ENERGY REINFORCED DIFFUSION FOR STRUCTURED REPRESENTATION FROM HETEROGENEOUS SPECTRAL OBSERVATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Heterogeneous spectral observations provide high dimensional molecular signals for knowledge discovery and structured evidence learning. However, global representations often obscure informative relationships among local peak groups, while insufficient structural guidance makes these dependencies difficult to preserve during evidence exploration and policy optimization. Consequently, learned representations may rely on isolated peaks or background patterns, rather than informative group dependencies with consistent structural semantics. To this end, we propose Knowledge-Energy Reinforced Diffusion (KERD) for structured evidence learning from heterogeneous data. KERD organizes peak clusters into a Peak Graph and learns structured relations among band groups through data driven consistency modeling. The learned knowledge energy guides diffusion sampling toward band group combinations with stable structural semantics. Next, we construct class conditional spectral prototypes from peak cluster relations to improve the stability of policy learning. When paired cohort variables are available, KERD optionally refines the structured state through a shared relation graph that strengthens the alignment between peak substructures and auxiliary factor groups. Finally, we evaluate KERD on more than 70,000 measurements from multicenter cohorts. KERD achieves consistent gains across the evaluated heterogeneous datasets covering different sources and biological sample types.

1 INTRODUCTION

Structured continuous observations provide a natural setting for studying how stable relational structure can be recovered from high dimensional data (Gui et al., 2024). Their informative content is often expressed through coordinated variations among local regions rather than isolated high magnitude responses (Wang et al., 2024). Raman spectroscopy captures molecular fingerprint information through a noninvasive and label free acquisition process (Zareno et al., 2025; Wang et al., 2025). It thus provides a representative validation scenario for relation aware representation learning under source and measurement heterogeneity. Recent advances in high throughput acquisition have enabled large scale analysis across increasingly diverse spectral collections (Li et al., 2025a; Sun et al., 2025). Learning based approaches make it possible to recover structural dependencies that remain useful across different observation conditions (Tian et al., 2021; Qian et al., 2024). These advances have supported stain free tissue characterization (Huang et al., 2023), minimally invasive molecular detection (Ishwar et al., 2022), and rapid microbial identification (Li et al., 2025c).

For heterogeneous spectral observations, reliable evidence learning depends strongly on the relations between local peak groups and latent functional structure (Alberts et al., 2024; Bushuiev et al., 2024; He et al., 2026). The organization of peak clusters captures microscopic vibrational behavior together with broader variation in the observed system (Hua et al., 2024). Accordingly, informative evidence often lies in coordinated changes across multiple regions rather than in isolated peak responses. This principle has motivated Raman based methods for molecular distribution analysis (Kamp et al., 2024) and high dimensional structural profiling (Sigle et al., 2023). However, conventional global embeddings often suppress the relative structure encoded among peak groups. Most existing methods summarize the complete spectrum with a single representation and do not explicitly retain local group interactions (Zhang et al., 2024b; Zhou et al., 2022b; Shi et al., 2022). The

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.

resulting search space receives only weak and sparse learning feedback for group configurations. Such strategies often fail to identify coordinated group dependencies that persist across heterogeneous conditions (Li et al., 2023). As a consequence, learned representations may capture source specific variation more readily than stable structural relations (Kumar et al., 2023; Yang et al., 2023).

Two methodological limitations help explain why existing methods struggle to preserve stable group level dependencies. First, conventional representations often compress local peak groups into global features, without explicitly preserving their structural relationships. As a result, relative intensity patterns and coordinated group dependencies may be weakened before evidence selection (Li et al., 2025b). Second, exploring informative group configurations remains challenging when relational constraints are not consistently incorporated into the learning process. Prediction driven exploration alone may favor incidental group combinations, while policy updates without relation consistent feedback may struggle to preserve coherent dependencies across successive decisions(Tang et al., 2024; Hu et al., 2025a). These limitations raise a central question:

How can structural priors be used to recover stable group level dependencies from high dimensional heterogeneous observations?

Building on the above analysis, we develop Knowledge-Energy Reinforced Diffusion (KERD) for robust spectral representation learning under data heterogeneity. In this work, structured evidence refers to coordinated group relations and prototype interactions. KERD first constructs a fixed group topology and then models data driven relational statistics over this structure. Knowledge Aware Energy Diffusion Policy converts this relational geometry into differentiable diffusion guidance, while relation consistent policy feedback reinforces coherent group configurations. When matched cohort variables are available, KERD further refines the shared structured state through an optional context graph. Our theoretical analysis characterizes an idealized regularized diffusion reference and motivates the local energy surrogate used for structured mask exploration.

In summary, our main contributions are four-fold:

- We propose Knowledge-Energy Reinforced Diffusion (KERD), a framework that uses relational knowledge to guide structured evidence exploration.
- We develop a structured relation learning methodology capturing coordinated peak group dynamics as higher order discriminative relationships beyond peak wise representations.
- We collect and curate a multicenter cohort that covers nails, blood, urine, and other sample types. The cohort contains 72,057 measurements and supports evaluation of structural priors under cohort and biological sample type heterogeneity.
- We systematically evaluate KERD across heterogeneous spectral cohorts and verify its effectiveness through comparative studies, ablation analyses, and robustness evaluations.

2 Related work

Spectral Representation Learning Spectral representation learning has shifted from handcrafted peak selection toward supervised and weakly supervised feature extraction (Georgiev et al., 2024; Zhu et al., 2025; Lu et al., 2024; Li et al., 2022). Spectral libraries and reproducible pipelines support automated analysis across acquisition devices (Ho et al., 2019; Jiang et al., 2022). Advances in sequence modeling have enabled hierarchical feature extraction (Shao et al., 2024), global and frequency aware modeling (Wu et al., 2021; Zhou et al., 2022a), and cross variable interaction learning (Liu et al., 2023b). Flat minima optimization has also been explored under acquisition variation (Zareno et al., 2025). Despite these advances, global representations do not necessarily preserve structural dependencies among local observation groups.

Peak Group Structure Recent structure aware models capture dependencies between local segments (Wang et al., 2025; Hu et al., 2025b; Zhang et al., 2024c). Patch based methods reduce attention cost while preserving local patterns (Nie, 2022), and multiscale architectures capture information across different resolutions (Naghashi et al., 2025). Graph filtering reduces irrelevant

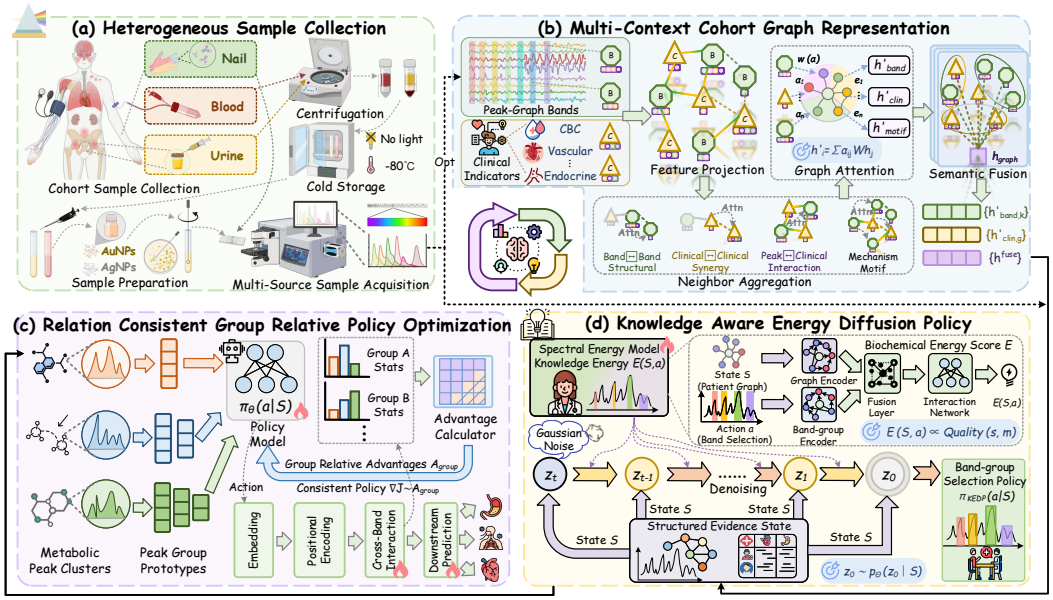


Figure 1: **Overall architecture of KERD.** KERD transforms heterogeneous spectral observations and cohort variables into a structured state, where Peak Graph relations preserve dependencies among band groups. The resulting state conditions an energy guided diffusion process that explores structurally coherent band group configurations, while prototype consistent group relative feedback reinforces trajectories that preserve stable relational structure.

interactions (Feng et al., 2026), while spectral pretraining introduces peak correlations through reconstruction and alignment objectives (Ren et al., 2025). These methods improve structural modeling, but predefined partitions and learned relations may change across sources and acquisition conditions.

Robust Evidence Learning under Heterogeneity Robust learning seeks to mitigate spurious correlations, incomplete observations, and distribution shifts (Wu et al., 2023; Deng et al., 2023). Existing solutions explore flat minima (Zhang et al., 2024a), concept interventions (Shao et al., 2023), progressive data expansion (Wang et al., 2023), and evidential alignment (Ye et al., 2025). Other approaches use multiple views or compensate for missing variables (Yu et al., 2025; Meng et al., 2026). However, these strategies often depend on additional annotations, multiple observations, or explicit missingness information. KERD instead integrates structural priors into the exploration and optimization of spectral evidence. Its learning process uses relational structure to guide diffusion and prototype consistent policy feedback, encouraging stable dependencies among observation groups.

3 Methodology

In this section, we introduce the Knowledge-Energy Reinforced Diffusion (KERD), as shown in Fig. 1. KERD uses shared relational structure to guide evidence exploration and policy learning.

3.1 Structured State Construction

Given an observation $x \in \mathbb{R}^L$ and optional context $c \in \mathbb{R}^D$, we partition the observation into B groups $\mathcal{G} = \{\mathcal{I}_1, \ldots, \mathcal{I}_B\}$ with topology $A \in \{0,1\}^{B\times B}$. The structural specification is $\mathcal{K} = (\mathcal{G}, A)$. For ordered observations, $\Delta = (\lambda_{\max} - \lambda_{\min})/B$, $\mathcal{I}_k = [\lambda_{\min} + (k-1)\Delta, \lambda_{\min} + k\Delta)$, with the final interval closed at $\lambda_{\max}$, and $A_{pq} = \mathbb{I}(|p-q| = 1)$. This topology forms the spectral Peak Graph. Here, knowledge comprises this topology and relation statistics estimated from training data, not external pathway annotations.

Local segments $X = \{x_i\}_{i=1}^N$ are embedded as $u_i = \Pi_\phi(x_i)$. We set $u_c = \Pi_c(c)$ or use a learnable null embedding when context is absent. The encoder yields $H = F_\theta([u_c, u_1, \ldots, u_N]) = [h_c, h_1, \ldots, h_N]$. We aggregate its output with $a_i = q^\top \tanh(W_h h_i + W_c h_c)$, $\alpha_i = \exp(a_i)/\sum_j \exp(a_j)$, and $z = \sum_i \alpha_i h_i$.

For each group, $b_k = |\mathcal{I}_k|^{-1} \sum_{\ell \in \mathcal{I}_k} x_\ell$ and $\mathcal{E} = \{(p,q) : A_{pq} > 0\}$. Multi Context Cohort Graph Representation (MCGR) combines these group features with K optional context groups in $H_0 = [b_1, \ldots, b_B, c_1, \ldots, c_K]^\top$. After $Z_g = H_0 W$, graph attention uses $e_{ij} = \mathrm{LeakyReLU}(a_{\mathrm{src}}^\top \zeta_i + a_{\mathrm{dst}}^\top \zeta_j)$ and $\alpha_{ij} = \exp(e_{ij})/\sum_{k \in \mathcal{N}(j)} \exp(e_{kj})$ to update $\zeta_j' = \sum_{i \in \mathcal{N}(j)} \alpha_{ij} \zeta_i$. The pooled graph state is fused with z. Without auxiliary context, MCGR reduces to the observation group graph.

3.2 Energy Guided Decision and Policy Optimization

Knowledge Aware Energy Diffusion Policy (KEDP) explores group configurations using an analytic structural energy. Relation Consistent Group Relative Policy Optimization (RCGRPO) uses the same relation geometry and prototypes for trajectory feedback.

Energy Shaped Diffusion Control Conditioned on $\mathrm{cond} = z$, KEDP defines a diffusion policy $q_\eta(m \mid \mathrm{cond})$ over continuous masks $m = \sigma(x_0) \in [0,1]^B$, with $x_0 \in \mathbb{R}^B$. Let $p_\omega(k \mid s, m) = \mathrm{softmax}(\mathrm{Head}_\omega(s, m))_k$ and $\bar{p}_k(s) = \mathrm{sg}[p_\omega(k \mid s, \mathbf{1})]$, where sg stops gradients. Predictive preservation is $Q_{\mathrm{pred}}(s, m) = -D_{\mathrm{KL}}(\bar{p}(s) \| p_\omega(\cdot \mid s, m))$.

