Contrastive Time-Series Representation Learning for Multiplexed Neurochemical Concentration Prediction

Na Dai[1,2], **Abdallah Daha**[1], **Joshua Labbe**[1], **Zhengxu Tang**[1], **Yunnuo Zhang**[1], **Jinxing Li**[1,3]
[1]Institute for Quantitative Health Science & Engineering, Michigan State University
[2]College of Osteopathic Medicine, Michigan State University
[3]Department of Biomedical Engineering, Michigan State University
daina@msu.edu jl@msu.edu

Abstract

Accurately characterizing the neurochemical environment is essential for advancing the understanding and treatment of neurological and psychiatric disorders. Fast-scan cyclic voltammetry (FSCV) enables high-temporal-resolution measurement of neurotransmitter dynamics, but predicting multiplex concentrations in complex fluids remains an open challenge. We identify three key obstacles: nonlinear analyte interactions, the limited utility of normalization, and pronounced batch effects stemming from sensor fabrication variability. To address these, we propose a metadata-aware contrastive representation learning framework that explicitly incorporates batch identity, scan sequence, and scan rate to model experimental variability. Preference-based ranking losses emphasize subtle yet discriminative features of voltammetry curves, while the learned representations are decoupled from downstream predictors. A convolutional neural network is then applied to capture non-linearity in concentration prediction. Preliminary results show improved accuracy over traditional baselines, highlighting a promising direction at the intersection of time-series representation learning and multiplexed biosensing.

1 Introduction

Neurotransmitters and hormones regulate essential physiological and psychological functions [9]. Neurotransmitters enable rapid brain communication, while hormones act as long-range messengers coordinating systemic physiology. Dopamine drives reward and motor control, serotonin shapes mood, norepinephrine modulates stress and attention, and melatonin regulates sleep–wake cycles. Real-time monitoring of their dynamics is crucial for understanding neural circuits, brain–body interactions, and for phenotyping psychiatric and neurological disorders.

Electrochemical sensing through cyclic voltammetry (CV) is the primary sensing method, recording redox currents of analytes [2]. While slow-scan CV yields high-quality signals, it lacks real-time resolution. Fast-scan CV (FSCV) improves temporal resolution, but discrimination between structurally similar monoamine compounds (e.g., dopamine vs. norepinephrine) remains challenging.

Despite substantial prior work on inferring multiplex concentrations from CV data [10, 14, 12, 19, 4, 6, 16], our analysis reveals three under-recognized obstacles that fundamentally limit current approaches. First, analyte–analyte interactions introduce non-linear coupling effects that degrade multiplex prediction accuracy. Second, contrary to common practice, we find that background-subtraction–based normalization [7, 8] does not improve concentration prediction. Third, CV curves exhibit substantial heterogeneity and batch- or sensor-specific variation, further complicating model generalization. We provide an expanded discussion of these observations in Appendix A.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Learning from Time-series for Health (TS4H).

We propose a metadata-aware contrastive learning framework that decouples representation learning from downstream prediction. Metadata such as batch identity, scan sequence, and scan rate guide robust representation, while preference-based ranking losses emphasize subtle concentration-dependent differences. A convolutional neural network then captures nonlinearities for prediction.

Our contributions are: (1) systematic characterization of CV-specific challenges, (2) a metadata-aware contrastive framework for robust representation learning, and (3) evidence of improved accuracy on unseen CV data.

