure  
 1068 of complex molecules that are physically larger than any sample seen during training.

1069 The results are presented in Table 11. While the performance on unseen larger molecules naturally  
 1070 dips compared to the standard random split due to the increased structural complexity, MSpecT-  
 1071 mol maintains a high F1-score of 0.925. Notably, our model consistently outperforms baselines in  
 1072 this challenging setting. This demonstrates that our PA-IB framework effectively learns intrinsic  
 1073 spectroscopic-structural correlations rather than simply memorizing dataset-specific patterns, con-  
 1074 firming its capability to generalize to more complex chemical spaces.

## 1075 L ANALYSIS OF PRIOR DISTRIBUTION CHOICE FOR LATENT VARIABLES

1078 In our Primary-Auxiliary Information Bottleneck (PA-IB) framework, the choice of the prior distri-  
 1079 butions for the latent bottleneck variables  $T_m$  and  $T_a$ , denoted as  $q(t_m)$  and  $q(t_a)$  respectively, is a  
 critical step that influences model performance. Specifically,  $q(t_m)$  regularizes the core information

1080  
 1081 **Table 11: Performance comparison (F1-score) on the stress test of the top 10% largest molecules**  
 1082 **versus the original random split.**

1083	Modality	Model	f1-score (Top 10% Large)	f1-score (Original Split)
1084	IR	1D-CNN	0.866	0.895
1085		Transformer	0.852	0.881
1086		Wu et al.	0.864	0.886
1087		Alberts et al.	0.874	0.891
1088		<b>MSpecTmol</b>	<b>0.900</b>	<b>0.920</b>
1089	<sup>13</sup> C-NMR	1D-CNN	0.623	0.674
1090		Transformer	0.845	0.913
1091		Wu et al.	0.873	0.914
1092		Alberts et al.	0.896	0.919
1093		<b>MSpecTmol</b>	<b>0.904</b>	<b>0.923</b>
1094	IR + <sup>13</sup> C-NMR + <sup>1</sup> H-NMR	1D-CNN	0.873	0.900
1095		Transformer	0.902	0.936
1096		Wu et al.	0.912	0.944
1097		Alberts et al.	0.916	0.947
1098		<b>MSpecTmol</b>	<b>0.925</b>	<b>0.959</b>

1100  
 1101 extracted from the primary spectrum, while  $q(t_a)$  regularizes the supplementary information from  
 1102 the auxiliary spectra. To ensure the scientific rigor and optimality of our model’s configuration, we  
 1103 systematically investigated the impact of different prior distributions on performance.  
 1104

1105 We designed a series of rigorous comparative experiments to evaluate three distinct prior distributions  
 1106 on the functional group classification task, applying them to both  $q(t_m)$  and  $q(t_a)$ :

1107

- 1108 **1. Gaussian Distribution:** The standard  $\mathcal{N}(0, I)$  distribution.
- 1109 **2. Laplace Distribution:** The standard  $Laplace(0, 1)$  distribution, which is known to effectively  
 1110 promote sparsity in the latent space.
- 1111 **3. Gamma Distribution:** The standard  $\Gamma(k = 1, \theta = 1)$  distribution, which constrains the latent  
 1112 variables to be non-negative.

1113  
 1114 Throughout these experiments, all other model hyperparameters (such as learning rate, batch size,  
 1115 and the trade-off coefficients  $\alpha$  and  $\beta$ ) were held strictly constant to ensure a fair comparison.

1116 The model trained with the **Gaussian prior** achieved the highest F1-score of 0.959, compared to  
 1117 0.951 for the Laplace prior and 0.946 for the Gamma prior. This superior performance suggests that  
 1118 assuming the compressed latent features of both primary and auxiliary spectra follow a Gaussian  
 1119 distribution provides an efficient and flexible representation space, allowing the model to optimally  
 1120 capture the complex relationships within the spectral data. Therefore, we selected the Gaussian  
 1121 distribution for our final model configuration, as its effectiveness is validated by these results.

## 1123 M MODEL PERFORMANCE STABILITY

1126 To provide a more detailed view of model stability, Figure 6 visualizes the root-mean-square de-  
 1127 viation (RMSD) distributions for conformations generated by different models. The MSpecTmol  
 1128 model exhibits a distribution with a notably lower median RMSD and smaller variance compared  
 1129 to all other baselines. This indicates that MSpecTmol not only generates conformations that are, on  
 1130 average, more accurate but also does so with higher consistency. In contrast, models relying on a  
 1131 single spectral modality (UV-Only, IR-Only, Raman-Only) and the attention-based ablation model  
 1132 all show higher median RMSDs and wider distributions. This demonstrates a greater variance in the  
 1133 quality of the generated conformations and underscores the effectiveness of our multi-modal fusion  
 strategy in achieving stable, high-fidelity results.

