FROM STATIC RANKING TO ACTIVE PLANNING: FEEDBACK-CONDITIONED CANDIDATE VERIFICATION FOR MOLECULE IDENTIFICATION WITH NEURAL SPECTRUM PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Molecule identification with neural spectrum prediction ranks candidate molecular structures by agreement between their predicted spectra and observed experimental spectra, then verifies candidates in this order. However, simply following this static ranking overlooks spectral similarity feedback obtained during candidate verification. We introduce ActiveMS, an active planning framework for MS/MS-based molecule identification that incorporates this feedback to determine which unverified candidate structures to verify next. ActiveMS-GP models feedback residuals with a Gaussian process, while ActiveMS-FCP learns an amortized feedback-conditioned posterior over candidate structures across verification rollouts. For evaluation, we construct ActiveMS Bench from the SpectraVerse and NIST spectral libraries as two separate benchmarks, with feedback supplied by an independent non-neural simulator. On ActiveMS Bench, both planners consistently outperform static-ranking verification with MassFormer and ICEBERG spectrum predictors. With either predictor, ActiveMS-GP and ActiveMS-FCP reduce mean verification rounds by 33.3–46.9% and 38.4–52.2% respectively. ActiveMS-FCP further raises identification success within 16 verification rounds: with MassFormer, from 50.9% to 77.7% on SpectraVerse and from 70.9% to 86.0% on NIST; with ICEBERG, from 87.2% to 93.7% on SpectraVerse and from 88.6% to 95.0% on NIST. These results indicate that feedback-conditioned planning and neural spectrum prediction improve identification along complementary axes: better predictions sharpen the initial ranking, while feedback during verification further reduces the rounds needed to reach the target.

1 INTRODUCTION

Tandem mass spectrometry (MS/MS) is a central technique for identifying unknown molecules in complex samples by matching their experimental spectra against spectral libraries of known structures. However, experimental spectral libraries cover only a small fraction of known molecular structures (Stein, 2012; Bittremieux et al., 2022). For molecules beyond library coverage, candidate structures sharing the target's molecular formula are retrieved from structure databases, and a spectrum predictor generates spectra for each candidate, ranking them by agreement between predicted and observed experimental spectra (Kind et al., 2018; Ruttkies et al., 2016; Wang et al., 2021; Hong et al., 2025). Recent advances in neural spectrum prediction have improved the fidelity of predicted spectra, with methods such as GrAFF-MS (Murphy et al., 2023), MassFormer (Young et al., 2024), 3DMolMS (Hong et al., 2023), SCARF (Goldman et al., 2023), ICEBERG (Goldman et al., 2024; Wang et al., 2025), and GLACIER (Wang et al., 2026). Despite these advances, identifying an unknown molecule still requires successive rounds of experimental verification, each testing a batch of ranked candidates by acquiring reference standards and measuring spectra under matching conditions. When a round does not identify the target, verification proceeds to the next batch in the ranking. Under this static ranking, information revealed by each round goes unexploited.

Verification returns more than a match or mismatch. Comparing a verified candidate's experimental spectrum with the target's provides spectral similarity feedback. For compounds sharing a molecular

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

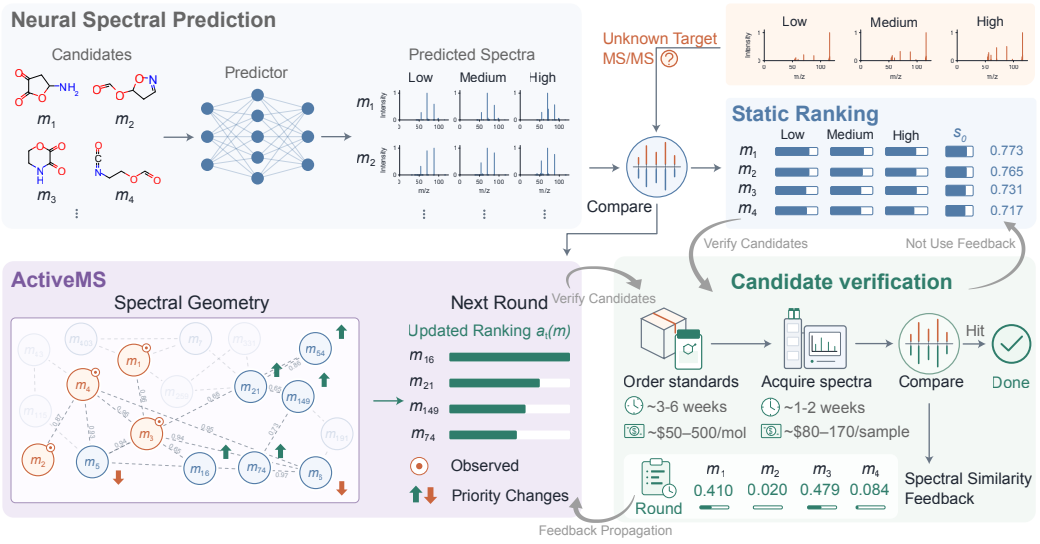


Figure 1: From static ranking to feedback-conditioned candidate verification. ActiveMS transfers spectral similarity feedback through the *in silico* spectral geometry to update verification priorities.

formula, high spectral similarity to the target is more common among close structural analogs but does not uniquely identify them (Wang et al., 2016; Uritboonthai et al., 2026). Feedback from each verified candidate therefore constrains but does not determine which remaining candidates are likely to be the target. This raises a question complementary to improving spectrum prediction: can feedback obtained during verification guide which candidates to verify next?

We address this question with ActiveMS, an active planning framework for candidate verification in MS/MS-based molecule identification, as illustrated in Figure 1. The central challenge is to transfer feedback from verified candidates to unverified ones that have no experimental spectra yet. The neural spectrum predictor, which supplies the initial ranking, also provides the basis for this transfer: its predicted spectra define an *in silico* spectral geometry that places all candidates in a shared space before any is verified. Over this geometry, ActiveMS-GP models the feedback with a Gaussian process, while ActiveMS-FCP learns a feedback-conditioned posterior over candidate structures, as illustrated in Figure 2. On ActiveMS Bench, built from SpectraVerse and NIST spectral libraries, both planners identify targets in fewer verification rounds than static ranking with MassFormer and ICEBERG spectrum predictors.

Our contributions are as follows:

- We introduce ActiveMS, extending static-ranking verification with feedback-conditioned active planning. We construct ActiveMS Bench from SpectraVerse and NIST, with spectral similarity feedback for each candidate simulated by an independent non-neural simulator.
- We propose two ActiveMS planners. ActiveMS-GP models the feedback by inferring feedback residuals with a Gaussian process. ActiveMS-FCP models the target by learning an amortized feedback-conditioned target posterior.
- With MassFormer and ICEBERG on both benchmarks, both planners reduce mean verification rounds by 33.3–52.2% over static ranking and consistently improve identification success rates. Experiments and further analysis show that the gains of ActiveMS-FCP rely on spectral geometry and correct feedback correspondence.

Related work is discussed in Appendix E.

2 ACTIVEMS: PROBLEM FORMULATION

2.1 PROBLEM SETUP

A molecule identification query q consists of observed experimental MS/MS spectra acquired at different collision energies from an unknown target molecule $m^\star$, and a candidate molecule set $\mathcal{C}_q = \{m_1, \ldots, m_K\}$ with $m^\star \in \mathcal{C}_q$. All candidates in $\mathcal{C}_q$ share the molecular formula of the target molecule $m^\star$. These observed spectra are denoted $\mathcal{S}_q = \{(S_j, e_j)\}_{j=1}^{J_q}$, where J_q is the number of experimental spectra and e_j is the collision energy associated with spectrum S_j.

Let $\mathcal{V}_{t-1}$ denote the candidates verified before round t, with $\mathcal{V}_0 = \varnothing$. The planner selects and verifies a batch of up to B candidates:

$$\mathcal{B}_t \subseteq \mathcal{C}_q \setminus \mathcal{V}_{t-1}, \qquad |\mathcal{B}_t| = \min\{B, |\mathcal{C}_q \setminus \mathcal{V}_{t-1}|\}.$$

Verification returns candidate–target spectral similarity feedback $y_q(m) \in [0, 1]$ for each selected $m \in \mathcal{B}_t$. If no candidate matches the target structure in the current round, this feedback guides candidate selection in the next round.

After each round, the verified set becomes $\mathcal{V}_t = \mathcal{V}_{t-1} \cup \mathcal{B}_t$. The sequence of selected batches and their spectral similarity feedback forms a verification trajectory, ending when $m^\star \in \mathcal{B}_t$ or a budget of $R_{\max}$ verification rounds is exhausted.

2.2 FORWARD PREDICTION

Trained on the training split of the corresponding benchmark, a neural spectrum predictor f_{sp} maps each candidate molecule m and collision energy e to a predicted spectrum $\hat{S}(m, e) = f_{\text{sp}}(m, e)$. At the query's collision energies $\mathcal{E}_q = \{e_1, \ldots, e_{J_q}\}$, matching these predictions to the experimental spectra gives energy-specific matching scores and an initial ranking score:

$$s_j(q, m) = \text{entropy}\left(S_j, f_{\text{sp}}(m, e_j)\right), \qquad s_0(q, m) = \frac{1}{J_q} \sum_{j=1}^{J_q} s_j(q, m),$$

where entropy denotes spectral-entropy similarity (Li et al., 2021). The static ranking verifies candidates in descending order of $s_0(q, m)$. The predicted spectra $\hat{S}(m, e)$, matching scores $s_j(q, m), s_0(q, m)$, and collision energies $\mathcal{E}_q$ together form the forward prediction $\mathcal{F}_q$.

