RISK AWARE COMPOSITIONAL EVIDENCE LEARNING FOR HETEROGENEOUS BIOMEDICAL OBSERVATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

High dimensional continuous observations provide rich local information for representation learning in biomedical and scientific data. Recent advances in patch based modeling, structured encoding, and feature selection have improved the ability to capture local dependencies. However, existing methods typically aggregate local observations into a single deterministic representation or retain one preferred subset, making it difficult to preserve multiple useful evidence combinations when local information is redundant, heterogeneous, and unequally reliable. Our key idea is to learn an input conditioned, risk aware distribution over multiple evidence compositions rather than searching for a single optimal representation. Based on this idea, we propose Quantized Structured Flow Representation (QSFR). Specifically, QSFR maps continuously varying local observations into reusable discrete evidence units, constructs alternative evidence compositions through flow based sequential sampling, and models composition level utility as a distribution. The estimated utility and utility dispersion jointly reshape probability mass toward informative and reliable evidence states. We evaluate QSFR across multiple heterogeneous cohorts spanning different biological sources and acquisition conditions. The results show consistently favorable representation quality across multiple evaluation criteria. Mechanism analyses reveal improved local stability, expanded compositional support, and lower risk evidence organization.

1 INTRODUCTION

High dimensional continuous observations are common in modern biomedical sensing and scientific data analysis, providing rich local information for representation learning (Yue et al., 2022; Goswami et al., 2024; Wang et al., 2025). Such observations usually consist of many correlated local patterns, with information also encoded in dependencies across different regions (Wu et al., 2021; Zhou et al., 2022). Recent advances in local patch representations, sequence models, and discrete encoding substantially expand the modeling capacity for continuous data (Dong et al., 2023; Razavi et al., 2019). These methods transform raw high dimensional inputs into compact latent representations and are widely applied to biomedical signals (Wang et al., 2024b). Therefore, modeling local structures and their compositional relations becomes an important direction in continuous representation learning (Jiang, 2025; He et al., 2026).

Given continuous observations, existing studies usually obtain predictive representations through fixed aggregation or salient feature selection (Yamada et al., 2020; Covert et al., 2023). The former directly integrates local features and produces a deterministic global vector (Lee et al., 2019). The latter first identifies informative regions, then constructs the final representation from the selected regions (Chen et al., 2018). However, either strategy usually retains only one final representation or one preferred feature set (Greenstein et al., 2025; Li et al., 2026). This design makes it difficult to capture cases where multiple local evidence combinations are simultaneously effective (Acikalin et al., 2025). Prior set based learning studies model interactions among multiple elements, but they generally aggregate these interactions into a task specific representation (Tokar & Sanner, 2026; Chen et al., 2026). Real biomedical observations also contain substantial redundancy, while different local regions may have different discriminative value and reliability (Wang et al., 2024b; 2025). Some methods consider feature selection or predictive uncertainty, but they usually treat these factors independently (Yamada et al., 2020; Liu et al., 2020; Charpentier et al., 2020). Therefore,

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.

the compositional dependencies among local evidence remain insufficiently modeled in a unified framework. Based on this observation, a key question arises:

When local evidence is redundant and heterogeneous, how can we learn a reliable distribution over compositional evidence from high dimensional continuous biomedical observations?

In this work, we study local evidence structures in high dimensional continuous observations. We analyze how redundancy, compositional dependence, and uncertainty jointly affect representation learning. Specifically, we view this problem as three sequential processes: local evidence abstraction, compositional relation modeling, and reliability assessment. Different local evidence can form multiple effective combinations (Bengio et al., 2021). A more appropriate objective is to learn an input conditioned distribution over compositional evidence, with probability mass favoring states with higher utility and lower risk (Zhang et al., 2022; Malkin et al., 2022). Based on this idea, we propose Quantized Structured Flow Representation (QSFR). QSFR first converts continuous observations into structured discrete local representations, then explores and samples multiple candidate evidence states. It then adjusts the relative probability of each combination according to predictive utility and utility dispersion, forming a risk aware distribution over compositional evidence. This process preserves multiple evidence combinations with complementary information instead of compressing all local information into a single deterministic representation. We systematically evaluate QSFR on heterogeneous Raman cohorts. We further use ECG and EEG to examine whether similar evidence organization behavior appears across biomedical signal domains. Experimental results support our motivation through predictive performance, compositional diversity, utility dispersion modeling, and changes in evidence distributions. In summary, our main contributions are fourfold:

- We introduce a new continuous to compositional evidence learning formulation. Its core learning target is an input conditioned distribution over evidence sets rather than a single deterministic representation.
- We propose QSFR, which allocates probability mass in the compositional space through quantized evidence prototypes and generative flow modeling.
- We model a conditional distribution of composition-level task utility and use its dispersion as a relative risk surrogate to guide risk-aware compositional sampling toward high-utility, low-dispersion evidence states.
- We systematically evaluate QSFR across eight Raman cohorts and cross condition robustness settings. We further use ECG and EEG as cross-signal mechanism probes to examine evidence distribution behavior.

2 Related work

Structured Representation Learning. Structured representation methods model local dependencies across regions, variables, and scales (He et al., 2026; Nie et al., 2022; Wu et al., 2022; Liu et al., 2024; Wang et al., 2024a). Discrete methods map continuous features to reusable codes through learned codebooks or scalar levels (Yu et al., 2021; Mentzer et al., 2024). Related spectral models also capture spatial and spectral dependencies to reduce redundancy (Wang et al., 2025). However, these approaches mainly improve local encoding and aggregation rather than learning an explicit distribution over alternative evidence compositions.

Generative Flow Modeling. Prior studies identify informative elements through instance dependent selection, differentiable subset optimization, stochastic gating, and sequential acquisition (Chen et al., 2018; Yoon et al., 2018; Balın et al., 2019; Yamada et al., 2020; Covert et al., 2023). Set models further capture interactions among selected elements (Lee et al., 2019), but typically optimize a subset or aggregate its information. Generative Flow Networks instead assign reward dependent probability to multiple compositional objects through sequential decisions (Bengio et al., 2021; Zhang et al., 2022; Malkin et al., 2022; Bengio et al., 2023). Later work extends this framework to richer compositional settings and generalization (Zhang et al., 2025).

Distributional Utility and Risk Modeling. Predictive reliability has been studied through evidential methods, distribution shift analysis, representation geometry, and density based posterior models (Sensoy et al., 2018; Ovadia et al., 2019; Liu et al., 2020; Amini et al., 2020; Charpentier

et al., 2020). Such uncertainty estimates can change under dataset shift. Related distributional reinforcement learning represents possible returns beyond their mean (Li et al., 2026). Quantile methods approximate return variability (Dabney et al., 2018; Yang et al., 2019), while risk sensitive methods use return distributions to guide decisions (Lim & Malik, 2022; Kastner et al., 2023). These approaches mainly assess uncertainty at the prediction, value, or return level rather than within individual evidence compositions.

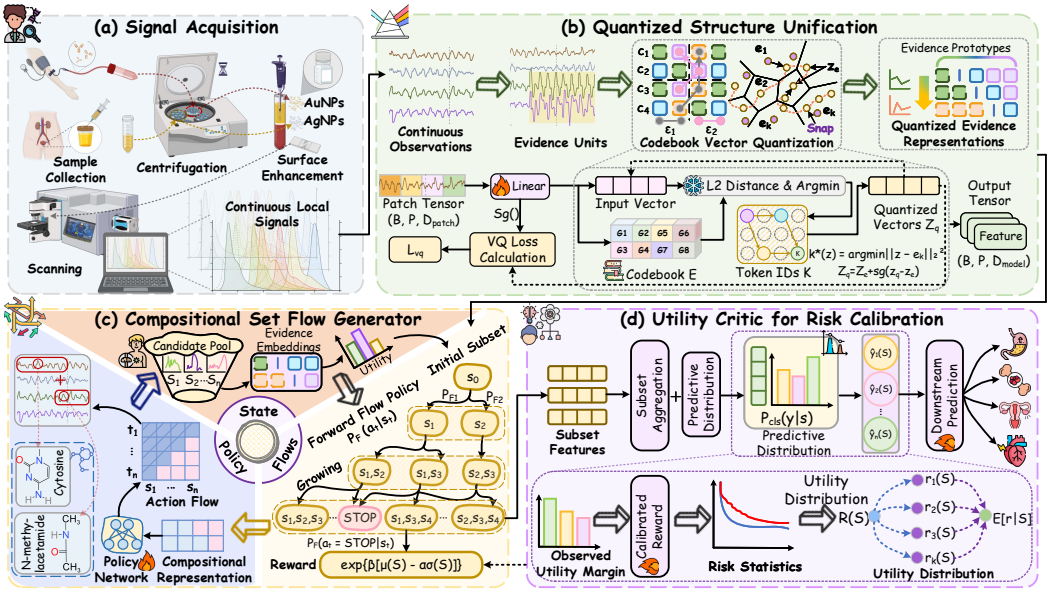


Figure 1: **The overall architecture of QSFR.** QSFR transforms heterogeneous continuous observations into reusable discrete evidence units and learns a distribution over alternative evidence compositions through sequential flow sampling. A distributional utility critic estimates the expected utility and dispersion of each composition to guide risk aware compositional sampling.

3 METHOD

In this section, we introduce the Quantized Structured Flow Representation (QSFR), as illustrated in Fig. 1. Given a high dimensional continuous observation $x \in \mathbb{R}^{C\times L}$, we seek to represent its prediction relevant information without collapsing all local observations into a single deterministic feature. Instead, QSFR learns a conditional probability law over compositional evidence,

$$\Pi(S \mid x), \qquad S \in \mathcal{S}_x, \tag{1}$$

where $\mathcal{S}_x$ is the admissible family of local evidence subsets. The central principle is that multiple evidence compositions may be simultaneously useful, while their reliability can differ substantially.

3.1 COMPOSITIONAL EVIDENCE REPRESENTATION

We decompose x into M local observations and encode them as $z_i = E_\theta(p_i) \in \mathbb{R}^d$. Continuous local representations often contain substantial redundancy. We therefore introduce a learned evidence alphabet $\mathcal{C} = \{c_k\}_{k=1}^K$ and define

$$\kappa_i = \arg\min_k \|z_i - c_k\|_2^2, \qquad q_i = c_{\kappa_i}, \qquad \bar{z}_i = z_i + \mathrm{sg}(q_i - z_i). \tag{2}$$

This operation, termed Quantized Structure Unification (QSU), maps continuously varying local observations onto reusable evidence identities. Its quantization objective is

$$\mathcal{L}_q = \frac{1}{M}\sum_{i=1}^{M}\left[\|\,\mathrm{sg}(z_i) - q_i\|_2^2 + \lambda_c\|z_i - \mathrm{sg}(q_i)\|_2^2\right]. \tag{3}$$

The quantized local representations are further encoded into evidence atoms $e_{1:M} = \mathcal{R}_\omega(\bar{z}_{1:M})$, where $\mathcal{R}_\omega$ models contextual interactions among the quantized local representations. These evidence atoms define the candidate space for compositional sampling.

For a partial subset S_t, we define its structural state as

$$g_{S_t} = \mathcal{A}_\zeta\left(\mathrm{Sel}_{S_t}(e_{1:M})\right), \tag{4}$$

where Sel_{S_t} selects the evidence atoms indexed by S_t. Rather than optimizing one subset directly, QSFR constructs a sequence $S_0 = \varnothing, S_1, \ldots, S_T$ through the Compositional Set Flow Generator (CSFG). For a trajectory $\tau = (S_0, a_0, \ldots, S_T)$,

$$P_\psi^F(\tau \mid x) = \prod_{t=0}^{T-1} P_\psi^F(a_t \mid S_t, x). \tag{5}$$

This construction converts combinatorial evidence selection into distribution learning over terminal subsets. For each terminal subset S, we aggregate the selected evidence atoms into a contextual subset representation

$$h_S = \mathcal{A}_\rho\left(\mathrm{Sel}_S(e_{1:M})\right), \tag{6}$$

where $\mathcal{A}_\rho$ denotes the subset aggregation operator. Since $e_{1:M}$ already contains contextual interactions produced by $\mathcal{R}_\omega$, h_S summarizes the selected composition without introducing an additional representation pathway. The resulting h_S is used for predictive and utility estimation.

3.2 Risk Aware Distributional Utility

A high utility evidence composition is not necessarily reliable. QSFR therefore models its utility as a distribution. The Utility Critic for Risk Calibration (UCRC) predicts $p_\phi(\xi_j \mid h_S)$ over fixed utility atoms $\{\xi_j\}_{j=1}^J$. Its first two moments are

$$\mu_\phi(S) = \sum_j p_\phi(\xi_j \mid h_S)\xi_j, \qquad \sigma_\phi^2(S) = \sum_j p_\phi(\xi_j \mid h_S)[\xi_j - \mu_\phi(S)]^2. \tag{7}$$

Here, $\sigma_\phi(S)$ measures the dispersion of the learned conditional utility distribution. We use this dispersion as a relative evidence risk surrogate. It is not interpreted as an identifiable estimate of epistemic or aleatoric uncertainty. We define the risk adjusted compositional potential

$$\mathcal{U}_{\phi,\alpha}(S) = \mu_\phi(S) - \alpha\sigma_\phi(S). \tag{8}$$

where $\alpha \geq 0$ controls sensitivity to utility dispersion and $\beta > 0$ controls distribution concentration.

The only positive terminal reward used by the compositional flow is

$$R_{\phi,\alpha}(S) = \exp\left\{\beta[\mu_\phi(S) - \alpha\sigma_\phi(S)]\right\}. \tag{9}$$

It induces the central QSFR evidence distribution

$$\Pi_{\phi,\alpha}(S \mid x) = \frac{\exp\{\beta[\mu_\phi(S) - \alpha\sigma_\phi(S)]\}}{\sum_{S' \in \mathcal{S}_x} \exp\{\beta[\mu_\phi(S') - \alpha\sigma_\phi(S')]\}}. \tag{10}$$

Eq. (10) is the principal learning object of QSFR. It replaces single subset optimization with a risk aware distribution over multiple evidence compositions. The observed supervised utility used to train UCRC is distinct from Eq. (9), as detailed in Appendix B.2.