For $b = (b_1, \ldots, b_B)$, define $\tilde{b}(m) = (m \odot b)/(\mathbf{1}^\top (m \odot b) + \epsilon_0)$. The shared relation feature preserves normalized group responses and their relative values:

$$\psi(b, m) = \left[\tilde{b}(m), \left\{ \log \frac{\tilde{b}_p(m) + \epsilon}{\tilde{b}_q(m) + \epsilon} \right\}_{(p,q) \in \mathcal{E}} \right]. \tag{1}$$

For each class k, relation prototypes are estimated once from the training subset $\mathcal{D}_k$ and then fixed:

$$\begin{aligned} \mu_k &= \frac{1}{|\mathcal{D}_k|} \sum_{j \in \mathcal{D}_k} \psi(b_j, \mathbf{1}), \\ \sigma_k^2 &= \frac{1}{|\mathcal{D}_k|} \sum_{j \in \mathcal{D}_k} (\psi(b_j, \mathbf{1}) - \mu_k)^2 + \epsilon. \end{aligned} \tag{2}$$

Their compatibility potential is

$$\Phi_k(b, m) = -\frac{1}{2} \sum_d \left[\frac{(\psi_d(b, m) - \mu_{k,d})^2}{\sigma_{k,d}^2} + \log \sigma_{k,d}^2 \right]. \tag{3}$$

KEDP uses posterior weighted compatibility $Q_{\mathrm{rel}}(s, b, m) = \sum_{k=1}^C \bar{p}_k(s) \Phi_k(b, m)$. Structural coherence is

$$Q_{\mathrm{str}}(m) = -\frac{\lambda_s}{B^2} \sum_{p,q} A_{pq} (m_p - m_q)^2 - \lambda_b \left(\frac{1}{B} \sum_{p=1}^B m_p - \rho_0 \right)^2. \tag{4}$$

The shared knowledge energy, used during optimization and deployment, is

$$Q(s, m) = w_q Q_{\mathrm{pred}}(s, m) + w_{\mathrm{rel}} Q_{\mathrm{rel}}(s, b, m) + Q_{\mathrm{str}}(m). \tag{5}$$

Theorem 3.1 (Energy tilted diffusion policy)**.** *Let* $\mathbb{P}^{\mathrm{pre}}(\cdot \mid \mathrm{cond})$ *and* $\mathbb{P}^g(\cdot \mid \mathrm{cond})$ *denote the unguided and guided reverse diffusion path measures. With* $\mathcal{R}(x_0) = \kappa Q(s, \sigma(x_0))$*, consider* $\mathcal{J}(g) = \mathbb{E}_{\mathbb{P}^g}[\mathcal{R}(x_0)] - \mathrm{KL}(\mathbb{P}^g \| \mathbb{P}^{\mathrm{pre}})$*. Its optimal path measure satisfies*

$$\mathbb{P}^{g^*}(\tau \mid \mathrm{cond}) \propto \mathbb{P}^{\mathrm{pre}}(\tau \mid \mathrm{cond}) \exp(\kappa Q(s, \sigma(x_0))). \tag{6}$$

The proof appears in Appendix A.1. The exact middle energy is

$$E_t^\star(x_t) = \log \mathbb{E}[\exp(\kappa Q(s, m)) \mid x_t, \mathrm{cond}]. \tag{7}$$

We approximate the decoded mask posterior locally by $m \approx \mathcal{N}(\tilde{m}, \mathrm{diag}(\nu_1, \ldots, \nu_B))$. With $g_j = \partial Q(s,m)/\partial m_j|_{m=\tilde{m}}$, its analytic surrogate is

$$\widehat{E}_t(x_t) \approx \kappa Q(s, \tilde{m}) + \frac{\kappa^2}{2} \sum_{j=1}^{B} \nu_j g_j^2. \tag{8}$$

To amortize this target, $E_\Theta = E_\Theta(x_t, t, s, \bar{p}; \mathcal{K})$ matches both its value and local score field:

$$\mathcal{L}_E = \mathbb{E}_t \big[(E_\Theta - \mathrm{sg}[\widehat{E}_t])^2 + \|\nabla_{x_t} E_\Theta - \mathrm{sg}[\nabla_{x_t} \widehat{E}_t]\|_2^2 \big]. \tag{9}$$

The resulting $\nabla_{x_t} E_\Theta$ guides reverse diffusion. It is an amortized local surrogate of $E_t^\star$, not an exact realization of the optimal drift.

Relation Consistent Group Relative Policy Optimization RCGRPO reuses Φ_k and the fixed prototypes. For observation x_j with label y_j, let $\ell_k(x_j, m_t^{(i)}) = \Phi_k(b_j, m_t^{(i)})$ and $\rho_{y_j} = \ell_{y_j} - \max_{k \neq y_j} \ell_k$, evaluated at $(x_j, m_t^{(i)})$. Its relation feedback is $r_{\mathrm{rel},t}^{(i)} = \beta_{\mathrm{rel}} \tanh(\rho_{y_j}/\tau)$. Given trajectory logits $o_t^{(i)} = \mathrm{Head}_\omega(h_t^{(i)})$, define the decision margin $\mathcal{M}_t^{(i)} = o_{t,y_j}^{(i)} - \max_{k \neq y_j} o_{t,k}^{(i)}$ and $\Delta\mathcal{M}_t^{(i)} = \mathcal{M}_t^{(i)} - \mathcal{M}_{t-1}^{(i)}$. For fixed horizon $T_{\max}$, the reward and return are

$$r_t^{(i)} = \Delta\mathcal{M}_t^{(i)} + r_{\mathrm{rel},t}^{(i)} - \lambda_{\mathrm{div}} s_t^{(i)}, \quad R^{(i)} = \sum_{t=1}^{T_{\max}} r_t^{(i)}. \tag{10}$$

Here $s_t^{(i)}$ penalizes repeated selection. The margin increments sum to the final margin change, so no additional terminal reward is needed.

For G trajectories from the same state, let $\bar{R} = G^{-1} \sum_i R^{(i)}$. Their sample standard deviation is $\sigma_R = [(G-1)^{-1} \sum_i (R^{(i)} - \bar{R})^2]^{1/2}$. The group relative advantage is

$$A^{(i)} = \frac{R^{(i)} - \bar{R}}{\mathrm{sg}(\sigma_R) + \epsilon}. \tag{11}$$

Theorem 3.2 (Full group centering of trajectory returns)**.** *Let $G > 1$ trajectories be sampled independently under the same policy and context. With $S^{(i)} = \sum_t \nabla_\eta \log \pi_\eta(m_t^{(i)} \mid s_t^{(i)})$ and $J(\eta) = \mathbb{E}_{\tau \sim \pi_\eta}[R(\tau)]$,*

$$\mathbb{E}\left[\frac{1}{G} \sum_{i=1}^{G} S^{(i)} (R^{(i)} - \bar{R}) \right] = \frac{G-1}{G} \nabla_\eta J(\eta). \tag{12}$$

Full group centering thus preserves the expected score function gradient direction up to a positive constant. The standard deviation in Eq. (11) *only normalizes the update scale and is not part of this expectation claim.*

The proof appears in Appendix A.2. KERD first trains the contextual encoder and prediction head, then fixes relation prototypes estimated from the training partition. RCGRPO trains the diffusion policy and sequential state updater, while Eq. (9) independently trains E_Θ. At inference, KEDP obtains $\bar{p}$ from the complete observation and uses $\nabla_{x_t} E_\Theta$ at each reverse step with the same sequential state updates. No trajectory reward is computed at deployment. The final mask after $T_{\max}$ steps defines the predictive representation.

4 Experiments

We conduct a systematic evaluation of KERD under source, medium, and acquisition heterogeneity. The study includes baseline comparisons, component ablations, robustness tests, evaluation across heterogeneous datasets, and structured relation analysis.

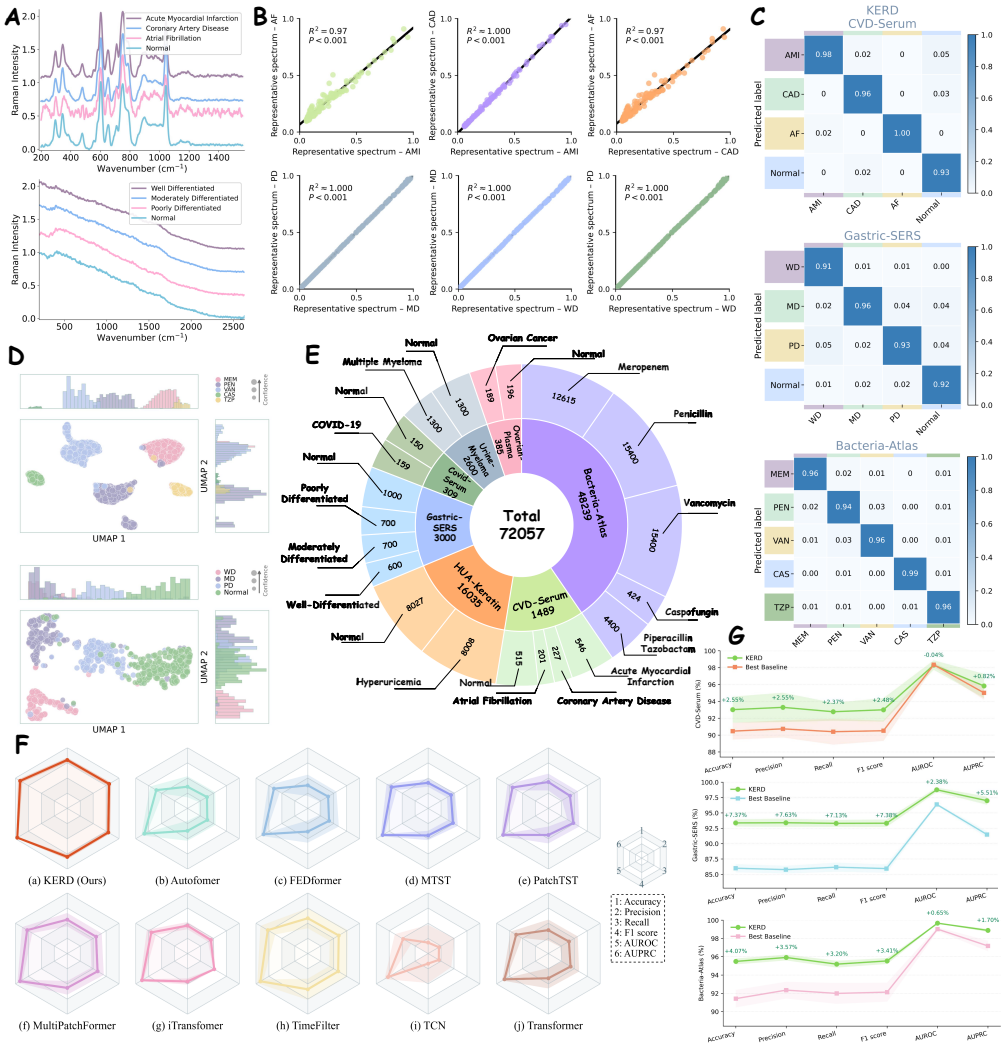


Figure 2: **Representation evaluation under heterogeneous observation conditions.** (A–B) Global profile similarity and local variation. (C–D) Prediction structure and latent-space organization. (E) Dataset composition. (F–G) Overall performance and metric-wise improvements of KERD over representative baselines across the primary evaluation tasks.

4.1 Implementation Details

Dataset Description Fig. 2E presents seven heterogeneous datasets comprising 72,057 molecular spectral observations from clinical and microbial sources. Entity level 80%/10%/10% splits separate subjects or microbial isolates across data partitions. Each dataset is trained and evaluated independently under this split. Accordingly, the multi dataset study evaluates consistency under heterogeneous conditions rather than transfer to an unseen source. More details are provided in Appendix C.

Baselines and Evaluation To evaluate structured relation learning, we compare KERD with nine general baselines: TimeFilter, MultiPatchFormer, iTransformer, MTST, FEDformer, PatchTST, Autoformer, TCN, and Transformer (Hu et al., 2025b; Naghashi et al., 2025; Liu et al., 2023b; Zhang et al., 2024c; Zhou et al., 2022a; Nie, 2022; Wu et al., 2021; Bai, 2018; Vaswani et al., 2017). They represent different mechanisms for modeling local, global, and variable dependencies. Detailed implementation settings, evaluation metrics, and hyperparameters are reported in Appendix D. All en-

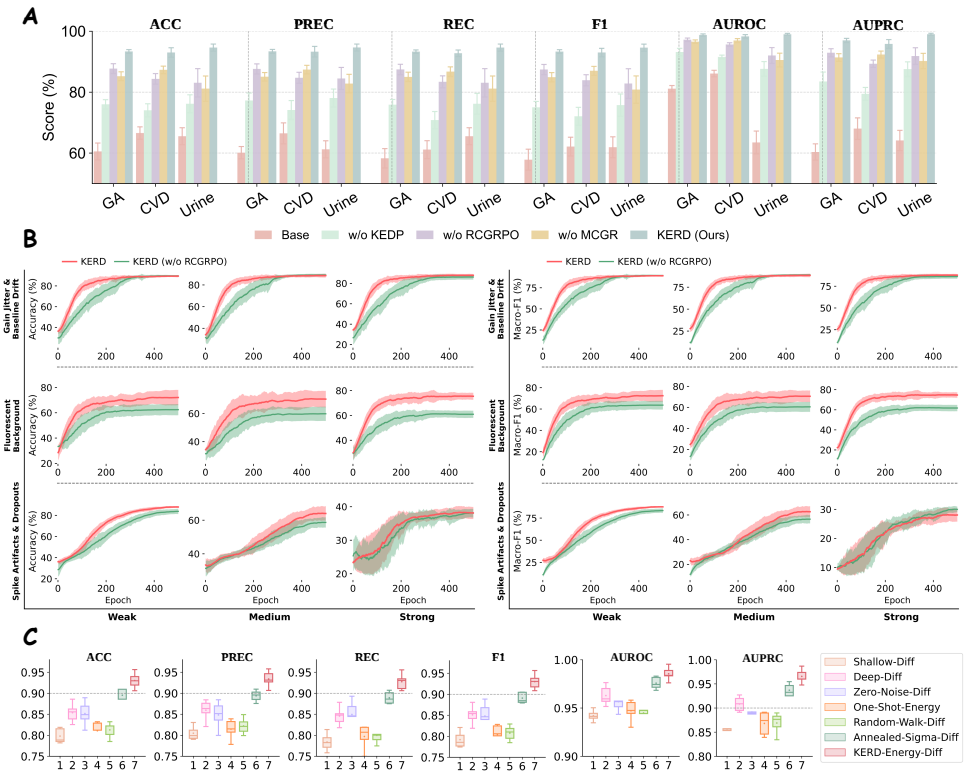


Figure 3: **Deep ablation analysis.** (A) Controlled variants quantify the roles of graph structured state construction, prototype aligned trajectory learning, and energy guided band group search. (B) Perturbation training curves show policy behavior under baseline drift, fluorescence background, and spike artifacts. (C) Diffusion strategy comparisons evaluate how noise design, sampling depth, and energy update frequency shape structured band group exploration.

ergy weights and reward coefficients are fixed across datasets and are used only to balance objective scales. We additionally compare six domain specific methods: DSCF, Objectosphere, DSCARNet, Raman Transformer, Raman-OSDL, and SANet (Xue et al., 2025; Balytskyi et al., 2024; Lu et al., 2024; Liu et al., 2023a; Zhu et al., 2025; Deng et al., 2021). These comparisons examine whether specialized feature learning can capture the dependencies targeted by KERD.

4.2 Structured Representation Evaluation under Heterogeneity

We investigate whether structural priors can expose stable dependencies that are otherwise obscured in high dimensional heterogeneous observations. In this section, we evaluate representation quality across datasets with different sources, media, acquisition settings, and label structures, providing a controlled test under substantial heterogeneity.

Fig. 2A illustrates the spectral profiles of different label groups, revealing similar global patterns alongside variations in local spectral responses. The representative comparisons in Fig. 2B further show high overall similarity. Consequently, the main challenge is not the extraction of large amplitude features, but the identification of weak structural deviations that are distributed across multiple local regions. Instead, informative differences may arise from coordinated variations across local regions, motivating the explicit modeling of group level structural dependencies.