2 Methods

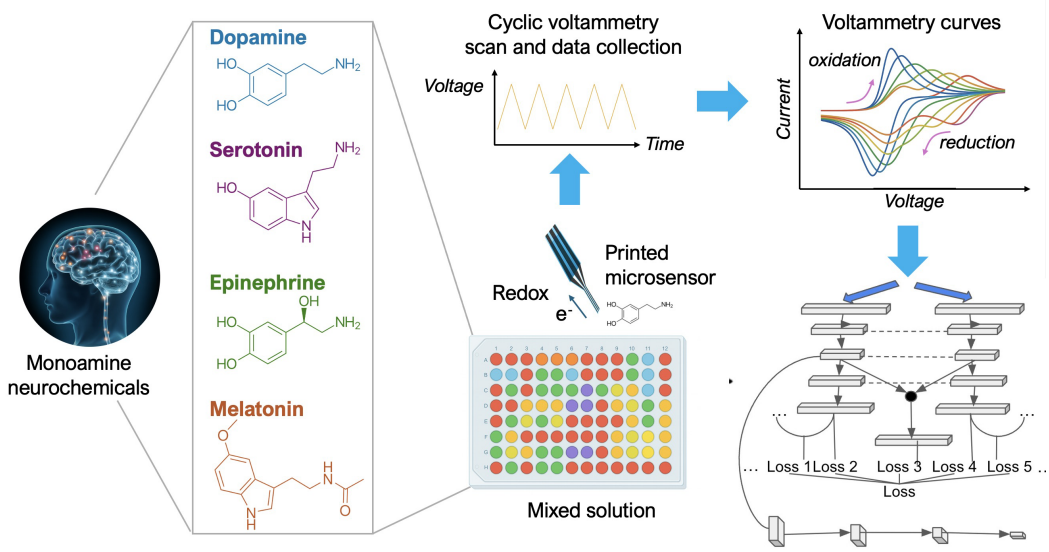

(a)

Figure 1: Overview of biosensing system empowered with AI for concentration prediction.

The system automatically predicts multiplex concentrations from neurochemical sensors using cyclic voltammetry. Figure 1 shows the AI-powered sensing setup for four target compounds—serotonin (5-HT), dopamine (DA), norepinephrine (NE), and melatonin (MEL). These compounds were dissolved in 1X phosphate buffered saline (PBS) to create solutions of known concentrations, which were placed into a neurochemical sensor array with PBS references at one end of each row.

Currents detected by the sensors and electrodes are transmitted to a cyclic voltammetry (CV) monitor, which outputs CV curves. This measurement is typically repeated multiple times, yielding a scan sequence of CV curves. For representation learning, features are extracted from pairs of successive curves within each sequence. The resulting representations are then fed into a downstream model for concentration prediction. The machine learning task is formulated to predict the concentrations of four target compounds from the current CV curve.

Data Process. Each CV curve was unfolded by horizontally flipping the reduction portion and concatenating it with the oxidation portion, yielding a continuous time series of electrical current. For each example, features were derived from the current CV curve and its preceding scan. Five feature modalities were constructed: (1) the average of the current and preceding curves, (2) the difference between the current and preceding curves, and the (3) first, (4) second, and (5) third derivatives of the time series derived from modality (1). Preliminary experiments showed that derivatives beyond the third order did not provide additional performance gains.

2.1 Meta-PairRank (MPR) Framework

Contrastive learning is a self-supervised framework that aligns similar examples while separating dissimilar ones, thereby preserving intrinsic structure in the embedding space and enhancing generalization. Pairwise ranking introduces preferential discrimination between examples with distinct