Since direct normalization over $\mathcal{S}_x$ is combinatorial, CSFG realizes Eq. (10) through sequential flow matching. Let P^B denote the normalized reverse transition. We optimize

$$\mathcal{L}_{\mathrm{TB}} = \mathbb{E}_\tau\Bigg[\log Z_\psi(x) + \sum_t \log P_\psi^F(a_t \mid S_t, x) - \sum_t \log P^B(S_t \mid S_{t+1}, x) - \mathrm{sg}[\log R_{\phi,\alpha}(S_T)]\Bigg]^2. \tag{11}$$

The stop gradient separates reward estimation from flow matching. Consequently, UCRC cannot change its predicted utility distribution merely to reduce the trajectory-balance loss: UCRC is optimized only against the supervised distributional utility target, whereas CSFG is optimized to match the resulting detached reward. This prevents a circular optimization path between the flow objective and the utility critic, although both components remain intentionally task conditioned.

3.3 Conditional Distributional Error Analysis

We provide a conditional error propagation analysis of how representation distortion, utility estimation error, and flow approximation affect the final compositional law. The quantities introduced below are analytical controls under explicit regularity and approximation assumptions.

For a fixed subset S, define the complete compositional representation map as

$$\mathcal{F}_S(v_{1:M}) = \mathcal{A}_\rho\left(\mathrm{Sel}_S\left(\mathcal{R}_\omega(v_{1:M})\right)\right). \tag{12}$$

Accordingly,

$$h_S^c = \mathcal{F}_S(z_{1:M}), \qquad h_S^q = \mathcal{F}_S(q_{1:M}). \tag{13}$$

For an observation x, let h_S^c and h_S^q denote the final contextual subset representations obtained from the continuous and quantized evidence paths, respectively, under the same downstream contextual mapping. Let the oracle risk utility be

$$u_\alpha^\star(S) = \mu^\star(h_S^c) - \alpha\sigma^\star(h_S^c), \tag{14}$$

which induces

$$\Pi_\alpha^\star(S \mid x) \propto \exp\{\beta u_\alpha^\star(S)\}. \tag{15}$$

QSFR instead uses the quantized representation h_S^q and the learned potential

$$\widehat{u}_{\phi,\alpha}(S) = \mu_\phi(h_S^q) - \alpha\sigma_\phi(h_S^q), \tag{16}$$

yielding the target law $\widetilde{\Pi}_{\phi,\alpha}(S \mid x) \propto e^{\beta\widehat{u}_{\phi,\alpha}(S)}$. Let $Q_\psi(S \mid x)$ denote the terminal distribution produced by CSFG. Assume that the complete mapping from local evidence representations to the final contextual subset representation is L_F Lipschitz, the average quantization distortion is bounded by δ_q, and the oracle utility mean and utility dispersion are respectively L_μ and L_σ Lipschitz. Let the critic approximation errors be bounded by ϵ_μ and ϵ_σ. Define

$$\Delta_\alpha = \epsilon_\mu + L_\mu L_F \delta_q + \alpha\left(\epsilon_\sigma + L_\sigma L_F \delta_q\right). \tag{17}$$

Theorem 3.1 (Distributional fidelity)**.** *Under the assumptions above,*

$$\sup_{S\in\mathcal{S}_x} |\widehat{u}_{\phi,\alpha}(S) - u_\alpha^\star(S)| \le \Delta_\alpha. \tag{18}$$

If the trajectory balance residual and partition estimation error are bounded by ϵ_{TB} and ϵ_Z, respectively, then

$$\|Q_\psi - \Pi_\alpha^\star\|_{\mathrm{TV}} \le \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}} + \sqrt{\beta\Delta_\alpha}. \tag{19}$$

Theorem 3.1 gives an end-to-end error decomposition. The quantities in the bound correspond directly to the three learning stages of QSFR. Specifically, δ_q characterizes the representation distortion controlled by the quantization objective $\mathcal{L}_q$, while ϵ_μ and ϵ_σ summarize approximation errors of the distributional critic trained by $\mathcal{L}_{\mathrm{dist}}$. The flow terms ϵ_{TB} and ϵ_Z characterize the residual mismatch of trajectory balance and partition estimation optimized through $\mathcal{L}_{\mathrm{TB}}$. These quantities are analytical error controls rather than additional model parameters, and the theorem states how errors from the three stages propagate to the final compositional distribution.

We further characterize when QSFR concentrates probability on reliable evidence. Let $\mathcal{G}_x$ and $\mathcal{B}_x$ denote reliable and unreliable evidence regions. Define

$$\Delta_\mu = \left[\sup_{S\in\mathcal{B}_x} \mu^\star(S) - \inf_{S\in\mathcal{G}_x} \mu^\star(S)\right]_+, \tag{20}$$

and

$$\Delta_\sigma = \inf_{S\in\mathcal{B}_x} \sigma^\star(S) - \sup_{S\in\mathcal{G}_x} \sigma^\star(S) > 0. \tag{21}$$

Let $N_{\mathcal{G}} = |\mathcal{G}_x|$ and $N_{\mathcal{B}} = |\mathcal{B}_x|$.

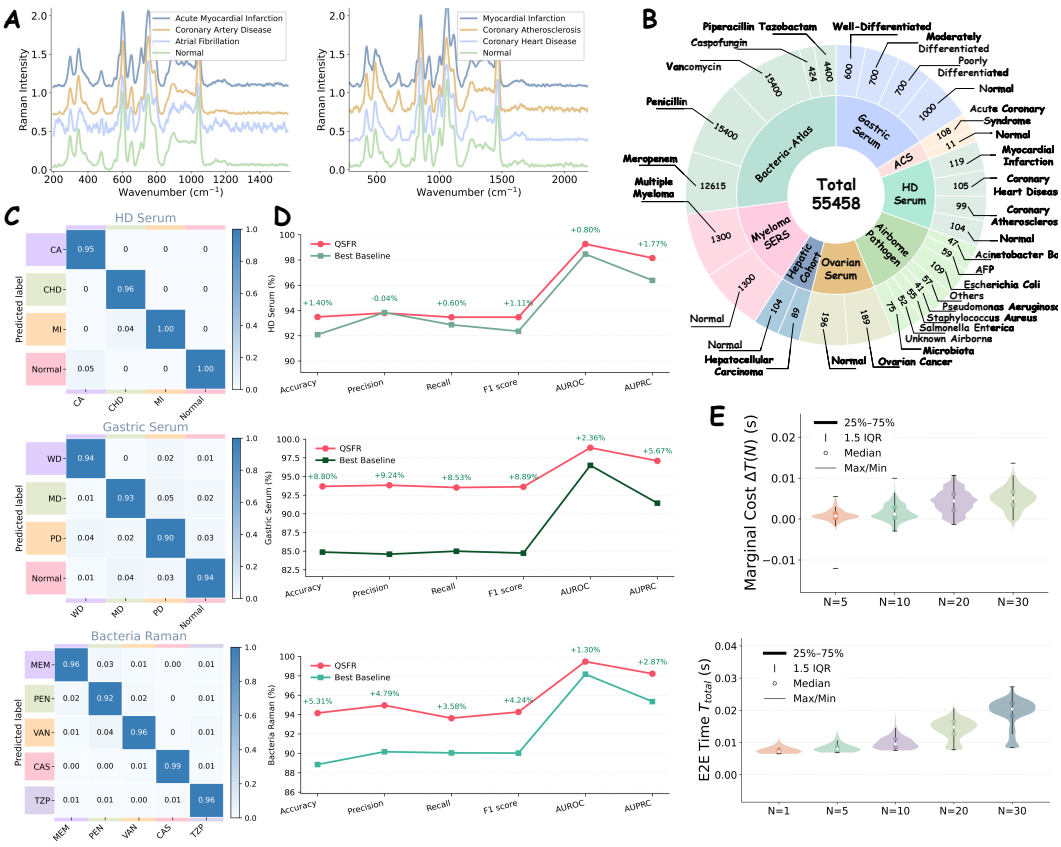


Figure 2: **Evaluating QSFR on redundant and heterogeneous observations.** (A, B) Representative spectral patterns and cohort composition. Representative observations with overlapping global structures and subtle local variations across groups. (C, D) Predictive performance across representative cohorts. (E) Computational cost under increasing observation scales.

Theorem 3.2 (Reliable evidence concentration)**.** *Define*

$$\Gamma_\alpha = \alpha\Delta_\sigma - \Delta_\mu - 2\Delta_\alpha. \tag{22}$$

If $\Gamma_\alpha > 0$*, then*

$$Q_\psi(\mathcal{B}_x \mid x) \le \frac{N_{\mathcal{B}} e^{-\beta\Gamma_\alpha}}{N_{\mathcal{G}} + N_{\mathcal{B}} e^{-\beta\Gamma_\alpha}} + \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}}. \tag{23}$$

Let

$$A_q = \epsilon_\mu + L_\mu L_F \delta_q, \qquad B_q = \epsilon_\sigma + L_\sigma L_F \delta_q.$$

If $\Delta_\sigma > 2B_q$*, a sufficient risk threshold is*

$$\alpha_c = \frac{\Delta_\mu + 2A_q}{\Delta_\sigma - 2B_q}. \tag{24}$$

For every $\alpha > \alpha_c$*, unreliable evidence mass decreases exponentially with* $\beta\Gamma_\alpha$*, up to the residual flow approximation error.*

Theorem 3.2 shows that dispersion aware risk adjustment changes the sampling distribution itself. Reliable evidence can dominate even when some unreliable states have higher expected utility, provided that their utility dispersion gap is sufficiently large. Proofs are provided in Appendix A.

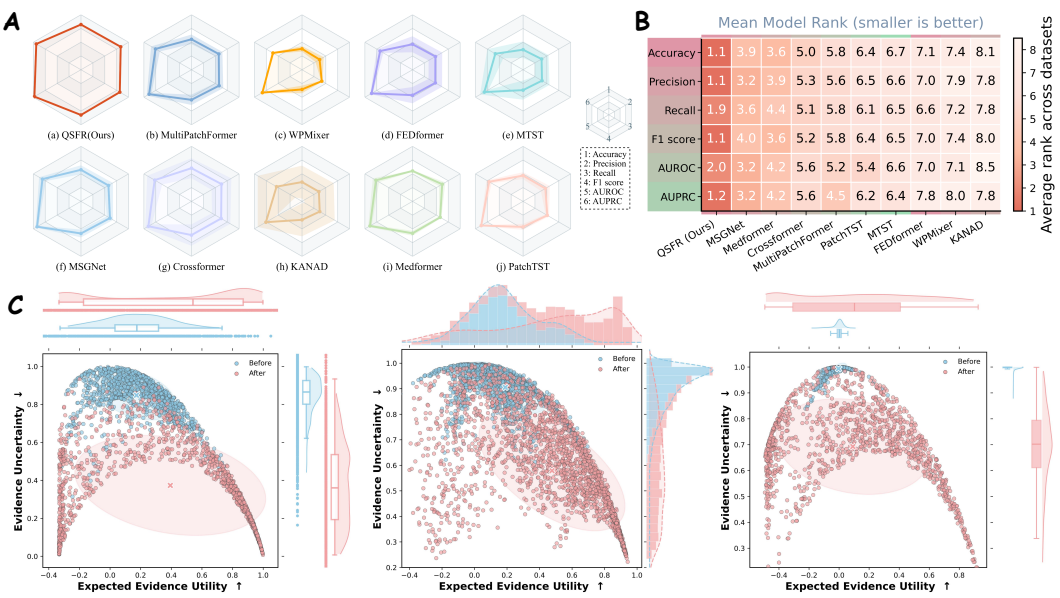


Figure 3: **From heterogeneous local observations to reliable compositional evidence.** (A) Performance profiles across complementary evaluation metrics. (B) Cross cohort model ranking. (C) Utility and utility dispersion distributions before and after QSFR modeling on Raman, ECG, and EEG. The results show reshaping toward higher utility, lower dispersion, or clearer differentiation.

4 Experiments

4.1 Experimental Setups

Evaluation Datasets. Fig. 2B summarizes the heterogeneous Raman cohorts. We analyze 55,458 continuous observations from eight independent cohorts, covering serum, urine, bacterial isolates, and airborne samples. These cohorts vary substantially in biological source, sample medium, acquisition condition, and spectral distribution, providing a test bed for representation learning under heterogeneous continuous observations. Each cohort uses a patient independent stratified 8:1:1 train, validation, and test split. Each patient's spectra are assigned to the same subset. For cohorts without patient identifiers, splitting occurs at the specimen or isolate level. Details appear in Appendix C.

Baselines and Evaluation. We evaluate QSFR against MultiPatchFormer (Naghashi et al., 2025), WPMixer (Murad et al., 2025), KANAD (Zhou et al., 2024), MTST (Zhang et al., 2024), Medformer (Wang et al., 2024b), MSGNet (Cai et al., 2024), PatchTST (Nie et al., 2023), Crossformer (Zhang & Yan, 2023), and FEDformer (Zhou et al., 2022). These models assess local structures, multiscale representations, and dependencies among spectral regions. We further include seven domain specific methods (Deng et al., 2021; Lu et al., 2024; Shin et al., 2023; Xue et al., 2025; Ibtehaz et al., 2023; Liu et al., 2023; Zhu et al., 2025) to extend the evaluation to specialized domain models. Their results are reported in Appendix F, with implementation details in Appendix D.

4.2 Compositional Evidence Learning in Heterogeneous Continuous Observations

We first examine whether QSFR can learn reliable representations when informative local observations are highly overlapping and redundant. Heterogeneous spectral observations provide a representative setting for this question. As shown in Fig. 2A, different groups share similar global profiles, while informative variations are distributed across weak and correlated local regions. Fig. 2B further shows substantial heterogeneity in sample sources and acquisition settings. We therefore use these cohorts to evaluate whether compositional evidence learning can preserve useful local structure without relying on a single dominant region.