The results in Fig. 2C–D support this interpretation. The confusion matrices show stable output structure, and the embeddings exhibit compact local organization with ambiguous observations concentrated near transition regions. The radar and gain plots in Fig. 2F–G indicate that KERD improves performance across complementary criteria, including accuracy and ranking quality. KERD exploits

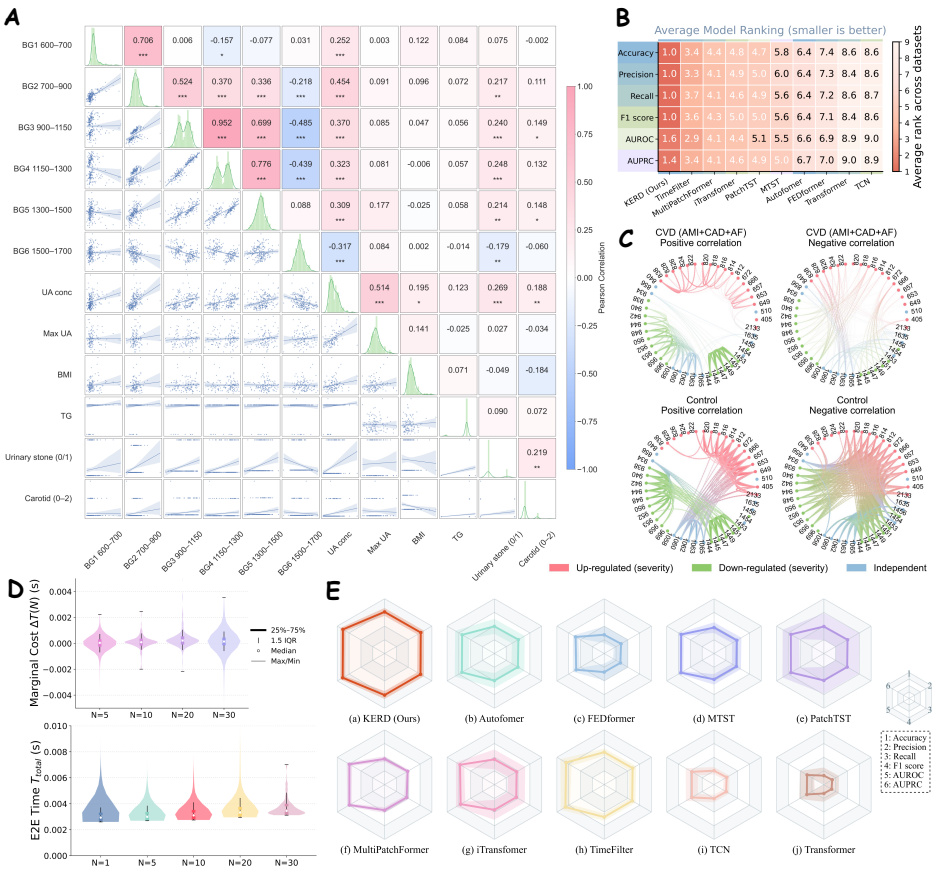


Figure 4: **Structured relation and multi source evaluation.** (A) Pairwise associations among spectral groups and auxiliary variables. (B) Average model ranks. (C) Signed dependency networks under different observation conditions. (D) Marginal inference cost and decision stage latency. (E) Comparison across additional heterogeneous source tasks.

coordinated deviations across multiple local regions instead of relying on isolated high-amplitude features. This mechanism provides a stable basis for representation learning under heterogeneity.

Complete comparisons covering all six evaluation metrics are provided in Appendix E.1. Table 2 reports the comparison with representative sequence models covering patch aggregation, multiscale encoding, frequency decomposition, and cross variable interaction. KERD obtains the best or tied best results on most metrics. While existing methods provide effective representations under their respective assumptions, they often encode local, global, or relational information through separate mechanisms. KERD instead uses a shared relational structure to guide evidence exploration and policy learning, encouraging coordinated dependencies among local observation groups.

4.3 Ablation Analysis

We examine how relational structure influences representation construction, exploration, and policy optimization. Fig. 3A shows that removing individual components reduces performance, indicating that structural modeling and guided exploration jointly support the learning of informative group dependencies. Fig. 3B indicates that relation consistent policy feedback improves stability under increasing observation perturbations. Fig. 3C further shows that successive energy guidance improves structured exploration compared with alternative diffusion strategies, supporting the learning of informative group dependencies. These observations suggest that relational priors integrated into energy guided exploration help maintain informative group dependencies during representation learning and policy optimization. Detailed ablation analyses are provided in Appendix D.1.

4.4 Multi Source Evaluation under Heterogeneity

We extend the evaluation to HUA Keratin, Urine Myeloma, Covid Serum, and Ovarian Plasma, which provide heterogeneous sources, media, and acquisition conditions. The average rank heatmap in Fig. 4B summarizes the overall performance across metrics. Detailed per-dataset results are reported in Appendix E. These results indicates that continuous energy guidance and relation consistent optimization favor group interactions that remain stable under heterogeneous conditions.

We further assess aggregate performance across seven cohorts using Friedman tests with Holm correction. Significant differences are observed for all six metrics, with KERD obtaining the lowest average rank in each case. However, pairwise differences against some strong baselines are not uniformly significant. Detailed statistical analyses are provided in Appendix E.

We additionally study efficiency and metric balance. In Fig. 4D, the marginal latency changes only slightly as the number of aggregated spectra increases from one to thirty. The reported latency covers only online inference. In deployment, encoded spectral features can be cached, so subsequent requests update only the cached state and decision module. Fig. 4E further demonstrates that KERD achieves consistently strong performance across multiple evaluation metrics on the additional heterogeneous cohorts. These findings provide further empirical evidence of KERD's performance and inference efficiency under the evaluated heterogeneous conditions.

4.5 Structured Relation Analysis

We investigate whether observable dependencies support the group level priors used by KERD. Fig. 4A analyzes six spectral groups spanning 600–1700 cm^{-1} and six auxiliary cohort variables through Pearson correlations, significance tests, marginal distributions, and pairwise regressions.

The strong positive correlations between adjacent spectral groups and negative associations across other regions reveal nonuniform relational patterns within the observations. Additional associations with auxiliary variables suggest that spectral and contextual information may provide complementary evidence. These findings motivate KERD's use of structural priors to organize local group interactions and relative intensity relationships. Such relational modeling provides a basis for exploring coordinated group configurations rather than treating spectral regions independently.

Fig. 4C examines how these dependencies vary across observation conditions. After SNV normalization and removal of a global brightness factor, we compute within condition Spearman correlations. Representative nodes are selected using group-wise tests with Benjamini Hochberg correction, while signed chord graphs visualize positive and negative dependencies. The resulting networks contain hubs and coherent subnetworks, but their connectivity varies across conditions. Some group relations persist, while others weaken or reorganize. These findings suggest that heterogeneity affects relational topology as well as individual feature magnitudes, motivating structural priors for learning stable group level dependencies from high dimensional heterogeneous observations.

4.6 Additional Experimental Evaluations

We conduct additional experiments to examine structural preservation, auxiliary context integration, and comparisons with domain specific methods. Appendix F evaluates spectral structure preservation under controlled observation perturbations. Appendix G examines whether structured auxiliary information improves representation learning. Appendix H compares KERD with domain specific methods to assess its performance against specialized representation architectures.

5 Conclusion

We introduce KERD to investigate how structural priors can support the learning of stable group dependencies from heterogeneous observations. KERD models dependencies among spectral peak structures in heterogeneous data and enables more robust policy optimization for relation driven evidence learning. Experiments on 72,057 Raman spectra from multicenter cohorts demonstrate strong performance against competitive baselines across the evaluated tasks. Future work will investigate causal and uncertainty aware extensions under stronger distribution shifts.

AI USE STATEMENT

Generative AI tools were used to assist with translation, language editing, and refinement of experimental result descriptions.

ETHICS STATEMENT

This multicenter study complied with relevant ethical standards and institutional data protection requirements throughout data collection and analysis. All clinical measurements were deidentified before analysis to safeguard participant privacy. The handling of clinical information followed applicable institutional policies to ensure responsible data use and confidentiality.

REPRODUCIBILITY STATEMENT

To support reproducibility, we document the proposed methodology, theoretical derivations, experimental configurations, and evaluation procedures in the manuscript and appendix. Additional experimental results and detailed implementation settings are provided to support verification of the reported findings. Due to institutional access policies and privacy requirements, the data and source code will be available upon approved request following publication.

REFERENCES

Marvin Alberts, Oliver Schilter, Federico Zipoli, Nina Hartrampf, and Teodoro Laino. Unraveling molecular structure: A multimodal spectroscopic dataset for chemistry. *Advances in Neural Information Processing Systems*, 37:125780–125808, 2024.

Shaojie Bai. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.

Yaroslav Balytskyi, Nataliia Kalashnyk, Inna Hubenko, Alina Balytska, and Kelly McNear. Enhancing open-world bacterial raman spectra identification by feature regularization for improved resilience against unknown classes. *Chemical & Biomedical Imaging*, 2(6):442, 2024.

Roman Bushuiev, Anton Bushuiev, Niek F de Jonge, Adamo Young, Fleming Kretschmer, Raman Samusevich, Janne Heirman, Fei Wang, Luke Zhang, Kai Dührkop, et al. Massspecgym: A benchmark for the discovery and identification of molecules. *Advances in Neural Information Processing Systems*, 37:110010–110027, 2024.

Lin Deng, Yuzhong Zhong, Maoning Wang, Xiujuan Zheng, and Jianwei Zhang. Scale-adaptive deep model for bacterial raman spectra identification. *IEEE Journal of Biomedical and Health Informatics*, 26(1):369–378, 2021.

Yihe Deng, Yu Yang, Baharan Mirzasoleiman, and Quanquan Gu. Robust learning with progressive data expansion against spurious correlation. *Advances in neural information processing systems*, 36:1390–1402, 2023.

Wenzheng Feng, Yuxiao Dong, Shaosheng Cao, Xu Cheng, and Jie Tang. Heterogeneous graph random neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 2026.

Dimitar Georgiev, Álvaro Fernández-Galiana, Simon Vilms Pedersen, Georgios Papadopoulos, Ruoxiao Xie, Molly M Stevens, and Mauricio Barahona. Hyperspectral unmixing for raman spectroscopy via physics-constrained autoencoders. *Proceedings of the National Academy of Sciences*, 121(45):e2407439121, 2024.

Jie Gui, Tuo Chen, Jing Zhang, Qiong Cao, Zhenan Sun, Hao Luo, and Dacheng Tao. A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9052–9071, 2024.

Yan-Lin He, Zi-Yang Lu, Yuan Xu, Qun-Xiong Zhu, and Xin Pan. A knowledge-data integrated graph convolutional network for fault diagnosis in industrial processes. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2026.

Chi-Sing Ho, Neal Jean, Catherine A Hogan, Lena Blackmon, Stefanie S Jeffrey, Mark Holodniy, Niaz Banaei, Amr AE Saleh, Stefano Ermon, and Jennifer Dionne. Rapid identification of pathogenic bacteria using raman spectroscopy and deep learning. *Nature communications*, 10(1): 4927, 2019.

Jifeng Hu, Sili Huang, Zhejian Yang, Shengchao Hu, Li Shen, Hechang Chen, Lichao Sun, Yi Chang, and Dacheng Tao. Analytic energy-guided policy optimization for offline reinforcement learning. *arXiv preprint arXiv:2505.01822*, 2025a.

Yifan Hu, Guibin Zhang, Peiyuan Liu, Disen Lan, Naiqi Li, Dawei Cheng, Tao Dai, Shu-Tao Xia, and Shirui Pan. Timefilter: Patch-specific spatial-temporal graph filtration for time series forecasting. *arXiv preprint arXiv:2501.13041*, 2025b.

Chenqing Hua, Bozitao Zhong, Sitao Luan, Liang Hong, Guy Wolf, Doina Precup, and Shuangjia Zheng. Reactzyme: A benchmark for enzyme-reaction prediction. *Advances in Neural Information Processing Systems*, 37:26415–26442, 2024.

Liping Huang, Hongwei Sun, Liangbin Sun, Keqing Shi, Yuzhe Chen, Xueqian Ren, Yuancai Ge, Danfeng Jiang, Xiaohu Liu, Wolfgang Knoll, et al. Rapid, label-free histopathological diagnosis of liver cancer based on raman spectroscopy and deep learning. *Nature Communications*, 14(1): 48, 2023.

Deeptha Ishwar, Rupa Haldavnekar, Krishnan Venkatakrishnan, and Bo Tan. Minimally invasive detection of cancer using metabolic changes in tumor-associated natural killer cells with oncoimmune probes. *Nature communications*, 13(1):4527, 2022.

Cheng Jiang, Asadur Chowdury, Xinhai Hou, Akhil Kondepudi, Christian Freudiger, Kyle Conway, Sandra Camelo-Piragua, Daniel Orringer, Honglak Lee, and Todd Hollon. Opensrh: optimizing brain tumor surgery using intraoperative stimulated raman histology. *Advances in neural information processing systems*, 35:28502–28516, 2022.

Marlous Kamp, Jakub Surmacki, Marc Segarra Mondejar, Tim Young, Karolina Chrabaszcz, Fadwa Joud, Vincent Zecchini, Alyson Speed, Christian Frezza, and Sarah E Bohndiek. Raman microspectroscopy reveals the spatial distribution of fumarate in cells and tissues. *Nature Communications*, 15(1):5386, 2024.

Abhinav Kumar, Amit Deshpande, and Amit Sharma. Causal effect regularization: Automated detection and removal of spurious correlations. *Advances in Neural Information Processing Systems*, 36:20942–20984, 2023.

Bo Li, Mikkel N Schmidt, and Tommy S Alstrøm. Raman spectrum matching with contrastive representation learning. *Analyst*, 147(10):2238–2246, 2022.

Honglin Li, Zhongyi Shui, Yunlong Zhang, Chenglu Zhu, and Lin Yang. Pathvq: Reforming computational pathology foundation model for whole slide image analysis via vector quantization. *arXiv preprint arXiv:2503.06482*, 2025a.

Jinning Li, Ruipeng Han, Jingying Zeng, Dachun Sun, Chenkai Sun, Hanghang Tong, ChengXiang Zhai, Boleslaw K Szymanski, and Tarek Abdelzaher. Learning to slice: Self-supervised interpretable hierarchical representation learning with graph auto-encoder tree. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 1388–1399, 2025b.

Yuetao Li, Jiabao Xu, Xiaofei Yi, Xiaobo Li, Yanjun Luo, Andrew Glidle, Phil Summersgill, Simon Allen, Tim Ryan, Xiaochen Liu, et al. Rapid culture-free diagnosis of clinical pathogens via integrated microfluidic-raman micro-spectroscopy. *Nature Communications*, 2025c.

Zhengxin Li, Feiping Nie, Danyang Wu, Zheng Wang, and Xuelong Li. Sparse trace ratio lda for supervised feature selection. *IEEE transactions on cybernetics*, 54(4):2420–2433, 2023.