2.3 PLANNING OBJECTIVE

Before round t, the verification history contains feedback from verified candidates $m^v \in \mathcal{V}_{t-1}$:

$$\mathcal{H}_{t-1} = \{(m^v, y_q(m^v)) : m^v \in \mathcal{V}_{t-1}\}.$$

A policy π conditions on the verification history $\mathcal{H}_{t-1}$ and the forward prediction $\mathcal{F}_q$ to select the next batch $\mathcal{B}_t^\pi \subseteq \mathcal{C}_q \setminus \mathcal{V}_{t-1}$ of up to B candidates.

Let R_q^π denote the number of verification rounds for policy π to select $m^\star$ for query q:

$$R_q^\pi = \min\{t : 1 \le t \le R_{\max},\ m^\star \in \mathcal{B}_t^\pi\},$$

with $R_q^\pi = R_{\max}$ if the target remains unverified. The objective is to reduce the mean number of verification rounds $\mathbb{E}_q[R_q^\pi]$ across all queries (**verification budgets**), while achieving $R_q^\pi \le r$ on as many queries as possible with $r \ll R_{\max}$ (**identification success rate**). Section 5.1 defines the evaluation metrics for these two objectives.

3 ACTIVEMS PLANNERS

ActiveMS planners combine the forward prediction $\mathcal{F}_q$ with spectral feedback $y_q(m^v)$ recorded in the verification history $\mathcal{H}_{t-1}$, assign each unverified candidate structure m an acquisition score $a_t(m)$, and select up to B highest-scoring candidates for verification in round t. Feedback Greedy serves as a diagnostic reference baseline through local feedback propagation in the *in silico* spectral space. We develop two complementary planners: ActiveMS-GP estimates the posterior feedback mean, while ActiveMS-FCP learns to estimate the probability that each unverified candidate is the target structure.

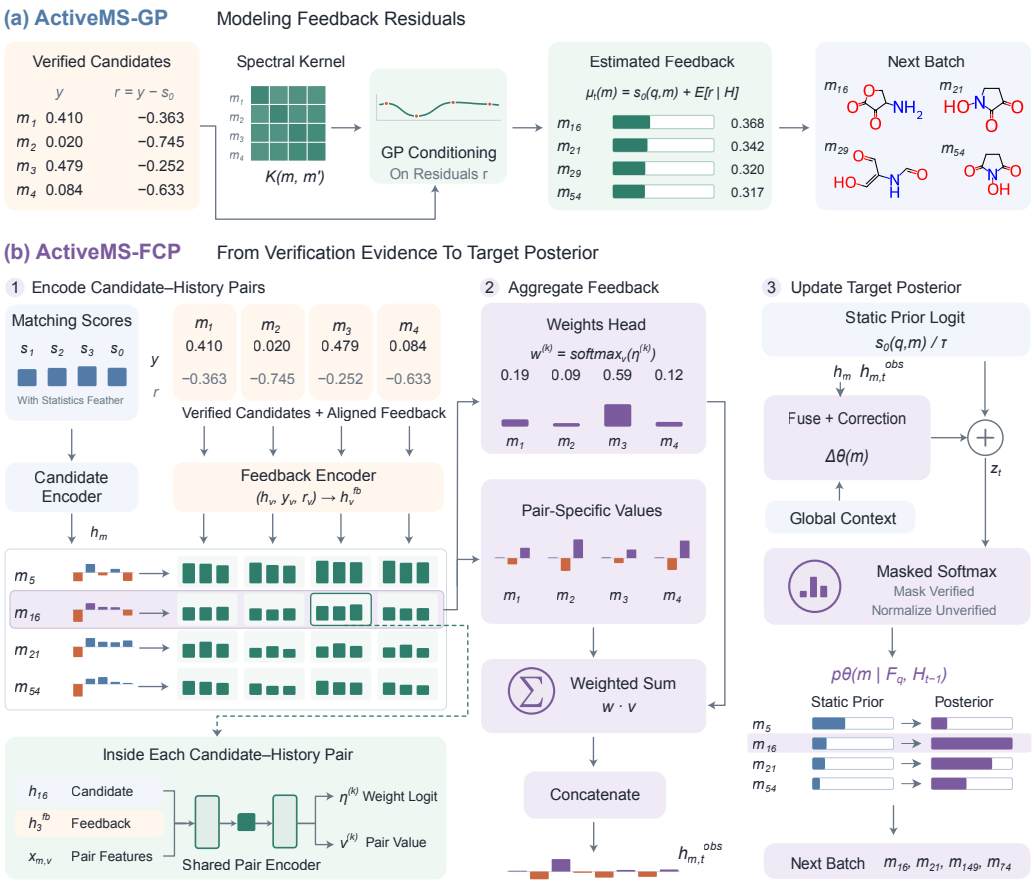


Figure 2: Feedback-conditioned candidate selection. (a) ActiveMS-GP conditions a Gaussian process on feedback residuals $r = y - s_0$ with a spectral kernel and selects by posterior mean. (b) ActiveMS-FCP encodes candidate–history pairs via spectral geometry and multi-head attention, then corrects a static prior to yield the target posterior over unverified candidates.

3.1 Feedback Greedy

Feedback Greedy propagates observed spectral feedback $y_q(m^v)$ from verified to unverified candidates according to the similarity of their predicted spectra. For candidate structures m_i and m_j, we average the spectral similarity sim between their binned predicted spectra across collision energies to obtain $k_{\mathrm{sim}}(m_i, m_j)$:

$$k_{\mathrm{sim}}(m_i, m_j) = \frac{1}{|\mathcal{E}_q|} \sum_{e \in \mathcal{E}_q} \mathrm{sim}\Big(\hat{S}(m_i, e),\, \hat{S}(m_j, e)\Big)\,.$$

We compute the feedback estimate $g_t(m)$ for each unverified candidate m as a similarity-weighted average of observed feedback $y_q(m^v)$:

$$g_t(m) = \frac{\sum_{m^v \in \mathcal{V}_{t-1}} k_{\mathrm{sim}}(m, m^v)\, y_q(m^v)}{\sum_{m^v \in \mathcal{V}_{t-1}} k_{\mathrm{sim}}(m, m^v) + \epsilon},$$

where $0 < \epsilon \ll 1$ stabilizes the denominator. The acquisition score combines the initial ranking score and the feedback estimate with a mixing weight $\alpha \in [0, 1]$:

$$a_t^{\mathrm{Greedy}}(m) = (1 - \alpha)\, s_0(q, m) + \alpha\, g_t(m).$$

3.2 ACTIVEMS-GP

The **G**aussian **P**rocess planner (ActiveMS-GP) estimates the posterior feedback mean $\mu_t(m)$ to guide candidate selection. It takes the initial ranking score $s_0(q, m)$ as the prior mean. A Gaussian process (Rasmussen & Williams, 2006) over the *in silico* spectral space then models the feedback residual $r(m^v) = y_q(m^v) - s_0(q, m^v)$ of each verified candidate m^v.

Kernel. To infer feedback residuals $r(m)$ for unverified candidates from those already observed, we assume that candidates with similar predicted spectra have correlated residuals. We encode this dependence with a radial basis function (RBF) kernel. Averaging energy-specific kernels across collision energies gives $k(m_i, m_j)$:

$$k(m_i, m_j) = \frac{1}{|\mathcal{E}_q|} \sum_{e \in \mathcal{E}_q} \exp\left(-\frac{d_e(i,j)^2}{2\ell^2}\right),$$

where $d_e(i, j)$ is the Euclidean distance between normalized binned predicted spectra at energy e. The length scale ℓ equals the median positive pairwise distance, pooled across candidate pairs and collision energies, scaled by $c_\ell > 0$.

Acquisition. Conditioning on the observed residuals $\mathbf{r}_{\mathcal{V}_{t-1}}$, ActiveMS-GP selects candidates by posterior mean (PM) acquisition:

$$a_t^{\mathrm{GP}}(m) = \mu_t(m) = s_0(q, m) + \mathbf{K}_{m\mathcal{V}_{t-1}} \left(\mathbf{K}_{\mathcal{V}_{t-1}\mathcal{V}_{t-1}} + \sigma_n^2 \mathbf{I}\right)^{-1} \mathbf{r}_{\mathcal{V}_{t-1}},$$

where the $\mathbf{K}$ blocks collect pairwise kernel values indicated by their subscripts, $\mathbf{r}_{\mathcal{V}_{t-1}}$ stacks observed residuals, and σ_n^2 is the observation-noise variance.

3.3 ACTIVEMS-FCP

The **F**eedback-**C**onditioned **P**osterior planner (ActiveMS-FCP) learns to update the probability that each unverified candidate m is the target structure $m^\star$ as verification feedback $y_q(m^v)$ accumulates. It combines a static prior derived from initial ranking scores $s_0(q, m)$ with candidate-specific corrections. Cross-attention computes these corrections over encoded spectral feedback from verified candidates $m^v \in \mathcal{V}_{t-1}$, conditioned on the *in silico* spectral geometry. The model is trained on verification rollouts for training-split queries to amortize posterior inference. At test time, a single forward pass updates candidate probabilities without retraining.