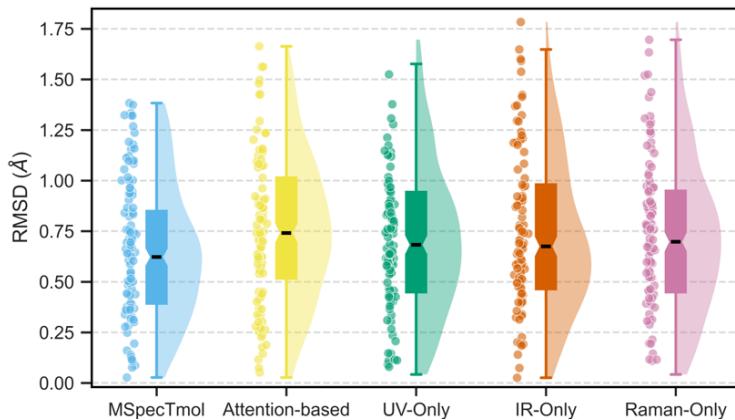


Figure 6: RMSD ( $\text{\AA}$ ) distributions for MSpecTmol compared to single-modality models and an attention-based ablation baseline. The black line indicates the median, the box represents the interquartile range, and the violin plot shows the probability density of the data.

## N COMPARISON WITH ALTERNATIVE MULTI-MODAL FUSION STRATEGIES

To validate the effectiveness of PA-IB as a multi-modal information fusion strategy, we conducted a comparative study against three standard fusion paradigms:

- **Early Fusion:** The features from all modalities are directly concatenated at the input level before being fed into the model, allowing the model to learn a joint representation from the raw data.
- **Mid-level Fusion:** Each modality is first processed through independent CNN encoders to extract latent features. These features are then concatenated and passed to a shared MLP for classification.
- **Late Fusion:** Each modality independently predicts the presence of functional groups through separate MLPs, and the final prediction is obtained by averaging the probability outputs from all modalities.

The results, presented in Table 12, show that MSpecTmol consistently outperforms all three baselines (in terms of F1-score). Early and Mid-level fusion strategies generally perform better than Late fusion, likely because they enable some degree of joint representation learning. However, they still fall short of MSpecTmol, as they fail to explicitly filter out redundant cross-modal information. Late fusion performs the worst, as it ignores inter-modal interactions entirely by processing each modality in isolation, leading to substantial information loss.

Table 12: Performance comparison (F1-score) of MSpecTmol against Early, Mid-level, and Late fusion strategies across different modality combinations.

Modality Configuration	MSpecTmol	Early Fusion	Mid-level Fusion	Late Fusion
IR + $^{13}\text{C}$ -NMR + $^1\text{H}$ -NMR	<b>0.959</b>	0.900	0.904	0.874
IR + MS/MS (Pos) + MS/MS (Neg)	<b>0.944</b>	0.887	0.895	0.854

## O INTERPRETABILITY ANALYSIS

To disentangle overlapping importance regions caused by functional group co-occurrence, we adopt a one-vs-all training strategy by training a dedicated model for each functional group. Each model receives the concatenation of all three spectral modalities as input. This design allows us to isolate the contribution of individual spectra to specific functional group predictions and analyze their region-wise importance, as shown in Figure 7.

1188 Different spectral modalities emphasize distinct molecular features. Different types of spectroscopy  
 1189 capture different aspects of molecular structures. Infrared (IR) spectroscopy is particularly important  
 1190 in identifying functional groups such as carbonyl (C=O), hydroxyl (-OH), and amine (-NH<sub>2</sub>). This  
 1191 is likely because IR spectroscopy primarily reflects the vibrational characteristics of polar functional  
 1192 groups, which exhibit strong absorption in the IR spectrum. In contrast, nuclear magnetic resonance  
 1193 (NMR) spectroscopy is more sensitive to structural motifs such as alkyl (-CH<sub>3</sub>, -CH<sub>2</sub>-), aromatic  
 1194 rings, and heterocycles. This is because NMR provides detailed insights into the electronic envi-  
 1195 ronment surrounding specific atomic nuclei, allowing for precise differentiation of these structural  
 1196 features. The complementary nature of these spectral modalities underscores the necessity of multi-  
 1197 modal approaches for comprehensive molecular characterization.