Bo Liu, Kunxiang Liu, Xiaoqing Qi, Weijia Zhang, and Bei Li. Classification of deep-sea cold seep bacteria by transformer combined with raman spectroscopy. *Scientific Reports*, 13(1):3240, 2023a.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. *arXiv preprint arXiv:2310.06625*, 2023b.

Songlin Lu, Yuanfang Huang, Wan Xiang Shen, Yu Lin Cao, Mengna Cai, Yan Chen, Ying Tan, Yu Yang Jiang, and Yu Zong Chen. Raman spectroscopic deep learning with signal aggregated representations for enhanced cell phenotype and signature identification. *PNAS nexus*, 3 (8):pgae268, 2024.

Jing Meng, Zhenfang Zhu, Qiang Lu, Hongli Pei, Kefeng Li, Yuzhi Ren, Meng Li, Xiaowen Sun, and Dawei Zhao. A prompt-driven framework for compensation and fusion in multimodal sentiment analysis with missing modalities. *Knowledge-Based Systems*, pp. 115438, 2026.

Vahid Naghashi, Mounir Boukadoum, and Abdoulaye Banire Diallo. A multiscale model for multivariate time series forecasting. *Scientific Reports*, 15(1):1565, 2025.

Y Nie. A time series is worth 64words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

Maria Paraskevaidi, Katherine M Ashton, Helen F Stringfellow, Nicholas J Wood, Patrick J Keating, Anthony W Rowbottom, Pierre L Martin-Hirsch, and Francis L Martin. Raman spectroscopic techniques to detect ovarian cancer biomarkers in blood plasma. *Talanta*, 189:281–288, 2018.

Quan Qian, Jun Luo, and Yi Qin. Heterogeneous federated domain generalization network with common representation learning for cross-load machinery fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(9):5704–5716, 2024.

Pengju Ren, Ri-gui Zhou, and Yaochong Li. A self-supervised learning method for raman spectroscopy based on masked autoencoders. *Expert Systems with Applications*, pp. 128576, 2025.

Xinyue Shao, Hongzhi Wang, Xiang Chen, Xiao Zhu, and Yan Zhang. Cube: causal intervention-based counterfactual explanation for prediction models. *IEEE Transactions on Knowledge and Data Engineering*, 36(6):2416–2429, 2023.

Zezhi Shao, Fei Wang, Yongjun Xu, Wei Wei, Chengqing Yu, Zhao Zhang, Di Yao, Tao Sun, Guangyin Jin, Xin Cao, et al. Exploring progress in multivariate time series forecasting: Comprehensive benchmarking and heterogeneity analysis. *IEEE Transactions on Knowledge and Data Engineering*, 37(1):291–305, 2024.

Yong Shi, Pei Quan, Yang Xiao, Minglong Lei, and Lingfeng Niu. Graph influence network. *IEEE transactions on cybernetics*, 53(10):6146–6159, 2022.

Manuel Sigle, Anne-Katrin Rohlfing, Martin Kenny, Sophia Scheuermann, Na Sun, Ulla Graeßner, Verena Haug, Jessica Sudmann, Christian M Seitz, David Heinzmann, et al. Translating genomic tools to raman spectroscopy analysis enables high-dimensional tissue characterization on molecular resolution. *Nature Communications*, 14(1):5799, 2023.

Yuxuan Sun, Yixuan Si, Chenglu Zhu, Xuan Gong, Kai Zhang, Pingyi Chen, Ye Zhang, Zhongyi Shui, Tao Lin, and Lin Yang. Cpath-omni: A unified multimodal foundation model for patch and whole slide image analysis in computational pathology. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 10360–10371, 2025.

Wenhao Tang, Fengtao Zhou, Sheng Huang, Xiang Zhu, Yi Zhang, and Bo Liu. Feature re-embedding: Towards foundation model-level performance in computational pathology. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11343–11352, 2024.

Qing Tian, Heyang Sun, Chuang Ma, Meng Cao, Yi Chu, and Songcan Chen. Heterogeneous domain adaptation with structure and classification space alignment. *IEEE Transactions on Cybernetics*, 52(10):10328–10338, 2021.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Changyue Wang, Changzhong Wang, Shuang An, Weiping Ding, and Yuhua Qian. Feature selection and classification based on directed fuzzy rough sets. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 55(1):699–711, 2024.

Haotian Wang, Kun Kuang, Long Lan, Zige Wang, Wanrong Huang, Fei Wu, and Wenjing Yang. Out-of-distribution generalization with causal feature separation. *IEEE Transactions on Knowledge and Data Engineering*, 36(4):1758–1772, 2023.

Liang Wang, Shaozhen Liu, Yu Rong, Deli Zhao, Qiang Liu, and Shu Wu. Molspectra: Pre-training 3d molecular representation with multi-modal energy spectra. *arXiv preprint arXiv:2502.16284*, 2025.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.

Shirley Wu, Mert Yuksekgonul, Linjun Zhang, and James Zou. Discover and cure: Concept-aware mitigation of spurious correlation. In *International Conference on Machine Learning*, pp. 37765–37786. PMLR, 2023.

Bingsen Xue, Xinyuan Bi, Zheyi Dong, Yunzhe Xu, Minghui Liang, Xin Fang, Yizhe Yuan, Ruoxi Wang, Shuyu Liu, Rushi Jiao, et al. Deep spectral component filtering as a foundation model for spectral analysis demonstrated in metabolic profiling. *Nature Machine Intelligence*, 7(5): 743–757, 2025.

Qisen Yang, Huanqian Wang, Mukun Tong, Wenjie Shi, Gao Huang, and Shiji Song. Leveraging reward consistency for interpretable feature discovery in reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(2):1014–1025, 2023.

Wenqian Ye, Guangtao Zheng, and Aidong Zhang. Improving group robustness on spurious correlation via evidential alignment. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 3610–3621, 2025.

Gang Yin, Lintao Li, Shun Lu, Yu Yin, Yuanzhang Su, Yilan Zeng, Mei Luo, Maohua Ma, Hongyan Zhou, Lucia Orlandini, et al. An efficient primary screening of covid-19 by serum raman spectroscopy. *Journal of Raman Spectroscopy*, 52(5):949–958, 2021.

Chengqing Yu, Fei Wang, Chuanguang Yang, Zezhi Shao, Tao Sun, Tangwen Qian, Wei Wei, Zhulin An, and Yongjun Xu. Merlin: Multi-view representation learning for robust multivariate time series forecasting with unfixed missing rates. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 3633–3644, 2025.

Kaitlin Zareno, Jarett Dewbury, Siamak K Sorooshyari, Hossein Mobahi, and Loza F Tadesse. Sharpness-aware minimization (sam) improves classification accuracy of bacterial raman spectral data enabling portable diagnostics. In *Optics and Biophotonics in Low-Resource Settings XI*, volume 13307, pp. 65–71. SPIE, 2025.

Jian Zhang, Lei Qi, Yinghuan Shi, and Yang Gao. Exploring flat minima for domain generalization with large learning rates. *IEEE Transactions on Knowledge and Data Engineering*, 36(11):6145–6158, 2024a.

Kexin Zhang, Qingsong Wen, Chaoli Zhang, Rongyao Cai, Ming Jin, Yong Liu, James Y Zhang, Yuxuan Liang, Guansong Pang, Dongjin Song, et al. Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. *IEEE transactions on pattern analysis and machine intelligence*, 46(10):6775–6794, 2024b.

Yitian Zhang, Liheng Ma, Soumyasundar Pal, Yingxue Zhang, and Mark Coates. Multi-resolution time-series transformer for long-term forecasting. In *International conference on artificial intelligence and statistics*, pp. 4222–4230. PMLR, 2024c.

Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pp. 27268–27286. PMLR, 2022a.

Zhichao Zhou, Yu Hu, Yue Zhang, Jiazhou Chen, and Hongmin Cai. Multiview deep graph infomax to achieve unsupervised graph embedding. *IEEE Transactions on Cybernetics*, 53(10):6329–6339, 2022b.

Longji Zhu, Yunan Yang, Fei Xu, Xinyu Lu, Mingrui Shuai, Zhulin An, Xiaomeng Chen, Hu Li, Francis L Martin, Peter J Vikesland, et al. Open-set deep learning–enabled single-cell raman spectroscopy for rapid identification of airborne pathogens in real-world environments. *Science Advances*, 11(2):eadp7991, 2025.

A Supplementary Theoretical Analysis

A.1 KEDP: Energy-Guided Diffusion as KL-Regularized Optimal Distribution

Compact notation Let $p_\omega(k \mid s, m)$ denote the predictive posterior under mask m, and define the complete observation reference posterior as

$$\bar{p}_k(s) = \mathrm{sg}\left[p_\omega(k \mid s, \mathbf{1})\right].$$

Let $\psi(b, m)$ denote the shared relation representation. For each class k, let (μ_k, σ_k^2) be estimated from the training partition using the complete group mask. We define the prototype compatibility

$$\Phi_k(b, m) = -\frac{1}{2} \sum_d \left[\frac{(\psi_d(b, m) - \mu_{k,d})^2}{\sigma_{k,d}^2} + \log \sigma_{k,d}^2 \right]. \tag{13}$$

The deployment consistent energy is

$$Q(s, m) = w_q Q_{\mathrm{pred}}(s, m) + w_{\mathrm{rel}} \sum_{k=1}^{C} \bar{p}_k(s) \Phi_k(b, m) + Q_{\mathrm{str}}(m). \tag{14}$$

The fixed Peak Graph provides only a local structural prior. Its contribution differs from the learned relation potential. The smoothness term constrains neighboring mask values, whereas the relation potential measures the compatibility between sample responses and class conditional prototypes. Both KEDP and RCGRPO reuse this relation potential for diffusion guidance and trajectory feedback, respectively. Thus, the relational guidance depends on the observed data and cannot be reduced to topology based smoothing alone. Our contribution lies in the shared use of spectral relation potentials across exploration and policy optimization. Specifically, KEDP uses posterior weighted prototype compatibility to guide diffusion sampling, while RCGRPO reuses the same relation potential to construct label conditioned trajectory feedback during training. Unlike topology based smoothing, this shared potential captures data dependent relations among spectral groups. It connects structured exploration with policy feedback without requiring true labels during inference.

The same analytic energy defines the target in both training and inference. A learnable model E_Θ amortizes its middle energy without introducing a separate energy objective.

For notational convenience, the appendix uses the forward control time index, where the terminal state x_T corresponds to the denoised endpoint denoted by x_0 under the reverse time convention in the main text.

Assumption (needed for Girsanov) We assume a Novikov-type condition so the path measures induced by the two SDEs below are mutually absolutely continuous.

A.1.1 KL between diffusion path measures (Girsanov chain)

Unguided (prior) diffusion:

$$\mathrm{d}x_t = f^{\mathrm{pre}}(t, x_t, \mathrm{cond})\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w_t, \qquad x_0 \sim \rho_0. \tag{15}$$

Guided diffusion:

$$\mathrm{d}x_t = g(t, x_t, \mathrm{cond})\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w_t, \qquad x_0 \sim \rho_0. \tag{16}$$

Let $\mathbb{P}^{\mathrm{pre}}(\cdot \mid \mathrm{cond})$ and $\mathbb{P}^{g}(\cdot \mid \mathrm{cond})$ denote the induced path measures of $x_{0:T}$.

Lemma A.1 (KL of controlled diffusion path measures)**.** *Under the Novikov-type condition,*

$$\mathrm{KL}(\mathbb{P}^g_{\mathrm{cond}} \| \mathbb{P}^{\mathrm{pre}}_{\mathrm{cond}}) = \mathbb{E}_{\tau \sim \mathbb{P}^g_{\mathrm{cond}}} \left[\int_0^T \frac{\|\Delta_t\|_2^2}{2\sigma^2(t)}\,\mathrm{d}t \right], \tag{17}$$

Proof. Define the normalized drift gap

$$u_t := \frac{g(t, x_t, \mathrm{cond}) - f^{\mathrm{pre}}(t, x_t, \mathrm{cond})}{\sigma(t)}. \tag{18}$$

By *Girsanov's theorem*,

$$\log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\mathbb{P}^{\mathrm{pre}}} = \int_0^T u_t^\top \, \mathrm{d}w_t^{\mathrm{pre}} - \frac{1}{2}\int_0^T \|u_t\|_2^2 \, \mathrm{d}t, \tag{19}$$

where w^{pre} is Brownian motion under $\mathbb{P}^{\mathrm{pre}}$. Under $\mathbb{P}^g$ define $w_t^g := w_t^{\mathrm{pre}} - \int_0^t u_s \, \mathrm{d}s$, then w^g is Brownian motion under $\mathbb{P}^g$ (Girsanov). Hence

$$\begin{aligned} \log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\mathbb{P}^{\mathrm{pre}}} &= \int_0^T u_t^\top \left(\mathrm{d}w_t^g + u_t \, \mathrm{d}t\right) - \frac{1}{2}\int_0^T \|u_t\|_2^2 \, \mathrm{d}t \\ &= \int_0^T u_t^\top \, \mathrm{d}w_t^g + \frac{1}{2}\int_0^T \|u_t\|_2^2 \, \mathrm{d}t. \end{aligned} \tag{20}$$

Therefore,

$$\begin{aligned} \mathrm{KL}(\mathbb{P}^g \| \mathbb{P}^{\mathrm{pre}}) &= \mathbb{E}_{\mathbb{P}^g}\left[\log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\mathbb{P}^{\mathrm{pre}}}\right] \\ &= \mathbb{E}_{\mathbb{P}^g}\left[\int_0^T u_t^\top \, \mathrm{d}w_t^g + \frac{1}{2}\int_0^T \|u_t\|_2^2 \, \mathrm{d}t\right] \\ &= \mathbb{E}_{\mathbb{P}^g}\left[\frac{1}{2}\int_0^T \|u_t\|_2^2 \, \mathrm{d}t\right] \\ &= \mathbb{E}_{\mathbb{P}^g}\left[\int_0^T \frac{\|\Delta_t\|_2^2}{2\sigma^2(t)} \, \mathrm{d}t\right], \end{aligned} \tag{21}$$

$$\Delta_t := g(t, x_t, \mathrm{cond}) - f^{\mathrm{pre}}(t, x_t, \mathrm{cond}).$$

where the third line uses that the Itô integral is a martingale with zero expectation. This equals Eq. (17). □

A.1.2 KL-REGULARIZED OBJECTIVE $\Longleftrightarrow$ BOLTZMANN-TILTED OPTIMAL PATH MEASURE

We apply terminal energy via $m_T = \sigma(x_T)$ and define the terminal utility

$$\mathcal{R}(x_T) := \kappa \, Q\big(s, \sigma(x_T)\big). \tag{22}$$

Define the log-partition function

$$\log Z(\mathrm{cond}) := \log \, \mathbb{E}_{x_{0:T} \sim \mathbb{P}^{\mathrm{pre}}(\cdot|\mathrm{cond})}\big[\exp(\mathcal{R}(x_T))\big]. \tag{23}$$

