

Agent-Guided Relational Concept Discovery: Toward Interpretable Surgical Margin Assessment

Nooshin Maghsoodi¹ *, Amoon Jamzad¹, Robert Policelli¹, Mohammad Farahmand¹, Dilakshan Srikanthan¹, Martin Kaufmann¹, Kevin Y. M. Ren¹, Shaila Merchant¹, Sonal Varma¹, Ross Walker¹, Doug McKay¹, John Rudan¹, Gabor Fichtinger¹, and Parvin Mousavi¹

Queen's University, Kingston, ON, Canada
nooshin.maghsoodi@queensu.ca

Abstract. Deep learning models can effectively use Rapid Evaporative Ionization Mass Spectrometry (REIMS) data for surgical margin assessment. However, their clinical adoption remains challenging due to limited generalization to operating room conditions. This difficulty arises because models are typically trained on labeled spectra collected from resected tissue samples, while they must operate on noisy, unlabeled data acquired directly during surgery. In addition, the black-box nature of deep learning models makes it difficult to understand and systematically improve their behavior. Concept-based learning offers a promising way to address these challenges by mapping raw measurements to human-understandable concepts. However, supervised concept-based approaches rely on concept annotations, which are difficult to obtain in complex mass spectrometry workflows. We propose Agent-Guided Concept Discovery, a framework that learns meaningful concepts directly from data without requiring predefined concept labels. During training, a reasoning agent refines semantic descriptions of the learned concepts and adaptively adjusts their weight based on diagnostic relevance. These concepts are further grounded using a biochemical knowledge graph to ensure consistency with known metabolic relationships. Across Skin and Breast Cancer datasets, our model improves balanced accuracy and sensitivity over the baseline. In a representative intraoperative case, it shows fewer false positives, indicating better generalization to surgical conditions.

Keywords: Agentic AI · Concept Learning · Mass Spectrometry

1 Introduction

Achieving clear surgical margins has a significant impact on cancer outcomes. Negative margins indicate that no malignant cells remain at the boundary of the excised tissue, suggesting that the tumor has been fully removed and reducing the risk of residual disease. Today, margin status is typically determined post-operatively by histopathology, leaving no opportunity for intraoperative adjustment. In contrast, real-time margin assessment allows surgeons to course-correct

* Corresponding author.

during resection, both avoiding cutting through the tumor and minimizing unnecessary removal of healthy tissue [10]. Rapid Evaporative Ionization Mass Spectrometry (REIMS) enables intraoperative feedback by providing instantaneous metabolomic profiles from cauterized tissue, which can be classified using machine learning to identify cancer at the point of incision [1].

Prior works have explored multiple strategies to improve REIMS-based margin assessment, including Bayesian neural networks for uncertainty estimation and image-based representations of mass spectra [8,6]. More recently, a foundation model pretrained on large-scale tandem mass spectrometry data, Deep Representations Empowering the Annotation of Mass Spectra (DreaMS) [5], has been proposed and later adapted for REIMS applications [7]. However, fine-tuning this foundation model on small, labeled ex vivo REIMS datasets risks overfitting and limits generalization to noisy intraoperative data. Moreover, despite strong performance, these models largely remain black boxes, offering limited insight into the biochemical drivers of their predictions.

Concept-based learning offers a promising alternative to purely end-to-end classification. Rather than mapping raw spectra directly to diagnostic labels, these models introduce an intermediate representation composed of human-understandable variables [13]. Related approaches further show how such high-level concepts can provide meaningful evidence for interpreting model behavior and predictions [12,17]. However, existing supervised frameworks depend on manually defined concept annotations, which are impractical in complex mass spectrometry workflows.