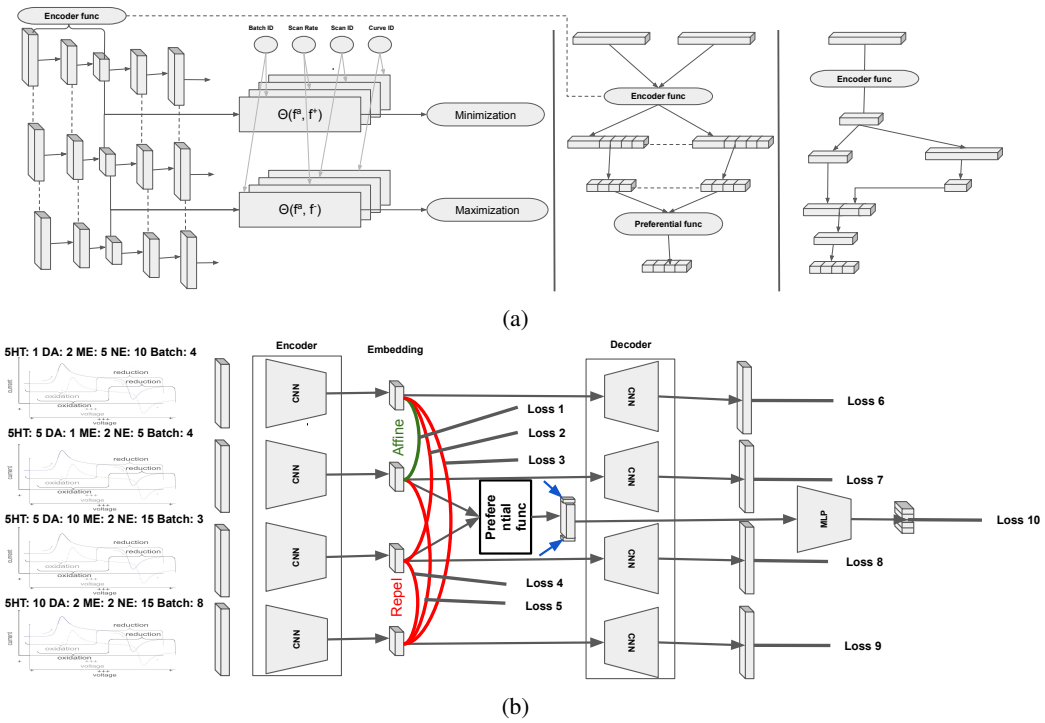

Figure 2: Overview of neurochemical sensing system and machine learning architecture. Machine learning system overview and contrastive loss schematic diagram.

labels. In this work, we assume: (1) CV curves obtained under identical experimental conditions are similar; and (2) CV curves with differing concentration ground truths are dissimilar.

Our representation framework consists of three components:

- Data augmentation ($Aug(\cdot)$): Each input CV curve is augmented using standard time-series techniques (e.g., Gaussian noise), producing jittered variants that share metadata with the original. Augmented and original samples jointly define positives and negatives for Meta-Contrast Loss, while PairRank Loss operates on explicitly constructed CV curve pairs.
- Encoder network ($Enc(\cdot)$): An encoder maps each CV curve to a k-dimensional representation ($k = 32$). Both original and augmented inputs undergo identical transformations, yielding embeddings used for downstream concentration prediction.
- Projection network ($Proj(\cdot)$): A projection head (MLP or linear layer) maps embeddings to a normalized high-dimensional space, where inter-sample distances are computed. As in prior work, $Proj(\cdot)$ is used only during representation learning and discarded during downstream tasks.

Additional architectural modules may be incorporated on top of $Enc(\cdot)$ or $Proj(\cdot)$ as complementary constraints to further regularize representation learning.

Meta-Contrast Loss. Using metadata (experiment batch, FSCV scan rate), we assume curves from the same batch or scan rate are positives, others negatives. The loss, adapted from supervised contrastive [11] and N-pair loss [15], is:

$$\mathcal{L}_{\text{meta}}(\{x_i, x_i^+\}; \text{Proj}) = \frac{1}{N}\sum_{i=1}^{N} \log\left(1 + \sum_{j\neq i} \exp(\text{Proj}(x_i)^T \text{Proj}(x_j^+) - \text{Proj}(x_i)^T \text{Proj}(x_i^+))\right)$$

where N is the number of classes. Anchors are encouraged to align with positives and diverge from negatives.

Pairwise Rank Loss. Learning-to-rank (LoR) methods compare item pairs to capture preference [3]. We adopt a pairwise formulation on CV examples to capture fine-grained discriminative features.

Following the pairing strategy in Meta-Contrast Loss, each anchor is compared with a positive sharing at least one metadata type, with labels defined by relative compound concentrations (1 if anchor > positive, else 0). A shared scoring network $Scr(\cdot)$ on top of $Proj(\cdot)$ outputs scores whose differences predict these preference labels. Multiple pairwise loss functions can be applied; in this work, we adopted logistic loss without loss of generality.

Reconstruction Loss. To preserve information in low-dimensional representations, we employ autoencoders with binary cross-entropy on normalized CV inputs:

$$\mathcal{L}_{\text{reconstr}}(x) = -\frac{1}{N}\sum_i [x_i \log(\hat{x}_i) + (1 - x_i)\log(1 - \hat{x}_i)],$$

where $\hat{x}_i$ is reconstructed from x_i via the decoder applied to $Enc(\cdot)$.

Loss Combination. The total loss is a weighted sum of meta-contrast, pairwise rank, and reconstruction losses:

$$\mathcal{L}_{\text{total}} = \frac{\alpha}{2}(\mathcal{L}_{\text{meta\_batch}} + \mathcal{L}_{\text{meta\_scan\_rate}}) + \beta\mathcal{L}_{\text{pair}} + (1 - \alpha - \beta)\mathcal{L}_{\text{reconstr}}$$

with $\alpha = \beta = \frac{1}{3}$.