1198

## 1199 P CONFUSION MATRIX ANALYSIS

1200

1201 To investigate functional group misclassification, we construct a confusion matrix based on co-  
 1202 occurring prediction errors, counting the instances where two groups are simultaneously mispre-  
 1203 dicted for each test sample. As shown in Figure 8, this reveals that certain groups—notably Ether,  
 1204 Haloalkane, and Sulfide—are frequently confused. To diagnose this, we visualized the model’s  
 1205 attention across the fused multi-modal spectra (Figure 7). The analysis demonstrates that these  
 1206 confusable groups exhibit significant overlapping attention, indicating that the model relies on shared  
 1207 features present across different spectral modalities for their identification. This finding highlights  
 1208 the inherent difficulty in distinguishing these groups, even when multiple sources of spectral infor-  
 1209 mation are available.

1210

1211 The observed spectral feature overlap is rooted in the intrinsic chemical properties of these functional  
 1212 groups. Similarities in their responses across various spectroscopic methods, such as shared absorp-  
 1213 tion bands or related electronegativity profiles, create highly correlated features that are challenging  
 1214 to disentangle. This inherent ambiguity confirms that no single spectral modality contains sufficient  
 1215 information for perfect discrimination. It therefore becomes critical to employ a framework that can  
 1216 synergistically fuse complementary information from multiple spectra. Our approach is designed to  
 1217 address this very challenge, resolving ambiguities by integrating diverse spectral evidence to achieve  
 1218 more accurate classification.

1219

1220

## 1221 Q IMPACT OF DATA AUGMENTATION

1222

1223 Real-world experimental spectra are often subject to variations from instrumental noise and cali-  
 1224 bration drift. This challenge is compounded by the scarcity of large-scale, multi-modal spectral  
 1225 datasets. Given that deep learning models, particularly those with convolutional neural network  
 1226 (CNN) architectures, are inherently data-hungry, data augmentation becomes a crucial technique.  
 1227 By synthetically expanding the training dataset to represent a wider range of experimental condi-  
 1228 tions, we can significantly enhance the model’s generalization, robustness, and overall predictive  
 1229 performance. All experiments were conducted on molecular data obtained from the SDBS database.

1230

1231 We implemented and tested several augmentation strategies:

1232

- 1233 • **Horizontal Shift:** A random horizontal shift of up to 10 pixels is applied to the spectrum’s data  
 1234 points.
- 1235 • **Vertical Noise:** Uniform random noise (up to a level of 0.05) is added to the intensity values, with  
 1236 the noise magnitude being inversely scaled by the signal intensity.
- 1237 • **Gaussian Smoothing:** A 1D Gaussian filter with a sigma value randomly chosen between 0.75  
 1238 and 1.25 is applied to the spectrum.
- 1239 • **Combined Strategies:** A horizontal shift or vertical noise is first applied, followed by the appli-  
 1240 cation of Gaussian smoothing.

1241

1242

1243

1244

1245

1246 As shown in Figure 10(a), all data augmentation strategies successfully improved the F1-score com-  
 1247 pared to the model trained on the original data (0.9134). Among these, the Horizontal Shift strategy  
 1248 was the most effective, achieving the highest F1-score of 0.9344. This suggests that teaching the  
 1249 model to be robust against positional variations in spectral peaks is highly beneficial. The addition of  
 1250 vertical noise also provided a substantial performance boost. Interestingly, while Gaussian smooth-

ing alone offered a modest improvement, combining it with other methods (e.g., vertical noise + smoothing) did not yield further gains and resulted in lower performance than the individual, more effective strategies. This indicates that while introducing variability is beneficial, excessive transformation can risk distorting the essential chemical information within the spectra, thereby creating a trade-off between robustness and signal fidelity.

## R IMPACT OF MISSING MODALITIES AND NOISE INJECTION

**Impact of Missing Modalities** To evaluate MSpecTmol’s robustness against incomplete data, which is a common real-world challenge, we masked individual spectral modalities in the test set and evaluated the pre-trained model’s performance on SDBS dataset. As presented in Table 13, we explicitly compared MSpecTmol against baseline models under these conditions. As shown in Figure 10(b) and Table 13, the model’s performance degrades gracefully rather than failing. The F1-score drops from 0.9134 with complete data to scores between 0.8974 and 0.8623 when a modality is absent. This resilience is a direct benefit of our PA-IB architecture. Notably, MSpecTmol with missing modalities still outperforms several baselines operating with complete data. The model’s ability to extract information from auxiliary spectra to supplement the primary modality allows it to maintain robust performance even when data is partially available, making it highly suitable for practical applications.