To address these challenges, we propose Agent-Guided Relational Concept Discovery, a framework for interpretable surgical margin assessment. We introduce a reasoning-agent-in-the-loop training framework that enables automatic discovery of discriminative concepts without requiring concept annotations. The agent analyzes shared spectral patterns, assigns semantic descriptions, and adaptively adjusts the relevance of concepts. These concepts are further grounded using a knowledge graph to ensure consistency with established metabolic pathways. Our main contributions are summarized as follows:

1. *Agent-Guided Concept Discovery without Supervision.* A reasoning agent is incorporated into the training loop to help the model discover and refine meaningful latent representations from embeddings without requiring manual annotations. This is achieved through two key contributions below.
2. *Relational Concept Modeling via Knowledge Graph Construction.* The agent receives the most informative spectral patterns associated with each learned concept and then queries metabolic databases to ground these concepts in biological knowledge. These links generate a knowledge graph connecting concepts to metabolites, pathways, and tissue states, which evolves during training to capture relationships and support more consistent agent feedback.
3. *Reasoning-Driven Concept Weight Optimization.* During training, the agent aggregates supporting evidence to provide feedback on the relevance of the learned concepts. This feedback is incorporated into an auxiliary alignment loss, which guides the learning of the concept weights.




Fig. 1. REIMS spectra are encoded to obtain latent embeddings, which are mapped to concept activations. Samples with high concept activation are analyzed to identify discriminative spectral features. A reasoning agent integrates this information with biochemical knowledge bases to ground the learned concepts, construct a relational knowledge graph, and iteratively refine concepts during training.

2 Materials and Methods

2.1 Data

Ex vivo data: Ex vivo REIMS spectra were collected from two surgical oncology cohorts at Kingston Health Sciences Center under institutional ethics approval, with written informed consent obtained from all participants. The basal cell carcinoma (BCC) dataset includes 693 annotated burns from 91 patients (252 tumors, 441 benign), acquired from freshly excised specimens. The breast cancer dataset comprises 144 ex vivo burns from 11 patients (41 tumors, 103 benign), sampled from malignant tissue and adjacent benign parenchyma using the same REIMS protocol.

Intraoperative data: Intraoperative data were collected during a breast-conserving surgery. Over a 27-minute procedure, 1,616 spectra were recorded at 1 Hz and labeled using intraoperative surgical annotations.

2.2 Proposed Model

Fig. 1 provides an overview of the proposed framework and its main components. The model is designed to jointly perform surgical margin classification while learning interpretable concepts based on a biochemical knowledge base. The overall pipeline consists of the following components.

Step 1: Embedding Generation via Foundation Model In the first step, we use the DreaMS foundation model [5] to process raw REIMS spectra. Due to the limited size of our labeled REIMS dataset, the DreaMS parameters are kept frozen and used solely as a feature extractor. For each input spectrum x_i , the model produces a d -dimensional embedding, z_i , which serves as the input to the concept embedding layer.

$$z_i = f_{\text{DreaMS}}(x_i) \in \mathbb{R}^d. \quad (1)$$

Step 2: Concept Embedding Layer The embedding z_i is mapped to a set of K latent concepts through a learnable linear projection layer, forming an explicit concept bottleneck. Each concept k is parameterized by a weight vector w_k and a bias term b_k . For sample i , the projection first produces an unbounded concept activation score s_{ik} :

$$s_{ik} = w_k^\top z_i + b_k, \quad p_{ik} = \sigma(s_{ik}). \quad (2)$$

Since s_{ik} is unbounded, its scale may vary across concepts and samples. Therefore, s_{ik} is passed through a sigmoid function to obtain the bounded activation value p_{ik} . These bounded values p_{ik} are only for ranking spectra for feature extraction, whereas the classifier uses the corresponding unbounded logits to preserve information.

To encourage separation between concept directions, the concept vectors $\{w_k\}_{k=1}^K$ are initialized using a cosine-similarity-based criterion. This initialization discourages highly similar concept vectors at the start of training, reducing redundancy in the initial concept space while allowing related concepts to become correlated during optimization when supported by the data.

Step 3: Reasoning Agent for Concept Refinement To convert concept activations into biologically meaningful representations, we integrate a reasoning agent into the training loop.

Discriminative spectral region identification. Each mass spectrum represents ion intensity over mass-to-charge ratios, denoted by m/z . Since cancer and normal tissues can differ in molecular composition, they may show distinct intensity patterns across specific m/z regions. The feature extraction module associates each learned concept with discriminative m/z ranges by contrasting spectra with high and low concept activation scores. Thus, concepts that were initially numerical parameters in Step 2 become linked to interpretable m/z regions.