2.2 Multiplex Concentration Prediction

Embeddings from $Enc(\cdot)$ are fed into a CNN for multiplex concentration prediction, with batch and scan rate metadata concatenated at a hidden layer. CNNs are selected for their capacity to capture nonlinear feature interactions and integrate with learned representations. To address unseen batches, the model jointly predicts batch IDs; the resulting probability distribution over batches is concatenated with the embeddings for final concentration prediction.

3 Experiments

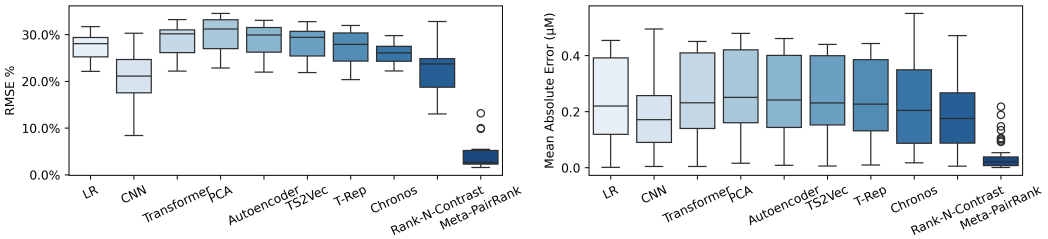

Figure 3: Performance summary on RMSE % (left) and MAE in μM (right).

Figure 3 compares model performance at a scan rate of 20, reported as RMSE (%) and mean absolute error (MAE, μM). Baselines include logistic regression (LR), convolutional neural network (CNN) [13], Transformer [17], PCA [12], Autoencoder [6], TS2Vec [20], T-Rep [5], Chronos [1], and Rank-N-Contrast [21]. Except for CNN, LR, and Meta-PairRank, all representation-based methods were paired with a multiple linear regression model. Most baselines cluster around 25–32% RMSE, with CNN exhibiting greater variability—occasionally reaching lower errors but with less stability. Rank-N-Contrast reduces RMSE to 22.9% with a narrower spread. Meta-PairRank significantly outperforms all baselines (pairwise t-tests against each baseline yield $p < 0.05$), achieving substantially lower RMSE (4.36%) with minimal variance. For MAE, Meta-PairRank again achieves the best results, with a median of 0.033 μM and exceptionally low variability.

Figure 4 shows PaCMAP projections of examples with identical multiplex concentrations. Each color denotes a unique multiplex concentration, with instances at a scan rate of 500 highlighted as large dots. In CNN, same-color dots form clusters that remain distinguishable. In Chronos, these clusters are more tightly packed, leading to greater overlap. In Meta-PairRank, same-color dots overlap substantially, indicating that its representations are more robust to noise affecting concentration prediction than those of the other methods.

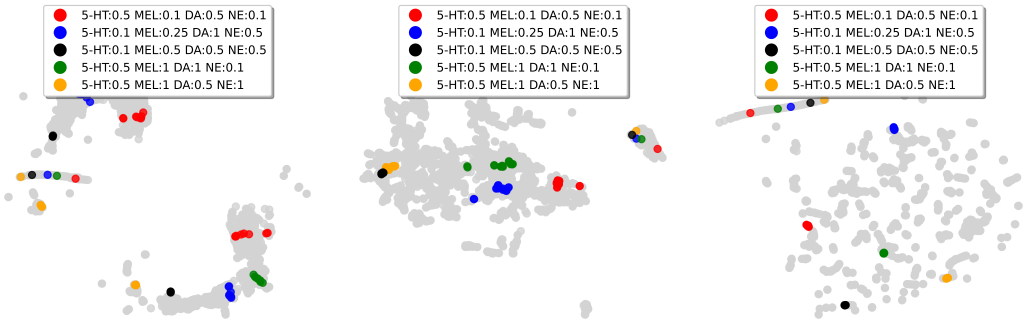

(a) (b) (c)

Figure 4: Visualization of examples with identical multiperplexed concentrations, projected into 2D using PaCMAP [18]. (a) CNN (b) Chronos (c) Meta-PairRank.