Table 1: Representation quality across heterogeneous Raman cohorts. Purple marks QSFR (Ours). The baselines span three groups: representation reconstruction, compositional collaborative flows, and robust biomedical signal modeling.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
HD-Serum	**MultiPatchFormer**	82.33±4.25	84.55±4.64	83.71±4.16	81.44±4.56	94.60±0.89	89.31±1.81
	WPMixer	79.07±5.50	80.42±5.26	79.86±5.72	78.70±5.90	93.57±2.21	87.36±3.64
	FEDformer	86.98±5.06	87.99±4.82	87.46±5.04	87.11±5.12	95.98±3.45	91.84±3.15
	MTST	70.47±1.85	71.76±1.90	71.43±2.41	69.64±1.90	89.19±2.77	77.41±2.73
	MSGNet	89.07±1.74	89.70±1.94	89.15±1.69	88.94±1.85	98.72±0.11	96.73±0.31
	Crossformer	92.09±1.71	**93.84±1.16**	92.87±1.57	92.36±1.68	98.47±0.86	96.39±1.97
	KANAD	87.21±3.53	88.17±3.17	87.54±3.67	87.41±3.43	97.27±0.58	93.92±1.31
	Medformer	83.72±1.95	85.89±1.98	84.50±1.94	83.78±2.07	94.56±1.49	89.26±3.07
	PatchTST	83.49±4.84	85.01±4.27	84.50±4.43	83.01±5.47	96.45±1.90	92.39±2.96
	QSFR (Ours)	**93.49±2.39**	93.80±2.18	**93.47±2.37**	**93.47±2.31**	**99.27±0.49**	**98.16±1.26**
Gastric-Serum	**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
	WPMixer	79.97±0.49	79.68±0.22	79.95±0.54	79.73±0.36	94.06±0.23	85.32±0.60
	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
	MSGNet	83.57±0.65	83.57±0.62	83.60±0.70	83.57±0.66	96.52±0.18	91.39±0.46
	Crossformer	77.40±1.19	77.03±1.36	77.45±1.19	76.97±1.20	93.37±0.48	84.33±0.96
	KANAD	62.87±1.13	62.58±1.10	62.60±1.33	62.51±1.21	83.23±0.74	64.20±1.43
	Medformer	84.87±0.69	84.59±0.80	84.99±0.70	84.73±0.72	96.49±0.21	91.43±0.42
	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
	QSFR (Ours)	**93.67±0.74**	**93.83±0.68**	**93.52±0.94**	**93.62±0.79**	**98.85±0.29**	**97.10±0.56**
Bacteria-Raman	**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
	WPMixer	75.88±6.67	73.70±6.52	78.78±5.25	74.16±5.21	94.73±2.37	86.36±5.48
	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
	MSGNet	82.61±2.15	82.33±3.09	86.14±1.53	83.78±2.54	96.54±0.68	91.81±1.60
	Crossformer	88.85±1.36	90.17±0.80	90.05±0.73	90.03±0.69	98.17±0.40	95.35±0.98
	KANAD	81.77±5.44	81.45±5.95	83.91±3.45	80.34±4.89	94.07±0.49	93.03±0.82
	Medformer	84.47±1.33	89.94±0.73	86.44±0.84	87.43±0.88	97.69±0.33	94.87±0.44
	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
	QSFR (Ours)	**94.16±0.43**	**94.96±0.43**	**93.63±0.49**	**94.27±0.45**	**99.47±0.08**	**98.22±0.25**

Table 1 reports the downstream predictive quality of the learned representations on three representative Raman cohorts. QSFR achieves the best result on 17 of the 18 reported metric and cohort combinations, while remaining close to the best result on the remaining comparison. The confusion matrices in Fig. 2C show that most predictions remain concentrated on the diagonal, while Fig. 2D confirms that the advantage persists across multiple evaluation criteria. Fig. 2E further shows that increasing the observation integration scale introduces only a modest marginal aggregation cost. We next examine whether the advantage is consistent across datasets and metrics. The radar profiles in Fig. 3A show that QSFR maintains a balanced performance profile. The average rank analysis in Fig. 3B provides a dataset level view and shows favorable ranking across evaluated criteria. The statistical significance of these performance differences is examined in Appendix H.

Finally, we test the central hypothesis of QSFR: whether redundant and uncertain local observations can be reorganized into reliable compositional evidence. We consider three distinct signal domains using Gastric Serum Raman data, PTB-XL ECG (Wagner et al., 2020), and an Alzheimer's disease, frontotemporal dementia, and healthy control EEG dataset (Miltiadous et al., 2023). Fig. 3C contrasts evidence states before and after QSFR modeling. Before corresponds to the predictive states prior to compositional evidence learning, while After corresponds to the states obtained after QSFR evidence composition and risk aware probability shaping. Each observation is characterized by task adjusted evidence utility and utility dispersion. Raman and ECG both show a population shift toward stronger utility and lower utility dispersion. EEG shows a different but consistent effect, with highly ambiguous states showing lower utility dispersion and clearer differentiation after modeling. Taken together, these results show a common pattern of systematic evidence reshaping across molecular, cardiac, and neural signals. We treat ECG and EEG as mechanism probes rather than evidence of cross modality predictive generalization. Their role is to test whether the evidence-distribution behavior observed in Raman also appears qualitatively in other biomedical signal structures. This observation supports our main premise that QSFR learns to organize heterogeneous local observations according to both predictive value and reliability.

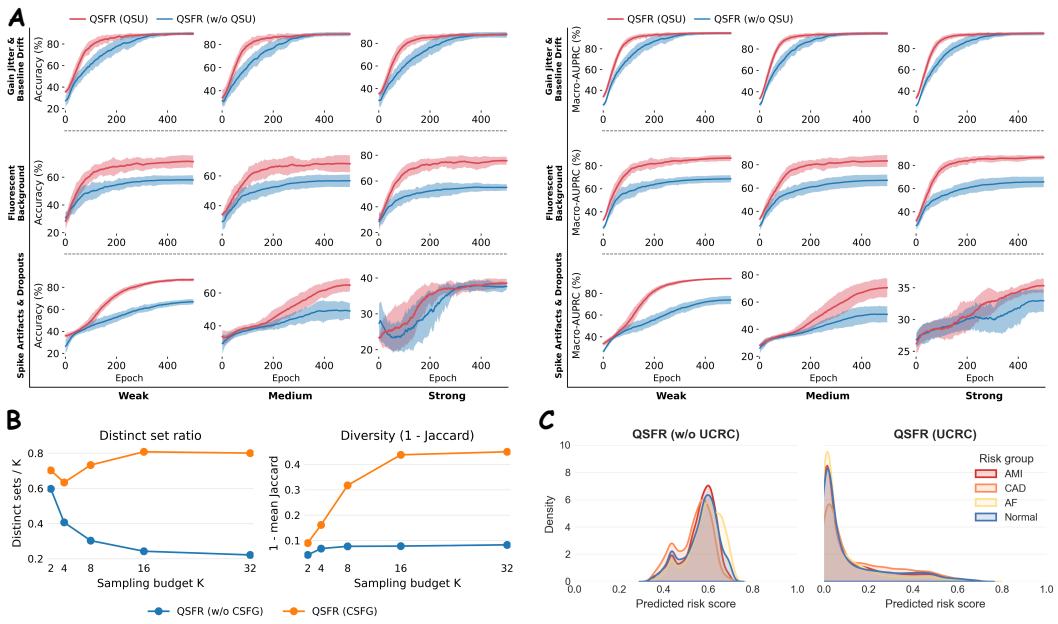


Figure 4: **From stable local evidence to reliable compositions.** (A) Discrete abstraction stabilizes local representations under structured perturbations. (B) Distributional exploration maintains multiple distinct evidence compositions across sampling budgets. (C) Reliability aware scoring separates low risk evidence from a smaller set of higher dispersion states.

4.3 Mechanism Ablation

We analyze mechanisms necessary for learning a reliable distribution over evidence compositions. Fig. 4A shows that prototype based discretization of local observations reduces representation instability under increasingly strong observation perturbations. Fig. 4B shows that flow based sampling over evidence subsets expands the effective support of the compositional distribution, yielding more distinct and less redundant compositions. Fig. 4C shows that dispersion-aware utility shaping changes the distribution of the defined relative risk score, shifting probability mass toward compositions with lower conditional utility dispersion while retaining a higher dispersion tail. These complementary effects suggest reliable compositional evidence requires stable local units, sufficient distributional support, and explicit reliability modeling. Complete analysis appears in Appendix J.

Additional Experiments. Additional experiments examine four aspects of QSFR: reliability of dispersion aware evidence scoring, robustness across extended heterogeneous cohorts, comparison with domain specific representation methods, and transfer to an independently held out target cohort. These analyses provide complementary evidence for relative risk alignment, robustness to heterogeneous observation conditions, competitiveness against domain specific methods, and transferability beyond the source cohorts. Complete results appear in Appendices I, K, F, and G.

5 Conclusion and Limitations

We formulate representation learning from heterogeneous continuous observations as risk aware compositional evidence learning. QSFR combines discrete evidence abstraction, sequential flow modeling, and distributional utility estimation to learn reliable evidence compositions. Experiments across heterogeneous Raman cohorts demonstrate its effectiveness, while ablation studies confirm the complementary roles of its components. Additional analyses on Raman, ECG, and EEG reveal related changes in evidence organization, providing mechanism evidence rather than cross modality predictive validation. Future work will explore broader modalities, stronger distribution shifts, and more efficient evidence encoders.

AI USE STATEMENT

Generative AI assisted with translation and linguistic revision.

ETHICS STATEMENT

This multicenter study followed applicable ethical standards and institutional data protection policies during data acquisition and analysis. Clinical measurements were deidentified before analysis to protect participant privacy. All clinical information was handled in accordance with institutional requirements for responsible and confidential data use.

REPRODUCIBILITY STATEMENT

We support reproducibility by providing detailed descriptions of our methodology, theoretical derivations, experimental configurations, and evaluation protocols. The manuscript and appendix also include implementation settings and supplementary experimental results to facilitate independent verification. Data and source code will be available upon approved request after publication, subject to institutional access and privacy requirements.

REFERENCES

Utku Umur Acikalin, Aaron Ferber, and Carla Gomes. Learning to explore and exploit with gnns for unsupervised combinatorial optimization. In *International Conference on Learning Representations*, volume 2025, pp. 93315–93333, 2025.

Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep evidential regression. *Advances in neural information processing systems*, 33:14927–14937, 2020.

Muhammed Fatih Balın, Abubakar Abid, and James Zou. Concrete autoencoders: Differentiable feature selection and reconstruction. In *International conference on machine learning*, pp. 444–453. PMLR, 2019.

Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in neural information processing systems*, 34:27381–27394, 2021.

Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

Wanlin Cai, Yuxuan Liang, Xianggen Liu, Jianshuai Feng, and Yuankai Wu. Msgnet: Learning multi-scale inter-series correlations for multivariate time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 11141–11149, 2024.

Bertrand Charpentier, Daniel Zügner, and Stephan Günnemann. Posterior network: Uncertainty estimation without ood samples via density-based pseudo-counts. *Advances in neural information processing systems*, 33:1356–1367, 2020.

Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *International conference on machine learning*, pp. 883–892. PMLR, 2018.

Yankai Chen, Xinni Zhang, Henry Peng Zou, Bowei He, Yangning Li, Philip Yu, Irwin King, and Xue Liu. Adversarial encoding perturbation and synthesis for set representation auxiliary learning. In *International Conference on Learning Representations*, volume 2026, pp. 18462–18487, 2026.

Ian Connick Covert, Wei Qiu, Mingyu Lu, Na Yoon Kim, Nathan J White, and Su-In Lee. Learning to maximize mutual information for dynamic feature selection. In *International Conference on Machine Learning*, pp. 6424–6447. PMLR, 2023.

Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*, pp. 1096–1105. PMLR, 2018.

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.

Jiaxiang Dong, Haixu Wu, Haoran Zhang, Li Zhang, Jianmin Wang, and Mingsheng Long. Simmtm: A simple pre-training framework for masked time-series modeling. *Advances in Neural Information Processing Systems*, 36:29996–30025, 2023.

Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski. Moment: A family of open time-series foundation models. *arXiv preprint arXiv:2402.03885*, 2024.

Dan Greenstein, Elazar Gershuni, Ilan Ben-Bassat, Yaroslav Fyodorov, Ran Moshe, Fiana Raiber, Alex Shtoff, Oren Somekh, and Nadav Hallak. A stochastic approach to the subset selection problem via mirror descent. In *International Conference on Learning Representations*, volume 2025, pp. 32394–32437, 2025.

Cheng He, Xu Huang, Gangwei Jiang, Zhaoyi Li, Defu Lian, Hong Xie, Enhong Chen, Patrick PC Lee, et al. Gtm: A general time-series model for enhanced representation learning of time-series data. In *International Conference on Learning Representations*, volume 2026, pp. 7280–7309, 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.

Nabil Ibtehaz, Muhammad EH Chowdhury, Amith Khandakar, Serkan Kiranyaz, M Sohel Rahman, and Susu M Zughaier. Ramannet: a generalized neural network architecture for raman spectrum analysis. *Neural Computing and Applications*, 35(25):18719–18735, 2023.

Yalong Jiang. Local patterns generalize better for novel anomalies. In *International Conference on Learning Representations*, volume 2025, pp. 49969–49995, 2025.

Tyler Kastner, Murat A Erdogdu, and Amir-massoud Farahmand. Distributional model equivalence for risk-sensitive reinforcement learning. *Advances in Neural Information Processing Systems*, 36:56531–56552, 2023.

Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pp. 3744–3753. PMLR, 2019.

Weida Li, Yaoliang Yu, and Bryan Kian Hsiang Low. Treegrad-ranker: feature ranking via $o(l)$-time gradients for decision trees. *arXiv preprint arXiv:2602.11623*, 2026.

Shiau Hong Lim and Ilyas Malik. Distributional reinforcement learning for risk-sensitive policies. *Advances in Neural Information Processing Systems*, 35:30977–30989, 2022.