Biochemical grounding. The agent analyzes the identified m/z regions by querying biochemical knowledge bases to retrieve candidate metabolites and pathways. We use RaMP [2,18] as the primary resource, which integrates metabolite identifiers, pathways, and chemical annotations across curated databases. The agent is additionally guided by text-based domain guidelines distilled from free biochemical and clinical resources and the literature [9,16,11,3].

Memory and knowledge graph construction. The agent maintains an external memory that records concept information across epochs, including associated

m/z regions, relevance feedback, and classification outcomes. For each candidate concept, overlapping concepts from previous epochs are retrieved and provided to the LLM to improve the consistency of feedback. The agent also updates a concept-level knowledge graph that links concepts to metabolites, pathways, and tissue pathological states, thereby supporting biological interpretation and concept refinement.

Description and relevance feedback. Based on this analysis, the agent provides (i) *textual descriptions* that associate concepts with putative metabolic classes, and (ii) *concept relevance* as directional feedback (*more relevant, less relevant*) based on biochemical consistency, historical trends from the memory board, and impact on classification.

Step 4: Agent-Guided Concept Alignment Loss As mentioned, the model discovers K meaningful concepts from the training data. However, in biochemical analysis, different spectral patterns or m/z regions contribute unequally to tissue characterization and diagnosis [11,3]. Motivated by this, we allow each learned concept to have a distinct weight on the final prediction.

To explicitly control the contribution of each concept, we associate it with a learnable scalar weight parameter α_k . To ensure interpretability and numerical stability, concept weight parameters are constrained to the range $[0, 1]$ using a sigmoid reparameterization. The gated concept signal, $\tilde{s}_{ik}$, is then defined as

$$\alpha_k = \sigma(a_k), \quad a_k \in \mathbb{R}, \quad \tilde{s}_{ik} = \alpha_k s_{ik}. \quad (3)$$

The classifier receives both the global embedding and the gated concept signals as input. Specifically, the predicted output for sample i is given by

$$\hat{y}_i = g([z_i \ \tilde{s}_{i1} \ \tilde{s}_{i2} \ \dots \ \tilde{s}_{iK}]), \quad (4)$$

where $g(\cdot)$ represents the classification head. To align the model’s use of concepts with the reasoning agent’s feedback, we introduce an alignment loss based on a perturbation-based relevance measure inspired by [14]. For each sample i and concept k , we measure the change in the predicted probability when the contribution of concept k is removed from the classifier input.

$$\Delta Y_{ik} = |\hat{y}_i - \hat{y}_{i(\tilde{s}_{ik} \rightarrow 0)}|. \quad (5)$$

A larger ΔY_{ik} indicates that the prediction for sample i depends more strongly on concept k .

Based on the reasoning agent’s feedback, concepts are grouped into high-relevance concepts $\mathcal{K}_H$ and low-relevance concepts $\mathcal{K}_L$. The sample-level alignment loss is defined as:

$$\mathcal{L}_{\text{align}}^{(i)} = \sum_{k \in \mathcal{K}_H} (1 - \Delta Y_{ik}) + \sum_{k \in \mathcal{K}_L} \Delta Y_{ik}. \quad (6)$$

The alignment loss, $\mathcal{L}_{\text{align}}$, is obtained by averaging $\mathcal{L}_{\text{align}}^{(i)}$ over the batch. The final objective combines classification loss with agent-guided alignment, where

Table 1. Classification performance (mean $\pm$ standard deviation over 30 runs) on REIMS datasets.