Candidate and feedback representations. A candidate encoder f_{cand} maps the energy-specific matching scores $s_j(q, m)$ and initial ranking score $s_0(q, m)$ to a candidate representation $\mathbf{h}_m$:

$$\mathbf{h}_m = f_{\mathrm{cand}}(s_1(q, m), \ldots, s_{J_q}(q, m),\ s_0(q, m)).$$

For each verified candidate $m^v \in \mathcal{V}_{t-1}$, a feedback encoder f_{fb} maps the candidate representation $\mathbf{h}_{m^v}$, observed feedback $y_q(m^v)$, and feedback residual $r(m^v)$ to a feedback representation $\mathbf{h}^{\mathrm{fb}}_{m^v}$:

$$\mathbf{h}^{\mathrm{fb}}_{m^v} = f_{\mathrm{fb}}(\mathbf{h}_{m^v},\ y_q(m^v),\ r(m^v)).$$

Geometry-aware feedback conditioning. Multi-head cross-attention (Vaswani et al., 2017) transfers feedback from verified candidates $m^v \in \mathcal{V}_{t-1}$ to each unverified candidate m, conditioned on the pairwise *in silico* spectral geometry. We represent this geometry by energy-specific spectral similarities between the binned predicted spectra of m and m^v:

$$\rho^e_{m,m^v} = \mathrm{sim}\left(\hat{S}(m, e),\ \hat{S}(m^v, e)\right), \qquad \boldsymbol{\rho}_{m,m^v} = (\rho^{e_j}_{m,m^v})_{j=1}^{J_q}.$$

A pairwise feature vector $\mathbf{x}_{m,m^v}$ augments this geometry with verification outcomes:

$$\mathbf{x}_{m,m^v} = \left[\boldsymbol{\rho}_{m,m^v},\ s_0(q, m),\ s_0(q, m^v),\ y_q(m^v),\ r(m^v),\ \bar{\rho}_{m,m^v}\, y_q(m^v),\ \bar{\rho}_{m,m^v}\, r(m^v)\right],$$

where $\bar{\rho}_{m,m^v}$ is the mean similarity across energies.

A shared pair encoder f_{pair} maps each candidate–verified pair to per-head attention logits and value vectors:

$$\left\{\left(\eta_{m,m^v}^{(k)},\, \mathbf{v}_{m,m^v}^{(k)}\right)\right\}_{k=1}^{N_{\text{head}}} = f_{\text{pair}}\left(\mathbf{h}_m,\, \mathbf{h}_{m^v}^{\text{fb}},\, \mathbf{x}_{m,m^v}\right).$$

Attention-weighted aggregation across verified candidates yields the feedback-conditioned representation $\mathbf{h}_{m,t}^{\text{obs}}$:

$$\mathbf{h}_{m,t}^{\text{obs}} = \underset{k=1}{\overset{N_{\text{head}}}{\text{Concat}}} \left(\sum_{m^v \in \mathcal{V}_{t-1}} w_{m,m^v,t}^{(k)}\, \mathbf{v}_{m,m^v}^{(k)} \right), \qquad w_{m,m^v,t}^{(k)} = \underset{m^v \in \mathcal{V}_{t-1}}{\text{softmax}}\left(\eta_{m,m^v}^{(k)}\right),$$

where N_{head} is the number of attention heads.

Posterior update and selection. A fusion network f_{fuse} combines the candidate representation $\mathbf{h}_m$, feedback-conditioned representation $\mathbf{h}_{m,t}^{\text{obs}}$, and a global context $\mathbf{h}_t^{\text{ctx}}$ summarizing the candidate set and verification history:

$$\mathbf{h}_{m,t}^{\text{joint}} = f_{\text{fuse}}(\mathbf{h}_m,\, \mathbf{h}_{m,t}^{\text{obs}},\, \mathbf{h}_t^{\text{ctx}}).$$

A linear head predicts a correction Δ_θ to the static prior, yielding the updated candidate score

$$z_t^{\text{FCP}}(m) = \frac{\tilde{s}_0(q,m)}{\tau} + \Delta_\theta(m \mid \mathcal{F}_q, \mathcal{H}_{t-1}),$$

where $\tilde{s}_0(q,m)$ is $s_0(q,m)$ standardized across $\mathcal{C}_q$, $\tau > 0$ is a learned temperature, and θ denotes all trainable parameters. The acquisition score is the target posterior over unverified candidates:

$$a_t^{\text{FCP}}(m) = p_\theta(m \mid \mathcal{F}_q, \mathcal{H}_{t-1}) = \frac{\exp(z_t^{\text{FCP}}(m))}{\sum_{m_i \in \mathcal{C}_q \setminus \mathcal{V}_{t-1}} \exp(z_t^{\text{FCP}}(m_i))}.$$

Training with bootstrapped rollouts. Each training example $(\mathcal{F}_q, \mathcal{H}_{t-1})$ pairs the forward prediction $\mathcal{F}_q$ with a verification history $\mathcal{H}_{t-1}$ before some round t where the target $m^\star$ remains unverified. Examples are pooled from three rollout sources: *prior rollouts* follow the static ranking $s_0(q,m)$; an FCP model trained on prior rollouts generates *greedy rollouts* that select candidates with the highest posterior p_θ; and a second model trained on both generates *stochastic rollouts* that sample candidate batches from p_θ. The final model is trained from scratch on examples pooled from prior, greedy, and stochastic rollouts.

The training objective combines a cross-entropy loss on the target posterior with an auxiliary residual loss $\mathcal{L}_{\text{res}}$:

$$\mathcal{L} = -\log p_\theta(m^\star \mid \mathcal{F}_q, \mathcal{H}_{t-1}) + \lambda_{\text{res}}\, \mathcal{L}_{\text{res}},$$

where $\lambda_{\text{res}} \geq 0$ controls the auxiliary contribution. The auxiliary head on $\mathbf{h}_{m,t}^{\text{joint}}$ predicts feedback residuals $r(m)$ for unverified non-target candidates under a feedback-weighted Smooth-L_1 loss. Network architecture, rollout construction, and auxiliary loss are detailed in Appendix B.

4 ACTIVEMS BENCH

We develop ActiveMS Bench with separate benchmarks built from the SpectraVerse (Gupta et al., 2026a) and NIST (National Institute of Standards and Technology, 2014) experimental spectral libraries, as reported in Table 1. Each query pairs three experimental MS/MS spectra of a target structure $m^\star$ at low, medium, and high collision energies with a candidate set $\mathcal{C}_q$ of up to 1,024 structures sharing the target's molecular formula. Candidates comprise PubChem (Kim et al., 2025) decoys and generated decoys sampled from a chemical language model (Bushuiev et al., 2024; Gupta et al., 2026b); queries with fewer than 128 candidates are filtered out. SpectraVerse follows a structure-based split with MCES distance (Kretschmer et al., 2025) at least 10 between targets in different splits; NIST follows a temporal split (NIST20 for training, NIST23 for validation and test). In both benchmarks, ECFP4 Tanimoto similarity (Rogers & Hahn, 2010) between targets across splits is below 0.7.

Table 1: Statistics of the SpectraVerse and NIST benchmarks.

Benchmark	Split	#Queries	#Molecules	#Spectra	Avg. #Cand.
SpectraVerse	Train	7,161	7,161	21,483	845.0
	Validation	895	895	2,685	770.8
	Test	895	895	2,685	788.4
NIST	Train	18,994	18,759	56,982	819.7
	Validation	6,673	6,672	20,019	775.3
	Test	6,674	6,671	20,022	775.7

Because experimental spectral libraries do not cover candidate structures, ActiveMS Bench generates verification feedback with CFM-ID 4.0, an independent non-neural simulator. Feedback $y_q(m)$ averages spectral-entropy similarity between the simulated spectra of candidate m and target $m^\star$ across three energy channels. The resulting feedback is an informative but probabilistic indicator of structural relatedness to the target: excluding target matches, mean within-set Spearman correlations with normalized MCES similarity are 0.208 on SpectraVerse and 0.276 on NIST. Benchmark construction and feedback diagnostics are detailed in Appendix C.

5 Experiments

5.1 Experimental Setup

Each query is evaluated with a budget of $R_{\max} = 64$ verification rounds and $B = 4$ candidate structures per round. Initial rankings are based on spectral-entropy similarity between predicted and experimental spectra. In the first round, Feedback Greedy and both ActiveMS planners verify the top-B candidates in this ranking.

Neural spectrum predictors. MassFormer and ICEBERG are trained separately on the training split of each benchmark and then frozen during candidate verification. On the test queries, MassFormer achieves Top-4 hit rates of 11.17% on SpectraVerse and 19.07% on NIST, while ICEBERG achieves 55.75% and 42.81%, respectively.

Baselines. *Random* verifies candidate structures in a random order. *Static Ranking* verifies candidates in descending order of their initial ranking scores $s_0(q, m)$. *Feedback Greedy* re-ranks candidates by mixing the initial ranking score $s_0(q, m)$ with similarity-weighted feedback estimates, as described in Section 3.1.

Evaluation metrics. *Mean verification rounds* (MVR) is the average number of rounds required to identify the target molecule across evaluation queries. *S@Rr* measures the percentage of queries for which the target molecule is identified within r verification rounds. *R@Pp* measures the fewest verification rounds needed to resolve at least $p\%$ of queries. If this proportion is not reached within $R_{\max}$ rounds, R@Pp is reported as $> R_{\max}$. We report S@Rr for $r \in \{4, 16\}$ and R@Pp for $p = 90$.

Implementation. Forward predictions $\mathcal{F}_q$ for ActiveMS-FCP training come from cross-fitting: the neural spectrum predictor is trained on each molecule-disjoint half of the training split and predicts candidate spectra for queries in the other half. CFM-ID supplies the feedback $y_q(m)$ along the resulting training rollouts. Cosine similarity between binned predicted spectra defines the *in silico* spectral geometry for Feedback Greedy and ActiveMS-FCP. For the stochastic methods (Random and ActiveMS-FCP), S@Rr and MVR are averaged over five seeds, and R@Pp is computed from the mean query-coverage curve. Hyperparameters and training details are provided in Appendix D.

5.2 Main Results

Table 2 shows that ActiveMS-GP and ActiveMS-FCP improve S@R4 and S@R16 over Static Ranking on both benchmarks with each neural spectrum predictor. With MassFormer, S@R16 on Spec-

traVerse increases from 50.95% for Static Ranking to 72.07% for ActiveMS-GP and 77.70% for ActiveMS-FCP. The same ordering holds on NIST, where ActiveMS-FCP again achieves the highest success rates. With ICEBERG, Static Ranking already reaches S@R16 of 87.15% on SpectraVerse and 88.63% on NIST. Both planners further improve identification success, with ActiveMS-FCP reaching 93.68% on SpectraVerse and 94.97% on NIST. The larger absolute gains with MassFormer are consistent with the greater room for improvement over its static ranking, while the gains with ICEBERG demonstrate that feedback-conditioned planning also complements more accurate initial rankings.