Table 1: Performance on metadata ablation.

Views	RMSE (%)	MAE (μM)
Scan rates, modality IDs	4.253 ± 0.295	0.033 ± 0.002
Batch IDs, modality IDs	3.268 ± 0.215	0.024 ± 0.002
Batch IDs, scan rates	3.369 ± 0.229	0.025 ± 0.002
Batch IDs, scan rates, modality IDs	$\mathbf{2.869 \pm 0.189}$	$\mathbf{0.022 \pm 0.001}$

Table 1 reports performance under different metadata ablation settings. Batch ID emerges as the most informative metadata, as its removal produces the largest increase in both RMSE and MAE, followed by scan rate and modality ID.

4 Conclusions

We presented Meta-PairRank, a metadata-aware contrastive framework for predicting multiplex neurochemical concentrations from cyclic voltammetry. By incorporating batch and scan information into representation learning, our approach achieves significantly lower errors and greater robustness than existing baselines, highlighting its potential to advance neurochemical sensing.

Acknowledgments

We thank Vittorio Mottini for contributing the initial version of the preliminary work. We also thank Anna Antoniak for assistance with data collection. This work was supported in part by the National Science Foundation under grants 2339495 and 2334134.

References

[1] A. F. Ansari, L. Stella, C. Turkmen, X. Zhang, P. Mercado, H. Shen, O. Shchur, S. S. Rangapuram, S. P. Arango, S. Kapoor, J. Zschiegner, D. C. Maddix, H. Wang, M. W. Mahoney, K. Torkkola, A. G. Wilson, M. Bohlke-Schneider, and Y. Wang. Chronos: Learning the language of time series, 2024.

[2] A. J. Bard and L. R. Faulkner. *Electrochemical Methods: Fundamentals and Applications*. Wiley, New York, 2nd edition, 2001.

[3] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, page 129–136, New York, NY, USA, 2007. Association for Computing Machinery.

[4] H. Choi, H. Shin, H. U. Cho, C. D. Blaha, M. L. Heien, Y. Oh, K. H. Lee, and D. P. Jang. Neurochemical concentration prediction using deep learning vs principal component regression in fast scan cyclic voltammetry: A comparison study. *ACS Chemical Neuroscience*, 13(15):2288–2297, 2022. PMID: 35876751.

[5] A. Fraikin, A. Bennetot, and S. Allassonnière. T-rep: Representation learning for time series using time-embeddings. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, 2024. arXiv preprint arXiv:2310.04486.

[6] A. Goyal, J. Yuen, S. Sinicrope, B. Winter, L. Randall, A. E. Rusheen, C. D. Blaha, K. E. Bennet, K. H. Lee, H. Shin, and Y. Oh. Resolution of tonic concentrations of highly similar neurotransmitters using voltammetry and deep learning. *Mol. Psychiatry*, Apr. 2024.

[7] M. A. Hayes, E. W. Kristensen, and W. G. Kuhr. Background-subtraction of fast-scan cyclic staircase voltammetry at protein-modified carbon-fiber electrodes. *Biosens. Bioelectron.*, 13(12):1297–1305, Dec. 1998.

[8] J. O. Howell, W. G. Kuhr, R. E. Ensman, and R. Mark Wightman. Background subtraction for rapid scan voltammetry. *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry*, 209(1):77–90, 1986.

[9] E. R. Kandel, J. D. Koester, S. H. Mack, and S. A. Siegelbaum. *Principles of Neural Science*. McGraw-Hill Education, New York, NY, 6 edition, 2021.

[10] R. B. Keithley and R. M. Wightman. Assessing principal component regression prediction of neurochemicals detected with fast-scan cyclic voltammetry. *ACS Chem. Neurosci.*, 2(9):514–525, June 2011.

[11] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan. Supervised contrastive learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc., 2020.

[12] J. Kim, Y. Oh, C. Park, Y. Kang, H. Shin, I. Y. Kim, and D. P. Jang. Comparison study of partial least squares regression analysis and principal component analysis in fast-scan cyclic voltammetry. *International Journal of Electrochemical Science*, 2019.

[13] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.

[14] D. R. Schuweiler, C. D. Howard, E. S. Ramsson, and P. A. Garris. Improving in situ electrode calibration with principal component regression for fast-scan cyclic voltammetry. *Analytical Chemistry*, 90(22):13434–13442, 2018. PMID: 30335966.