Data	Method	Bal. Acc.	Sens.	Spec.	AUROC
Skin Cancer	Transformer	0.74 $\pm$ 0.027	0.72 $\pm$ 0.111	0.76 $\pm$ 0.084	0.85 $\pm$ 0.015
	DreaMS [5]	0.71 $\pm$ 0.034	0.59 $\pm$ 0.086	0.83 $\pm$ 0.054	0.83 $\pm$ 0.014
	Our method	0.76 $\pm$ 0.026	0.70 $\pm$ 0.043	0.82 $\pm$ 0.031	0.86 $\pm$ 0.013
Breast Cancer	Transformer	0.82 $\pm$ 0.035	0.78 $\pm$ 0.216	0.86 $\pm$ 0.062	0.93 $\pm$ 0.021
	DreaMS [5]	0.84 $\pm$ 0.028	0.80 $\pm$ 0.051	0.89 $\pm$ 0.026	0.96 $\pm$ 0.012
	Our method	0.87 $\pm$ 0.018	0.81 $\pm$ 0.044	0.93 $\pm$ 0.024	0.98 $\pm$ 0.013

$\mathcal{L}_{\text{class}}$ is binary cross-entropy and $\lambda \geq 0$ controls the alignment strength.

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{class}} + \lambda \mathcal{L}_{\text{align}}. \quad (7)$$

2.3 Experiments

Our experiments consist of the following evaluations: (i) *Classification performance comparison*. We compare the proposed framework against a DreaMS-based baseline that performs classification on foundation model embeddings, on two REIMS mass spectrometry datasets. (ii) *Concept quality evaluation*. We assess interpretability and biological relevance by examining how concept activations vary with hormone receptor status and align with established metabolic pathways. (iii) *Generalization to intraoperative data*. We assess by evaluating performance on an in vivo REIMS case acquired during surgery. (iv) *Ablation studies*. We analyze the contribution of key components by ablating biochemical knowledge base integration and knowledge graph grounding.

Implementation Details. The reasoning agent uses **Qwen2.5-7B-Instruct** as the backbone LLM. Datasets are split into training, validation, and test sets using patient-wise stratification, ensuring no patient overlap between training and test samples. The model learns $K = 8$ concepts and is trained with Adam for up to 40 epochs, using a learning rate of 10^{-3} and a batch size of 32. We set $\lambda = 0.5$ and update the agent every three epochs. Hyperparameters are selected on the validation set. All experiments are repeated 30 times, with results reported as mean and standard deviation. The code and configuration files used in this study are publicly available at: <https://github.com/nooshinmaghsoodi/cem-qwen-reims>.

3 Results and Discussion

Classification performance comparison The classification results are shown in Table 1. Overall, the proposed model outperforms the DreaMS baseline across both datasets. This indicates that incorporating agent-guided concept embeddings, aimed at improving generalization and interpretability, does not sacrifice classification performance and can even improve it. On the Skin Cancer dataset,




Fig. 2. (a) Breast cancer knowledge graph linking m/z ranges, metabolites, and pathways. (b) Zoomed-in selected pathways for some m/z ranges in (a). (c) Concept activations across tumor samples with their Progesterone Receptor (PR) status.

we observe an approximate 7% relative improvement in balanced accuracy, along with higher AUROC and a noticeable increase in sensitivity, suggesting improved detection of tumor tissue. Similarly, on the Breast Cancer dataset, our method achieves higher balanced accuracy (0.87) while also improving both sensitivity and specificity.

Clinical Evaluation of Concepts Fig. 2 illustrates how the learned concepts are grounded in biochemical knowledge. Fig. 2(a) shows a portion of the Breast Cancer knowledge graph linking informative m/z ranges to candidate metabolites and pathways discovered during training. A zoomed-in view from this knowledge graph is shown in Fig. 2(b), highlighting selected m/z ranges and some of their associated pathways, enabling assessment of clinical relevance. We evaluated this relevance by analyzing concept activation scores for tumor samples in the Breast Cancer test set. Fig. 2(c) shows the activation of selected concepts together with progesterone receptor (PR) status. Prior studies [15,4] report that PR-negative breast cancers are associated with glutamine metabolism and increased amino acids. Consistent with these findings, Fig. 2(b) highlights amino acid- and glutamine-related pathways linked to the m/z subbands 130–135, 211–222, and 892–894, while Fig. 2(c) shows higher activation of these subbands in PR-negative samples. This consistency supports the biological relevance of the learned concepts and demonstrates the interpretability enabled by the generated knowledge graph.