These improvements in identification success are accompanied by lower verification costs. Relative to Static Ranking, MVR decreases by 33.3–46.9% with ActiveMS-GP and by 38.4–52.2% with ActiveMS-FCP. On SpectraVerse with ICEBERG, ActiveMS-FCP reaches 90% identification coverage in 9 rounds, compared with 10 rounds for ActiveMS-GP and 29 rounds for Static Ranking. Feedback Greedy also improves upon Static Ranking, but both ActiveMS planners achieve higher success rates and lower MVR, extending the gains beyond simple feedback propagation. ActiveMS-FCP outperforms ActiveMS-GP on S@R4, S@R16, MVR, and R@P90 in all four settings, with larger advantages under MassFormer. Across-seed variation and paired statistical comparisons are reported in Appendix A.1.

Table 2: Verification budgets and identification success rates on the SpectraVerse and NIST benchmarks. > 64 indicates that 90% query coverage is not reached within 64 rounds. Bold marks the best result for each benchmark and neural spectrum predictor.

	MassFormer				ICEBERG			
Method	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
SpectraVerse								
Random	50.56	2.41	10.86	> 64	50.56	2.41	10.86	> 64
Static Ranking	24.88	27.49	50.95	> 64	8.44	73.74	87.15	29
Feedback Greedy	18.66	33.85	63.46	64	5.48	77.88	91.73	13
ActiveMS-GP	14.78	35.87	72.07	42	4.49	80.67	93.18	10
ActiveMS-FCP	**11.89**	**39.33**	**77.70**	**30**	**4.37**	**82.46**	**93.68**	**9**
NIST								
Random	49.90	2.92	11.82	> 64	49.90	2.92	11.82	> 64
Static Ranking	15.40	42.36	70.90	52	7.06	69.49	88.63	19
Feedback Greedy	12.87	47.00	76.10	40	6.16	72.40	90.73	15
ActiveMS-GP	9.70	52.02	83.19	25	4.71	76.40	93.90	10
ActiveMS-FCP	**8.69**	**54.02**	**85.99**	**22**	**4.35**	**78.07**	**94.97**	**9**

5.3 Ablation Study

Table 3 compares mixed-rollout training with three single-source variants, each trained on prior, greedy, or stochastic rollouts alone. To our surprise, each single-source variant already retains most of the gains over Static Ranking, suggesting that limited rollout diversity suffices to train an effective planner. On NIST with MassFormer, single-source variants reach S@R16 of 84.25–85.67% and MVR of 8.75–9.28 rounds, compared with 85.99% and 8.69 rounds for mixed training. With ICEBERG, their MVR of 4.36–4.45 rounds nearly matches the 4.35 rounds of mixed training. Mixed training attains the lowest MVR in all four settings, with the largest differences on SpectraVerse with MassFormer. The SpectraVerse ablation is reported in Appendix A.2.

To isolate the contribution of spectral geometry, Zeroed Geometry in Table 3 sets all *in silico* spectral geometry inputs to zero at inference while retaining initial ranking evidence and observed spectral feedback. Across all four settings, S@R16 decreases and MVR increases, approaching Static Ranking performance. On NIST with MassFormer, S@R16 falls from 85.99% to 70.90%, matching Static Ranking. These results indicate that the trained model relies on spectral geometry to turn verification feedback into improved candidate selection.

Table 3: Ablation of ActiveMS-FCP on the NIST benchmark.

	MassFormer				ICEBERG			
Variant	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
Full ActiveMS-FCP	**8.69**	**54.02**	**85.99**	**22**	**4.35**	**78.07**	**94.97**	**9**
w/o Mixed Rollouts								
Prior rollouts only	8.94	53.61	85.26	23	4.36	77.69	94.89	10
Greedy rollouts only	9.28	52.61	84.25	24	4.40	77.49	94.82	10
Stochastic rollouts only	8.75	53.40	85.67	**22**	4.45	77.28	94.79	10
w/o *In Silico* Spectral Geometry								
Zeroed Geometry	15.41	42.71	70.90	53	6.95	70.03	88.95	18

5.4 Feedback Correspondence Analysis

Shuffled Feedback examines whether ActiveMS-FCP's gains depend on the correct correspondence between verified candidates and their spectral similarity feedback. At each planning update, observed feedback values are randomly permuted among verified candidates, while the trained model, initial ranking evidence, and *in silico* spectral geometry remain unchanged. As Table 4 shows, shuffling on NIST lowers S@R4 and increases MVR with each neural spectrum predictor, and SpectraVerse shows the same pattern. With MassFormer on NIST, S@R4 falls from 54.02% to 42.31%, while MVR increases from 8.69 to 15.35. The SpectraVerse results are reported in Appendix A.3.

After shuffling, performance approaches Static Ranking across all four settings: S@R4 differs by at most 0.69 percentage points and MVR by at most 0.83 rounds. Because the trained model is unchanged, this degradation indicates that feedback obtained during verification guides candidate selection. The planning gains of ActiveMS-FCP depend on correctly associating observed spectral feedback with verified candidate structures.

Table 4: ActiveMS-FCP feedback correspondence analysis on the NIST benchmark. Results are five-seed averages. Δ reports the mean $\pm$ std of Shuffled minus the reference.

	MVR↓			S@R4 (%)↑		
Predictor	Shuffled	Δ vs. Original	Δ vs. Static	Shuffled	Δ vs. Original	Δ vs. Static
MassFormer	15.35	$+6.67_{\pm 0.34}$	$-0.04_{\pm 0.22}$	42.31	$-11.71_{\pm 0.70}$	$-0.05_{\pm 0.53}$
ICEBERG	6.76	$+2.41_{\pm 0.14}$	$-0.30_{\pm 0.07}$	70.09	$-7.98_{\pm 0.50}$	$+0.60_{\pm 0.29}$

6 Conclusion

ActiveMS formulates candidate verification after neural spectrum prediction as active planning, moving from static ranking to decisions conditioned on the spectral similarity feedback that each verification returns. Predicted spectra serve two roles in this formulation: they supply the initial ranking and define the *in silico* spectral geometry through which feedback from verified candidates transfers to unverified ones. Over this geometry, ActiveMS-GP models the feedback and ActiveMS-FCP models the target. On ActiveMS Bench, built from two spectral libraries with an independent non-neural simulator for verification feedback, both planners identify targets in fewer verification rounds than static ranking with MassFormer and ICEBERG. Ablation and shuffling experiments trace the gains of ActiveMS-FCP to spectral geometry and correct feedback correspondence, while revealing robustness to the choice of rollout source. With ICEBERG, the gains persist even when the initial ranking is already strong, indicating that feedback-conditioned planning complements advances in spectrum prediction. Two directions remain open: planning with experimental feedback measured on reference standards rather than simulated by CFM-ID, and candidate sets that may not contain the target, where a planner must decide when to stop verifying.

Reproducibility Statement

Reproducing our results requires the planners, the benchmarks, and the experimental protocol. The planners are specified in Section 3, and the ActiveMS-FCP architecture and training procedure are detailed in Appendix B. Benchmark construction, including query selection, candidate sets, data splits, and CFM-ID verification feedback, is documented in Section 4 and Appendix C. The evaluation protocol, predictor training, planner configurations, and random seeds are described in Section 5.1 and Appendix D. Upon acceptance, we will release source code, trained model weights, the SpectraVerse portion of ActiveMS Bench, and CFM-ID verification feedback for both benchmarks, along with scripts to rebuild the NIST portion from licensed NIST20 and NIST23 exports.

AI Use Statement

Generative AI tools assisted with language polishing of the authors' original draft and code debugging. The authors reviewed and verified the suggested textual revisions and code changes, and take responsibility for the content of this paper.

References

Wout Bittremieux, Mingxun Wang, and Pieter C. Dorrestein. The critical role that spectral libraries play in capturing the metabolomics community knowledge. *Metabolomics*, 18:94, 2022. doi: 10.1007/s11306-022-01947-y. URL https://doi.org/10.1007/s11306-022-01947-y.

Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010. URL https://arxiv.org/abs/1012.2599.

Roman Bushuiev, Anton Bushuiev, Niek F. de Jonge, Adamo Young, Fleming Kretschmer, Raman Samusevich, Janne Heirman, Fei Wang, Luke Zhang, Kai Dührkop, Marcus Ludwig, Nils A. Haupt, Apurva Kalia, Corinna Brungs, Robin Schmid, Russell Greiner, Bo Wang, David S. Wishart, Li-Ping Liu, Juho Rousu, Wout Bittremieux, Hannes Röst, Tytus D. Mak, Soha Hassoun, Florian Huber, Justin J. J. van der Hooft, Michael A. Stravs, Sebastian Böcker, Josef Sivic, and Tomáš Pluskal. MassSpecGym: A benchmark for the discovery and identification of molecules. In *Advances in Neural Information Processing Systems*, volume 37, pp. 110010–110027, 2024. doi: 10.52202/079017-3491. URL https://doi.org/10.52202/079017-3491.

Marta Garnelo, Dan Rosenbaum, Christopher J. Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo J. Rezende, and S. M. Ali Eslami. Conditional neural processes. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1704–1713, 2018. URL https://proceedings.mlr.press/v80/garnelo18a.html.

Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. In *Proceedings of the 29th International Conference on Machine Learning*, pp. 843–850. Omnipress, 2012. URL https://icml.cc/2012/papers/113.pdf.

Samuel Goldman, John Bradshaw, Jiayi Xin, and Connor W. Coley. Prefix-tree decoding for predicting mass spectra from molecules. In *Advances in Neural Information Processing Systems*, volume 36, pp. 48548–48572, 2023. URL https://arxiv.org/abs/2303.06470.