[15] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[16] T. Twomey, L. Barbosa, T. Lohrenz, and P. R. Montague. Deep learning architectures for fscv, a comparison, 2022.

[17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[18] Y. Wang, H. Huang, C. Rudin, and Y. Shaposhnik. Understanding how dimension reduction tools work: An empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization. *Journal of Machine Learning Research*, 22(201):1–73, 2021.

[19] Y. Xue, W. Ji, Y. Jiang, P. Yu, and L. Mao. Deep learning for voltammetric sensing in a living animal brain. *Angew. Chem. Int. Ed Engl.*, 60(44):23777–23783, Oct. 2021.

[20] Z. Yue, Y. Wang, J. Duan, T. Yang, C. Huang, Y. Tong, and B. Xu. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8980–8987. AAAI Press, 2022.

[21] K. Zha, P. Cao, J. Son, Y. Yang, and D. Katabi. Rank-n-contrast: learning continuous representations for regression. *Advances in Neural Information Processing Systems*, 36:17882–17903, 2023.

A Challenges of Multiplex Concentration Prediction

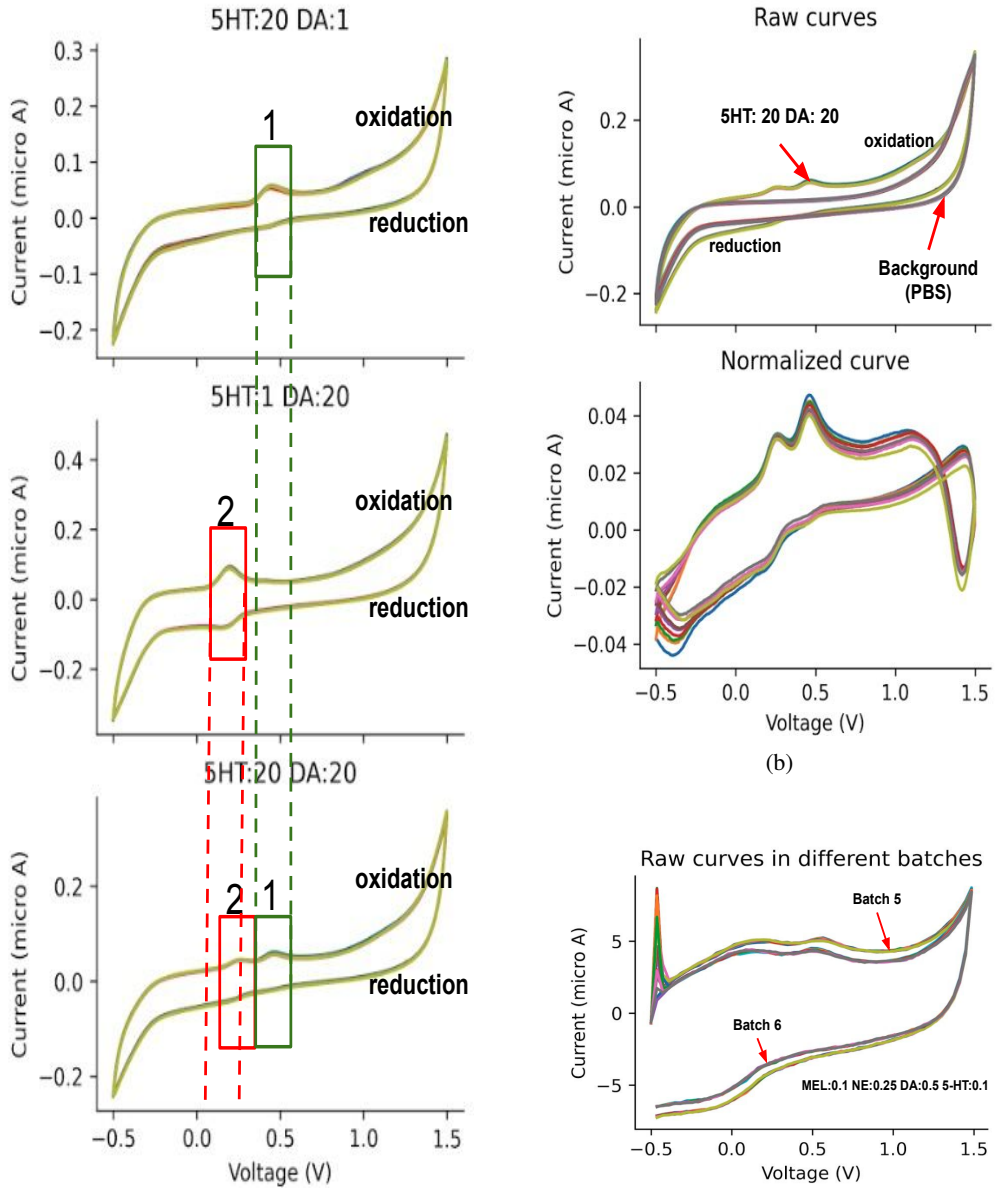

Figure 1: Data examples explaining why traditional machine learning for CV data may fail. 5-HT: serotonin. DA: dopamine. MEL: melatonin. NE: norepinephrine. Concentration unit is μM. (a) Slow scan CV curves varying concentration of analytes. Ranges 1 (green) and 2 (red) are two salient CV portions discriminating the solution predominated by 5-HT and DA. Top: serotonin dominant. Middle: dopamine dominant. Bottom: co-dominant. (b) Slow scan CV curves and their normalized ones by subtracting the background measured from PBS. PBS: phosphate-buffered saline. Top: the raw curves from 20μM serotonin+20μM dopamine, and from PBS. Bottom: raw curve substracting background. (c) FSCV curves collected in two batches from the solution with the same concentration.

Kang et al. [4] and Bond et al. [1] are among the few studies that have addressed the challenges and potential solutions to predict multiplex concentration. The former suggested that CV curves might be inadequate for training, while the latter proposed strategies to improve electrochemical sensing data

for machine learning. Nevertheless, prior research has not systematically examined the fundamental causes underlying these shortcomings. We summarized our discovery as follows.

A.1 Interactions among analytes affect multiplex concentration prediction non-linearly.