Table 13: Performance comparison (F1-score) under missing modality conditions.

Input Configuration	MSpecTmol	1D-CNN	Trans.	Wu et al.	Alberts et al.
Full (MS+ <sup>13</sup> C+ <sup>1</sup> H)	<b>0.913</b>	0.847	0.858	0.872	0.881
w/o MS	<b>0.862</b>	0.811	0.826	0.831	0.847
w/o <sup>1</sup> H-NMR	<b>0.897</b>	0.823	0.834	0.848	0.866
w/o <sup>13</sup> C-NMR	<b>0.878</b>	0.819	0.818	0.832	0.851

**Impact of Noise** Noise is an unavoidable component of experimental spectra, making a model’s performance under such conditions a key indicator of its practical utility. To assess this, we introduced varying levels of Gaussian noise to all spectra in the test set. The standard deviation of the noise was scaled proportionally to the maximum intensity of each spectrum, ensuring a consistent signal-to-noise ratio for the evaluation. As illustrated in Figure 10(c) and Table 14, the model’s F1-score exhibits a steady and predictable decline as the noise level increases, decreasing from 0.9134 on clean data to 0.7469 at the highest noise level of 0.1. Importantly, the performance does not suffer a catastrophic collapse but rather degrades gracefully. Comparatively, as detailed in Table 14, baseline models suffer more severe degradation. At the highest noise level ( $\sigma = 0.10$ ), MSpecTmol (0.747) significantly outperforms 1D-CNN (0.635) and Transformer (0.704). This demonstrates that MSpecTmol can effectively discern core spectral features from background noise, further confirming its robustness for real-world applications where data quality is variable.

Table 14: Performance comparison (F1-score) under varying levels of Gaussian noise.

Noise ( $\sigma$ )	MSpecTmol	1D-CNN	Trans.	Wu et al.	Alberts et al.
0.00 (Clean)	<b>0.913</b>	0.847	0.858	0.872	0.881
0.02	<b>0.882</b>	0.815	0.825	0.838	0.846
0.05	<b>0.825</b>	0.745	0.761	0.767	0.772
0.10	<b>0.747</b>	0.635	0.704	0.711	0.703

## S ALGORITHMIC PROCEDURE OF CONFORMATION GENERATION

In this work, we propose a dual-encoder diffusion framework that generates molecular conformations by conditioning a geometric diffusion model on spectroscopic information.

1296

1297

Table 15: Performance comparison of different models.

1298

1299

1300

1301

1302

1303

1304

## S.1 PROBLEM FORMULATION

1305

1306

1307

1308

Given a molecular graph  $G = (\mathcal{V}, \mathcal{E})$  with atom types  $\mathbf{z} \in \mathbb{Z}^{|\mathcal{V}|}$  and spectroscopic measurements  $\mathbf{s} = [\mathbf{s}_{uv}, \mathbf{s}_{ir}, \mathbf{s}_{raman}] \in \mathbb{R}^{d_s}$ , we aim to generate 3D molecular conformations  $\mathbf{x} \in \mathbb{R}^{3|\mathcal{V}|}$  that are consistent with both the molecular connectivity and observed spectra.

1309

1310

## S.2 SPECTROSCOPY-TO-SMILES PREDICTION

1311

1312

1313

1314

1315

1316

1317

1318

1319

While it is technically feasible to generate complete 3D conformations directly from spectra, direct generation involves a huge potential space. An effective approach is to first obtain molecular samples based on spectra and then perform conformation generation. Since spectroscopy-to-SMILES prediction is already a well-established research direction, we conducted an additional feasibility study by replacing our predictor layer (MLP) with a Transformer-based decoder to predict SMILES strings directly from spectra. This modification does not alter our underlying PA-IB framework. To evaluate performance, we trained the model on the QM9S dataset—consistent with the conformation generation task—and calculated the Top-1 and Top-5 accuracy of the generated results, as well as the MCES score (normalized by the number of heavy atoms). The results are as Table 15:

1320

1321

1322

These results show that our model achieves high accuracy on the dataset. Even in cases where predictions are not perfectly accurate, the MCES scores close to 1.0 demonstrate that the model is capable of reconstructing molecular graph structures with high similarity.