Samuel Goldman, Janet Li, and Connor W. Coley. Generating molecular fragmentation graphs with autoregressive neural networks. *Analytical Chemistry*, 96(8):3419–3428, 2024. doi: 10.1021/acs.analchem.3c04654. URL https://doi.org/10.1021/acs.analchem.3c04654.

Vishu Gupta, Hantao Qiang, Hsin Hsiang Chung, Ehud Herbst, and Michael A. Skinnider. Comprehensive curation and harmonization of small-molecule MS/MS libraries in Spectraverse. *Analytical Chemistry*, 98(5):3934–3943, 2026a. doi: 10.1021/acs.analchem.5c06256. URL https://doi.org/10.1021/acs.analchem.5c06256.

Vishu Gupta, Chencheng Xu, Ehud Herbst, Fei Wang, David S Wishart, and Michael A Skinnider. Confronting spurious evaluations of computational methods in small molecule mass spectrometry. *bioRxiv*, pp. 2026–05, 2026b.

Stephen R. Heller, Alan McNaught, Igor Pletnev, Stephen Stein, and Dmitrii Tchekhovskoi. InChI, the IUPAC international chemical identifier. *Journal of Cheminformatics*, 7:23, 2015. doi: 10.1186/s13321-015-0068-4. URL https://doi.org/10.1186/s13321-015-0068-4.

Yuhui Hong, Sujun Li, Christopher J. Welch, Shane Tichy, Yuzhen Ye, and Haixu Tang. 3DMolMS: Prediction of tandem mass spectra from 3D molecular conformations. *Bioinformatics*, 39(6): btad354, 2023. doi: 10.1093/bioinformatics/btad354. URL https://doi.org/10.1093/bioinformatics/btad354.

Yuhui Hong, Yuzhen Ye, and Haixu Tang. Machine learning in small-molecule mass spectrometry. *Annual Review of Analytical Chemistry*, 18:193–215, 2025. doi: 10.1146/annurev-anchem-071224-082157. URL https://doi.org/10.1146/annurev-anchem-071224-082157.

Florian Huber, Lars Ridder, Stefan Verhoeven, Jurriaan H. Spaaks, Faruk Diblen, Simon Rogers, and Justin J. J. van der Hooft. Spec2Vec: Improved mass spectral similarity scoring through learning of structural relationships. *PLoS Computational Biology*, 17(2):e1008724, 2021a. doi: 10.1371/journal.pcbi.1008724. URL https://doi.org/10.1371/journal.pcbi.1008724.

Florian Huber, Sven van der Burg, Justin J. J. van der Hooft, and Lars Ridder. MS2DeepScore: A novel deep learning similarity measure to compare tandem mass spectra. *Journal of Cheminformatics*, 13:84, 2021b. doi: 10.1186/s13321-021-00558-4. URL https://doi.org/10.1186/s13321-021-00558-4.

Shali Jiang, Gustavo Malkomes, Geoff Converse, Alyssa Shofner, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic active search. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1714–1723, 2017. URL https://proceedings.mlr.press/v70/jiang17d.html.

Shali Jiang, Gustavo Malkomes, Matthew Abbott, Benjamin Moseley, and Roman Garnett. Efficient nonmyopic batch active search. In *Advances in Neural Information Processing Systems*, volume 31, pp. 1099–1109, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/4c27cea8526af8cfee3be5e183ac9605-Abstract.html.

Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, S. M. Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=SkE6PjC9KX.

Sunghwan Kim, Jie Chen, Tiejun Cheng, Asta Gindulyte, Jia He, Siqian He, Qingliang Li, Benjamin A. Shoemaker, Paul A. Thiessen, Bo Yu, Leonid Zaslavsky, Jian Zhang, and Evan E. Bolton. PubChem 2025 update. *Nucleic Acids Research*, 53(D1):D1516–D1525, 2025. doi: 10.1093/nar/gkae1059. URL https://doi.org/10.1093/nar/gkae1059.

Tobias Kind, Hiroshi Tsugawa, Tomas Cajka, Yan Ma, Zijuan Lai, Sajjan S. Mehta, Gert Wohlgemuth, Dinesh Kumar Barupal, Megan R. Showalter, Masanori Arita, and Oliver Fiehn. Identification of small molecules using accurate mass MS/MS search. *Mass Spectrometry Reviews*, 37(4):513–532, 2018. doi: 10.1002/mas.21535. URL https://doi.org/10.1002/mas.21535.

Wouter Kool, Herke van Hoof, and Max Welling. Stochastic beams and where to find them: The Gumbel-Top-k trick for sampling sequences without replacement. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3499–3508, 2019. URL https://proceedings.mlr.press/v97/kool19a.html.

Fleming Kretschmer, Jan Seipp, Marcus Ludwig, Gunnar W. Klau, and Sebastian Böcker. Coverage bias in small molecule machine learning. *Nature Communications*, 16:554, 2025. doi: 10.1038/s41467-024-55462-w. URL https://doi.org/10.1038/s41467-024-55462-w.

Yuanyue Li, Tobias Kind, Jacob Folz, Arpana Vaniya, Sajjan Singh Mehta, and Oliver Fiehn. Spectral entropy outperforms MS/MS dot product similarity for small-molecule compound identification. *Nature Methods*, 18(12):1524–1531, 2021. doi: 10.1038/s41592-021-01331-z. URL https://doi.org/10.1038/s41592-021-01331-z.

Michael Murphy, Stefanie Jegelka, Ernest Fraenkel, Tobias Kind, David Healey, and Thomas Butler. Efficiently predicting high resolution mass spectra with graph neural networks. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 25549–25562, 2023. URL https://proceedings.mlr.press/v202/murphy23a.html.

National Institute of Standards and Technology. NIST/EPA/NIH mass spectral library (NIST 14). NIST Standard Reference Database 1A, 2014. URL https://www.nist.gov/system/files/documents/srd/NIST1aVer22Man.pdf.

Rıza Özçelik, Sarah de Ruiter, Emanuele Criscuolo, and Francesca Grisoni. Chemical language modeling with structured state space sequence models. *Nature Communications*, 15(1):6176, 2024.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006. ISBN 978-0-262-18253-9. URL https://gaussianprocess.org/gpml/.

David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010. doi: 10.1021/ci100050t. URL https://doi.org/10.1021/ci100050t.

Daniel J. Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on thompson sampling. *Foundations and Trends in Machine Learning*, 11(1):1–96, 2018. doi: 10.1561/2200000070. URL https://doi.org/10.1561/2200000070.

Christoph Ruttkies, Emma L. Schymanski, Sebastian Wolf, Juliane Hollender, and Steffen Neumann. MetFrag relaunched: Incorporating strategies beyond in silico fragmentation. *Journal of Cheminformatics*, 8:3, 2016. doi: 10.1186/s13321-016-0115-9. URL https://doi.org/10.1186/s13321-016-0115-9.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. doi: 10.1109/JPROC.2015.2494218. URL https://doi.org/10.1109/JPROC.2015.2494218.

Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 1015–1022, 2010. URL https://proceedings.mlr.press/v9/srinivas10a.html.

Stephen E. Stein. Mass spectral reference libraries: An ever-expanding resource for chemical identification. *Analytical Chemistry*, 84(17):7274–7282, 2012. doi: 10.1021/ac301205z. URL https://doi.org/10.1021/ac301205z.

Winnie Uritboonthai, Linh Hoang, and Gary Siuzdak. Structure–spectrum fidelity in tandem mass spectrometry: Quantifying the limits of similarity-based structural inference. *Analytical Chemistry*, 2026. doi: 10.1021/acs.analchem.6c04898. URL https://doi.org/10.1021/acs.analchem.6c04898. Advance online publication.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008, 2017.

Fei Wang, Jaanus Liigand, Siyang Tian, David Arndt, Russell Greiner, and David S. Wishart. CFM-ID 4.0: More accurate ESI-MS/MS spectral prediction and compound identification. *Analytical Chemistry*, 93(34):11692–11700, 2021. doi: 10.1021/acs.analchem.1c01465. URL https://doi.org/10.1021/acs.analchem.1c01465.

Mingxun Wang, Jeremy J. Carver, Vanessa V. Phelan, Laura M. Sanchez, Neha Garg, Yao Peng, Don D. Nguyen, Jeramie Watrous, Clifford A. Kapono, Tal Luzzatto-Knaan, et al. Sharing and community curation of mass spectrometry data with global natural products social molecular networking. *Nature Biotechnology*, 34(8):828–837, 2016. doi: 10.1038/nbt.3597. URL https://doi.org/10.1038/nbt.3597.

Rui-Xi Wang, Runzhong Wang, and Connor W. Coley. GLACIER: Rethinking mass spectrum prediction as an object detection problem. *arXiv preprint arXiv:2606.29161*, 2026.

Runzhong Wang, Mrunali Manjrekar, Babak Mahjour, Julian Avila-Pacheco, Joules Provenzano, Erin E. Reynolds, Magdalena Lederbauer, Eivgeni Mashin, Samuel L. Goldman, Mingxun Wang, Jing-Ke Weng, Desirée L. Plata, Clary B. Clish, and Connor W. Coley. Neural spectral prediction for structure elucidation with tandem mass spectrometry. *bioRxiv*, 2025. doi: 10.1101/2025.05.28.656653. URL https://doi.org/10.1101/2025.05.28.656653. Version 2 posted January 22, 2026.

Adamo Young, Hannes Röst, and Bo Wang. Tandem mass spectrum prediction for small molecules using graph transformers. *Nature Machine Intelligence*, 6(4):404–416, 2024. doi: 10.1038/s42256-024-00816-8. URL https://doi.org/10.1038/s42256-024-00816-8.