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, 2023.

Jeremiah Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax Weiss, and Balaji Lakshminarayanan. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *Advances in neural information processing systems*, 33:7498–7512, 2020.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. In *International conference on learning representations*, volume 2024, pp. 11116–11140, 2024.

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.

Nikolay Malkin, Moksh Jain, Emmanuel Bengio, Chen Sun, and Yoshua Bengio. Trajectory balance: Improved credit assignment in gflownets. *Advances in Neural Information Processing Systems*, 35:5955–5967, 2022.

Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: Vq-vae made simple. In *International Conference on Learning Representations*, volume 2024, pp. 51772–51783, 2024.

Andreas Miltiadous, Katerina D Tzimourta, Theodora Afrantou, Panagiotis Ioannidis, Nikolaos Grigoriadis, Dimitrios G Tsalikakis, Pantelis Angelidis, Markos G Tsipouras, Euripidis Glavas, Nikolaos Giannakeas, et al. A dataset of scalp eeg recordings of alzheimer's disease, frontotemporal dementia and healthy subjects from routine eeg. *Data*, 8(6):95, 2023.

Md Mahmuddun Nabi Murad, Mehmet Aktukmak, and Yasin Yilmaz. Wpmixer: Efficient multiresolution mixing for long-term time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 19581–19588, 2025.

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

Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*, 2023.

Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.

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.

Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.

Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31, 2018.

Hyunku Shin, Byeong Hyeon Choi, On Shim, Jihee Kim, Yong Park, Suk Ki Cho, Hyun Koo Kim, and Yeonho Choi. Single test-based diagnosis of multiple cancer types using exosome-sers-ai for early stage cancers. *Nature communications*, 14(1):1644, 2023.

Tomas Tokar and Scott Sanner. Improving set function approximation with quasi-arithmetic neural networks. *arXiv preprint arXiv:2602.04941*, 2026.

Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):154, 2020.

Di Wang, Meiqi Hu, Yao Jin, Yuchun Miao, Jiaqi Yang, Yichu Xu, Xiaolei Qin, Jiaqi Ma, Lingyu Sun, Chenxing Li, et al. Hypersigma: Hyperspectral intelligence comprehension foundation model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Zhang, and Jun Zhou. Timemixer: Decomposable multiscale mixing for time series forecasting. In *International conference on learning representations*, volume 2024, pp. 38626–38652, 2024a.

Yihe Wang, Nan Huang, Taida Li, Yujun Yan, and Xiang Zhang. Medformer: A multi-granularity patching transformer for medical time-series classification. *Advances in Neural Information Processing Systems*, 37:36314–36341, 2024b.

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.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186*, 2022.

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.

Yutaro Yamada, Ofir Lindenbaum, Sahand Negahban, and Yuval Kluger. Feature selection using stochastic gates. In *International conference on machine learning*, pp. 10648–10659. PMLR, 2020.

Derek Yang, Li Zhao, Zichuan Lin, Tao Qin, Jiang Bian, and Tie-Yan Liu. Fully parameterized quantile function for distributional reinforcement learning. *Advances in neural information processing systems*, 32, 2019.

Jinsung Yoon, James Jordon, and Mihaela Van der Schaar. Invase: Instance-wise variable selection using neural networks. In *International conference on learning representations*, 2018.

Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021.

Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 8980–8987, 2022.

Dinghuai Zhang, Nikolay Malkin, Zhen Liu, Alexandra Volokhova, Aaron Courville, and Yoshua Bengio. Generative flow networks for discrete probabilistic modeling. In *International Conference on Machine Learning*, pp. 26412–26428. PMLR, 2022.

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, 2024.

Yudong Zhang, Xuan Yu, Xu Wang, Zhaoyang Sun, Chen Zhang, Pengkun Wang, and Yang Wang. Coflownet: Conservative constraints on flows enable high-quality candidate generation. In *International Conference on Learning Representations*, volume 2025, pp. 9427–9441, 2025.

Yunhao Zhang and Junchi Yan. Crossformer: Transformer utilizing cross-dimension dependency for multivariate time series forecasting. In *The Eleventh International Conference on Learning Representations*, 2023.

Quan Zhou, Changhua Pei, Fei Sun, Jing Han, Zhengwei Gao, Dan Pei, Haiming Zhang, Gaogang Xie, and Jianhui Li. Kan-ad: time series anomaly detection with kolmogorov-arnold networks. *arXiv preprint arXiv:2411.00278*, 2024.

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, 2022.

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.

Andrey Zyubin, Elizaveta Demishkevich, Karina Matveeva, Yana Byichkova, Anna Tsibulnicova, Vladimir Rafalskiy, Alina Tsapkova, and Ilia Samusev. Dataset of human platelets sers spectra from individuals with acute coronary syndrome. *Data in Brief*, 54:110259, 2024.

A Proofs of the End-to-End Distributional Guarantees

A.1 Preliminaries

For a fixed observation x, let $\mathcal{S}_x$ denote the finite family of admissible terminal evidence compositions.

We first make the representation mapping explicit. For any local representation sequence $v_{1:M} = \{v_i\}_{i=1}^M$ and terminal subset $S \in \mathcal{S}_x$, define

$$\mathcal{F}_S(v_{1:M}) = \mathcal{A}_\rho\left(\mathrm{Sel}_S\left(\mathcal{R}_\omega(v_{1:M})\right)\right), \tag{25}$$

where $\mathcal{R}_\omega$ produces contextual evidence atoms, Sel_S retains the atoms indexed by S, and $\mathcal{A}_\rho$ aggregates the selected evidence into a terminal subset representation.

Accordingly, the continuous and quantized representation paths are

$$h_S^c = \mathcal{F}_S(z_{1:M}), \qquad h_S^q = \mathcal{F}_S(q_{1:M}). \tag{26}$$

We assume that $\mathcal{F}_S$ is uniformly L_F Lipschitz over $S \in \mathcal{S}_x$ with respect to the average local representation distance. These assumptions are used to characterize stability rather than to introduce quantities that must be estimated during training. In particular, L_F, L_μ, and L_σ describe regularity of the representation and utility mappings, while δ_q, ϵ_μ, ϵ_σ, and the trajectory balance terms describe the three approximation sources introduced by quantization, critic learning, and sequential flow realization. Specifically,

$$\|\mathcal{F}_S(v_{1:M}) - \mathcal{F}_S(v'_{1:M})\|_2 \le L_F \frac{1}{M}\sum_{i=1}^M \|v_i - v'_i\|_2. \tag{27}$$

The average quantization distortion is assumed to satisfy

$$\frac{1}{M}\sum_{i=1}^M \|z_i - q_i\|_2 \le \delta_q. \tag{28}$$

We distinguish three distributions. The oracle law is

$$\Pi_\alpha^\star(S \mid x) = \frac{e^{\beta u_\alpha^\star(S)}}{Z_\alpha^\star}. \tag{29}$$

where

$$u_\alpha^\star(S) = \mu^\star(h_S^c) - \alpha\sigma^\star(h_S^c). \tag{30}$$

The learned reward induced law is

$$\widetilde{\Pi}_{\phi,\alpha}(S \mid x) = \frac{e^{\beta \widehat{u}_{\phi,\alpha}(S)}}{\widehat{Z}_{\phi,\alpha}}, \tag{31}$$

with

$$\widehat{u}_{\phi,\alpha}(S) = \mu_\phi(h_S^q) - \alpha\sigma_\phi(h_S^q). \tag{32}$$

Finally, $Q_\psi(S \mid x)$ denotes the marginal terminal distribution produced by the learned forward policy.

A.2 Proof of Theorem 3.1

We first propagate the local quantization error through the complete compositional representation.

For every fixed $S \in \mathcal{S}_x$, Eqs. (25) and (26) give

$$\begin{aligned}\|h_S^q - h_S^c\|_2 &= \|\mathcal{F}_S(q_{1:M}) - \mathcal{F}_S(z_{1:M})\|_2 \\ &\le L_F \frac{1}{M}\sum_{i=1}^M \|q_i - z_i\|_2 \\ &\le L_F \delta_q. \end{aligned} \tag{33}$$

We now decompose the learned mean error. By adding and subtracting $\mu^\star(h_S^q)$,

$$\begin{aligned}
&|\mu_\phi(h_S^q) - \mu^\star(h_S^c)| \\
&= |\mu_\phi(h_S^q) - \mu^\star(h_S^q) + \mu^\star(h_S^q) - \mu^\star(h_S^c)| \\
&\le |\mu_\phi(h_S^q) - \mu^\star(h_S^q)| + |\mu^\star(h_S^q) - \mu^\star(h_S^c)| \\
&\le \epsilon_\mu + L_\mu \|h_S^q - h_S^c\|_2 \\
&\le \epsilon_\mu + L_\mu L_F \delta_q \\
&= A_q. \qquad (34)
\end{aligned}$$

Using the same argument for the utility dispersion term,

$$\begin{aligned}
&|\sigma_\phi(h_S^q) - \sigma^\star(h_S^c)| \\
&\le |\sigma_\phi(h_S^q) - \sigma^\star(h_S^q)| + |\sigma^\star(h_S^q) - \sigma^\star(h_S^c)| \\
&\le \epsilon_\sigma + L_\sigma \|h_S^q - h_S^c\|_2 \\
&\le \epsilon_\sigma + L_\sigma L_F \delta_q \\
&= B_q. \qquad (35)
\end{aligned}$$

Therefore,

$$\begin{aligned}
&|\widehat{u}_{\phi,\alpha}(S) - u_\alpha^\star(S)| \\
&= |\mu_\phi(h_S^q) - \mu^\star(h_S^c) - \alpha\,[\sigma_\phi(h_S^q) - \sigma^\star(h_S^c)]| \\
&\le |\mu_\phi(h_S^q) - \mu^\star(h_S^c)| + \alpha\,|\sigma_\phi(h_S^q) - \sigma^\star(h_S^c)| \\
&\le A_q + \alpha B_q \\
&= \Delta_\alpha. \qquad (36)
\end{aligned}$$

This proves Eq. (18).

We next transfer this perturbation from the potential to the normalized distribution. Eq. (36) implies

$$-\Delta_\alpha \le \widehat{u}_{\phi,\alpha}(S) - u_\alpha^\star(S) \le \Delta_\alpha. \qquad (37)$$

Multiplying by β and exponentiating gives

$$\begin{aligned}
&e^{-\beta\Delta_\alpha} e^{\beta u_\alpha^\star(S)} \\
&\qquad \le e^{\beta \widehat{u}_{\phi,\alpha}(S)} \\
&\qquad \le e^{\beta\Delta_\alpha} e^{\beta u_\alpha^\star(S)}. \qquad (38)
\end{aligned}$$

Summing over all $S \in \mathcal{S}_x$ yields

$$e^{-\beta\Delta_\alpha} Z_\alpha^\star \le \widehat{Z}_{\phi,\alpha} \le e^{\beta\Delta_\alpha} Z_\alpha^\star. \qquad (39)$$

Hence,

$$\left| \log \frac{\widehat{Z}_{\phi,\alpha}}{Z_\alpha^\star} \right| \le \beta\Delta_\alpha. \qquad (40)$$

For every terminal state,

$$\begin{aligned}
&\log \frac{\widetilde{\Pi}_{\phi,\alpha}(S \mid x)}{\Pi_\alpha^\star(S \mid x)} \\
&= \log \frac{e^{\beta \widehat{u}_{\phi,\alpha}(S)}}{\widehat{Z}_{\phi,\alpha}} - \log \frac{e^{\beta u_\alpha^\star(S)}}{Z_\alpha^\star} \\
&= \beta\,[\widehat{u}_{\phi,\alpha}(S) - u_\alpha^\star(S)] + \log \frac{Z_\alpha^\star}{\widehat{Z}_{\phi,\alpha}}. \qquad (41)
\end{aligned}$$

Thus,

$$\begin{aligned}
&\left|\log \frac{\widetilde{\Pi}_{\phi,\alpha}(S \mid x)}{\Pi^{\star}_{\alpha}(S \mid x)}\right| \\
&\leq \beta \left|\widehat{u}_{\phi,\alpha}(S) - u^{\star}_{\alpha}(S)\right| + \left|\log \frac{Z^{\star}_{\alpha}}{\widehat{Z}_{\phi,\alpha}}\right| \\
&\leq 2\beta\Delta_{\alpha}. \qquad (42)
\end{aligned}$$

Consequently,

$$\begin{aligned}
&D_{\mathrm{KL}}\left(\widetilde{\Pi}_{\phi,\alpha} \| \Pi^{\star}_{\alpha}\right) \\
&= \sum_{S \in \mathcal{S}_x} \widetilde{\Pi}_{\phi,\alpha}(S \mid x) \log \frac{\widetilde{\Pi}_{\phi,\alpha}(S \mid x)}{\Pi^{\star}_{\alpha}(S \mid x)} \\
&\leq 2\beta\Delta_{\alpha} \sum_{S \in \mathcal{S}_x} \widetilde{\Pi}_{\phi,\alpha}(S \mid x) \\
&= 2\beta\Delta_{\alpha}. \qquad (43)
\end{aligned}$$

Pinsker's inequality then gives

$$\left\|\widetilde{\Pi}_{\phi,\alpha} - \Pi^{\star}_{\alpha}\right\|_{\mathrm{TV}} \leq \sqrt{\beta\Delta_{\alpha}}. \qquad (44)$$

We now incorporate the sequential flow approximation. For a complete trajectory $\tau = (S_0, a_0, \ldots, S_T)$, define the normalized backward trajectory probability as

$$P^{B}(\tau^{-} \mid S_T, x) = \prod_{t=0}^{T-1} P^{B}(S_t \mid S_{t+1}, x). \qquad (45)$$

Define the target trajectory distribution

$$\widetilde{P}(\tau \mid x) = \widetilde{\Pi}_{\phi,\alpha}(S_T \mid x) P^{B}(\tau^{-} \mid S_T, x). \qquad (46)$$

Because the backward trajectory probability is normalized for every terminal state,

$$\begin{aligned}
\sum_{\tau} \widetilde{P}(\tau \mid x) &= \sum_{S \in \mathcal{S}_x} \widetilde{\Pi}_{\phi,\alpha}(S \mid x) \sum_{\tau : S_T = S} P^{B}(\tau^{-} \mid S, x) \\
&= \sum_{S \in \mathcal{S}_x} \widetilde{\Pi}_{\phi,\alpha}(S \mid x) \\
&= 1. \qquad (47)
\end{aligned}$$

Hence $\widetilde{P}$ is a valid trajectory distribution.