1323

1324

## S.3 SPECTRUM-CONDITIONED DIFFUSION PROCESS

1325

1326

1327

We formulate the generation process as a conditional diffusion model operating in the coordinate space. The forward diffusion process adds Gaussian noise to the true conformation  $\mathbf{x}_0$ :

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I}) \quad (29)$$

1328

1329

where  $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$  and  $\{\beta_i\}$  follows a predefined noise schedule.

1330

1331

The reverse process is parameterized by a neural network  $\epsilon_\theta$  that predicts the noise conditioned on the spectrum:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(\mathbf{x}_t, t, \mathbf{s}, G) \right) + \sigma_t \boldsymbol{\eta} \quad (30)$$

1332

where  $\boldsymbol{\eta} \sim \mathcal{N}(0, \mathbf{I})$  and  $\sigma_t$  is the posterior variance.

1333

1334

## S.4 DUAL-ENCODER ARCHITECTURE

1335

1336

Our model consists of three key components:

1337

1338

**Spectrum Encoder:** We design a MSpecTmol encoder to process multi-modal spectroscopic data. The spectrum data is encoded using the PA-IB-based method described previously. The features are fused through a gated mechanism:

1339

1340

$$\mathbf{h}_s = \text{MLP}(\mathbf{h}_{spec} \oplus \mathbf{t}_{emb}) \quad (31)$$

1341

1342

where  $\mathbf{h}_{spec}$  is the spectrum embedding,  $\mathbf{t}_{emb}$  is the timestep embedding, and  $\oplus$  denotes concatenation.

1343

1344

1345

1346

1347

1348

1349

**Dual Geometric Encoders:** We employ two complementary graph encoders: (1) a SchNet-based global encoder that captures long-range interactions through radius graphs, and (2) a GIN-based local encoder focusing on chemical bond structures. Both encoders incorporate the spectrum condition  $\mathbf{h}_s$  into node representations.

1350  
 1351 **Distance-based Denoising:** In this step, we predict noise in the distance space and transform back  
 1352 to coordinates. The training objective combines global and local distance predictions:  
 1353

$$1354 \quad \mathcal{L} = \mathbb{E}_{t,\epsilon} \left[ \lambda_g \|\mathbf{d}_g - \hat{\mathbf{d}}_g\|_2^2 + \lambda_l \|\mathbf{d}_l - \hat{\mathbf{d}}_l\|_2^2 \right] \quad (32)$$

1355 where  $\mathbf{d}_g, \mathbf{d}_l$  are target distances for global and local edges respectively, and  $\lambda_g, \lambda_l$  are weighting  
 1356 factors.

## 1358 T HYPERPARAMETER EXPERIMENTS AND ABLATION STUDY

1360 To assess the impact of hyperparameter choices on model performance, we conduct a series of  
 1361 experiments by varying the information bottleneck trade-off coefficients  $\alpha$  and  $\beta$  from  $1 \times 10^{-7}$  to  
 1362  $1 \times 10^{-3}$ . The best results are achieved when both  $\alpha$  and  $\beta$  are set to  $1 \times 10^{-6}$ . Besides, we present  
 1363 the implementation details of the ablation settings and illustrate how the loss function changes when  
 1364 specific components are removed. The corresponding results are shown in Figure 11.