A Additional Results and Analysis

A.1 Full Coverage Curves

Figure 3 shows query coverage as a function of the verification round. With MassFormer, ActiveMS-FCP reaches 98.2% coverage on each benchmark by round 64, compared with 82.2% on SpectraVerse and 92.6% on NIST for Static Ranking. With ICEBERG, both planners exceed 99% by round 64. Across five training seeds, ActiveMS-FCP's standard deviations are 0.14–0.32 pp for S@R4, 0.35–0.63 pp for S@R16, and 0.11–0.23 rounds for MVR. Paired molecule-level bootstrap with 10,000 replicates, additionally resampling training seeds, yields positive 95% confidence intervals for both planners' gains over Static Ranking in all four settings.

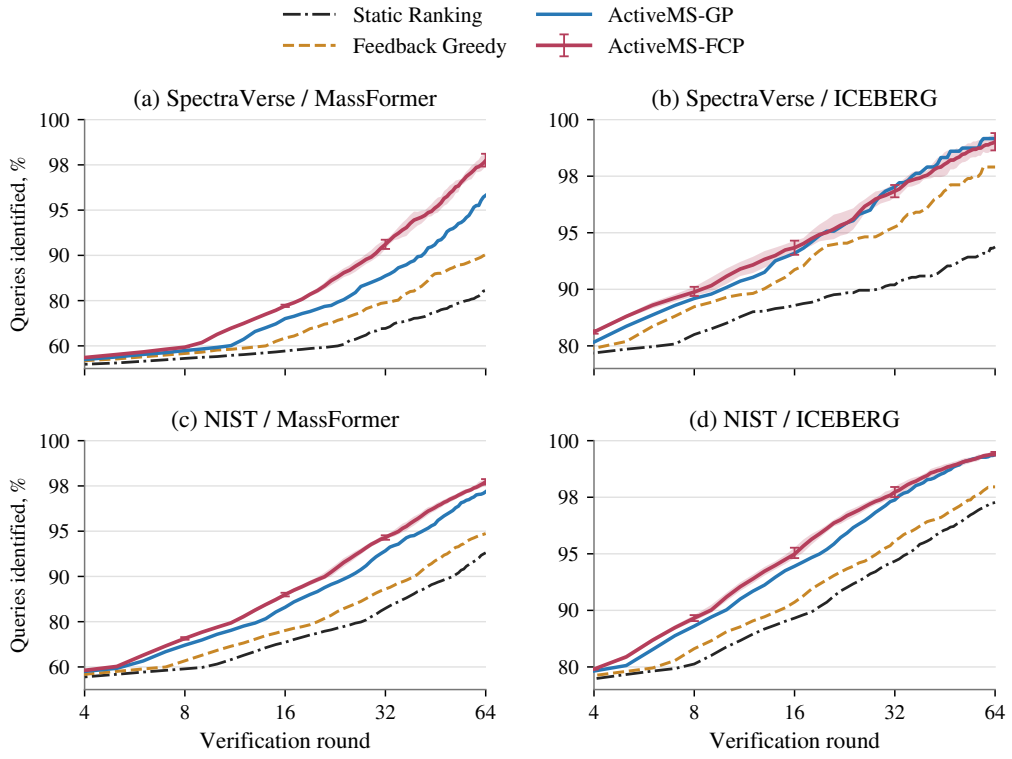


Figure 3: Query coverage as a function of the verification round. (a, b) SpectraVerse; (c, d) NIST. Left column: MassFormer; right column: ICEBERG. Horizontal axes are logarithmic. ActiveMS-FCP curves show five-seed means with shaded bands indicating ±1 std.

A.2 ActiveMS-FCP Ablation Study on SpectraVerse

Table 5: Ablation of ActiveMS-FCP on the SpectraVerse benchmark.

	MassFormer				ICEBERG			
Variant	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
Full ActiveMS-FCP	**11.89**	39.33	**77.70**	**30**	**4.37**	**82.46**	93.68	**9**
w/o Mixed Rollouts								
Prior rollouts only	13.22	38.86	74.70	35	4.43	81.81	**93.74**	**9**
Greedy rollouts only	13.19	**39.35**	74.84	36	4.55	81.97	93.27	10
Stochastic rollouts only	13.03	38.84	75.20	35	4.54	81.83	93.09	10
w/o *In Silico* Spectral Geometry								
Zeroed Geometry	24.91	27.55	50.79	> 64	8.53	73.90	86.97	28

Table 5 reports the SpectraVerse ablation. With MassFormer, single-source variants reach S@R16 of 74.70–75.20% and MVR of 13.03–13.22 rounds, compared with 77.70% and 11.89 rounds for mixed training. With ICEBERG, single-source and mixed training yield similar performance. Zeroed Geometry brings performance close to Static Ranking; with MassFormer, S@R16 falls to 50.79%, compared with 50.95% for Static Ranking.

A.3 ActiveMS-FCP Feedback Correspondence on SpectraVerse

Table 6 reports the SpectraVerse feedback correspondence analysis. Shuffling increases MVR and lowers S@R4 with each predictor, approaching Static Ranking performance.

Table 6: ActiveMS-FCP feedback correspondence analysis on the SpectraVerse benchmark. Results are five-seed averages. Δ reports the mean $\pm$ std of Shuffled minus the reference.

	MVR↓			S@R4 (%)↑		
Predictor	Shuffled	Δ vs. Original	Δ vs. Static	Shuffled	Δ vs. Original	Δ vs. Static
MassFormer	24.05	$+12.16_{\pm 0.85}$	$-0.83_{\pm 0.68}$	26.93	$-12.40_{\pm 0.89}$	$-0.56_{\pm 0.91}$
ICEBERG	8.39	$+4.02_{\pm 0.20}$	$-0.05_{\pm 0.20}$	74.44	$-8.02_{\pm 0.76}$	$+0.69_{\pm 0.62}$

A.4 ActiveMS-GP Ablation Study on Kernels

Table 7 compares cosine, RBF, Matérn-$\frac{1}{2}$, Matérn-$\frac{3}{2}$, and Matérn-$\frac{5}{2}$ kernels with Posterior Mean acquisition. RBF is the main configuration. Each covariance is averaged over the same collision energies. The observation noise and batch-correlation threshold remain fixed, and the radial kernels share the same length-scale multiplier without retuning.

Table 7: ActiveMS-GP kernel ablation with Posterior Mean acquisition.

	MassFormer				ICEBERG			
Kernel	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
SpectraVerse								
Static Ranking	24.88	27.49	50.95	> 64	8.44	73.74	87.15	29
Cosine	15.68	35.87	69.05	45	5.03	80.11	91.96	12
RBF	14.78	35.87	**72.07**	42	4.49	80.67	93.18	**10**
Matérn-$\frac{1}{2}$	15.09	36.20	68.60	**40**	4.58	80.34	93.07	11
Matérn-$\frac{3}{2}$	14.83	**36.42**	71.28	42	4.47	**81.34**	93.30	11
Matérn-$\frac{5}{2}$	**14.71**	36.20	70.95	41	**4.42**	80.78	**93.52**	11
NIST								
Static Ranking	15.40	42.36	70.90	52	7.06	69.49	88.63	19
Cosine	10.15	50.66	82.20	28	4.87	76.37	93.53	11
RBF	9.70	52.02	83.19	25	4.71	76.40	93.90	**10**
Matérn-$\frac{1}{2}$	**9.42**	51.57	**84.18**	**24**	4.66	75.95	**94.11**	11
Matérn-$\frac{3}{2}$	9.60	**52.26**	83.50	26	4.66	**76.72**	94.04	11
Matérn-$\frac{5}{2}$	9.65	51.80	83.53	26	**4.63**	76.55	94.05	**10**

Every kernel improves S@R4, S@R16, and MVR over Static Ranking in all four settings. Matérn-$\frac{3}{2}$ attains the highest S@R4 throughout, while the best S@R16 and MVR depend on the benchmark and spectrum predictor. RBF attains the highest S@R16 on SpectraVerse with MassFormer, at 72.07%. On NIST with MassFormer, Matérn-$\frac{1}{2}$ reduces MVR from 9.70 rounds with RBF to 9.42 rounds. The gains over Static Ranking therefore persist across covariance families.

A.5 ActiveMS-GP Ablation Study on Acquisition Rules

Table 8 compares five acquisition rules with the RBF kernel, GP parameters, and batch-selection settings held fixed: posterior mean (PM), probability of improvement (PI) and expected improvement (EI) (Brochu et al., 2010), upper confidence bound (UCB) (Srinivas et al., 2010), and Thompson

sampling (TS) (Russo et al., 2018). PM provides a competitive reference, achieving higher S@R16 and lower MVR than PI, EI, and TS across all four settings. UCB attains the highest S@R16 and the lowest MVR in every setting. On SpectraVerse with MassFormer, UCB reduces MVR from 14.78 to 14.19 and R@P90 from 42 to 39 rounds relative to PM.

Table 8: ActiveMS-GP acquisition rule ablation with a fixed RBF kernel.

	MassFormer				ICEBERG			
Acquisition	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
SpectraVerse								
Thompson Sampling	39.54	13.54	23.64	> 64	17.62	57.94	65.32	> 64
Expected Improvement	26.88	18.55	44.02	> 64	9.24	70.95	83.69	36
Probability of Improvement	18.84	29.83	63.13	60	5.22	78.44	92.51	12
Posterior Mean	14.78	35.87	72.07	42	4.49	80.67	93.18	**10**
Upper Confidence Bound	**14.19**	**37.65**	**72.51**	**39**	**4.18**	**80.89**	**94.75**	**10**
NIST								
Thompson Sampling	29.29	23.40	37.53	> 64	29.61	44.32	47.81	> 64
Expected Improvement	21.51	28.81	54.90	> 64	21.50	50.90	58.86	> 64
Probability of Improvement	14.74	44.01	72.45	49	8.71	71.01	86.84	28
Posterior Mean	9.70	**52.02**	83.19	25	4.71	76.40	93.90	**10**
Upper Confidence Bound	**9.47**	51.65	**83.86**	**24**	**4.64**	**76.58**	**94.13**	**10**

A.6 Feedback Greedy Ablation Study on Similarity Metric and Propagation Weight

Table 9: Feedback Greedy ablation on similarity metric and propagation weight α.