Since

$$R_{\phi,\alpha}(S) = e^{\beta \widehat{u}_{\phi,\alpha}(S)}, \qquad (48)$$

and

$$\widetilde{\Pi}_{\phi,\alpha}(S \mid x) = \frac{R_{\phi,\alpha}(S)}{\widehat{Z}_{\phi,\alpha}}, \qquad (49)$$

define the trajectory balance residual as

$$\begin{aligned}
\rho_{\psi}(\tau) = &\log Z_{\psi}(x) + \sum_{t=0}^{T-1} \log P^{F}_{\psi}(a_t \mid S_t, x) \\
&- \sum_{t=0}^{T-1} \log P^{B}(S_t \mid S_{t+1}, x) - \log R_{\phi,\alpha}(S_T). \qquad (50)
\end{aligned}$$

Assume

$$\left(\mathbb{E}_{\tau\sim P_\psi^F}[\rho_\psi(\tau)^2]\right)^{1/2} \le \epsilon_{\mathrm{TB}}, \tag{51}$$

and

$$\left|\log Z_\psi(x) - \log \widehat{Z}_{\phi,\alpha}\right| \le \epsilon_Z. \tag{52}$$

Let

$$\eta_Z = \log Z_\psi(x) - \log \widehat{Z}_{\phi,\alpha}. \tag{53}$$

Using

$$P_\psi^F(\tau \mid x) = \prod_{t=0}^{T-1} P_\psi^F(a_t \mid S_t, x),$$

we obtain

$$\begin{aligned}
&\log \frac{P_\psi^F(\tau \mid x)}{\widetilde{P}(\tau \mid x)} \\
&= \log P_\psi^F(\tau \mid x) - \log R_{\phi,\alpha}(S_T) + \log \widehat{Z}_{\phi,\alpha} - \log P^B(\tau^- \mid S_T, x) \\
&= \rho_\psi(\tau) - \eta_Z.
\end{aligned} \tag{54}$$

Therefore,

$$\begin{aligned}
&D_{\mathrm{KL}}\left(P_\psi^F \| \widetilde{P}\right) \\
&= \mathbb{E}_{\tau\sim P_\psi^F}\left[\log \frac{P_\psi^F(\tau \mid x)}{\widetilde{P}(\tau \mid x)}\right] \\
&= \mathbb{E}_{P_\psi^F}[\rho_\psi(\tau)] - \eta_Z \\
&\le \left|\mathbb{E}_{P_\psi^F}[\rho_\psi(\tau)]\right| + |\eta_Z| \\
&\le \left(\mathbb{E}_{P_\psi^F}[\rho_\psi(\tau)^2]\right)^{1/2} + \epsilon_Z \\
&\le \epsilon_{\mathrm{TB}} + \epsilon_Z.
\end{aligned} \tag{55}$$

The penultimate inequality follows from the Cauchy Schwarz inequality.

The terminal state is a deterministic measurable function of the complete trajectory. Therefore, the data processing inequality for KL divergence yields

$$D_{\mathrm{KL}}\left(Q_\psi \| \widetilde{\Pi}_{\phi,\alpha}\right) \le D_{\mathrm{KL}}\left(P_\psi^F \| \widetilde{P}\right) \le \epsilon_{\mathrm{TB}} + \epsilon_Z. \tag{56}$$

Applying Pinsker's inequality again,

$$\left\|Q_\psi - \widetilde{\Pi}_{\phi,\alpha}\right\|_{\mathrm{TV}} \le \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}}. \tag{57}$$

Finally, the triangle inequality for total variation gives

$$\begin{aligned}
&\|Q_\psi - \Pi_\alpha^\star\|_{\mathrm{TV}} \\
&\le \left\|Q_\psi - \widetilde{\Pi}_{\phi,\alpha}\right\|_{\mathrm{TV}} + \left\|\widetilde{\Pi}_{\phi,\alpha} - \Pi_\alpha^\star\right\|_{\mathrm{TV}} \\
&\le \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}} + \sqrt{\beta \Delta_\alpha}.
\end{aligned} \tag{58}$$

This completes the proof of Theorem 3.1. The resulting bound should be interpreted as an error propagation guarantee under the stated regularity and approximation assumptions. Its role is to characterize how representation, critic, and flow errors jointly affect the learned compositional distribution, rather than to provide a dataset specific numerical certificate.

Corollary A.1. For any bounded statistic $f : \mathcal{S}_x \to \mathbb{R}$ with $\|f\|_\infty \le M_f$,

$$\begin{aligned}&\left|\mathbb{E}_{Q_\psi}[f(S)] - \mathbb{E}_{\Pi^\star_\alpha}[f(S)]\right| \\ &\le 2M_f \left\|Q_\psi - \Pi^\star_\alpha\right\|_{\mathrm{TV}} \\ &\le 2M_f \left[\sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}} + \sqrt{\beta\Delta_\alpha}\right].\end{aligned} \tag{59}$$

Thus, the same end to end error controls every bounded statistic of the sampled compositional evidence distribution.

A.3 Proof of Theorem 3.2

We next analyze how the learned distribution allocates probability between reliable and unreliable evidence compositions.

Let $\mathcal{G}_x$ and $\mathcal{B}_x$ form a disjoint partition of the admissible evidence space,

$$\mathcal{S}_x = \mathcal{G}_x \,\dot\cup\, \mathcal{B}_x. \tag{60}$$

Here, reliability is defined operationally with respect to the oracle conditional utility distribution: $\mathcal{B}_x$ denotes evidence states with relatively larger utility dispersion. It does not refer to an externally identifiable epistemic or aleatoric uncertainty label.

For compactness, define

$$\mu^\star(S) = \mu^\star(h^c_S), \qquad \sigma^\star(S) = \sigma^\star(h^c_S). \tag{61}$$

From Eqs. (34) and (35), for every $S \in \mathcal{S}_x$,

$$|\widehat{u}_{\phi,\alpha}(S) - u^\star_\alpha(S)| \le \Delta_\alpha. \tag{62}$$

Consider any $S_g \in \mathcal{G}_x$ and $S_b \in \mathcal{B}_x$. Then

$$\begin{aligned}&\widehat{u}_{\phi,\alpha}(S_g) - \widehat{u}_{\phi,\alpha}(S_b) \\ &\ge u^\star_\alpha(S_g) - u^\star_\alpha(S_b) - 2\Delta_\alpha \\ &= \mu^\star(S_g) - \mu^\star(S_b) - \alpha\left[\sigma^\star(S_g) - \sigma^\star(S_b)\right] - 2\Delta_\alpha.\end{aligned} \tag{63}$$

By the definition of Δ_μ,

$$\mu^\star(S_g) - \mu^\star(S_b) \ge -\Delta_\mu. \tag{64}$$

By the definition of Δ_σ,

$$\sigma^\star(S_b) - \sigma^\star(S_g) \ge \Delta_\sigma, \tag{65}$$

or equivalently

$$-\alpha\left[\sigma^\star(S_g) - \sigma^\star(S_b)\right] \ge \alpha\Delta_\sigma. \tag{66}$$

Substituting Eqs. (64) and (66) into Eq. (63) gives

$$\begin{aligned}&\widehat{u}_{\phi,\alpha}(S_g) - \widehat{u}_{\phi,\alpha}(S_b) \\ &\ge \alpha\Delta_\sigma - \Delta_\mu - 2\Delta_\alpha \\ &= \Gamma_\alpha.\end{aligned} \tag{67}$$

Since this inequality holds for every pair $(S_g, S_b) \in \mathcal{G}_x \times \mathcal{B}_x$,

$$\inf_{S\in\mathcal{G}_x} \widehat{u}_{\phi,\alpha}(S) - \sup_{S\in\mathcal{B}_x} \widehat{u}_{\phi,\alpha}(S) \ge \Gamma_\alpha. \tag{68}$$

Let

$$W_{\mathcal{G}} = \sum_{S\in\mathcal{G}_x} e^{\beta\widehat{u}_{\phi,\alpha}(S)}, \tag{69}$$

$$W_{\mathcal{B}} = \sum_{S\in\mathcal{B}_x} e^{\beta\widehat{u}_{\phi,\alpha}(S)}. \tag{70}$$

Using Eq. (68),

$$W_{\mathcal{G}} \geq \sum_{S \in \mathcal{G}_x} \exp\left\{\beta \inf_{S' \in \mathcal{G}_x} \widehat{u}_{\phi,\alpha}(S')\right\} = N_{\mathcal{G}} \exp\left\{\beta \inf_{S' \in \mathcal{G}_x} \widehat{u}_{\phi,\alpha}(S')\right\}, \tag{71}$$

whereas

$$W_{\mathcal{B}} \leq \sum_{S \in \mathcal{B}_x} \exp\left\{\beta \sup_{S' \in \mathcal{B}_x} \widehat{u}_{\phi,\alpha}(S')\right\} = N_{\mathcal{B}} \exp\left\{\beta \sup_{S' \in \mathcal{B}_x} \widehat{u}_{\phi,\alpha}(S')\right\}. \tag{72}$$

Therefore,

$$\frac{W_{\mathcal{B}}}{W_{\mathcal{G}}} \leq \frac{N_{\mathcal{B}}}{N_{\mathcal{G}}} \exp\left\{-\beta\left[\inf_{\mathcal{G}_x} \widehat{u}_{\phi,\alpha} - \sup_{\mathcal{B}_x} \widehat{u}_{\phi,\alpha}\right]\right\} \leq \frac{N_{\mathcal{B}}}{N_{\mathcal{G}}} e^{-\beta\Gamma_\alpha}. \tag{73}$$

Since $\mathcal{G}_x$ and $\mathcal{B}_x$ partition $\mathcal{S}_x$, the learned target probability of the unreliable region is

$$\begin{aligned}\widetilde{\Pi}_{\phi,\alpha}(\mathcal{B}_x \mid x) &= \frac{W_{\mathcal{B}}}{W_{\mathcal{G}} + W_{\mathcal{B}}} \\ &= \frac{W_{\mathcal{B}}/W_{\mathcal{G}}}{1 + W_{\mathcal{B}}/W_{\mathcal{G}}} \\ &\leq \frac{(N_{\mathcal{B}}/N_{\mathcal{G}})e^{-\beta\Gamma_\alpha}}{1 + (N_{\mathcal{B}}/N_{\mathcal{G}})e^{-\beta\Gamma_\alpha}} \\ &= \frac{N_{\mathcal{B}}e^{-\beta\Gamma_\alpha}}{N_{\mathcal{G}} + N_{\mathcal{B}}e^{-\beta\Gamma_\alpha}}.\end{aligned} \tag{74}$$

This establishes the concentration bound for the learned target law.

For any event $\mathcal{E} \subseteq \mathcal{S}_x$, the definition of total variation gives

$$\left|Q_\psi(\mathcal{E} \mid x) - \widetilde{\Pi}_{\phi,\alpha}(\mathcal{E} \mid x)\right| \leq \left\|Q_\psi - \widetilde{\Pi}_{\phi,\alpha}\right\|_{\mathrm{TV}}. \tag{75}$$

Setting $\mathcal{E} = \mathcal{B}_x$ and using Eq. (57),

$$\begin{aligned}Q_\psi(\mathcal{B}_x \mid x) &\leq \widetilde{\Pi}_{\phi,\alpha}(\mathcal{B}_x \mid x) + \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}} \\ &\leq \frac{N_{\mathcal{B}}e^{-\beta\Gamma_\alpha}}{N_{\mathcal{G}} + N_{\mathcal{B}}e^{-\beta\Gamma_\alpha}} + \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}}.\end{aligned} \tag{76}$$

This establishes Eq. (23).

It remains to characterize when Γ_α is strictly positive. Since

$$\Delta_\alpha = A_q + \alpha B_q, \tag{77}$$

we have

$$\begin{aligned}\Gamma_\alpha &= \alpha\Delta_\sigma - \Delta_\mu - 2(A_q + \alpha B_q) \\ &= \alpha\left(\Delta_\sigma - 2B_q\right) - \left(\Delta_\mu + 2A_q\right).\end{aligned} \tag{78}$$

Suppose

$$\Delta_\sigma - 2B_q > 0. \tag{79}$$

Then $\Gamma_\alpha > 0$ is equivalent to

$$\begin{aligned}&\alpha\left(\Delta_\sigma - 2B_q\right) > \Delta_\mu + 2A_q\\ \iff &\alpha > \frac{\Delta_\mu + 2A_q}{\Delta_\sigma - 2B_q}\\ &= \alpha_c.\end{aligned} \tag{80}$$

Hence, whenever $\alpha > \alpha_c$, the reliable and unreliable regions remain separated after accounting for quantization and critic approximation errors.

Substituting Eq. (78) into Eq. (76) gives

$$\begin{aligned}&Q_\psi(\mathcal{B}_x \mid x)\\ &\le \frac{N_{\mathcal{B}} \exp\left\{-\beta\left[\alpha(\Delta_\sigma - 2B_q) - (\Delta_\mu + 2A_q)\right]\right\}}{N_{\mathcal{G}} + N_{\mathcal{B}} \exp\left\{-\beta\left[\alpha(\Delta_\sigma - 2B_q) - (\Delta_\mu + 2A_q)\right]\right\}}\\ &\quad + \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}}.\end{aligned} \tag{81}$$

This completes the proof of Theorem 3.2.