Theorem A.2 (KL-regularized energy guidance yields a Boltzmann-tilted path distribution)**.** *Consider the KL-regularized objective*

$$\mathrm{obj}(g) := \mathbb{E}_{\mathbb{P}^g}\big[\mathcal{R}(x_T)\big] - \mathrm{KL}(\mathbb{P}^g(\cdot \mid \mathrm{cond}) \,\|\, \mathbb{P}^{\mathrm{pre}}(\cdot \mid \mathrm{cond}))\,. \tag{24}$$

Define the tilted measure on trajectories $\tau = x_{0:T}$:

$$\widetilde{\mathbb{P}}(\tau \mid \mathrm{cond}) := \frac{\mathbb{P}^{\mathrm{pre}}(\tau \mid \mathrm{cond}) \, \exp(\mathcal{R}(x_T))}{Z(\mathrm{cond})}. \tag{25}$$

Then

$$\mathrm{obj}(g) = \log Z(\mathrm{cond}) - \mathrm{KL}\Big(\mathbb{P}^g(\cdot \mid \mathrm{cond}) \,\Big\|\, \widetilde{\mathbb{P}}(\cdot \mid \mathrm{cond})\Big)\,. \tag{26}$$

Consequently, $\sup_g \mathrm{obj}(g) = \log Z(\mathrm{cond})$ *and any optimizer satisfies*

$$\mathbb{P}^{g^*}(\cdot \mid \mathrm{cond}) = \widetilde{\mathbb{P}}(\cdot \mid \mathrm{cond}) \Longleftrightarrow \; \mathbb{P}^{g^*}(\tau \mid \mathrm{cond}) \propto \mathbb{P}^{\mathrm{pre}}(\tau \mid \mathrm{cond}) \cdot \exp\big(\kappa Q(s, \sigma(x_T))\big). \tag{27}$$

Proof. Write the objective as a chain of equalities:

$$\begin{aligned}
\mathrm{obj}(g) &= \mathbb{E}_{\mathbb{P}^g}\big[\mathcal{R}(x_T)\big] - \mathbb{E}_{\mathbb{P}^g}\left[\log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\mathbb{P}^{\mathrm{pre}}}\right] \\
&= \mathbb{E}_{\mathbb{P}^g}\left[\mathcal{R}(x_T) - \log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\mathbb{P}^{\mathrm{pre}}}\right] \\
&= \mathbb{E}_{\mathbb{P}^g}\left[\log\left(\frac{\exp(\mathcal{R}(x_T))\,\mathrm{d}\mathbb{P}^{\mathrm{pre}}}{\mathrm{d}\mathbb{P}^g}\right)\right] \\
&= \mathbb{E}_{\mathbb{P}^g}\left[\log\left(\frac{Z(\mathrm{cond})\,\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\mathbb{P}^g}\right)\right] \\
&= \log Z(\mathrm{cond}) - \mathbb{E}_{\mathbb{P}^g}\left[\log \frac{\mathrm{d}\mathbb{P}^g}{\mathrm{d}\widetilde{\mathbb{P}}}\right] \\
&= \log Z(\mathrm{cond}) - \mathrm{KL}\Big(\mathbb{P}^g \,\Big\|\, \widetilde{\mathbb{P}}\Big),
\end{aligned} \tag{28}$$

which is Eq. (26). Nonnegativity of KL implies $\mathrm{obj}(g) \le \log Z(\mathrm{cond})$ with equality iff $\mathbb{P}^g = \widetilde{\mathbb{P}}$, giving Eq. (27). □

A.1.3 Optimal drift form (HJB → exponential transform → Feynman–Kac)

Define the optimal value function (fixed (s, cond)):

$$\begin{aligned}
v_t^*(x) := \sup_g\ \mathbb{E}\Bigg[&\mathcal{R}(x_T) - \int_t^T \frac{\|\Delta_\tau\|_2^2}{2\sigma^2(\tau)}\,\mathrm{d}\tau \\
&\Big|\ x_t = x, \mathrm{cond}\Bigg], \\
\Delta_\tau := g(\tau, x_\tau, \mathrm{cond}) &- f^{\mathrm{pre}}(\tau, x_\tau, \mathrm{cond}).
\end{aligned} \tag{29}$$

with terminal condition $v_T^*(x) = \mathcal{R}(x)$.

Lemma A.3 (Equation and optimal drift)**.** *The optimizer satisfies*

$$g^*(t, x, \mathrm{cond}) = f^{\mathrm{pre}}(t, x, \mathrm{cond}) + \sigma^2(t)\,\nabla_x v_t^*(x). \tag{30}$$

Moreover v_t^* solves

$$\begin{aligned}
&\partial_t v_t^* + f^{\mathrm{pre}}(t, x, \mathrm{cond})^\top \nabla_x v_t^* + \frac{\sigma^2(t)}{2}\Delta_x v_t^* + \frac{\sigma^2(t)}{2}\|\nabla_x v_t^*\|_2^2 = 0, \\
&v_T^*(x) = \mathcal{R}(x).
\end{aligned} \tag{31}$$

Proof. Complete the square:

$$\begin{aligned}
g^\top \nabla v - \frac{1}{2\sigma^2}\|g - f^{\mathrm{pre}}\|_2^2 = &-\frac{1}{2\sigma^2}\Big\|g - \big(f^{\mathrm{pre}} + \sigma^2 \nabla v\big)\Big\|_2^2 \\
&+ \frac{\sigma^2}{2}\|\nabla v\|_2^2 + f^{\mathrm{pre}\top}\nabla v,
\end{aligned} \tag{32}$$

so the maximizer is Eq. (30) and substitution yields Eq. (31). □

Lemma A.4 (Exponential transform and Feynman–Kac representation)**.** *Let $\psi_t(x) := \exp(v_t^*(x))$. Then Eq.* (31) *is equivalent to*

$$\begin{aligned}
&\partial_t \psi_t(x) + f^{\mathrm{pre}}(t, x, \mathrm{cond})^\top \nabla_x \psi_t(x) \\
&\qquad + \frac{\sigma^2(t)}{2}\Delta_x \psi_t(x) = 0, \\
&\psi_T(x) = \exp(\mathcal{R}(x)).
\end{aligned} \tag{33}$$

Moreover,

$$\begin{aligned}
&\psi_t(x) = \mathbb{E}_{x_{t:T}\sim\mathbb{P}^{\mathrm{pre}}(\cdot|x_t=x,\mathrm{cond})}\big[\exp(\mathcal{R}(x_T))\big], \\
&v_t^*(x) = \log \psi_t(x).
\end{aligned} \tag{34}$$

Proof. Use $\partial_t \psi = \psi\, \partial_t v$, $\nabla \psi = \psi\, \nabla v$, $\Delta \psi = \psi(\Delta v + \|\nabla v\|_2^2)$ to obtain Eq. (33). Eq. (34) follows from the *Feynman–Kac formula* for Eq. (15). □

Corollary A.5 (Principled guidance direction)**.** *The optimal drift can be written as a drift correction*

$$g^*(t, x, \text{cond}) = f^{\text{pre}}(t, x, \text{cond}) + \sigma^2(t)\, \nabla_x \log \psi_t(x). \tag{35}$$

A.1.4 Analytic middle-energy approximation

The exact energy to go is

$$E_t^\star(x_t) = \log \psi_t(x_t) = \log \mathbb{E}\left[\exp\left(\kappa Q(s, m)\right) \mid x_t, \text{cond}\right]. \tag{36}$$

In the discrete sampler, we first approximate $E_t^\star$ locally using the posterior moments of the decoded mask. Assume at a reverse step the decoded mask uncertainty is approximately diagonal Gaussian:

$$m \approx \mathcal{N}\big(\tilde{m}, \text{diag}(\nu_1, \ldots, \nu_B)\big), \qquad \nu_j = \text{Var}(m_j). \tag{37}$$

Let $b_j := \frac{\partial Q(s,m)}{\partial m_j}\Big|_{m=\tilde{m}}$.

Lemma A.6 (Gaussian MGF)**.** *Using the first-order expansion* $Q(s, m) \approx Q(s, \tilde{m}) + \sum_{j=1}^B b_j(m_j - \tilde{m}_j)$,

$$\begin{aligned}
\log \mathbb{E}\big[\exp(\kappa Q(s, m))\big] &\approx \log \mathbb{E}\left[\exp\left(\kappa Q(s, \tilde{m}) + \kappa \sum_{j=1}^B b_j(m_j - \tilde{m}_j)\right)\right] \\
&= \kappa Q(s, \tilde{m}) + \log \mathbb{E}\left[\exp\left(\kappa \sum_{j=1}^B b_j(m_j - \tilde{m}_j)\right)\right] \\
&= \kappa Q(s, \tilde{m}) + \frac{\kappa^2}{2} \sum_{j=1}^B \nu_j\, b_j^2.
\end{aligned} \tag{38}$$

Proof. $\sum_j b_j(m_j - \tilde{m}_j)$ is Gaussian with variance $\sum_j \nu_j b_j^2$. For Gaussian Z, $\log \mathbb{E}[\exp(\kappa Z)] = \frac{\kappa^2}{2}\text{Var}(Z)$. □

Let $\widehat{E}_t$ denote the right-hand side of Eq. (38). In implementation, the amortized model

$$E_\Theta = E_\Theta(x_t, t, s, \bar{p}; \mathcal{K})$$

matches both $\widehat{E}_t$ and its local guidance field,

$$\mathcal{L}_E = \Big(E_\Theta - \text{sg}[\widehat{E}_t]\Big)^2 + \Big\|\nabla_{x_t} E_\Theta - \text{sg}[\nabla_{x_t} \widehat{E}_t]\Big\|_2^2. \tag{39}$$

Practical guidance therefore uses $\nabla_{x_t} E_\Theta$. If Q is twice differentiable and $\|\nabla_m^2 Q\|_2 \le L_Q$ in the local posterior region, the linearization residual satisfies

$$|Q(s, m) - Q_{\text{lin}}(s, m)| \le \frac{L_Q}{2}\|m - \tilde{m}\|_2^2. \tag{40}$$

Thus, the implementation follows $E_t^\star \to \widehat{E}_t \to E_\Theta$, and we do not claim exact recovery of the optimal drift outside this local regime.

A.2 Group Relative Advantage with Class Conditional Spectral Consistency

RCGRPO reuses the same relation representation ψ and training prototypes $\{(\mu_k, \sigma_k^2)\}_{k=1}^C$ used by KEDP.

A.2.1 TASK TRANSFERABLE SPECTRAL RELATION FEATURES

Let x_j be an input spectrum under the sampled band-group mask. We extract its Peak-Graph band-group feature as

$$g_j = g(x_j) \in \mathbb{R}_+^B, \tag{41}$$

where B is the number of band groups. We use a positive normalized band feature

$$\tilde{g}_j = \frac{g_j}{\mathbf{1}^\top g_j + \epsilon_0}, \qquad \epsilon_0 > 0. \tag{42}$$

Using the fixed edge set $\mathcal{E} = \{(p, q) : A_{pq} > 0\}$ defined by the deterministic group adjacency, we define the relation feature

$$\psi_j = \left[\tilde{g}_j, \left\{ \log \frac{\tilde{g}_{j,p} + \epsilon}{\tilde{g}_{j,q} + \epsilon} \right\}_{(p,q)\in\mathcal{E}} \right], \qquad \epsilon > 0. \tag{43}$$

This feature describes band group distributions and relative band relations. It contains no class label and no hand crafted class rule. The same relation representation is used by KEDP and RCGRPO. Complete masks are used to estimate fixed prototypes, while sampled masks are used to evaluate trajectory feedback.

Lemma A.7 (Scale stability of transferable spectral relation features)**.** *For any multiplicative intensity drift $g_j \mapsto \alpha g_j$ with $\alpha > 0$, the normalized band feature satisfies*

$$\tilde{g}_j(\alpha) = \frac{\alpha g_j}{\alpha \mathbf{1}^\top g_j + \epsilon_0} = \frac{g_j}{\mathbf{1}^\top g_j + \epsilon_0/\alpha} \approx \tilde{g}_j(1). \tag{44}$$

For each $(p, q) \in \mathcal{E}$, the log ratio term satisfies

$$\begin{aligned} \chi_{pq}(\alpha) &= \log \frac{\tilde{g}_{j,p}(\alpha) + \epsilon}{\tilde{g}_{j,q}(\alpha) + \epsilon} \\ &\approx \log \frac{\tilde{g}_{j,p}(1) + \epsilon}{\tilde{g}_{j,q}(1) + \epsilon} = \chi_{pq}(1). \end{aligned} \tag{45}$$

The invariance is exact when $\epsilon_0 = 0$ and $\epsilon = 0$, provided all denominators are positive.

Proof. Eq. (44) follows by substituting αg_j into Eq. (42). When ϵ_0 is small relative to $\alpha \mathbf{1}^\top g_j$, the normalized feature is stable under multiplicative scaling. Substituting this result into the log ratio definition gives Eq. (45). If $\epsilon_0 = 0$ and $\epsilon = 0$, the common factor α cancels exactly. □

A.2.2 CLASS CONDITIONAL SPECTRAL PROTOTYPE CONSISTENCY

Let $\mathcal{D}_k$ denote the training subset associated with class k. For each class k, we estimate its fixed relation prototype from the training partition, using complete mask features $\psi_j = \psi(b_j, \mathbf{1})$:

$$\mu_k = \frac{1}{|\mathcal{D}_k|} \sum_{j:y_j=k} \psi_j, \tag{46}$$

and

$$\sigma_k^2 = \frac{1}{|\mathcal{D}_k|} \sum_{j:y_j=k} (\psi_j - \mu_k)^2 + \epsilon. \tag{47}$$

The variance is computed elementwise. It controls the reliability of each spectral relation dimension. Training labels are used to estimate the class conditional prototypes and to identify the target prototype during policy optimization. They are not encoded in ψ_j and are not required at inference.

For an input spectrum under the sampled band-group mask x_j, we define its consistency score with prototype k as a diagonal Gaussian log likelihood up to a constant:

$$\ell_k(x_j) = -\frac{1}{2} \sum_r \left[\frac{(\psi_{j,r} - \mu_{k,r})^2}{\sigma_{k,r}^2} + \log \sigma_{k,r}^2 \right]. \tag{48}$$

For a training sample with label y_j, we define the target aligned prototype margin as

$$\rho_{y_j}(x_j) = \ell_{y_j}(x_j) - \max_{k \neq y_j} \ell_k(x_j). \tag{49}$$

The relation consistency reward is

$$r^{(i)}_{\mathrm{spc},t} = \beta_{\mathrm{spc}} \tanh\left(\frac{\rho_{y_j}((x_j, m^{(i)}_t))}{\tau}\right). \tag{50}$$

This reward is used only during policy optimization, where the training label is available.