One branch of previous work used principle component regression (PCR) [5, 9, 6] because machine learning trained from CV data are prone to over-fitting. CV curve features are high-dimensional, but obtaining measurements from diverse concentrations to support training is difficult. Since PCR is linear, its underlying assumption is that analytes are independently contribute to multiplex concentration prediction and CV curves correlate to analyte concentration. However, our examples in Figure 1a fails to support it. Adding dopamine into serotonin dominant solution delayed dopamine oxidation and reduced its feature saliency during reduction. Dopamine CV features are robust to the addition of serotonin though.

A.2 Normalizing CV curves by subtracting background does not improve prediction.

The background-subtraction technique on voltammetry was first proposed in 1980s [3], and was subsequently used in [2] to maximize sensitivity and selectivity of voltammetry signals. It aims to increase the signal-to-noise ratio or help visualize small faradaic currents (tens of nanoamperes or less) produced by neurotransmitter release during biological stimulus events [8]. However, our preliminary experiment by using this data process approach failed to perform better than using raw CV curves. Mean absolute error on slow scan CV data increases from 1.51 ± 0.04 to $2.36\pm0.20\mu$M based on 5-fold cross validation, suggesting that subtracting background may incorporate more noise than producing sensitive and selective signals from the machine learning perspective. Figure 1b demonstrates that CV data subtracting background fall into a much narrow range and lead to larger relative inconsistency among multiple CV scans from the same solution, even though they seem more sensitive and selective. Movassaghi et al. [8] also discussed concerns of subtracting background for voltammetry data normalization, although from another perspective.

A.3 CV curves can be heterogeneous and inconsistent across trials.

CV curves can be heterogeneous and inconsistent across experiment batches, even though the measured solution has the same analyte concentration. Figure 1c shows one example from FSCV data. Electric current during oxidation and reduction does not well overlap between the data collected in Batch 5 and 6 performed on two days. While various work used computational approaches to predict neurochemical concentration from CV curves, none of them discovered and emphasized this issue, which negatively affects the performance of machine learning for multiplex concentration prediction.