Corollary A.2. If CSFG is exactly flow matched, $\epsilon_{\mathrm{TB}} = \epsilon_Z = 0$, then for every $\alpha > \alpha_c$,

$$Q_\psi(\mathcal{B}_x \mid x) \le \frac{N_{\mathcal{B}} e^{-\beta\Gamma_\alpha}}{N_{\mathcal{G}} + N_{\mathcal{B}} e^{-\beta\Gamma_\alpha}}. \tag{82}$$

Thus,

$$\lim_{\beta\Gamma_\alpha \to \infty} Q_\psi(\mathcal{B}_x \mid x) = 0. \tag{83}$$

Corollary A.3. For arbitrary flow error, the irreducible upper limit satisfies

$$\limsup_{\beta\Gamma_\alpha \to \infty} Q_\psi(\mathcal{B}_x \mid x) \le \sqrt{\frac{\epsilon_{\mathrm{TB}} + \epsilon_Z}{2}}. \tag{84}$$

Therefore, increasing risk separation alone cannot compensate for an arbitrarily inaccurate sequential flow. The remaining flow approximation error defines the limiting floor of reliable evidence concentration.

B Additional Methodological Details

B.1 Evidence Quantization and Structural Encoding

Let the local observations be $\{p_i\}_{i=1}^M$. The encoder first produces $z_i = E_\theta(p_i)$. QSU maintains a learnable prototype dictionary $\mathcal{C} = \{c_k\}_{k=1}^K$ and performs nearest prototype assignment according to Eq. (2).

The straight through representation $\bar{z}_i = z_i + \mathrm{sg}(q_i - z_i)$ preserves gradients with respect to the continuous encoder while the forward computation uses the discrete prototype. The quantized local representations are further encoded as

$$e_{1:M} = \mathcal{R}_\omega(\bar{z}_{1:M}), \tag{85}$$

where $\mathcal{R}_\omega$ maps the quantized representations into structured evidence atoms. The resulting evidence atoms constitute the candidate space used by CSFG.

For a partial subset S, we use a permutation stable aggregation operator to obtain its structural representation

$$g_S = \mathcal{A}_\zeta\left(\mathrm{Sel}_S(e_{1:M})\right), \qquad g_\varnothing = 0. \tag{86}$$

This representation provides the state information required during compositional flow construction. Concretely, $\mathcal{A}_\zeta$ is implemented by a GRU state aggregator, while $\mathcal{A}_\rho$ uses mean pooling over the selected contextual evidence atoms. These operators are fixed across all cohorts.

B.2 Distributional Utility Supervision

For a terminal subset S, we use the same subset representation as in Eq. (6),

$$h_S = \mathcal{A}_\rho\left(\mathrm{Sel}_S(e_{1:M})\right). \tag{87}$$

The representation h_S is shared by the prediction head and UCRC.

During supervised training, the observed utility of a terminal subset is measured using the prediction margin

$$r_{\mathrm{obs}}(S) = p_{\mathrm{pred}}(y \mid h_S) - \max_{c \neq y} p_{\mathrm{pred}}(c \mid h_S). \tag{88}$$

This quantity is used only as a supervised utility target for UCRC. It is label conditioned and is not used as an inference time confidence score. UCRC does not take y, $r_{\mathrm{obs}}(S)$, or $p_{\mathrm{pred}}(\cdot \mid h_S)$ as input. Instead, it receives h_S and learns the conditional utility distribution $p_\phi(\xi \mid h_S)$.

Let $\Xi = \{\xi_1, \ldots, \xi_J\}$ be the support of the predicted utility distribution. UCRC produces

$$p_\phi(\xi_j \mid h_S) = \frac{\exp(g_\phi(h_S)_j)}{\sum_{\ell=1}^{J} \exp(g_\phi(h_S)_\ell)}. \tag{89}$$

We convert the observed scalar utility into a smooth distributional target

$$\widetilde{p}_j(S) = \frac{\exp\left[-\frac{(\xi_j - r_{\mathrm{obs}}(S))^2}{2\tau_r^2}\right]}{\sum_{\ell=1}^{J} \exp\left[-\frac{(\xi_\ell - r_{\mathrm{obs}}(S))^2}{2\tau_r^2}\right]}, \tag{90}$$

and optimize

$$\mathcal{L}_{\mathrm{dist}} = -\mathbb{E}_S \sum_{j=1}^{J} \widetilde{p}_j(S) \log p_\phi(\xi_j \mid h_S). \tag{91}$$

The fixed scale τ_r only smooths the scalar utility target. Ignoring discretization and support boundary effects, the learned variance satisfies $\mathrm{Var}\,\phi(\xi \mid h_S) \approx \mathrm{Var}(r\mathrm{obs} \mid h_S) + \tau_r^2$. Because τ_r is shared by all states, it adds a common variance floor and changes the scale but not the ordering induced by conditional utility dispersion at the population optimum. The randomness here comes from observations and sampled compositions that induce similar h_S. We therefore use $\sigma_\phi(S)$ only as a relative utility dispersion score and do not interpret it as calibrated epistemic or aleatoric uncertainty.

The resulting moments are exactly those in Eq. (7). Accordingly, three quantities have distinct roles: $r_{\mathrm{obs}}(S)$ is the supervised utility observation, $\mu_\phi(S) - \alpha\sigma_\phi(S)$ is the learned risk adjusted utility, and $R_{\phi,\alpha}(S)$ is the unique positive terminal reward used by CSFG.

Scope of the dispersion measure. Throughout this work, "risk" refers specifically to relative composition-level utility dispersion within the learned evidence space. Accordingly, $\sigma_\phi(S)$ is not interpreted as a calibrated probability of prediction failure, an OOD score, a measure of clinical decision risk, or an identifiable epistemic or aleatoric uncertainty estimate. Reliability claims therefore concern the relative ordering of candidate evidence compositions under the defined conditional-utility criterion, rather than general predictive or clinical risk calibration.

Relation to scalar utility and prediction confidence. UCRC and the prediction head share h_S, so their outputs are not assumed to be statistically independent. The label conditioned margin in Eq. (88) supervises UCRC but is not provided to it at inference. At the population optimum of the smoothed distributional objective, the first two moments approximate $\mathbb{E}[r_{\mathrm{obs}} \mid h_S]$ and $\mathrm{Var}(r_{\mathrm{obs}} \mid h_S) + \tau_r^2$, subject to discretization and model approximation. This distinction matters when two compositions have the same estimated mean but different utility dispersion. For $\mu_\phi(S_a) = \mu_\phi(S_b)$ and $\sigma_\phi(S_a) < \sigma_\phi(S_b)$, Eq. (9) yields $R_{\phi,\alpha}(S_a)/R_{\phi,\alpha}(S_b) = \exp\{\beta\alpha[\sigma_\phi(S_b) - \sigma_\phi(S_a)]\} > 1$ when $\alpha > 0$, whereas mean only scoring gives equal rewards. This establishes a difference in the reward mechanism, not an empirical advantage over a capacity matched scalar critic or an independent estimate of predictive uncertainty. Appendix I examines whether the learned dispersion tracks held out utility residuals.

C SUPPLEMENTARY DATASET DETAILS

This section complements the main text from a data perspective. It provides detailed descriptions of clinical background, acquisition settings, and characteristics of the cohorts used in this study.

(1) **Myeloma SERS**. The multiple myeloma cohort focuses on surface enhanced Raman signals from urine. During acquisition we add silver nanoparticles AgNP to the urine samples to enhance peaks associated with small molecule metabolites and protein components. We collect all data in 2025 at a tertiary grade A hospital and obtain 2600 spectra from 26 urine specimens. Based on this cohort we construct a biofluid Raman library that is used to test the ability of the model to recognize disease in non blood biofluid media.

(2) **Gastric Serum**. To study pathological grading in gastric cancer, we collect serum from gastric cancer patients at a tertiary grade A hospital in 2025 and perform surface enhanced Raman measurements with a Raman spectrometer ATR3110 from Aopu Tiancheng. All samples follow a unified SERS preparation protocol. For each sample we take $100, \mu$L of serum and mix it with $20, \mu$L of gold nanocolloid before spotting. To improve spectral stability and repeatability, each subject provides four independent serum drops that are measured under identical instrument settings with laser power fixed at 100 mW and integration time of 3s. On this basis we build a serum Raman dataset for gastric cancer grading that includes poorly differentiated PD, moderately differentiated MD, and well differentiated WD samples. In total we obtain 3000 enhanced Raman spectra, which provide a reliable source of data for analyzing spectral features related to gastric cancer differentiation.

(3) **Hepatic Cohort**. The liver cancer cohort consists of ultrafiltered serum Raman spectra from patients with hepatocellular carcinoma HCC. All samples are collected in 2021 at a tertiary grade A hospital. The cohort includes 65 patients and 193 Raman spectral sequences. We use this cohort as an independent evaluation cohort to test band recognition in tumor related biofluids and to evaluate the transferability and generalization of the model across different disease types.

(4) **HD Serum**. The heart disease HD cohort is composed of ultrafiltered serum Raman spectra that reflect spectral characteristics in a real clinical environment. The samples are collected in 2021 at a tertiary grade A hospital and include 128 patients with 427 Raman spectral sequences. Clinical subtypes include myocardial infarction MI, coronary heart disease CHD, and coronary atherosclerosis CA. This independent dataset is used to assess how well the model supports fine grained risk stratification.

(5) **Ovarian Serum**. The plasma ovarian cancer dataset evaluates the discriminative power of blood Raman spectra for ovarian cancer diagnosis (Paraskevaidi et al., 2018). The study enrolls 27 patients with ovarian cancer and 28 patients with benign gynecologic disease. We collect EDTA anticoagulated whole blood from each subject and finally obtain 385 spectra, each containing 1480 sampling points. Whole blood is centrifuged at 2000 revolutions per minute for 10 minutes to separate plasma, which is stored under low temperature conditions. For measurement we pipette 50 microliters of plasma onto an aluminum foil slide and allow it to dry, and we also mix plasma with diluted silver nanoparticles at a ratio of one to one to prepare SERS samples that contain nanoparticles.

(6) **Bacteria Raman**. The bacterial Raman dataset is built from the Bacteria ID resource and targets automatic identification of multiple clinically common bacteria and yeasts whose strains almost cover inpatient infection cases at Stanford Hospital in 2016 to 2017 (Ho et al., 2019). During sample preparation the researchers first culture target strains and clinical isolates on blood agar plates, then scrape about 0.6 mg of biomass from a single colony and suspend it in 10 microliters of sterile water. They deposit 3 microliters of the suspension onto a gold coated silicon slide to form the test region. Spectra are acquired on a Horiba LabRAM HR Evolution micro Raman platform with a 633 nm laser, power of about 13.17 mW, a 300 lines per millimeter grating, and a 100 times objective lens with numerical aperture 0.9. In mapping mode the system measures a 45×45 grid on the sample surface with a step size of about 3 micrometers. The integration time of each spectrum is 1 or 2 seconds, and the spectral range covers 382 to 1792 cm^{-1}. All spectra are processed with polynomial baseline subtraction and intensity normalization. In total the dataset contains 48239 spectra, each with 1000 sampling points, and it serves as the base data for fast bacterial identification experiments.

(7) **Platelet SERS**. The acute coronary syndrome ACS dataset consists of surface enhanced Raman spectra from patient platelets (Zyubin et al., 2024). Venous blood is collected from patients, and platelets are enriched by repeated low speed and medium speed centrifugation, after which the su-

pernatant is further centrifuged to obtain the target platelet cells. The platelets are then dropped onto a titanium substrate that has been anodized to form nanostructures, and spherical gold nanoparticles on the surface provide signal enhancement. Finally a Centaur U Raman analyzer with a 532 nm solid state laser records 119 platelet spectra in microscopic mode over the range 400 to 1800 cm^{-1}, and these spectra are used as the data source for model evaluation.

(8) **Airborne Pathogen**. The airborne pathogen single cell Raman dataset supports rapid identification of airborne pathogens in complex real air environments (Zhu et al., 2025). A microbial aerosol generation and collection system atomizes pathogen cultures into aerosols, then collects them in sterile air. We then acquire single cell Raman spectra to build a training library. The library includes five airborne pathogens: Acinetobacter baylyi AB, Escherichia coli EC, Pseudomonas aeruginosa Pae, Staphylococcus aureus SA, and Salmonella enterica SE. We additionally include an others class and an unknown class composed of nonpathogenic airborne microbes and real indoor and outdoor air samples to represent background communities, with 495 spectra in total. Raman measurements are performed on a single cell micro Raman system over 600 to 1800 cm^{-1}, and spectra are baseline corrected, intensity normalized, and discretized into about 225 feature points.

D Implementation Details and Hyperparameters

The implementation details and hyperparameter settings used in this study are listed in Table 2. During training, we use the AdamW optimizer with a learning rate of 4×10^{-5} and a batch size of 32. The hidden dimension is 128, and the feedforward dimension is 512. We train the model for a maximum of 500 epochs and set the dropout rate to 0.1. We use early stopping to terminate training when the validation Macro-F1 does not improve for 30 consecutive epochs. We adopt six evaluation metrics, including Accuracy, Precision, Recall, F1 score, AUROC, and AUPRC. For data splitting, we use a patient-independent stratified split for each dataset. We assign 80%, 10%, and 10% of patients to the training, validation, and test sets, respectively. The validation set is used only for early stopping and hyperparameter selection. All reported test statistics are computed after model selection on the held out test partition. In particular, the UCRC risk alignment statistics in Table 5 are computed on this test partition and are not used to optimize UCRC, CSFG, or the prediction head. All spectra, technical replicates, and augmented samples from the same patient are assigned to only one subset. This protocol avoids spectrum level leakage and ensures that the test set evaluates patient level generalization. For non patient cohorts, such as bacterial isolates and airborne samples, we apply the same rule using the specimen or isolate identifier. The contextual evidence encoder and subset aggregation modules use the model dimensions reported in Table 2. Training is performed with Python 3.11 and CUDA 12.0. To ensure reliability, we repeat all experiments with five random seeds and report the mean and standard deviation.