1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403

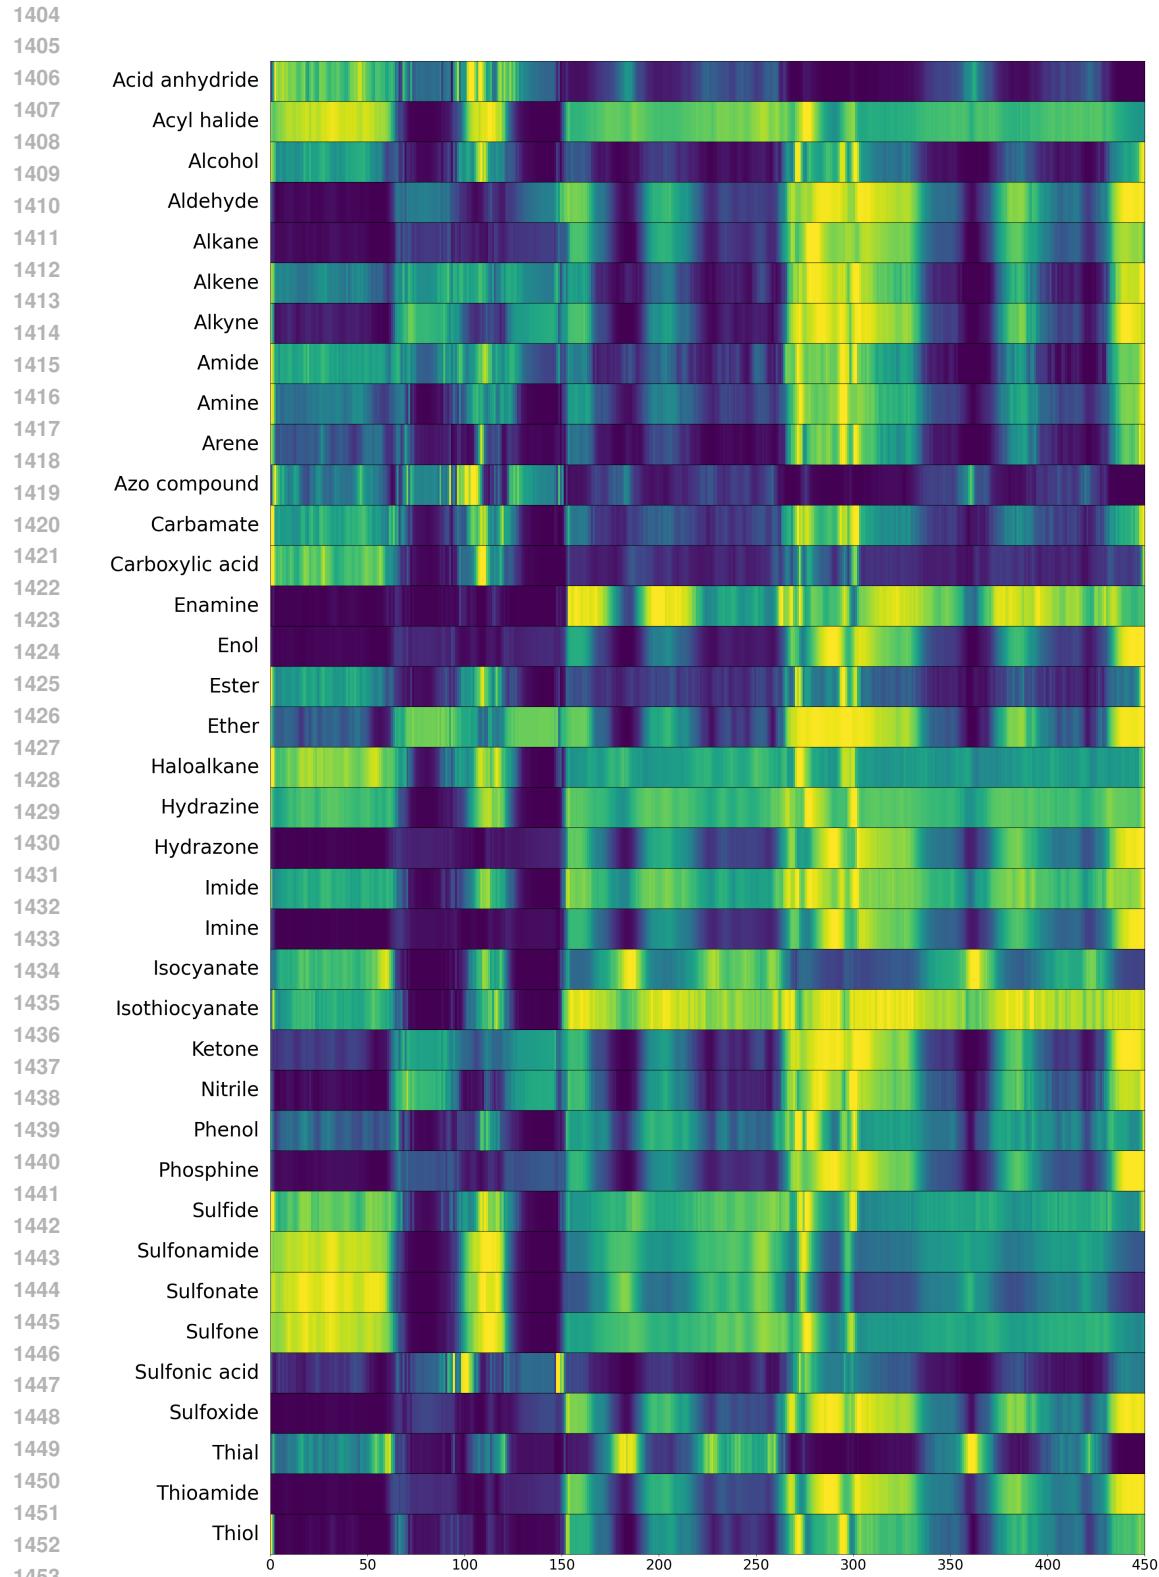


Figure 7: Illustration of the importance of spectral regions. The input spectrum is partitioned as follows: [0, 150] corresponds to IR spectra, [151, 300] to  $^1\text{H}$ -NMR spectra, and [301, 450] to  $^{13}\text{C}$ -NMR spectra. Warmer colors indicate crucial (high-importance) information, while cooler colors represent redundant (low-importance) information.