Fig. 3. (a) Intraoperative breast cancer example. The x-axis denotes surgery time. Striped red regions indicate surgeon-noted tumors, while dotted purple regions correspond to skin cuts. (b) Ablation study on both datasets, as metabolic database querying (DB) and knowledge graph reasoning (KG) are added to support the agent + clinical guidelines (GL). Values indicate relative performance gains.

Generalization to intraoperative data Fig. 3(a) shows model predictions on a representative intraoperative breast cancer case with surgeon-provided call-out labels. Stripped red overlaid regions denote the tumor, and dotted purple regions denote the skin cuts during incision. While both models achieve perfect tumor sensitivity, the concept-enhanced model produces fewer false positives than the baseline, particularly in skin regions with tumor-like signatures. This suggests that incorporating biologically aware concepts helps the model better distinguish tumor tissue from skin.

Ablation Study Ablation study results are shown in Fig. 3(b). In this experiment, components that support the reasoning agent are incrementally added to analyze the effects of the metabolic database and the knowledge graph. Incorporating the metabolic database improves classification performance, with additional gains observed when the knowledge graph is introduced. These effects are more observable on the skin dataset, where the largest relative improvements are observed in balanced accuracy, indicating improved sensitivity to tumor-related spectral patterns. We believe that adding the guideline component alone primarily helps reduce hallucinated semantic explanations generated by the agent.

4 Conclusion

In our proposed framework, an agent operates within the training loop to help the model discover meaningful concepts without explicit concept annotations, aiming to improve interpretability and strengthen generalization. This approach

extends models trained on ex vivo data to challenging intraoperative settings while making their decisions easier to understand. The model learns high-level metabolic concepts and directly integrates them into the prediction process, grounding decisions in patterns that remain stable under intraoperative noise. In addition, a dynamically constructed, data-specific knowledge graph reveals underlying biochemical structure and supports iterative interpretation and concept refinement beyond what fixed rule-based systems can provide. A current limitation is the reliance on general metabolic database queries; future work will explore more targeted querying strategies and the integration of gene-level information to further enhance biological specificity and interpretability.

Acknowledgments. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC); the Canadian Institutes of Health Research (CIHR); and the Vector Institute. The work of Parvin Mousavi was supported in part by a Canada CIFAR AI Chair and a Canada Research Chair.

Disclosure of Interests. The authors have no competing interests.

References

1. Balog, J., Sasi-Szabó, L., Kinross, J., Lewis, M.R., Muirhead, L.J., Veselkov, K., Mirnezami, R., Dezső, B., Damjanovich, L., Darzi, A., Nicholson, J.K.: Intraoperative tissue identification using rapid evaporative ionization mass spectrometry. *Science Translational Medicine* **5**(194), 194ra93–194ra93 (2013)
2. Braisted, J., Patt, A., Tindall, C., Sheils, T., Neyra, J., Spencer, K., Eicher, T., Mathé, E.A.: RaMP-DB 2.0: a renovated knowledgebase for deriving biological and chemical insight from metabolites, proteins, and genes. *Bioinformatics* **39**(1), btac726 (2023)
3. Brorsen, L.F., McKenzie, J.S., Pinto, F.E., Glud, M., Hansen, H.S., Haedersdal, M., Takats, Z., Janfelt, C., Lerche, C.M.: Metabolomic profiling and accurate diagnosis of basal cell carcinoma by maldi imaging and machine learning. *Experimental Dermatology* **33**(7), e15141 (2024)
4. Budczies, J., Pfitzner, B.M., Györfy, B., Winzer, K.J., Radke, C., Dietel, M., Fiehn, O., Denkert, C.: Glutamate enrichment as new diagnostic opportunity in breast cancer. *International Journal of Cancer* **136**(7), 1619–1628 (2015)
5. Bushuiev, R., Bushuiev, A., Samusevich, R., Brungs, C., Sivic, J., Pluskal, T.: Self-supervised learning of molecular representations from millions of tandem mass spectra using DreaMS. *Nature Biotechnology* **44**(4), 630–640 (2026)
6. Connolly, L., Fooladgar, F., Jamzad, A., Kaufmann, M., Syeda, A., Ren, K., Abolmaesumi, P., Rudan, J.F., McKay, D., Fichtinger, G., Mousavi, P.: ImSpect: image-driven self-supervised learning for surgical margin evaluation with mass spectrometry. *International Journal of Computer Assisted Radiology and Surgery* **19**(6), 1129–1136 (2024)
7. Farahmand, M., Jamzad, A., Fooladgar, F., Connolly, L., Kaufmann, M., Ren, K.Y.M., Rudan, J., McKay, D., Fichtinger, G., Mousavi, P.: FACT: Foundation model for assessing cancer tissue margins with mass spectrometry. *International Journal of Computer Assisted Radiology and Surgery* **20**(6), 1097–1104 (2025)