	MassFormer				ICEBERG			
Configuration	MVR↓	S@R4↑	S@R16↑	R@P90↓	MVR↓	S@R4↑	S@R16↑	R@P90↓
SpectraVerse								
Static Ranking	24.88	27.49	50.95	> 64	8.44	73.74	87.15	29
Cosine ($\alpha = 0.25$)	23.49	29.27	53.52	> 64	6.75	75.98	89.16	17
Cosine ($\alpha = 0.50$)	21.50	30.61	57.09	> 64	6.21	77.09	90.50	15
Cosine ($\alpha = 0.75$)	18.66	**33.85**	63.46	64	**5.48**	**77.88**	91.73	**13**
Cosine ($\alpha = 1.00$)	21.52	21.45	55.42	61	8.48	64.47	84.25	26
Entropy ($\alpha = 0.25$)	23.40	29.61	53.63	> 64	6.79	76.31	89.27	19
Entropy ($\alpha = 0.50$)	21.28	30.84	57.77	> 64	6.27	77.32	90.61	15
Entropy ($\alpha = 0.75$)	**18.49**	**33.85**	**63.58**	62	5.56	77.43	**91.84**	**13**
Entropy ($\alpha = 1.00$)	21.25	23.46	55.08	**59**	8.37	64.25	86.26	26
NIST								
Static Ranking	15.40	42.36	70.90	52	7.06	69.49	88.63	19
Cosine ($\alpha = 0.25$)	14.88	43.51	72.00	49	6.89	70.26	88.99	18
Cosine ($\alpha = 0.50$)	14.15	45.07	73.27	46	6.65	71.11	89.51	17
Cosine ($\alpha = 0.75$)	12.87	47.00	76.10	40	6.16	72.40	90.73	**15**
Cosine ($\alpha = 1.00$)	18.18	30.64	62.23	55	12.39	49.60	75.92	39
Entropy ($\alpha = 0.25$)	14.82	43.68	72.10	49	6.87	70.29	89.06	18
Entropy ($\alpha = 0.50$)	14.02	45.22	73.55	45	6.60	71.23	89.66	17
Entropy ($\alpha = 0.75$)	**12.63**	**47.38**	**76.66**	**39**	**6.08**	**72.55**	**90.85**	**15**
Entropy ($\alpha = 1.00$)	17.77	31.51	63.59	53	12.01	49.52	76.79	37

Table 9 reports the complete test grid with entropy-based initial scores $s_0(q, m)$ and CFM-ID feedback $y_q(m)$ held fixed. Every configuration starts with the Static top-four batch. Both propagation metrics improve S@R4, S@R16, and MVR over Static Ranking in all four settings for $\alpha \in \{0.25, 0.5, 0.75\}$. At $\alpha = 0.75$, entropy gives slightly higher S@R16 throughout, while MVR favors entropy in three settings and cosine in one. At $\alpha = 1$, subsequent selections discard the initial scores, and S@R4 falls below Static Ranking in every setting with either metric. This grid is a descriptive sensitivity analysis; the main configuration remains cosine propagation with $\alpha = 0.75$.

B ACTIVEMS-FCP DETAILS

B.1 NETWORK ARCHITECTURE

The candidate encoder receives 12 features: the three energy-specific matching scores and their mean $s_0(q, m)$, their four standardized values across the candidate set, and the standard deviation, range, maximum, and minimum of the three energy-specific scores. The static prior is the standardized $s_0(q, m)$ divided by the learned temperature τ, which is constrained to $[0.1, 10]$.

The pairwise feature vector $\mathbf{x}_{m,m^v}$ contains energy-specific spectral similarities $\boldsymbol{\rho}_{m,m^v}$ and their mean, minimum, and maximum; initial ranking scores $s_0(q, m)$, $s_0(q, m^v)$, and their difference; and observed feedback $y_q(m^v)$ and residual $r(m^v)$. The products $\bar{\rho}_{m,m^v} y_q(m^v)$ and $\bar{\rho}_{m,m^v} r(m^v)$ encode geometry–feedback interactions, where $\bar{\rho}_{m,m^v}$ is the mean similarity across energies.

The global context $\mathbf{h}_t^{\mathrm{ctx}}$ pools candidate representations $\mathbf{h}_m$ and verified feedback representations $\mathbf{h}_{m^v}^{\mathrm{fb}}$:

$$\mathbf{h}_t^{\mathrm{ctx}} = f_{\mathrm{ctx}}\big(\mathrm{Concat}\big(\mathrm{Pool}\{\mathbf{h}_m : m \in \mathcal{C}_q\},\, \mathrm{Pool}\{\mathbf{h}_{m^v}^{\mathrm{fb}} : m^v \in \mathcal{V}_{t-1}\},\, \mathbf{u}_t\big)\big)\,,$$

where f_{ctx} is a learned context encoder and Pool concatenates mean and element-wise max pooling. The auxiliary feature vector $\mathbf{u}_t$ contains the log candidate-set size $\log K$, verified fraction $|\mathcal{V}_{t-1}|/K$, and normalized collision energies from $\mathcal{E}_q$. The resulting context $\mathbf{h}_t^{\mathrm{ctx}}$ is shared across candidates in round t.

A fusion network f_{fuse} forms the joint representation $\mathbf{h}_{m,t}^{\mathrm{joint}}$, from which a linear head predicts the correction:

$$\mathbf{h}_{m,t}^{\mathrm{joint}} = f_{\mathrm{fuse}}\big(\mathrm{Concat}\big(\mathbf{h}_m,\, \mathbf{h}_{m,t}^{\mathrm{obs}},\, \mathbf{h}_t^{\mathrm{ctx}}\big)\big)\,,$$

$$\Delta_\theta(m \mid \mathcal{F}_q, \mathcal{H}_{t-1}) = \mathbf{w}_\Delta^\top \mathbf{h}_{m,t}^{\mathrm{joint}} + b_\Delta,$$

where $\mathbf{w}_\Delta$ and b_Δ are the weight vector and bias of the correction head.

All encoding and fusion networks have two linear layers with a hidden dimension of 64, SiLU activation, and dropout of 0.1. Geometry-aware cross-attention has $N_{\mathrm{head}} = 4$ heads with 16-dimensional value vectors, yielding the 64-dimensional feedback-conditioned representation $\mathbf{h}_{m,t}^{\mathrm{obs}}$. The correction head's weights $\mathbf{w}_\Delta$ and bias b_Δ are initialized to zero, and the learned temperature τ starts at 1, so the initial model preserves the static ranking.

B.2 ROLLOUT CONSTRUCTION

Rollout collection proceeds in stages. An intermediate ActiveMS-FCP model π_0, trained on prior-rollout examples $\mathcal{D}_{\mathrm{prior}}$, collects greedy-rollout examples $\mathcal{D}_{\mathrm{greedy}}$. A second model π_1, trained on their union, collects stochastic-rollout examples $\mathcal{D}_{\mathrm{stoch}}$ with Gumbel-top-B sampling (Kool et al., 2019) at temperature 1.0. The final training set is

$$\mathcal{D}_{\mathrm{train}} = \mathrm{Dedup}(\mathcal{D}_{\mathrm{prior}} \cup \mathcal{D}_{\mathrm{greedy}} \cup \mathcal{D}_{\mathrm{stoch}})\,,$$

where examples are deduplicated by $(q, \mathcal{V}_{t-1})$. Examples are collected for decisions in rounds 2–16 while the target remains unverified. The final model is trained from scratch on this combined set.

B.3 AUXILIARY RESIDUAL SUPERVISION

Let $\mathcal{U}_{q,t} = \mathcal{C}_q \setminus (\mathcal{V}_{t-1} \cup \{m^\star\})$ denote the unverified non-target candidates in a training example. The auxiliary head predicts their feedback residuals $r(m) = y_q(m) - s_0(q, m)$ from the joint representation $\mathbf{h}_{m,t}^{\mathrm{joint}}$:

$$\hat{r}_\theta(m \mid \mathcal{F}_q, \mathcal{H}_{t-1}) = \tanh\Big(\mathbf{w}_{\mathrm{res}}^\top \mathbf{h}_{m,t}^{\mathrm{joint}} + b_{\mathrm{res}}\Big)\,,$$

where $\mathbf{w}_{\mathrm{res}}$ and b_{res} are the weight and bias of the residual head. The feedback-weighted residual loss is

$$\mathcal{L}_{\mathrm{res}} = \frac{\sum_{m \in \mathcal{U}_{q,t}} w(m)\, \ell_{\mathrm{SL1}}(\hat{r}_\theta(m \mid \mathcal{F}_q, \mathcal{H}_{t-1}) - r(m))}{\sum_{m \in \mathcal{U}_{q,t}} w(m)},$$

$$w(m) = 1 + 2y_q(m).$$

Here, ℓ_{SL1} is the Smooth-L_1 loss with transition parameter 0.1. The residual head participates only in training.

C ACTIVEMS BENCH DETAILS

C.1 QUERY CONSTRUCTION

Queries are constructed from structure-annotated records in SpectraVerse 1.0.2 and NIST20/23. Both benchmarks retain positive-ion Orbitrap MS/MS spectra with $[\mathrm{M} + \mathrm{H}]^+$ adducts and scalar NCE annotations. Spectra are grouped by acquisition conditions and InChIKey-14 (the first 14 characters of the InChIKey that disregard stereochemistry) (Heller et al., 2015). Each group must contain nonempty spectra at three or more distinct collision energies.