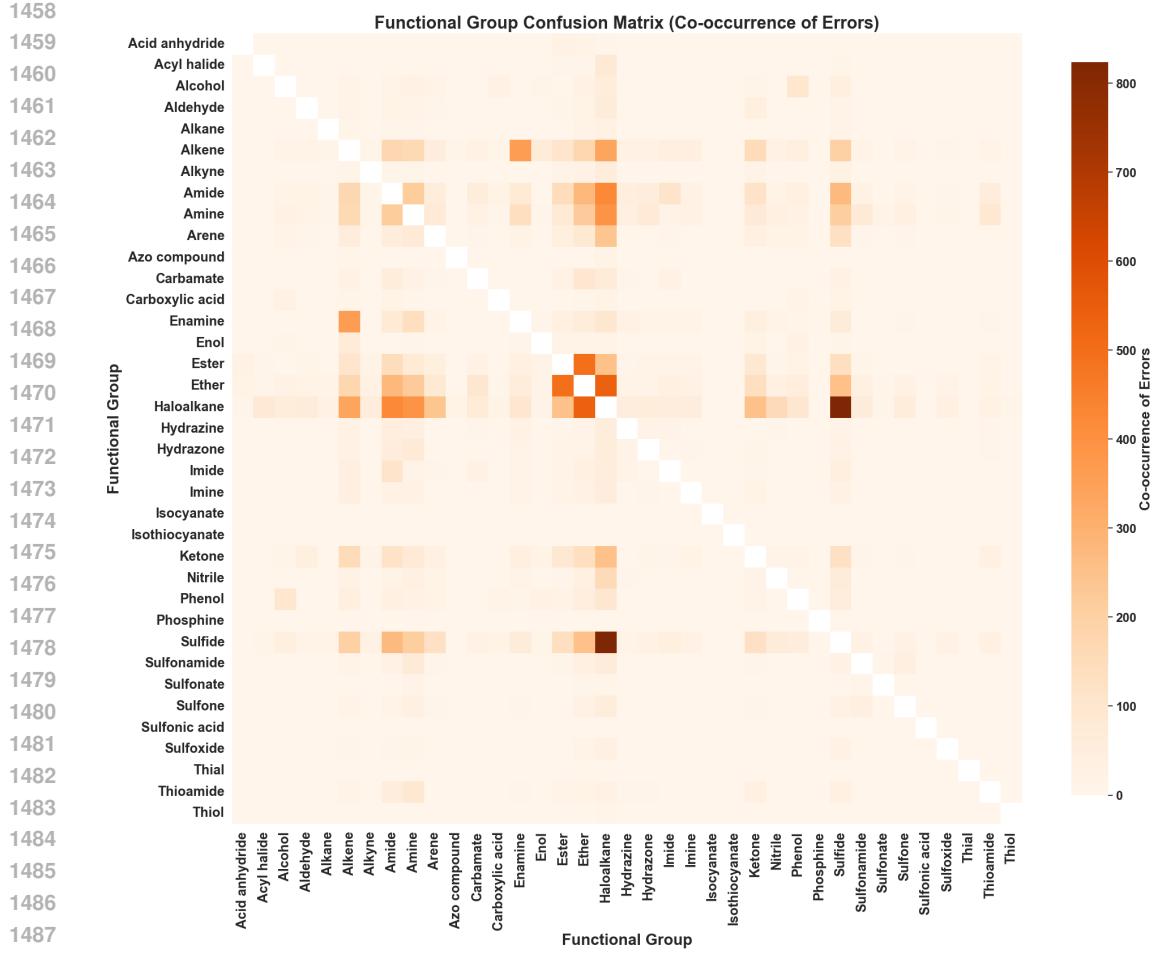


Figure 8: The confusion matrix between functional groups: the darker the color in the blocks, the higher the number of samples where the two functional groups were predicted incorrectly simultaneously.