A.2.3 Group Relative Advantage: Centering and Scale Normalization

Policy Distribution under Energy Guidance During each RCGRPO update, the energy model E_Θ is held fixed. We define $\pi_\eta(a_t \mid s_t)$ as the conditional distribution of the actual guided reverse transition, including the energy correction from $\nabla_{x_t} E_\Theta$. Here, a_t denotes the sampled reverse action, while m_t is the mask obtained from this transition. RCGRPO evaluates the log probability under this guided distribution rather than the unguided diffusion policy. Conditional on fixed guidance, independently sampled trajectories follow the same policy. Therefore, the score function identity and group centering result apply to the guided sampling process.

Let π_η induce a fixed horizon trajectory $\tau = (s_0, a_0, \ldots, s_{T_{\max}})$ with

$$R(\tau) = \sum_{t=1}^{T_{\max}} r_t.$$

Define

$$J(\eta) := \mathbb{E}_{\tau \sim \pi_\eta}[R(\tau)]. \tag{51}$$

Lemma A.8 (REINFORCE score function identity)**.**

$$\nabla_\eta J(\eta) = \mathbb{E}_{\tau \sim \pi_\eta}\left[\sum_{t=1}^{T_{\max}} \nabla_\eta \log \pi_\eta(a_t \mid s_t) R(\tau)\right]. \tag{52}$$

Proof. Using the score function identity,

$$\nabla_\eta \mathbb{E}_{\tau \sim \pi_\eta}[R(\tau)] = \mathbb{E}_{\tau \sim \pi_\eta}\left[R(\tau) \nabla_\eta \log p_\eta(\tau)\right]. \tag{53}$$

Since $p_\eta(\tau)$ contains the policy factors $\prod_t \pi_\eta(a_t \mid s_t)$, we have

$$\nabla_\eta \log p_\eta(\tau) = \sum_{t=1}^{T_{\max}} \nabla_\eta \log \pi_\eta(a_t \mid s_t). \tag{54}$$

Substituting Eq. (54) into Eq. (53) gives Eq. (52). □

Lemma A.9 (Baseline identity)**.** *Let b be any random variable independent of the current trajectory's actions. Then*

$$\mathbb{E}\left[\nabla_\eta \log \pi_\eta(a_t \mid s_t) b | s_t\right] = 0. \tag{55}$$

Therefore, replacing $R(\tau)$ by $R(\tau) - b$ in Eq. (52) *leaves the expected policy gradient unchanged.*

Proof. For a discrete action space,

$$\begin{aligned}
\mathbb{E}\left[\nabla_\eta \log \pi_\eta(a_t \mid s_t) b | s_t\right] &= b \sum_{a_t} \pi_\eta(a_t \mid s_t) \nabla_\eta \log \pi_\eta(a_t \mid s_t) \\
&= b \sum_{a_t} \nabla_\eta \pi_\eta(a_t \mid s_t) \\
&= b \nabla_\eta \sum_{a_t} \pi_\eta(a_t \mid s_t) = 0.
\end{aligned} \tag{56}$$

The same result holds for continuous actions by replacing the summation with an integral. □

Full group centering For one training context, sample $G > 1$ independent trajectories $\{\tau^{(i)}\}_{i=1}^{G}$ under the same policy and context, and let $R^{(i)} = R(\tau^{(i)})$. Define

$$\bar{R} = \frac{1}{G}\sum_{i=1}^{G} R^{(i)} \tag{57}$$

and the trajectory score

$$S^{(i)} = \sum_{t} \nabla_\eta \log \pi_\eta \left(m_t^{(i)} \mid s_t^{(i)}\right). \tag{58}$$

Theorem A.10 (Full group centering)**.** *The centered score estimator satisfies*

$$\mathbb{E}\left[\frac{1}{G}\sum_{i=1}^{G} S^{(i)}\left(R^{(i)} - \bar{R}\right)\right] = \frac{G-1}{G}\nabla_\eta J(\eta). \tag{59}$$

Thus, full group centering preserves the expected policy gradient direction up to a positive constant.

Proof. Using the score function identity, $\mathbb{E}[S^{(i)}] = 0$. Since different trajectories are sampled independently under the same context,

$$\mathbb{E}\left[S^{(i)}\bar{R}\right] = \frac{1}{G}\mathbb{E}\left[S^{(i)}R^{(i)}\right]. \tag{60}$$

Therefore,

$$\mathbb{E}\left[S^{(i)}\left(R^{(i)} - \bar{R}\right)\right] = \frac{G-1}{G}\nabla_\eta J(\eta). \tag{61}$$

Averaging over $i = 1, \ldots, G$ gives Eq. (59). □

In implementation, the centered returns are additionally normalized by the full group sample standard deviation

$$\sigma_R = \sqrt{\frac{1}{G-1}\sum_{i=1}^{G}\left(R^{(i)} - \bar{R}\right)^2}. \tag{62}$$

The quantity σ_R is treated as fixed within each sampled group. It multiplies the centered group estimator by one positive scalar and therefore preserves the realized update direction while stabilizing its scale. Since this scalar is random across groups, the expectation identity above is stated only for the centered estimator.

B Pseudocode of Knowledge-Energy Reinforced Diffusion Framework (KERD)

The pseudocode for KERD is shown in Algorithm 1.

C Multi Source Cohorts

We describe the heterogeneous spectral datasets used in this study and summarize their acquisition settings, sample scale, and clinical annotations. All clinical variables were de-identified before analysis.

(1) HUA Keratin: We construct a nail based Raman dataset to describe the pathology of hyperuricemia in a noninvasive way. Samples come from more than two hundred subjects enrolled at Anonymized Center A in 2025. For each subject, multiple nail surface sites were measured, and 16,035 quality-controlled nail Raman spectra were retained. Each spectrum contains 800 sampling points, which describe local nail spectral structure in detail. These spectra support nail based representation and modeling of hyperuricemia. In addition, we collect clinical variables for all subjects. These variables include uric acid, lipid metabolism markers, inflammatory cell counts, and cardiovascular risk factors, covering more than 40 clinical feature categories. These quantities provide complementary context for Raman bands related to proteins, lipids, nucleic acids, and metabolites.

Algorithm 1: Knowledge Energy Reinforced Diffusion (KERD)

Input: Training bags $\{(\mathcal{X}_i, c_i, y_i)\}_{i=1}^{M}$, structural specification $\mathcal{K} = (\mathcal{G}, A)$, group size G, and horizon $T_{\max}$

Output: f_θ, g_ϕ, G_γ, E_Θ, μ_η, A_ξ, and Head_ω

Initialize $\theta, \phi, \gamma, \Theta, \eta, \xi, \omega$

for *each supervised training epoch* **do**

  for *each mini batch* (x, c, y) **do**

    Partition x by $\mathcal{G}$ and compute $b_k \leftarrow |\mathcal{I}_k|^{-1} \sum_{\ell \in \mathcal{I}_k} x_\ell$

    $v \leftarrow f_\theta(x)$, $u \leftarrow g_\phi(c)$, $h \leftarrow \mathrm{Fuse}\big(v, u, G_\gamma(b, u; \mathcal{K})\big)$

    $L_{\mathrm{ce}} \leftarrow \mathrm{CE}\big(\mathrm{Head}_\omega(h), y\big)$

    Update $(\theta, \phi, \gamma, \omega)$ using ∇L_{ce}

  end

end

Estimate $\{(\mu_k, \sigma_k)\}_{k=1}^{C}$ from $\{\psi(b_j, \mathbf{1})\}_{j \in \mathcal{D}_{\mathrm{train}}}$

for *each policy training epoch* **do**

  for *each training bag* $(\mathcal{X}_i, c_i, y_i)$ **do**

    Encode $\mathcal{X}_i$ into $\{v_{i,j}, b_{i,j}\}_{j=1}^{N_i}$ and compute $u_i \leftarrow g_\phi(c_i)$

    Construct the complete state h_i^{full} and set

$$\bar{p}_i \leftarrow \mathrm{sg}\left[\mathrm{softmax}\left(\mathrm{Head}_\omega(h_i^{\mathrm{full}})\right)\right]$$

    for $g = 1$ **to** G **do**

      $h_0^{(g)} \leftarrow \mathbf{0}$, $\mathcal{T}^{(g)} \leftarrow \emptyset$

      for $t = 1$ **to** $T_{\max}$ **do**

        Compute $\widehat{E}_t^{(g)}$ by Eq. (38) and generate $m_t^{(g)}$ from μ_η using $\nabla_{x_t} E_\Theta$

        Aggregate the selected groups to obtain $(\bar{v}_t^{(g)}, \bar{b}_t^{(g)})$

        $h_t^{(g)} \leftarrow A_\xi(h_{t-1}^{(g)}, \bar{v}_t^{(g)})$

        $\Delta\mathcal{M}_t^{(g)} \leftarrow \mathcal{M}(h_t^{(g)}, y_i) - \mathcal{M}(h_{t-1}^{(g)}, y_i)$

        $r_{\mathrm{rel},t}^{(g)} \leftarrow \beta_{\mathrm{rel}} \tanh\left(\frac{\Phi_{y_i}(\bar{b}_t^{(g)}, m_t^{(g)}) - \max_{k \neq y_i} \Phi_k(\bar{b}_t^{(g)}, m_t^{(g)})}{\tau}\right)$

        $r_t^{(g)} \leftarrow \Delta\mathcal{M}_t^{(g)} + r_{\mathrm{rel},t}^{(g)} - \lambda_{\mathrm{div}} s_t^{(g)}$

        Store $(h_{t-1}^{(g)}, m_t^{(g)}, \log \pi_\eta(m_t^{(g)}), r_t^{(g)})$ in $\mathcal{T}^{(g)}$

      end

      $R^{(g)} \leftarrow \sum_{t=1}^{T_{\max}} r_t^{(g)}$

    end

    $\bar{R} \leftarrow \frac{1}{G} \sum_{g=1}^{G} R^{(g)}, \quad \sigma_R \leftarrow \sqrt{\frac{1}{G-1} \sum_{g=1}^{G} (R^{(g)} - \bar{R})^2}$

    $\hat{A}^{(g)} \leftarrow \frac{R^{(g)} - \bar{R}}{\mathrm{sg}(\sigma_R) + \epsilon}, \quad g = 1, \ldots, G$

    Update (η, ξ) by RCGRPO and update Θ by Eq. (39)

  end

end

At inference, compute $\bar{p}$ from the complete observation and run $T_{\max}$ reverse updates using μ_η with $\nabla_{x_t} E_\Theta$ and A_ξ

return $\arg\max \mathrm{Head}_\omega(h_{T_{\max}})$

These variables cover many laboratory tests and follow up indicators to enable joint Raman clinical modeling.

(2) CVD Serum: We construct a surface enhanced ultrafiltration serum Raman dataset for cardiovascular disease. All spectra come from clinical sampling at Anonymized Center A. The dataset provides 1489 Raman spectral sequences. It covers major cardiovascular subtypes such as atrial fibrillation, coronary artery disease, acute myocardial infarction. We acquire spectra on a WITec alpha300R Raman microscope with a 532 nm laser. Laser power stays between 10 and 14 mW. We use a 1200 g per mm grating and a one hundred times objective with numerical aperture 0.9. Each spectrum has an integration time from 16 to 20 seconds. The acquisition range spans from 280 to 2187 cm^{-1} with 1024 sampling points. This dataset stably captures serum Raman signatures of each cardiovascular subtype. It provides reliable support for noninvasive cardiovascular diagnosis models. We also collect structured clinical variables for these subjects. These variables include percutaneous coronary intervention, essential hypertension, diabetes mellitus, acute cerebral infarction, smoking related status. They supply systematic clinical context for serum spectral modeling.

(3) Gastric SERS: We construct a surface enhanced serum Raman dataset from patients with gastric cancer. All serum samples come from Anonymized Center B in 2025. We use a Raman spectrometer ATR3110 for spectral acquisition. The dataset covers pathological grades of gastric cancer such as poorly differentiated, moderately differentiated, well differentiated. It describes spectral differences across these differentiation states. Sample preparation follows a serum surface enhancement protocol. For each serum sample we mix 100 microliters with 20 microliters gold colloid then drop the mixture on the substrate. Each subject provides four serum drops under identical acquisition parameters. Laser power is fixed at 100 milliwatts with an integration time of 3 seconds. We finally obtain 3000 serum Raman spectra with 2048 sampling points each. These data support systematic analysis of spectral patterns related to gastric cancer grading. We also collect extensive clinical information in parallel. Examples include monocyte percentage MONO percent, white blood cell count WBC, neutrophil percentage NEUT percent, lymphocyte percentage LYM. In total we record more than twenty clinical indices to build a feature space for joint analysis.

(4) Urine Myeloma: We focus on urine samples from patients with multiple myeloma. We use silver nanoparticles AgNP as a surface enhanced substrate to strengthen metabolite related and protein related vibrational signals. All samples are collected in 2025 at Anonymized Center B. We obtain 2600 urine SERS spectra in total. Measurements for each subject are finished within a single batch to reduce within batch bias and between batch bias. The dataset forms a reproducible urine spectral library for body fluid scenarios beyond blood tumors. It enables evaluation of model band recognition in fluid media outside hematologic disease.

(5) Covid Serum: We use a serum Raman dataset on coronavirus disease 2019 to study early screening performance (Yin et al., 2021). The dataset contains 309 spectra from 103 distinct subjects. We include this independently collected cohort as an additional heterogeneous evaluation task.

(6) Ovarian Plasma: We use the ovarian plasma dataset to evaluate blood based spectra for ovarian cancer (Paraskevaidi et al., 2018). We collect EDTA anticoagulated whole blood from 27 ovarian cancer patients and 28 benign patients. The dataset includes 385 spectra with 1480 sampling points each. We separate plasma by centrifugation at 2000 revolutions per minute for 10 minutes then store it at low temperature. We place 50 microliters plasma on aluminum foil slides and let the drops dry. We also mix plasma with diluted silver nanoparticles in a one to one ratio to prepare SERS samples.

(7) Bacteria Atlas: The bacteria atlas dataset derives from the Bacteria ID collection (Ho et al., 2019). It targets identification of common pathogenic bacteria and yeast strains. These strains cover most inpatient infection cases at Stanford Hospital from 2016 to 2017. We culture reference strains and clinical isolates on blood agar plates. We scrape about 0.6 milligrams biomass from a single colony and add 10 microliters sterile water to form a suspension. We then place 3 microliters of this suspension on a gold coated silicon slide. Raman acquisition uses a Horiba LabRAM HR Evolution microscope with a 633 nanometer laser. We set laser power near 13.17 milliwatts. We use a 300 line per millimeter grating with a one hundred times objective of numerical aperture 0.9. In mapping mode we define a 45 by 45 grid on the sample surface with spot spacing about 3 micrometers. The integration time is 1 or 2 seconds for each spectrum. We record cellular spectra from 382 to 1792 cm^{-1}. We perform polynomial baseline correction followed by intensity normalization.