Within each eligible group, an ordered triple of distinct energies is selected to minimize the total absolute deviation from the NCE anchors 20, 30, and 40. One experimental spectrum is retained at each selected energy, and the resulting spectrum–energy pairs form the query observation $\mathcal{S}_q$.

C.2 CANDIDATE SET CONSTRUCTION

For each target structure $m^\star$, up to 512 PubChem decoys are selected from structures sharing its molecular formula, ranked by ECFP4 Tanimoto similarity to the target. S4 chemical language models (Özçelik et al., 2024) trained separately on the molecular structures of each benchmark provide the generated structure pools, which are filtered by the target's molecular formula. After removing target matches and structures already selected from PubChem, up to 256 generated decoys are selected by the same similarity criterion, followed by up to 255 sampled uniformly without replacement from the remaining pool.

Candidates are restricted to valid, neutral, single-component structures and deduplicated by InChIKey-14. The target is retained once alongside the selected decoys, yielding a candidate set $\mathcal{C}_q$ of up to 1,024 structures. When a source cannot fill its quota, the available candidates are retained; queries with fewer than 128 candidates are filtered out. Candidate order is randomized before evaluation.

C.3 DATA SPLITS

SpectraVerse follows a component-based structural split. Target molecules form the nodes of a graph, with an edge connecting two molecules if their MCES distance is below 10 or their ECFP4 Tanimoto similarity is at least 0.7. Each connected component is assigned as a whole to training, validation, or test, with proportions of approximately 80%, 10%, and 10% respectively. Component assignment accounts for molecular weight and heavy-atom count to balance molecular characteristics across splits. All queries associated with a target molecule inherit its split assignment.

NIST follows a temporal split, with NIST20 target molecules forming the training set. Molecules newly appearing in NIST23 are filtered to retain targets whose maximum ECFP4 Tanimoto similarity to any training target is below 0.7. Among the retained targets, those connected by similarities of at least 0.7 are grouped into connected components, which are assigned as a whole to two approximately equal validation and test splits, balancing target and query counts as well as molecular and acquisition attributes.

C.4 Simulated Feedback and Structural Relatedness

ActiveMS propagates verification feedback to guide subsequent candidate selection. Because the target structure is unknown during identification, structural similarity to it cannot be computed directly. Spectral feedback provides an accessible alternative: candidate spectra can be compared with the query spectrum without knowing the target. This is useful only if spectral feedback $y_q(m)$ correlates with structural similarity between candidate m and target $m^\star$. The following analysis tests this assumption on both benchmarks.

Structural relatedness is measured by ECFP4 Tanimoto and normalized MCES similarity,

$$s_{\mathrm{MCES}}(m^\star, m) = \max\left\{0,\, 1 - \frac{d_{\mathrm{MCES}}(m^\star, m)}{n_{\mathrm{bond}}(m^\star) + n_{\mathrm{bond}}(m)}\right\},$$

where d_{MCES} is the MCES distance and $n_{\mathrm{bond}}(m)$ counts the bonds in molecule m. Spectral feedback is computed from CFM-ID spectra projected onto a 15,000-bin grid spanning m/z 0–1,500, normalized to unit total intensity, and compared via spectral-entropy similarity.

Per-query Spearman correlation ρ measures whether candidates with higher feedback tend to be structurally closer to the target. Table 10 reports results for the 895 SpectraVerse and 6,671 NIST test queries. Positive MCES correlations occur in 86.48% and 92.88% of queries respectively, with positive correlations also predominating under ECFP4. Structural enrichment is the mean structural similarity of the top 5% of non-target candidates ranked by feedback minus the query-level average. This quantity is positive under both metrics on both benchmarks. The association is nonetheless probabilistic: within the top-5% subsets, on average 24.24% of candidates on SpectraVerse and 20.62% on NIST have MCES similarity at or below the query median. These results characterize feedback as a graded indicator of molecular relatedness, consistent with the association observed by Uritboonthai et al. (2026) in experimental METLIN spectra.

Table 10: Structural information in non-target candidate feedback on the SpectraVerse and NIST test splits.

Benchmark	Structural measure	Mean ρ	$\rho > 0$ queries (%)	Top-5% enrichment
SpectraVerse	MCES	0.208	86.48	0.081
	ECFP4	0.194	82.01	0.092
NIST	MCES	0.276	92.88	0.103
	ECFP4	0.249	90.05	0.109

C.5 Experimental Spectral Coverage

Experimental spectral feedback requires reference spectra for candidate structures as well as query targets. NIST23 is among the largest tandem MS libraries, containing 2.4 million spectra for 51,501 compounds.[1] In the NIST23 records analyzed here, same-formula coverage is sparse: after deduplicating by InChIKey-14, excluding training molecules, and restricting to positive-ion Orbitrap/HCD spectra with $[\mathrm{M} + \mathrm{H}]^+$ adducts at three or more collision energies, the largest same-formula group contains only 29 structures, well below the 128-candidate lower bound of ActiveMS Bench. Constructing candidate sets from experimental spectra alone is therefore infeasible at benchmark scale.

D Experimental Details

D.1 Neural Spectrum Predictors

Training. Both predictors are trained separately on the training spectra of SpectraVerse and NIST. MassFormer is initialized with PCQM4Mv2-pretrained graph weights, while ICEBERG follows its fragment-generation, intensity-prediction, and contrastive fine-tuning procedure. Training employs Adam with batch sizes of 128 for MassFormer and 32 for ICEBERG. Initial learning rates are

[1] https://www.nist.gov/programs-projects/nist23-updates-nist-tandem-and-electron-ionization-spectral-libraries

1.4×10^{-4} for MassFormer, 9.96×10^{-4} for ICEBERG fragment generation, and 7.36×10^{-4} for intensity prediction. Checkpoints are selected by minimum validation loss. Predictors trained on the full training split provide predicted candidate spectra for benchmark evaluation.

Prediction and scoring. For each query, candidate spectra $\hat{S}(m, e)$ are predicted at the three selected collision energies in $\mathcal{E}_q$. Spectral representations span m/z 0–1,500 with 15,000 bins. MassFormer retains the 100 highest-intensity output bins, whereas ICEBERG retains up to 100 highest-intensity predicted peaks before binning. Spectral-entropy matching against the experimental spectra yields energy-specific scores $s_j(q, m)$ and initial ranking scores $s_0(q, m)$. For the *in silico* spectral geometry, binned intensities are square-root transformed and L_2-normalized. The resulting forward predictions $\mathcal{F}_q$ are cached and shared across planners evaluated with the same predictor.

D.2 Planner Training and Configuration

Feedback Greedy uses $\alpha = 0.75$ across all settings. ActiveMS-GP and ActiveMS-FCP hyperparameters are listed in Table 11. ActiveMS-FCP training states are resampled each epoch, balancing queries and the round ranges 2–4, 5–8, and 9–16; AdamW is used with early stopping based on internal validation loss. Models are trained with seeds 11, 23, 37, 53, and 71 in each setting.

Table 11: Planner configurations for each benchmark and neural spectrum predictor.

	SpectraVerse		NIST	
Parameter	MassFormer	ICEBERG	MassFormer	ICEBERG
ActiveMS-GP				
Length-scale multiplier c_ℓ	2.0	2.0	2.0	1.0
Noise standard deviation σ_n	0.25	0.10	0.10	0.25
Correlation threshold	1.00	1.00	0.98	0.98
ActiveMS-FCP				
Training states per epoch	12,000	12,000	12,000	8,000
Internal validation states	1,500	1,500	1,500	1,000
Auxiliary loss weight λ_{res}	0.20	0.20	0.20	0.10

E Related Work

Forward MS/MS prediction provides a practical route to molecule identification beyond the coverage of experimental spectral libraries. By predicting spectra for candidate molecular structures and comparing them with observed experimental spectra, these methods rank candidates by spectral agreement. Neural spectrum prediction spans direct prediction from molecular representations and explicit fragmentation modeling, with representative methods including GrAFF-MS (Murphy et al., 2023), MassFormer (Young et al., 2024), 3DMolMS (Hong et al., 2023), SCARF (Goldman et al., 2023), ICEBERG (Goldman et al., 2024; Wang et al., 2025), and GLACIER (Wang et al., 2026). ActiveMS builds on spectra generated by frozen neural spectrum predictors to initialize candidate rankings and define an *in silico* spectral geometry over candidate structures.

Candidate verification can provide spectral similarity feedback by comparing experimental spectra of the candidate and the unknown target. The structural information conveyed by spectral similarity underpins molecular networking (Wang et al., 2016), as well as learned similarity measures such as Spec2Vec (Huber et al., 2021a) and MS2DeepScore (Huber et al., 2021b). Structure–spectrum fidelity analyses show that high spectral similarity is enriched among structurally related compounds, although closely related structures can produce substantially different spectra (Uritboonthai et al., 2026). This motivates treating spectral similarity feedback as informative but probabilistic evidence to guide further candidate verification. ActiveMS investigates how feedback from unmatched verified candidates can revise the assessment of unverified candidates within the *in silico* spectral space.

This formulation is closely related to active search, which allocates a limited query budget to discover members of a target class (Garnett et al., 2012), with extensions to nonmyopic planning (Jiang et al., 2017) and batch selection (Jiang et al., 2018). ActiveMS seeks a single target molecular

structure, with continuous spectral similarity feedback accompanying each identity check. Bayesian optimization provides a related framework for selecting queries from posterior predictions and uncertainty (Shahriari et al., 2016). ActiveMS-GP models continuous feedback residuals with a Gaussian process and derives acquisition scores from the resulting posterior. ActiveMS-FCP models the target structure with a neural posterior over candidate structures and selects candidates by posterior probability. Its rollout-based training amortizes posterior updates conditioned on frozen forward predictions and observed verification feedback, connecting to conditional neural processes (Garnelo et al., 2018) and their attentive extensions (Kim et al., 2019).