E Pseudocode of QSFR

The pseudocode for QSFR is shown in Algorithm 1. Algorithm 1 uses the same subset representation as Eq. (6) in both training and inference. UCRC affects test time prediction only through the CSFG policy learned during training and does not introduce a separate prediction fusion rule. At inference, neither the label y, the supervised utility r_{obs}, nor a sample specific confidence penalty is used. The learned CSFG policy selects the evidence composition, and the final prediction is produced by the prediction head from the selected representation.

F Comparison with Domain-Specific Representation Baselines

To examine whether the gains of QSFR depend on comparisons with general sequence models, we further evaluate seven domain-specific learning methods under the same data splits and evaluation protocol. The comparison includes SANet, DSCARNet, Exosome-SERS-AI, DSCF, RamanNet, Raman Transformer, and Raman-OSDL, covering multiscale spectral modeling, aggregated representations, spectral foundation modeling, attention based encoding, and open set learning.

Fig. 5 summarizes the results across six evaluation metrics. QSFR exhibits the broadest and most balanced performance profile, while the domain-specific baselines show larger variation across individual criteria. In particular, methods with specialized spectral representations remain competitive

Table 2: Core implementation hyperparameters of QSFR.

Hyperparameter	**Value**
Label smoothing	0.10
Weight decay	1×10^{-2}
Batch size	32
Early stopping patience	30
Model dimension d_{model}	128
Feed forward dimension d_{ff}	512
Patch length / stride	8 / 4
VQ codebook size	1024
VQ commitment cost	0.25
Peak graph encoder dimension	16
Band actor hidden dimension	128
UCRC hidden dimension	8
Reward atom number	51
Reward support range	$[-1, 1]$
Risk aversion coefficient α	0.50
Reward temperature β	1.00
QSU loss weight λ_q	0.10
UCRC loss weight λ_r	0.10
Flow loss weight λ_f	0.05
Inference decoding	Greedy until $T_{\max}$

on several metrics, but their gains are less uniform across the complete evaluation profile. These results suggest that the advantage of QSFR is not explained solely by the absence of Raman-specific inductive biases in the main baselines. Instead, learning multiple evidence compositions together with reliability-aware weighting provides a complementary representation principle.

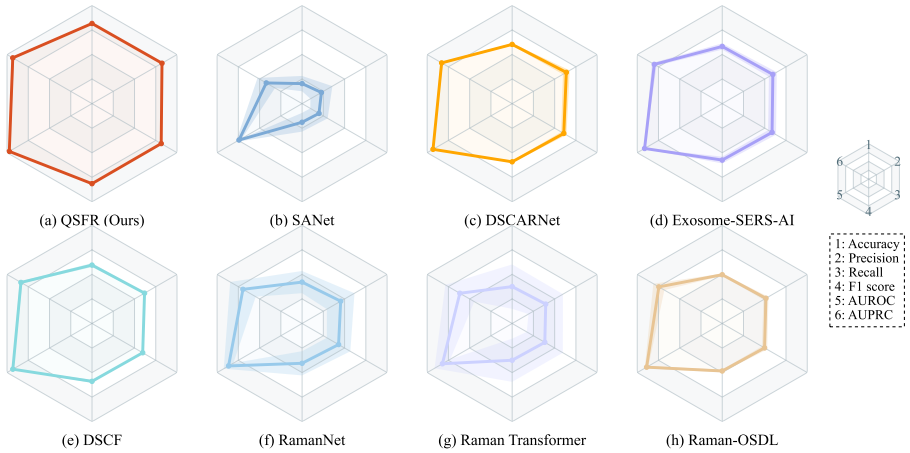


Figure 5: **Domain-specific controls for compositional evidence learning.** Radar profiles summarize six evaluation metrics under the same protocol. QSFR shows a comparatively broad and balanced performance profile.

G Real Source-to-Unseen-Target Transfer

We further evaluate whether the learned representation transfers to an independently held-out target cohort. Gastric Serum, Hepatic Cohort, HD Serum, and Bacteria Raman are used as source

Algorithm 1 Training and Inference of QSFR

Input: continuous observation x, label y
Components: evidence encoder $\mathcal{R}_\omega$, QSU, CSFG, subset aggregator $\mathcal{A}_\rho$, and UCRC
Training phase
for each mini batch (x, y) **do**
  $p \leftarrow \mathrm{Patchify}(x)$
  $(\bar{z}, L_q) \leftarrow \mathrm{QSU}(p)$
  $e_{1:M} \leftarrow \mathcal{R}_\omega(\bar{z}_{1:M})$
  $(S, \tau) \leftarrow \mathrm{CSFGSample}(e_{1:M})$
  $h_S \leftarrow \mathcal{A}_\rho\left(\mathrm{Sel}_S(e_{1:M})\right)$
  $\hat{y}_S \leftarrow \mathrm{Head}(h_S)$
  $L_{\mathrm{task}} \leftarrow \mathrm{CE}(\hat{y}_S, y)$
  $r_{\mathrm{obs}}(S) \leftarrow p_{\mathrm{pred}}(y \mid h_S) - \max_{c \neq y} p_{\mathrm{pred}}(c \mid h_S)$
  $p_\phi(\xi \mid h_S) \leftarrow \mathrm{UCRC}(h_S)$
  $L_{\mathrm{dist}} \leftarrow \mathrm{DistributionLoss}\left(p_\phi(\xi \mid h_S), r_{\mathrm{obs}}(S)\right)$
  Compute $\mu_\phi(S)$ and $\sigma_\phi(S)$
  $R_{\phi,\alpha}(S) \leftarrow \exp\left\{\beta\left[\mu_\phi(S) - \alpha\sigma_\phi(S)\right]\right\}$
  $L_{\mathrm{TB}} \leftarrow \mathrm{TrajectoryBalanceLoss}\left(\tau, \mathrm{sg}\left[R_{\phi,\alpha}(S)\right]\right)$
  Optimize $L_{\mathrm{task}} + \lambda_q L_q + \lambda_r L_{\mathrm{dist}} + \lambda_f L_{\mathrm{TB}}$
end for
Inference phase
$(\bar{z}, \cdot) \leftarrow \mathrm{QSU}(\mathrm{Patchify}(x))$
$e_{1:M} \leftarrow \mathcal{R}_\omega(\bar{z}_{1:M})$
$S \leftarrow \mathrm{CSFGSample}(e_{1:M})$
$h_S \leftarrow \mathcal{A}_\rho\left(\mathrm{Sel}_S(e_{1:M})\right)$
$\hat{y} \leftarrow \mathrm{Head}(h_S)$
Output $\arg\max \hat{y}$

Table 3: Performance on an unseen target cohort under four transfer protocols.

Setting	Accuracy	Macro-F1	AUROC
Source-only zero-shot	82.42 ± 2.18	81.60 ± 2.36	88.43 ± 1.03
Target-only	88.52 ± 1.24	88.12 ± 1.02	97.12 ± 0.76
Frozen + linear probe	89.06 ± 1.71	88.79 ± 1.89	97.10 ± 0.72
Full adaptation	$\mathbf{92.15 \pm 2.19}$	$\mathbf{91.85 \pm 2.24}$	$\mathbf{98.37 \pm 0.70}$

cohorts, while Cardiovascular Spectra is completely excluded from source training, validation, preprocessing-statistic estimation, and model selection.

We consider four settings that separate zero-shot generalization, target-only learning, frozen representation transfer, and full adaptation. As shown in Table 3, the source-only model reaches 82.42% Accuracy and 81.60% Macro-F1 on the unseen target cohort without using target labels. Frozen transfer performs comparably to target-only training, indicating that the source-trained representation remains useful on the new cohort. Full adaptation further improves Accuracy and Macro-F1 by 3.63 and 3.73 percentage points over target-only training, respectively. These results provide direct evidence that the learned compositional representation can transfer across real cohort and acquisition differences, while target supervision remains beneficial for further adaptation.

H Statistical Analysis

We assess cross-cohort performance consistency using Friedman tests over eight Raman cohorts, with cohorts treated as blocks and methods as treatments. Holm correction is applied across the six evaluation metrics. Table 4 shows that all corrected tests remain significant, while QSFR obtains the lowest average rank on every metric. This provides statistical support for the aggregate cross-cohort ranking observed in the main experiments.

Table 4: Cross-cohort Friedman analysis across eight Raman cohorts. Lower average rank indicates better overall ranking.

Metric	Friedman χ^2	Raw p	Holm-adjusted p	QSFR Avg. Rank ↓
Accuracy	35.952	4.0×10^{-5}	1.60×10^{-4}	**1.0625**
Precision	37.044	2.6×10^{-5}	1.30×10^{-4}	**1.1250**
Recall	24.977	2.996×10^{-3}	2.996×10^{-3}	**1.8750**
F1	33.873	9.4×10^{-5}	2.82×10^{-4}	**1.1250**
AUROC	29.782	4.78×10^{-4}	9.56×10^{-4}	**2.0000**
AUPRC	37.882	1.8×10^{-5}	1.08×10^{-4}	**1.2500**

I Evaluation of Utility Dispersion and Risk Alignment

In this subsection, we evaluate whether the utility dispersion produced by UCRC provides meaningful relative risk information. All statistics are computed only after model selection, with all model parameters frozen, on the held-out test partition. Test labels are used solely to instantiate r_{obs} for retrospective evaluation and never used to optimize UCRC, CSFG, the prediction head, or model selection. UCRC predicts a conditional utility distribution $p_\phi(z \mid h)$, extracting its mean $\mu(h)$ and dispersion $\sigma(h)$. The risk adjusted utility is $\mu(h) - \alpha\sigma(h)$. Consistent with Appendix B.2, $\sigma(h)$ is evaluated as a relative utility dispersion score, not as calibrated epistemic or aleatoric uncertainty.

Based on this formulation, we use two metrics that match the intended role of UCRC. The first is $\mathrm{Corr}(\sigma, |r_{\text{obs}} - \mu|)$. Importantly, this metric does not correlate σ with the raw prediction margin. It measures whether the learned dispersion tracks the residual utility error after the conditional mean μ has been accounted for. A positive correlation therefore indicates that the learned dispersion is aligned with residual utility error after the conditional mean has been accounted for. This metric evaluates residual risk ordering, not statistical independence from the prediction head. The second metric is $\mathrm{Coverage}(\mu - \alpha\sigma)$, which measures the empirical coverage of the risk adjusted utility at the same α used during training. This quantity evaluates the behavior of the chosen risk score at one operating point and should not be interpreted as a general calibration curve.

Specifically, UCRC first outputs the reward distribution $p_\phi(z \mid h)$ and extracts $\mu(h)$ and $\sigma(h)$ from it. At the same time, it uses the observed label conditioned utility margin as r_{obs} during evaluation. Next, we compute $\mathrm{Corr}(\sigma, |r_{\text{obs}} - \mu|)$ to measure the alignment between learned utility dispersion and utility prediction error. We compute $\mathrm{Coverage}(\mu - \alpha\sigma)$, which is the proportion of observed label conditioned utility margin values r_{obs} above the risk aware score $\mu - \alpha\sigma$. Here, α is fixed to the same risk coefficient used during training. The results are reported in Table 5.

From the results, $\mathrm{Corr}(\sigma, |r_{\text{obs}} - \mu|)$ reaches a mean of 0.8915 across three representative datasets. This result shows strong alignment between the learned utility dispersion and the observed utility prediction error. $\mathrm{Coverage}(\mu - \alpha\sigma)$ reaches a mean of 0.9471, showing high one-sided empirical coverage at the operating value of α used in QSFR. This statistic is reported only as an empirical property of the risk-adjusted utility at this operating point and is not intended as a calibrated coverage guarantee. On Gastric Serum, Corr is 0.9066 ± 0.0047 and Coverage is 0.9850 ± 0.0060, showing that UCRC provides informative reward level risk ordering under subtle class differences. On HD Serum, Coverage remains 0.9043 ± 0.0288, suggesting that the risk adjusted utility remains empirically stable under finer subtype boundaries. These results support the use of σ as a relative utility dispersion score for risk aware ranking, rather than as a calibrated epistemic or aleatoric uncertainty estimate. Theorem 3.2 provides a conditional guarantee under dispersion separation and bounded approximation errors. Our experiments evaluate empirical risk alignment at the chosen operating point rather than directly verifying these theoretical conditions.

J Ablation and Analysis of the Learned Evidence Distribution

As shown in Table 6, we conduct a controlled ablation study to separate the effects of the main components in QSFR. The reference model is a standard GFlowNet that shares the same encoder, prediction head, candidate evidence space, subset budget, sampling budget, and training protocol with QSFR. The standard GFlowNet retains the same trajectory balance objective and sequential subset

Table 5: Utility dispersion alignment and empirical risk coverage on held out test partitions of representative datasets.