Table 1: Core hyperparameters of KERD.

Hyperparameter	**Value**
Model dimension d_{model}	128
Feature dimension d_{ff}	512
Patch length / stride	8 / 4
Auxiliary context embedding dimension	128
Auxiliary context dropout	0.10
Diffusion noise scale	0.10
Energy guidance type	Amortized middle energy, hidden 256
Guidance scale	1.00
Guidance rescale	0.70
Mask threshold	0.50
Energy weight	0.20
Smoothness weights	0.05
Target mask	0.30
Peak-Graph band groups	16
Graph attention slope	0.20
Diversity penalty	0.30

The dataset contains 48239 usable spectra with 1000 sampling points per spectrum. These spectra support construction of fast identification models for bacterial species.

Each dataset is trained and evaluated independently under this protocol. Accordingly, the multi dataset evaluation measures consistency under heterogeneous conditions rather than transfer to an unseen source.

D Experimental Settings and Hyperparameters

The hyperparameters used in our method are shown in Table 1. All major weights are kept consistent across datasets. They are used only to balance the numerical scales of energy and reward terms, and do not encode disease specific decision boundaries. To evaluate model performance, we use accuracy (ACC), precision (PREC), recall, F1 score, AUROC and AUPRC. During training, we use AdamW with a learning rate of 3×10^{-4}. We set batch size to 32, model dimension to 128, hidden feedforward dimension to 512. We set the auxiliary-context dropout to 0.1. We train each model for up to 500 epochs. Early stopping is applied when the validation F1 score does not improve for 30 consecutive epochs. All experiments are conducted on a server running Ubuntu 20.04.6 LTS. The server has Intel Xeon Gold 5220 CPUs with 144 logical cores. We perform entity level stratified partitioning for each dataset, with 80% of entities for training, 10% for validation, and 10% for testing. For subject annotated cohorts, the entity denotes a subject. For microbial cohorts, the entity denotes an isolate or strain when such metadata are available. We split the data by entity identity to avoid overlap between the training and test sets. The spectral interaction encoder is initialized from scratch. It uses a model dimension of 128 and a feed forward dimension of 512. We train the model on A40 GPUs with Python 3.11 and CUDA 12.0. To assess result stability, we report mean and standard deviation over five random seeds during training.

We use an auxiliary reconstruction head to visualize spectral recovery in the perturbation analysis. For a perturbed spectrum $\tilde{x}$, KERD generates a robust fused representation $z_{\text{rob}} \in \mathbb{R}^d$. The reconstruction head then projects z_{rob} back to the spectral domain and produces $\hat{x} \in \mathbb{R}^L$. The head is implemented as a lightweight MLP with LayerNorm, GELU activation, and a final linear layer. Its output length is the same as the original spectrum, so $\hat{x}$ can be compared with the clean spectrum x at each spectral position. This reconstruction head is used only for analysis in the perturbation experiment. It helps assess whether the fused representation retains structural information that can be recovered after spectral corruption.

For Gastric SERS, CVD Serum, and HUA Keratin, KERD and all nine general baselines use identical clinical variables under each shared context setting. The six domain specific baselines use the same clinical inputs on Gastric SERS and CVD Serum. Models without native context inputs receive the same MCGR interface. Direct label indicators and clinical fields used to define the target diagnosis are excluded from prediction for every method.

MCGR Construction Each available auxiliary variable is represented as one context node after training set normalization. Continuous variables are standardized with training set statistics. Missing continuous values are imputed with the corresponding training set median and accompanied by a missingness indicator. The MCGR candidate graph contains the fixed observation group edges defined by A, context to context edges, and all observation to context edges. Self loops are included. Therefore, $\mathcal{N}(j)$ is the fixed set of incoming neighbors under these edge types, while graph attention learns their data dependent weights. All preprocessing statistics are estimated from the training partition only.

Diffusion and Policy Implementation The diffusion policy adopts a conditional noise predictor with a linear variance schedule. At each reverse step, the denoised latent state and its uncertainty provide the posterior moments for analytic energy approximation.

The resulting structural energy guides diffusion sampling, while RCGRPO optimizes the policy using group relative trajectory feedback. The energy model is optimized separately and held fixed during each policy update. Detailed configurations are reported in Table 1.

D.1 Ablation Analysis

The ablation study investigates how structural priors affect different stages of representation construction, exploration, and optimization. Fig. 3A decomposes KERD along three stages of the representation process. MCGR introduces relational context into the latent state, KEDP uses structural energy to explore group configurations, and RCGRPO reinforces configurations that remain consistent relative to group prototypes. The base model removes these structured constraints, whereas each ablated variant isolates one source of inductive bias.

The complete model consistently outperforms all variants. The reduction caused by removing KEDP shows that unconstrained exploration does not reliably recover coordinated group patterns, even when a strong sequence encoder is available. The degradation without RCGRPO indicates that sample-wise rewards alone provide an unstable learning signal and may amplify incidental dependencies. MCGR yields additional improvements when external context is available, suggesting that auxiliary information is most useful when it refines an existing structured state rather than forming a separate representation pathway. These observations support a unified view in which robust evidence emerges from the repeated preservation of relational geometry across encoding, exploration, and optimization.

Stability under heterogeneous perturbations To examine the stability of relational geometry under heterogeneous conditions, Fig. 3B compares learning behavior under global drift, background interference, and localized corruption at varying severity levels. KERD maintains more stable learning curves than the model without RCGRPO, especially under shifts that alter global magnitude while preserving relative local structure. This behavior suggests that prototype-consistent updates encourage invariance to nuisance variation without discarding informative group interactions.

Effect of diffusion design Fig. 3C studies how the diffusion process affects structural recovery. The alternatives vary the number of reverse steps, the presence of stochastic noise, and whether energy is applied once or throughout the trajectory. KERD-Energy-Diff provides the strongest and most balanced results. Removing stochasticity reduces coverage of plausible configurations, whereas random-walk sampling lacks a task-relevant direction. One-shot energy fails to correct later deviations, and overly deep diffusion accumulates unnecessary noise. By applying energy guidance at successive reverse steps, KEDP retains exploration capacity while continuously aligning sampled configurations with the learned relational structure. The comparison therefore validates continuous energy shaping as a key mechanism for recovering stable group level dependencies.

Table 2: Performance comparison on representative tasks. Purple marks KERD (Ours). The baseline methods under comparison are divided into three groups: sequence oriented global dependency modeling, variable driven cross channel correlation modeling, and structure driven general modeling.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
CVD-Serum (shared auxiliary-context setting)	**Autoformer**	87.85±0.91	87.96±1.28	88.14±0.85	87.93±1.06	97.31±0.34	93.06±0.81
	FEDformer	87.52±0.58	87.71±0.51	89.16±0.51	87.96±0.52	97.54±0.05	94.16±0.28
	MTST	85.17±2.02	85.68±1.50	87.82±1.28	85.88±1.72	97.77±0.23	94.60±0.67
	PatchTST	88.46±1.43	88.27±2.14	89.33±1.02	88.67±1.55	97.81±0.33	94.80±0.84
	MultiPatchFormer	89.53±1.13	89.44±1.28	90.56±0.98	89.87±1.18	98.18±0.15	94.55±0.40
	iTransformer	81.68±0.81	82.27±1.18	85.05±0.94	82.55±0.95	95.10±0.23	87.63±0.57
	TimeFilter	90.47±0.99	90.74±1.02	90.40±1.53	90.52±1.21	**98.35±0.25**	95.01±0.81
	TCN	84.03±2.31	85.74±2.63	83.38±1.81	84.40±2.12	95.06±0.97	89.14±1.83
	Transformer	88.93±0.60	88.97±0.78	89.48±0.62	89.11±0.58	97.11±0.23	92.93±0.53
	KERD (Ours)	**93.02±1.55**	**93.29±1.74**	**92.77±1.07**	**93.00±1.39**	98.31±0.57	**95.83±1.44**
Gastric-SERS (shared auxiliary-context setting)	**Autoformer**	78.57±2.91	78.64±2.79	78.68±2.77	78.52±2.91	93.51±1.40	84.76±2.91
	FEDformer	82.03±0.87	81.66±0.87	82.39±0.93	81.87±0.90	95.50±0.30	88.35±0.73
	MTST	80.50±1.34	80.17±1.30	80.84±1.25	80.37±1.29	94.91±0.84	87.53±2.15
	PatchTST	82.20±0.76	81.97±0.67	82.39±0.91	82.09±0.76	95.79±0.42	89.51±0.86
	MultiPatchFormer	82.17±1.05	82.12±1.06	82.55±1.00	82.26±1.04	95.20±0.18	89.11±0.25
	iTransformer	86.00±0.67	85.77±0.63	86.17±0.75	85.94±0.68	96.38±0.18	91.48±0.51
	TimeFilter	80.53±1.44	80.58±1.30	80.29±1.40	80.34±1.39	94.91±0.49	87.28±0.99
	TCN	78.80±1.55	78.94±1.69	78.85±1.41	78.63±1.54	94.13±0.84	86.27±2.02
	Transformer	77.60±1.00	77.22±0.89	77.69±1.05	77.31±0.95	94.06±0.21	85.37±0.51
	KERD (Ours)	**93.37±0.59**	**93.40±0.66**	**93.30±0.57**	**93.32±0.61**	**98.76±0.37**	**96.99±0.66**
Bacteria-Atlas	**Autoformer**	75.75±2.46	77.10±2.97	77.51±2.53	76.79±2.80	94.18±0.38	86.43±1.19
	FEDformer	74.63±1.66	75.81±3.34	75.33±2.58	74.53±3.00	93.79±0.49	86.22±1.06
	MTST	82.31±2.42	83.94±2.63	84.46±2.23	83.90±2.40	96.28±0.90	91.18±1.84
	PatchTST	78.47±2.00	77.21±2.81	81.59±2.54	78.67±2.51	94.93±0.72	87.94±1.46
	MultiPatchFormer	88.63±2.18	88.49±3.20	89.53±1.63	88.75±2.57	98.55±0.26	95.95±0.69
	iTransformer	83.78±1.65	82.38±1.91	85.14±1.48	82.88±1.79	97.63±0.80	93.76±1.58
	TimeFilter	91.42±0.95	92.35±0.91	91.99±1.13	92.13±1.03	99.02±0.20	97.18±0.56
	TCN	63.50±2.43	64.17±1.74	66.39±2.15	61.90±2.56	90.48±0.79	77.97±1.76
	Transformer	78.25±4.29	79.52±3.91	81.34±2.42	80.17±3.17	95.11±1.25	87.63±2.62
	KERD (Ours)	**95.49±0.32**	**95.92±0.32**	**95.19±0.37**	**95.54±0.33**	**99.67±0.05**	**98.88±0.16**

E Detailed Performance and Statistical Analysis

E.1 Detailed Baseline Comparisons

Table 2 provides complete baseline comparisons across the three representative tasks.

E.2 Extended Cohort Results

The complete results on the four additional cohorts are reported in Table 6.

E.3 Statistical Significance Analysis

To assess aggregate performance, we conduct Friedman tests across seven cohorts, treating cohorts as blocks and the ten methods as treatments. Each method is represented by its mean performance over five random seeds, with ties assigned average ranks.

As shown in Table 3, significant differences are observed for all six metrics, and all tests remain significant after Holm correction. KERD obtains the lowest average rank for every metric.

Nemenyi post hoc comparisons further identify significant differences against several baselines, including TCN and Transformer on all six metrics ($p < 3 \times 10^{-4}$). However, differences from some strong competitors are not uniformly significant. These results support an aggregate performance advantage rather than uniform pairwise superiority on every individual dataset.

Table 3: Friedman tests across heterogeneous cohorts. Lower average rank is better.

Metric	χ^2	Raw p	Holm p	Rank
Accuracy	38.646	1.33×10^{-5}	5.63×10^{-5}	**1.000**
Precision	38.158	1.63×10^{-5}	5.63×10^{-5}	**1.000**
Recall	38.193	1.61×10^{-5}	5.63×10^{-5}	**1.000**
F1	36.288	3.52×10^{-5}	5.63×10^{-5}	**1.000**
AUROC	39.494	9.38×10^{-6}	5.63×10^{-5}	**1.571**
AUPRC	39.312	1.01×10^{-5}	5.63×10^{-5}	**1.429**

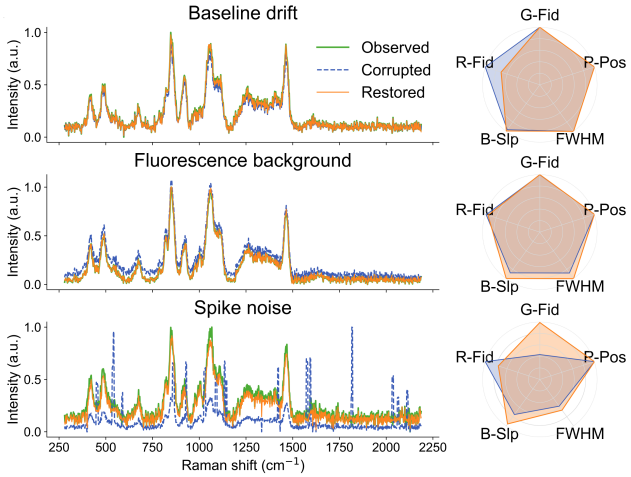


Figure 5: **Perturbed spectral reconstruction and statistical consistency.** Recovery of informative spectral structures. The left panels compare observed, corrupted, and restored spectra under baseline drift, fluorescence background, and spike noise. The radar plots summarize spectral fidelity, peak position, FWHM, baseline slope, and peak ratio fidelity. KERD restores the global spectral profile while preserving key Raman peak structures.

F Robust Spectral Structure Restoration

The core difficulty in heterogeneous spectral evidence learning comes not only from sample heterogeneity, but also from nonbiological acquisition disturbances. Baseline drift changes low frequency trends. Fluorescence background elevates broadband responses. Spike noise introduces nonstructural abrupt changes. These factors may cause the model to rely on spurious correlations. They can reduce generalization across scenarios. Therefore, this section further validates the sequence structure recovery ability of KERD. Specifically, this experiment constructs three typical corruption scenarios. It trains the model to recover the target clean spectrum from the corrupted spectrum. The model constrains the global spectral shape, local gradients and key peak regions. This design preserves Raman peak positions, Full Width at Half Maximum (FWHM), baseline trends and peak ratio relations during recovery.

Fig. 5 shows that the restored spectra closely match the clean spectra under baseline drift and fluorescence background. The model effectively suppresses low frequency background and global intensity shift. Under spike noise, the model weakens abrupt abnormal peaks while preserving major decision relevant peaks. The radar plots on the right also show higher global spectral fidelity after recovery. They also show higher peak position fidelity, FWHM fidelity, baseline slope fidelity and peak ratio fidelity than corrupted spectra. This result shows that KERD does not simply smooth noise. It recovers reliable decision relevant signals while preserving molecular vibrational structures. KERD tends to recover stable structures related to molecular vibrations. This provides more reliable input representations.