B Experiment Setup

B.1 Data sets

Slow-scan cyclic voltammetry (general) was used to measure solutions at millimolar (mM) concentrations. Serotonin, dopamine, melatonin, and norepinephrine were prepared at 1, 2, 5, and 10 mM, yielding 256 unique multiplex concentrations. For each concentration, ten CV curves were collected, resulting in 2,560 curves in total.

Fast-scan cyclic voltammetry (general) was used to measure solutions at micromolar (μM) concentrations. As in slow-scan CV, serotonin, dopamine, melatonin, and norepinephrine were prepared at 1, 5, 10, and 20 μM, yielding 256 unique multiplex concentrations. FSCV was conducted at scan rates of 20, 50, 100, 200, 300, and 500 per second. For each scan rate and concentration, twenty CV curves were collected, producing 30,720 curves across 16 experimental batches.

Fast-scan cyclic voltammetry (random) consisted of CVs from 16 randomly generated multiplex concentrations, centered at 0.5, 0.3, 0.03, and 0.2 μM for melatonin, norepinephrine, dopamine, and serotonin, respectively. These values approximate estimated levels in the human brain. For each compound, concentrations were sampled from a normal distribution $N(\text{avg}, 1)$, constrained to positive values.

In addition, CVs on pure PBS were collected after every four unique multiplex concentrations, resulting in approximately 25% more PBS curves alongside the multiplex measurements.

The first two datasets were combined to train the representation model and local convolutional neural network. Fine-tuning was performed on PBS curves from the final dataset, with evaluation conducted on solute-containing solutions from the same set.

B.2 Tissue-like Neurotransmitter Sensors

This work builds on recent advances in tissue-like neurotransmitter sensors engineered for soft, complex, and dynamically moving organs [7]. These devices provide an elastic, conformable biosensing interface with promising potential for future in vivo applications. Such designs open opportunities to investigate neurotransmitter dynamics in contexts such as gut–microbiota interactions and brain–gut communication.

References

[1] A. M. Bond, J. Zhang, L. Gundry, and G. F. Kennedy. Opportunities and challenges in applying machine learning to voltammetric mechanistic studies. *Current Opinion in Electrochemistry*, 34:101009, 2022.

[2] M. A. Hayes, E. W. Kristensen, and W. G. Kuhr. Background-subtraction of fast-scan cyclic staircase voltammetry at protein-modified carbon-fiber electrodes. *Biosens. Bioelectron.*, 13(12):1297–1305, Dec. 1998.

[3] J. O. Howell, W. G. Kuhr, R. E. Ensman, and R. Mark Wightman. Background subtraction for rapid scan voltammetry. *Journal of Electroanalytical Chemistry and Interfacial Electrochemistry*, 209(1):77–90, 1986.

[4] M. Kang, D. Kim, J. Kim, N. Kim, and S. Lee. Strategies to enrich electrochemical sensing data with analytical relevance for machine learning applications: A focused review. *Sensors*, 24(12), 2024.

[5] R. B. Keithley and R. M. Wightman. Assessing principal component regression prediction of neurochemicals detected with fast-scan cyclic voltammetry. *ACS Chem. Neurosci.*, 2(9):514–525, June 2011.

[6] J. Kim, Y. Oh, C. Park, Y. Kang, H. Shin, I. Y. Kim, and D. P. Jang. Comparison study of partial least squares regression analysis and principal component analysis in fast-scan cyclic voltammetry. *International Journal of Electrochemical Science*, 2019.

[7] J. Li, Y. Liu, L. Yuan, B. Zhang, E. S. Bishop, K. Wang, J. Tang, Y.-Q. Zheng, W. Xu, S. Niu, et al. A tissue-like neurotransmitter sensor for the brain and gut. *Nature*, 606(7912):94–101, 2022.

[8] C. S. Movassaghi, M. Alcañiz Fillol, K. T. Kishida, G. McCarty, L. A. Sombers, K. M. Wassum, and A. M. Andrews. Maximizing electrochemical information: A perspective on background-inclusive fast voltammetry. *Analytical Chemistry*, 96(16):6097–6105, 2024. PMID: 38597398.

[9] D. R. Schuweiler, C. D. Howard, E. S. Ramsson, and P. A. Garris. Improving in situ electrode calibration with principal component regression for fast-scan cyclic voltammetry. *Analytical Chemistry*, 90(22):13434–13442, 2018. PMID: 30335966.