Dataset	Corr $(\sigma, \lvert r_{\text{obs}} - \mu \rvert) \uparrow$	Coverage$(\mu - \alpha\sigma) \uparrow$
Myeloma SERS	0.8773 ± 0.0431	0.9521 ± 0.0314
Gastric Serum	0.9066 ± 0.0047	0.9850 ± 0.0060
HD Serum	0.8906 ± 0.0194	0.9043 ± 0.0288
Mean	0.8915	0.9471

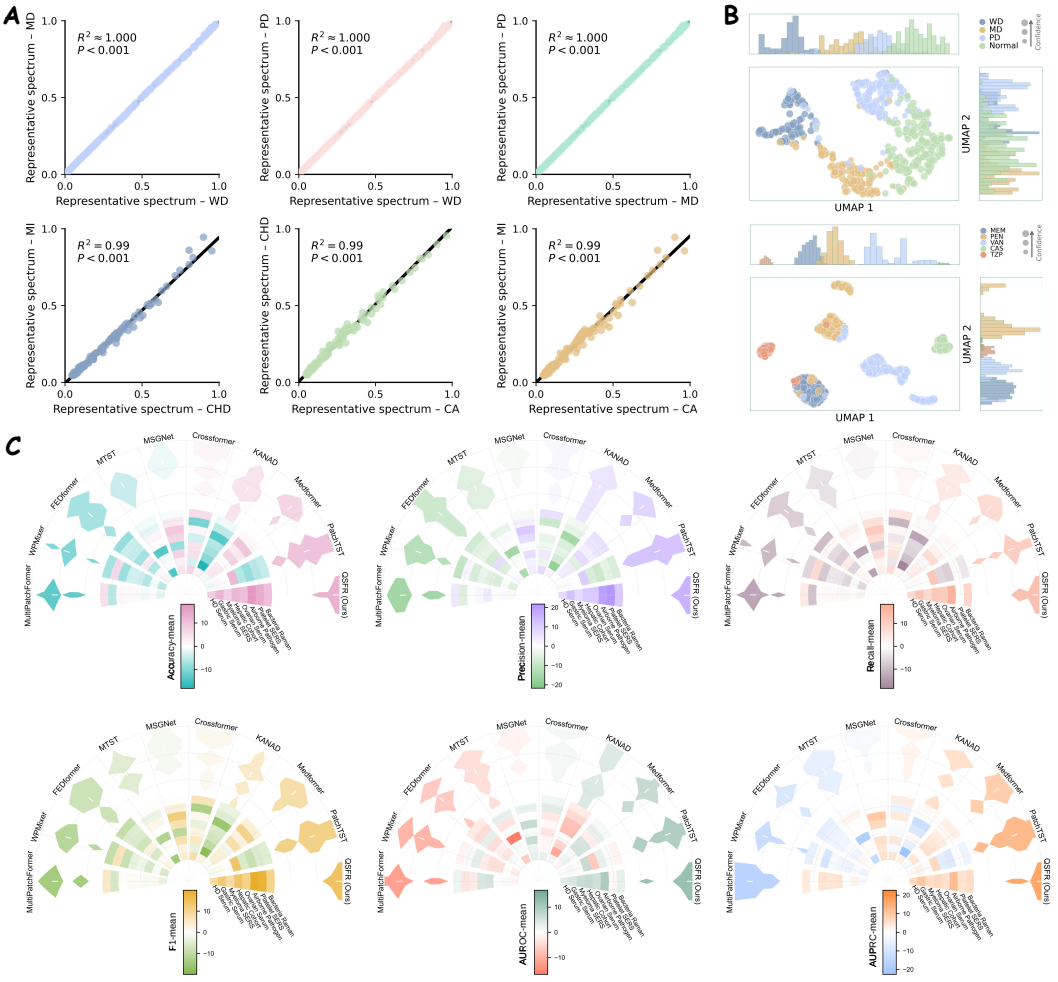


Figure 6: **Representation analysis under heterogeneous observation conditions.** (A) Redundancy in representative raw observations. (B) Geometry and confidence of learned representations. (C) Multi metric comparison across extended observation settings.

construction, but uses the risk neutral predictive utility as its terminal reward. Starting from the full framework, we further remove each component individually to evaluate its contribution under otherwise matched conditions. The results show consistent gains from structured evidence abstraction, compositional flow modeling, and reliability aware utility estimation. Their combination produces the strongest overall representation quality across the evaluated datasets.

We further analyze how each design choice affects a different property of the learned compositional evidence distribution. We focus on local representation stability, support coverage in the combinatorial space, and reliability aware redistribution.

Table 6: Ablation Study of QSFR Model Configurations.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
Myeloma SERS	**Base**	64.19±0.92	63.30±1.37	64.19±0.92	61.44±1.36	61.41±1.70	61.80±1.27
	w/o UCRC	74.15±1.41	75.21±1.40	74.15±1.41	73.88±1.53	82.73±2.00	82.62±2.40
	w/o CSFG	86.50±1.38	86.89±1.32	86.50±1.38	86.46±1.40	94.17±1.00	94.33±1.04
	w/o QSU	84.08±1.32	84.66±0.96	84.08±1.32	84.00±1.37	92.75±0.88	93.02±0.88
	QSFR	**94.46±0.66**	**94.49±0.65**	**94.46±0.66**	**94.46±0.64**	**98.66±0.28**	**98.69±0.29**
Gastric Serum	**Base**	59.93±1.26	59.12±1.81	57.82±1.17	57.03±1.33	81.23±0.80	60.04±1.56
	w/o UCRC	74.13±1.38	73.72±1.47	73.60±1.48	73.43±1.47	91.20±0.70	78.67±2.06
	w/o CSFG	82.53±1.37	82.37±1.30	81.98±1.30	82.00±1.29	95.29±0.54	87.91±1.65
	w/o QSU	87.17±1.26	87.11±1.22	86.78±1.30	86.90±1.26	97.12±0.42	92.39±1.23
	QSFR	**93.67±0.74**	**93.83±0.68**	**93.52±0.94**	**93.62±0.79**	**98.85±0.29**	**97.10±0.56**
Airborne Pathogen	**Base**	57.98±2.90	58.50±4.52	54.97±3.84	54.74±4.69	86.46±2.67	60.30±4.93
	w/o UCRC	66.06±2.13	68.19±3.17	63.34±4.66	63.88±2.62	89.09±2.61	68.63±4.57
	w/o CSFG	74.14±2.68	76.38±4.02	71.65±2.92	72.45±3.24	92.46±2.16	77.65±4.60
	w/o QSU	82.83±2.71	85.10±2.63	81.00±3.43	82.17±3.49	95.80±1.36	86.10±3.69
	QSFR	**85.86±2.86**	**87.51±2.45**	**84.10±4.24**	**84.84±3.83**	**96.11±0.93**	**89.14±2.28**

Stability of local evidence units. Reliable compositional learning first requires local evidence to remain stable under observation level variation. We therefore evaluate prototype based discretization under baseline drift, fluorescence background, and spike artifacts with increasing perturbation strength. Fig. 4A shows that discretized evidence representations generally converge faster and retain higher Accuracy and Macro AUPRC than their continuous counterparts. The advantage is more pronounced in several medium and strong perturbation settings. This suggests that mapping locally varying observations to reusable prototypes reduces representation drift before compositional sampling is performed.

Support coverage of the compositional distribution. A second requirement is that the learned distribution should not degenerate to one repeatedly selected evidence set. We examine this property by varying the sampling budget K. Fig. 4B shows that structured flow sampling maintains a high distinct set ratio while substantially increasing pairwise subset diversity. In contrast, the alternative sampler increasingly revisits similar compositions as more samples are drawn. These observations suggest that flow based sampling expands the effective support of the learned evidence distribution and preserves multiple candidate solutions.

Reliability aware probability shaping. The final requirement is to distinguish useful evidence from evidence with high utility dispersion. Distributional utility estimation provides both an expected utility and a dispersion estimate for each composition. Fig. 4C shows that dispersion aware scoring changes the empirical risk landscape from a narrow intermediate region to a strongly concentrated low risk mode with a smaller high dispersion tail. This pattern indicates that relative risk information affects the preference among candidate evidence states. It is consistent with the risk aware formulation, which favors evidence with high predicted utility when its utility dispersion remains controlled. These analyses mirror the three approximation sources isolated: quantization distortion, critic approximation, and sequential flow mismatch. They provide empirical evidence for the behavior of the corresponding mechanisms rather than direct numerical estimates of the theoretical error terms.

K Generalization of Compositional Evidence Representations

We examine whether the compositional evidence representation remains effective across additional Raman observation settings. The evaluation includes hepatic serum, ovarian serum, platelet SERS, myeloma SERS, and airborne single cell cohorts, which introduce substantial variation in biological

Table 7: Results on expanded cohorts. Purple marks QSFR.
The baselines span three groups: representation reconstruction, compositional collaborative flows, and robust medical detection.

Datasets	Models	Accuracy	Precision	Recall	F1 score	AUROC	AUPRC
Hepatic Cohort	**MultiPatchFormer**	85.13±1.62	90.38±0.82	85.56±1.50	83.33±1.93	96.52±0.63	96.57±0.58
	WPMixer	80.51±1.57	80.63±1.63	80.48±1.61	80.24±1.62	87.30±1.35	88.40±1.18
	FEDformer	87.69±1.95	92.76±0.96	88.49±1.81	85.28±2.43	89.68±1.88	90.12±1.79
	MTST	70.40±2.43	70.58±2.69	70.40±2.43	70.35±2.37	74.67±2.12	73.04±2.82
	MSGNet	75.90±2.04	86.11±0.98	76.90±1.84	71.73±2.52	91.96±0.91	91.52±1.03
	Crossformer	88.72±0.66	83.38±0.37	87.94±0.19	84.94±0.25	88.20±0.22	89.62±0.34
	KANAD	76.92±1.68	81.87±1.31	77.86±1.56	75.30±1.86	92.22±6.06	92.53±5.89
	Medformer	84.62±0.16	80.94±0.27	83.33±0.18	80.50±0.23	96.51±3.70	96.82±3.38
	PatchTST	96.41±3.84	96.38±3.86	**96.43±3.87**	**96.39±3.86**	98.20±2.63	98.29±2.51
	QSFR(Ours)	**96.41±1.26**	**96.66±1.19**	96.27±1.34	96.37±1.27	**98.31±1.76**	**98.45±1.60**
Ovarian Serum	**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
	WPMixer	76.62±5.26	76.95±5.69	76.59±5.23	76.57±5.22	86.86±3.00	87.81±2.69
	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
	MSGNet	87.53±2.41	88.24±2.67	87.46±2.38	87.46±2.39	**96.22±0.87**	**96.49±0.82**
	Crossformer	76.62±1.84	76.95±1.98	76.55±1.83	76.52±1.82	88.50±2.74	89.24±2.83
	KANAD	74.29±5.77	74.45±5.74	74.28±5.79	74.23±5.80	81.32±5.56	81.25±5.21
	Medformer	87.79±3.03	88.22±2.77	87.77±3.06	87.74±3.08	94.36±1.33	94.12±2.17
	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
	QSFR(Ours)	**91.17±2.23**	**91.84±1.88**	**91.15±2.24**	**91.12±2.26**	95.14±1.59	95.15±1.83
Airborne Pathogen	**MultiPatchFormer**	69.49±2.42	67.45±4.06	68.48±1.53	65.06±2.24	83.73±0.51	82.16±1.43
	WPMixer	60.00±3.17	60.96±4.07	59.49±3.25	56.49±2.32	89.51±1.23	67.60±1.98
	FEDformer	63.43±2.25	63.46±4.59	62.33±2.66	59.27±4.52	89.37±0.37	73.68±4.03
	MTST	69.90±2.81	68.69±3.55	69.05±2.12	66.80±2.96	82.75±0.35	77.52±0.83
	MSGNet	77.58±1.18	78.58±1.85	78.68±1.75	77.03±1.48	86.14±0.65	87.09±1.98
	Crossformer	69.90±4.30	67.77±5.84	69.59±3.96	66.46±4.66	83.52±0.64	81.04±1.92
	KANAD	54.34±3.47	52.37±1.54	54.38±1.78	48.99±2.88	78.65±1.53	64.52±5.77
	Medformer	71.52±1.18	70.97±1.79	70.30±1.27	68.00±1.35	84.09±0.25	80.45±1.85
	PatchTST	59.60±2.21	58.80±5.77	60.53±2.24	56.39±2.51	88.95±0.82	68.77±3.47
	QSFR(Ours)	**85.86±2.86**	**87.51±2.45**	**84.10±4.24**	**84.84±3.83**	**96.11±0.93**	**89.14±2.28**
Platelet SERS	**MultiPatchFormer**	74.17±1.67	62.22±0.56	75.91±0.91	61.43±1.35	78.22±0.95	80.21±0.33
	WPMixer	86.67±4.08	69.86±3.39	**92.73±2.23**	74.30±4.85	82.68±2.57	80.78±3.55
	FEDformer	75.83±1.67	62.86±0.71	86.82±0.91	62.85±1.48	77.28±2.37	75.23±2.33
	MTST	78.33±3.12	64.05±1.53	88.18±1.70	65.17±2.89	84.38±2.15	81.76±1.68
	MSGNet	84.17±4.08	70.86±3.69	91.36±2.23	71.35±4.07	84.13±2.11	80.48±3.12
	Crossformer	69.17±1.55	59.20±0.95	74.09±1.74	54.80±1.41	**85.00±1.33**	75.85±1.98
	KANAD	75.00±1.87	63.56±3.33	86.36±5.93	62.92±3.64	76.89±3.49	79.23±4.56
	Medformer	90.00±2.04	73.00±2.45	84.55±1.11	78.54±2.96	84.23±2.76	81.29±1.71
	PatchTST	74.17±3.19	63.54±3.90	85.91±3.49	62.44±3.67	82.51±2.89	79.73±2.44
	QSFR(Ours)	**94.31±0.92**	**89.65±2.40**	85.00±0.20	**86.49±1.96**	80.71±2.38	**85.22±1.95**
Myeloma SERS	**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
	WPMixer	80.27±4.05	80.63±4.28	80.27±4.05	80.22±4.03	88.13±3.35	88.16±3.13
	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
	MSGNet	93.88±1.63	94.06±1.54	93.88±1.63	93.88±1.63	**98.72±0.59**	**98.73±0.65**
	Crossformer	91.92±4.38	92.11±4.24	91.92±4.38	91.91±4.40	91.12±2.85	97.19±2.81
	KANAD	77.00±6.30	79.72±3.97	77.00±6.30	76.17±4.78	86.98±2.20	87.02±2.43
	Medformer	89.77±7.31	90.86±5.88	89.77±7.31	89.60±7.56	93.45±7.51	93.18±7.66
	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
	QSFR(Ours)	**94.46±0.66**	**94.49±0.65**	**94.46±0.66**	**94.46±0.64**	98.66±0.28	98.69±0.29

source, spectral background, and acquisition protocol. Table 7 shows that QSFR retains competitive and generally strong performance across these heterogeneous settings.

The analyses in Fig. 6 further characterize the learned representation. Panel A shows strong global similarity between observations from different groups, indicating that useful information cannot always be attributed to a single dominant region. Panel B shows that the learned feature space organizes observations into more structured local neighborhoods, while ambiguous samples remain closer to transition regions. Panel C provides a cohort level summary across six metrics and shows that the representation remains comparatively stable across different observation conditions. These results suggest that the compositional formulation can organize weak and distributed local evidence beyond the main evaluation cohorts.