Table 4: Optional Auxiliary Context Integration.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
Gastric SERS	**Spectral only**	85.23 ± 1.45	85.06 ± 1.40	85.02 ± 1.53	84.90 ± 1.53	96.58 ± 0.57	91.38 ± 1.27
	Clinical Concat	85.90 ± 1.20	85.76 ± 1.04	85.84 ± 1.15	85.64 ± 1.08	96.91 ± 0.49	92.36 ± 1.42
	KERD+MCGR	$\mathbf{93.37 \pm 0.59}$	$\mathbf{93.40 \pm 0.66}$	$\mathbf{93.30 \pm 0.57}$	$\mathbf{93.32 \pm 0.61}$	$\mathbf{98.76 \pm 0.37}$	$\mathbf{96.99 \pm 0.66}$
CVD Serum	**Spectral only**	87.32 ± 1.24	87.41 ± 1.41	86.74 ± 1.57	87.00 ± 1.48	96.96 ± 0.63	92.35 ± 1.14
	Clinical Concat	89.33 ± 1.59	89.19 ± 1.88	89.32 ± 1.60	89.21 ± 1.73	97.72 ± 0.62	93.98 ± 1.40
	KERD+MCGR	$\mathbf{93.02 \pm 1.55}$	$\mathbf{93.29 \pm 1.74}$	$\mathbf{92.77 \pm 1.07}$	$\mathbf{93.00 \pm 1.39}$	$\mathbf{98.31 \pm 0.57}$	$\mathbf{95.83 \pm 1.44}$
HUA Keratin	**Spectral only**	85.49 ± 0.48	85.52 ± 0.49	85.49 ± 0.48	85.48 ± 0.48	92.25 ± 0.28	91.85 ± 0.30
	Clinical Concat	85.58 ± 0.64	85.61 ± 0.64	85.58 ± 0.64	85.57 ± 0.64	92.15 ± 0.39	91.83 ± 0.45
	KERD+MCGR	$\mathbf{90.39 \pm 0.30}$	$\mathbf{90.40 \pm 0.30}$	$\mathbf{90.39 \pm 0.30}$	$\mathbf{90.39 \pm 0.30}$	$\mathbf{96.28 \pm 0.20}$	$\mathbf{96.30 \pm 0.17}$

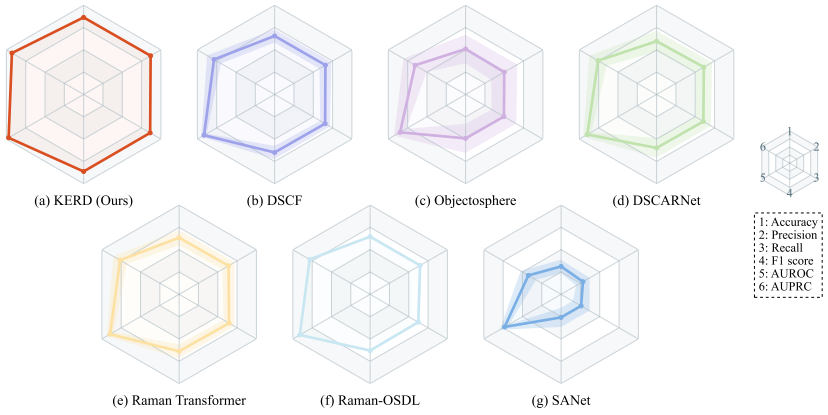


Figure 6: **Performance comparison with domain specific baselines.**

G Auxiliary Context Extension Study

For datasets with paired clinical variables, we further include clinical information as optional external evidence, rather than as required input for the main experiments. Raman spectra mainly reflect the molecular vibrational state of samples at acquisition time. They reveal spectral differences related to proteins, lipids, nucleic acids, and metabolites. However, spectral patterns may be affected by cohort level factors such as metabolic level, inflammatory state, and cardiovascular indicators. Including clinical variables further tests whether this external physiological information can complement Raman representations. The spectral only setting evaluates whether KERD can learn structured spectral evidence without external cohort variables. The naive concatenation row tests whether direct clinical feature input brings effective gains. KERD with MCGR further evaluates whether structured clinical evidence modeling improves over ordinary feature merging. It also tests whether clinical information can regulate spectral representations more properly.

As shown in Table 4, direct concatenation of clinical features brings some gains on several datasets, but the overall improvement remains limited. Auxiliary indicators describe cohort level context, while spectral peaks describe sample level molecular states. MCGR provides more stable gains. MCGR places clinical variables and spectral bands into the same heterogeneous graph for message passing. Spectral nodes retain metabolic and pathological patterns in the spectra, while clinical nodes provide patient level external evidence.

The additional baseline extension experiment provides more direct evidence for fairness. We provide the same clinical fusion channel for representative baselines. We compare spectral only, naive concatenation, and MCGR enhanced versions. As shown in Table 5, MCGR also improves different baseline models. When matched cohort variables are available, MCGR provides structured auxiliary evidence for robust evidence learning. This result also shows that the performance gain comes from controlled evidence transfer. It does not come from exclusive access to extra input information by one model.

Table 5: Auxiliary Context Extension for Baselines.

Dataset	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
HUA Keratin	**MultiPatchFormer-Base**	80.12 ± 1.58	81.11 ± 1.56	80.42 ± 1.57	80.05 ± 1.60	88.40 ± 1.51	88.68 ± 1.42
	MultiPatchFormer+Concat	81.50 ± 1.56	82.46 ± 1.47	81.79 ± 1.54	81.44 ± 1.58	89.95 ± 1.36	90.18 ± 1.25
	MultiPatchFormer+MCGR	$\mathbf{85.04 \pm 0.40}$	$\mathbf{85.05 \pm 0.40}$	$\mathbf{85.04 \pm 0.40}$	$\mathbf{85.04 \pm 0.40}$	$\mathbf{92.97 \pm 0.32}$	$\mathbf{93.11 \pm 0.32}$
	MTST-Base	79.03 ± 1.57	79.07 ± 1.56	79.03 ± 1.50	79.00 ± 1.55	86.66 ± 0.62	86.14 ± 0.66
	MTST+Concat	80.73 ± 2.21	80.80 ± 2.26	80.69 ± 2.14	80.69 ± 2.18	87.31 ± 1.06	86.53 ± 0.92
	MTST+MCGR	$\mathbf{82.63 \pm 0.29}$	$\mathbf{82.90 \pm 0.22}$	$\mathbf{82.63 \pm 0.29}$	$\mathbf{82.59 \pm 0.31}$	$\mathbf{90.69 \pm 0.17}$	$\mathbf{90.34 \pm 0.20}$

H Comparison with Domain-Specific Baselines

We further strengthen this evaluation by introducing six spectrum-specific baselines on the Gastric-SERS and CVD-Serum cohorts: DSCF (Xue et al., 2025), Objectosphere (Balytskyi et al., 2024), DSCARNet (Lu et al., 2024), Raman Transformer (Liu et al., 2023a), Raman-OSDL (Zhu et al., 2025), and SANet (Deng et al., 2021). The selected methods cover several complementary modeling principles encountered in modern Raman analysis. Specifically, the comparison includes spectral component and pretrained representation learning through DSCF, Raman-specific signal aggregation through DSCARNet, attention-based long-range spectral modeling through Raman Transformer, scale-adaptive feature extraction through SANet, and open-set representation learning through Objectosphere and Raman-OSDL. This extension therefore evaluates whether KERD remains competitive when compared directly with approaches whose inductive biases were designed around spectral characteristics rather than generic sequence structures. For both Gastric SERS and CVD Serum, all domain specific baselines use the same entity partitions and the same auxiliary clinical variables as KERD. Models without native clinical inputs use the same MCGR context refinement interface, while their original spectral backbones are kept unchanged.

Fig. 6 summarizes the comparison using six evaluation dimensions. For each model, the polygon is constructed from the metric means averaged across Gastric-SERS and CVD-Serum. The surrounding shaded band measures the cross-dataset standard deviation of these metric means and hence characterizes how strongly a method varies between the two clinical spectral tasks. This representation provides a compact view of both predictive quality and stability: a method with strong but task-dependent performance may reach the outer region along individual dimensions, whereas a consistently effective model is expected to retain a large and balanced profile across all axes.

The results show a clear separation between KERD and the spectrum-specific alternatives. KERD forms the outermost radar profile on every evaluation dimension and maintains consistently high performance on both cohorts. Its Accuracy/F1 scores are 93.37%/93.32% on Gastric-SERS and 93.02%/93.00% on CVD-Serum, substantially exceeding the strongest competing results of 86.53%/86.18% and 83.49%/82.77%, respectively. Similar trends are observed for AUROC and AUPRC. This observation suggests that explicitly organizing and optimizing relations among informative peak groups provides a complementary source of discriminative structure.

Table 6: Multi source screening results using heterogeneous biofluid spectral signals. Purple marks KERD (Ours).
The baseline methods under comparison are divided into three groups: sequence oriented global dependency modeling, variable driven cross channel correlation modeling, and structure driven general modeling.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
HUA Keratin (shared auxiliary-context setting)	**Autoformer**	79.28±1.03	79.40±1.09	79.27±1.03	79.26±1.02	87.60±0.92	87.55±1.04
	FEDformer	82.62±0.20	82.68±0.20	82.62±0.20	82.61±0.20	90.71±0.28	90.60±0.32
	MTST	82.63±0.29	82.90±0.22	82.63±0.29	82.59±0.31	90.69±0.17	90.34±0.20
	PatchTST	83.70±0.12	83.86±0.09	83.70±0.12	83.68±0.13	91.60±0.15	91.44±0.17
	MultiPatchFormer	85.04±0.40	85.05±0.40	85.04±0.40	85.04±0.40	92.97±0.32	93.11±0.32
	iTransformer	85.92±0.14	85.94±0.14	85.92±0.14	85.92±0.14	94.04±0.21	94.25±0.21
	TimeFilter	84.74±0.38	84.77±0.39	84.74±0.38	84.74±0.38	92.62±0.29	92.70±0.34
	TCN	76.78±0.65	76.91±0.63	76.78±0.65	76.75±0.65	84.13±0.27	83.88±0.26
	Transformer	74.72±0.54	75.00±0.57	74.72±0.54	74.65±0.54	81.93±0.57	81.19±0.72
	KERD (Ours)	**90.39±0.30**	**90.40±0.30**	**90.39±0.30**	**90.39±0.30**	**96.28±0.20**	**96.30±0.17**
Urine Myeloma	**Autoformer**	87.38±1.91	87.49±1.91	87.38±1.91	87.38±1.91	93.80±1.44	93.36±1.80
	FEDformer	80.27±2.58	81.21±2.64	80.27±2.58	80.12±2.60	88.87±3.79	88.66±4.57
	MTST	85.54±2.10	85.92±2.18	85.54±2.10	85.50±2.11	93.38±1.16	93.55±1.21
	PatchTST	86.23±2.54	86.72±2.40	86.23±2.54	86.18±2.57	91.65±2.83	91.27±3.73
	MultiPatchFormer	85.15±1.38	86.09±1.24	85.15±1.38	85.05±1.42	92.05±1.62	92.06±1.68
	iTransformer	85.54±3.55	86.19±3.38	85.54±3.55	85.46±3.60	93.05±1.84	93.34±1.72
	TimeFilter	89.38±0.70	89.74±0.52	89.38±0.70	89.36±0.72	96.84±0.51	96.90±0.51
	TCN	80.88±0.62	81.02±0.45	80.88±0.62	80.85±0.66	89.38±0.88	89.81±0.68
	Transformer	78.85±5.69	79.05±5.86	78.85±5.69	78.82±5.68	86.22±5.93	85.82±6.21
	KERD (Ours)	**94.62±1.17**	**94.65±1.15**	**94.62±1.17**	**94.61±1.17**	**99.00±0.31**	**99.03±0.30**
Covid Serum	**Autoformer**	82.90±0.73	83.80±0.60	83.13±0.74	82.81±0.80	85.02±0.50	84.90±0.11
	FEDformer	85.16±1.15	85.49±1.35	85.27±1.27	85.13±1.20	86.31±1.30	87.09±0.25
	MTST	88.06±4.03	88.30±4.17	88.08±4.16	88.03±4.06	89.81±3.21	89.39±3.91
	PatchTST	94.84±0.65	95.19±0.53	95.00±0.63	94.84±0.65	**97.58±1.17**	**97.67±1.18**
	MultiPatchFormer	84.52±2.19	84.62±2.15	84.44±2.24	84.47±2.21	90.19±1.01	89.76±0.97
	iTransformer	89.03±3.73	89.23±3.80	89.15±3.75	89.03±3.73	91.96±1.46	91.93±1.45
	TimeFilter	87.10±3.06	87.22±3.08	87.08±3.16	87.07±3.09	90.92±0.69	89.18±1.05
	TCN	75.48±2.77	76.34±3.08	75.69±2.76	75.37±2.79	78.71±3.86	78.85±4.36
	Transformer	78.39±0.81	78.69±0.94	78.48±0.86	78.35±0.82	81.98±0.65	81.94±0.84
	KERD (Ours)	**95.16±1.02**	**95.22±1.03**	**95.21±1.02**	**95.16±1.02**	96.29±1.80	95.62±2.27
Ovarian Plasma	**Autoformer**	89.87±2.77	90.06±2.72	89.86±2.79	89.85±2.78	96.02±1.86	95.97±2.21
	FEDformer	72.99±0.52	73.28±0.47	72.91±0.54	72.86±0.57	85.20±0.53	86.16±0.53
	MTST	80.52±2.96	80.59±2.96	80.51±2.99	80.50±2.97	89.92±1.81	90.08±1.97
	PatchTST	74.29±2.52	74.41±2.62	74.26±2.49	74.24±2.48	82.58±2.24	83.21±2.00
	MultiPatchFormer	81.56±2.65	82.75±2.23	81.44±2.68	81.33±2.78	94.01±0.99	94.26±1.00
	iTransformer	74.81±2.11	74.94±2.04	74.76±2.13	74.74±2.15	86.52±1.98	87.55±2.11
	TimeFilter	89.87±2.23	90.43±1.48	89.82±2.29	89.81±2.32	**97.92±0.65**	**98.04±0.61**
	TCN	77.40±2.26	77.70±2.11	77.40±2.30	77.33±2.33	83.95±1.23	83.69±1.38
	Transformer	67.79±0.77	68.50±0.83	67.69±0.77	67.50±0.76	72.91±0.79	73.56±0.66
	KERD (Ours)	**90.91±1.42**	**91.45±1.22**	**90.89±1.46**	**90.87±1.45**	94.60±1.72	94.58±1.94