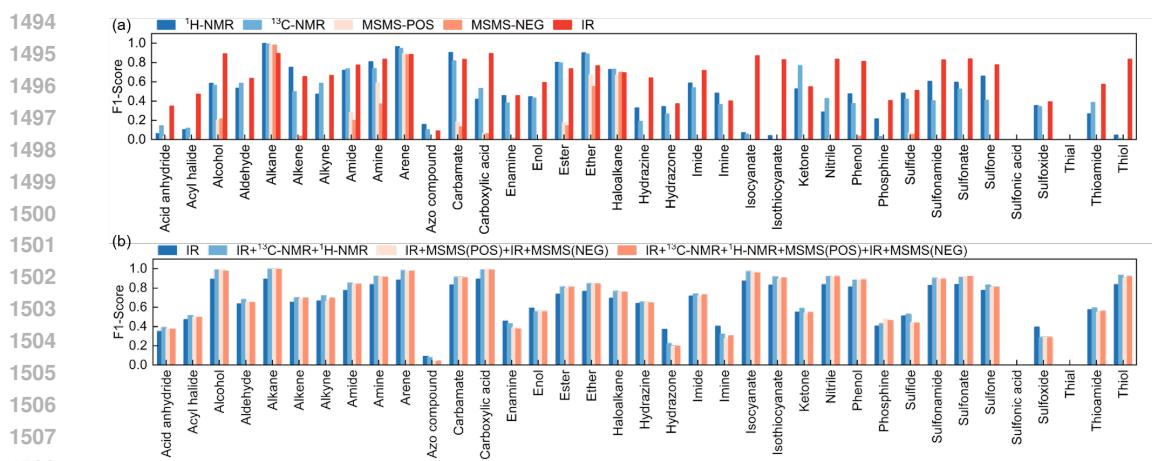


Figure 9: Performance of the model under unimodal and multimodal settings. (a) Results using a single spectrum as input (unimodal); (b) Results under multimodal fusion of spectra.

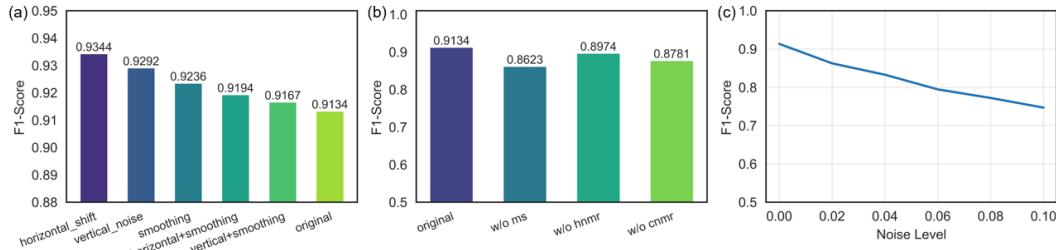


Figure 10: (a) Comparison of F1-scores for different data augmentation strategies applied to the training data. (b) Model performance with missing spectral modalities (MS,  $^1\text{H}$ -NMR, and  $^{13}\text{C}$ -NMR) in the test set. (c) The impact of increasing levels of Gaussian noise on the final F1-score.

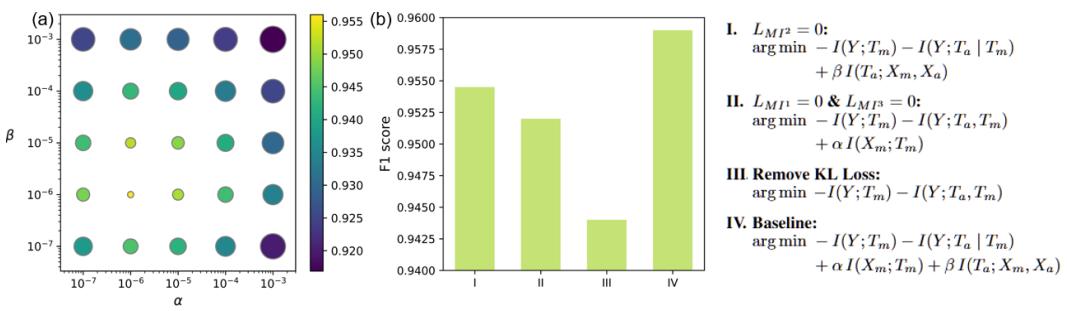


Figure 11: (a) Hyperparameter Experiments on functional group classification task. The circle size is proportional to the magnitude of the error. (b) Ablation study: by selectively removing different KL divergence terms, we adjust the optimization objectives of the model. The left panel shows the F1 scores of the prediction results, while the right panel illustrates the minimized objectives of the ablated models.