8. Fooladgar, F., Jamzad, A., Connolly, L., Santilli, A., Kaufmann, M., Ren, K., Abolmaesumi, P., Rudan, J.F., McKay, D., Fichtinger, G., Mousavi, P.: Uncertainty estimation for margin detection in cancer surgery using mass spectrometry. *International Journal of Computer Assisted Radiology and Surgery* **17**(12), 2305–2313 (2022)
9. Fu, Y., Zou, T., Shen, X., Nelson, P.J., Li, J., Wu, C., Yang, J., Zheng, Y., Bruns, C., Zhao, Y., Qin, L., Dong, Q.: Lipid metabolism in cancer progression and therapeutic strategies. *MedComm* **2**(1), 27–59 (2021)
10. Jamzad, A., Sedghi, A., Santilli, A.M.L., Janssen, N.N.Y., Kaufmann, M., Ren, K.Y.M., Vanderbeck, K., Wang, A., McKay, D., Rudan, J.F., Fichtinger, G., Mousavi, P.: Improved resection margins in surgical oncology using intraoperative mass spectrometry. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020. Lecture Notes in Computer Science*, vol. 12263, pp. 44–53. Springer (2020)
11. Kaufmann, M., Vaysse, P.M., Savage, A., Kooreman, L.F.S., Janssen, N., Varma, S., Ren, K.Y.M., Merchant, S., Engel, C.J., Olde Damink, S.W.M., Smidt, M.L., Shousha, S., Chauhan, H., Karali, E., Kazanc, E., Poulogiannis, G., Fichtinger, G., Tauber, B., Leff, D.R., Pringle, S.D., Rudan, J.F., Heeren, R.M.A., Porta Siegel, T., Takáts, Z., Balog, J.: Testing of rapid evaporative mass spectrometry for histological tissue classification and molecular diagnostics in a multi-site study. *British Journal of Cancer* **131**, 1298–1308 (2024)
12. Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., Sayres, R.: Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In: *Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 80, pp. 2668–2677. PMLR (2018)
13. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept bottleneck models. In: *Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 119, pp. 5338–5348. PMLR (2020)
14. Pang, W., Ke, X., Tsutsui, S., Wen, B.: Integrating clinical knowledge into concept bottleneck models. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. Lecture Notes in Computer Science*, vol. 15004, pp. 243–253. Springer (2024)
15. Shahnazari, P., Kavousi, K., Khorram Khorshid, H.R., Minuchehr, Z., Goliaei, B., M. Salek, R.: Uncovering subtype-specific metabolic signatures in breast cancer through multimodal integration, attention-based deep learning, and self-organizing maps. *Scientific Reports* **15**(1), 21775 (2025)
16. Thomsen, A.A., Rabjerg, M., Marcussen, N., Lund, L., Jensen, O.N.: Mass spectrometry and machine learning for classification and molecular phenotyping of renal cell carcinoma and benign tumors. *medRxiv* (2026). <https://doi.org/10.64898/2026.01.13.26343935>
17. Yuksekgonul, M., Wang, M., Zou, J.: Post-hoc concept bottleneck models. In: *The Eleventh International Conference on Learning Representations* (2023)
18. Zhang, B., Hu, S., Baskin, E., Patt, A., Siddiqui, J.K., Mathé, E.A.: RaMP: A comprehensive relational database of metabolomics pathways for pathway enrichment analysis of genes and metabolites. *Metabolites* **8**(1), 16 (2018)