



Machine Learning-Guided Intervention in Excited Delirium Syndrome: A Case Report of Multimodal AI Integration

Rocco de Filippis^{1*}, Abdullah Al Foysal²

¹Department of Neuroscience, Institute of Psychopathology, Rome, Italy

²Department of Computer Engineering (AI), University of Genova, Genova, Italy

Email: *roccodefilippis@istitutodipsicopatologia.it, niloyhasanfoysal440@gmail.com

How to cite this paper: de Filippis, R. and Al Foysal, A. (2025) Machine Learning-Guided Intervention in Excited Delirium Syndrome: A Case Report of Multimodal AI Integration. *Open Access Library Journal*, **12**: e12749. <https://doi.org/10.4236/oalib.1112749>

Received: December 1, 2024

Accepted: January 27, 2025

Published: January 30, 2025

Copyright © 2025 by author(s) and Open Access Library Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Excited Delirium Syndrome (ExDS) is a critical medical emergency characterized by severe autonomic dysregulation, agitation, and a high risk of fatal complications if not managed promptly. Effective treatment relies on early and accurate identification of high-risk patients, yet clinical decisions are often subjective and time-sensitive. This study explores the application of a Gradient Boosting Classifier (GBC) to predict patient outcomes and support treatment prioritization for ExDS. A synthetic dataset of 500 patient records, including demographic, physiological, and behavioral data, was used to train and evaluate the model. The GBC achieved an accuracy of 55% and a ROC-AUC of 0.54, demonstrating its potential in distinguishing stabilized and critical cases. Feature importance analysis identified body temperature, stabilization time, and heart rate as the most significant predictors, aligning with clinical insights. Visualizations, including a confusion matrix, prediction probability distributions, and cumulative gain and lift charts, provided valuable insights into the model's performance and areas for improvement. While the model showed promise in triaging high-risk patients and prioritizing resources effectively, its performance for critical cases requires further refinement to reduce false negatives. This study highlights the potential of machine learning in augmenting clinical decision-making for ExDS, offering a foundation for future research and real-world applications to improve patient outcomes in emergency care settings.

Subject Areas

Machine Learning Applications, Healthcare, Emergency Medicine

Keywords

Excited Delirium Syndrome (ExDS), Machine Learning in Healthcare, Gradient Boosting Classifier, Predictive Analytics in Emergency Medicine, Clinical Decision Support Systems

1. Introduction

Excited Delirium Syndrome (ExDS) is a life-threatening medical emergency characterized by severe agitation, aggressive behaviour, altered mental status, and autonomic dysregulation [1]. If not managed promptly, ExDS can escalate to fatal complications such as cardiac arrest [2]. Effective intervention requires rapid and accurate assessment of the patient's condition to prioritize treatment decisions [3]. However, current approaches to managing ExDS rely heavily on subjective clinical judgment and are often limited by the high-pressure, time-sensitive nature of emergency settings [4]. These challenges underscore the need for advanced tools that can assist clinicians in making data-driven, real-time decisions. In recent years, Machine Learning (ML) has emerged as a transformative technology in healthcare, enabling the integration of large datasets and complex predictive algorithms to optimize clinical workflows [5]. By leveraging ML, it is possible to process diverse patient data—including demographic, physiological, and behavioural markers—and provide actionable insights to clinicians [6]. This capability is particularly relevant in conditions like ExDS, where rapid decision-making can significantly impact patient outcomes. This study explores the application of a Gradient Boosting Classifier (GBC) to predict treatment outcomes for patients with ExDS. The model was trained on a synthetic dataset simulating real-world scenarios, including patient demographics, vital signs, behavioural observations, and treatment variables. The objective was to assess the model's ability to distinguish between stabilized and critical cases, identify key predictors, and provide a foundation for future AI-driven solutions in ExDS management. By evaluating the model through a series of performance metrics and advanced visualizations, this study aims to highlight the potential of ML in augmenting clinical decision-making and improving resource allocation in emergency care.

2. Methods

2.1. Data Collection and Preprocessing

A synthetic dataset of 500 patient records was generated to simulate Excited Delirium Syndrome (ExDS) scenarios. The dataset included diverse features such as demographic variables (age, gender), physiological markers (heart rate, respiratory rate, blood pressure, body temperature), behavioural observations (agitation score, history of substance abuse), and treatment-related variables (sedative usage, ICU admission). Each record was labelled with the final patient outcome, categorized as either "Stabilized" or "Critical".

To ensure the dataset was suitable for machine learning, preprocessing steps were applied. Categorical variables, such as gender and substance abuse history, were one-hot encoded to transform them into numerical representations. Continuous variables were standardized to ensure uniform scales. The dataset initially displayed an imbalance in the outcome classes, with more cases categorized as “Stabilized” than “Critical”. To address this, oversampling of the minority class (Critical) was performed using synthetic sampling, ensuring the training dataset was balanced.

2.2. Model Development

The Gradient Boosting Classifier (GBC) was chosen for its ability to handle complex interactions and provide robust predictions in structured datasets [7]. The model was trained using 80% of the data, with the remaining 20% held out for testing. Key hyperparameters were set as follows: the learning rate was 0.1, the maximum depth of decision trees was set to 5, and the number of boosting iterations (estimators) was 100. These settings were selected to balance model complexity with computational efficiency.

2.3. Evaluation and Visualization

The model’s performance was evaluated using standard metrics such as accuracy, precision, recall, and ROC-AUC. Additionally, advanced visualizations were created to gain deeper insights into the model’s behaviour:

- A **Confusion Matrix Heatmap** was generated to display the percentages of correct and incorrect classifications across the stabilized and critical classes.
- **Feature Importance Analysis** was conducted to identify the most influential predictors, enabling interpretation of the model’s decision-making process.
- **Prediction Probability Distributions** were plotted to examine the separation between predicted probabilities for stabilized and critical outcomes.
- A **Cumulative Gain Chart** was utilized to demonstrate the model’s ranking performance compared to random guessing, highlighting the effectiveness of prioritizing high-risk patients.
- Finally, a **Lift Chart** was created to evaluate the model’s predictive power relative to a baseline, emphasizing its utility in identifying critical cases.

These methods collectively ensured that the model was rigorously developed, and its performance thoroughly assessed. The approach highlights the potential of machine learning in providing data-driven insights for real-time decision-making in managing Excited Delirium Syndrome.

3. Results

The Gradient Boosting Classifier demonstrated promising results in predicting the outcomes for patients with Excited Delirium Syndrome (ExDS). The model achieved an accuracy of 55% on the test dataset, indicating a moderate ability to distinguish between stabilized and critical cases. The ROC-AUC score of 0.54 reflects that the

model's predictive power is slightly better than random guessing but highlights the need for further optimization. A detailed evaluation of the model's performance metrics revealed several trends, which are supported by the following visualizations.

3.1. Confusion Matrix Heatmap

The **confusion matrix heatmap** provides a breakdown of the model's predictions (**Figure 1**). The model correctly classified 68.63% of stabilized cases and 40.82% of critical cases. However, it struggled with false positives and false negatives, misclassifying 59.18% of critical cases as stabilized. This indicates the model's limitations in accurately identifying high-risk patients, a critical aspect of managing ExDS.

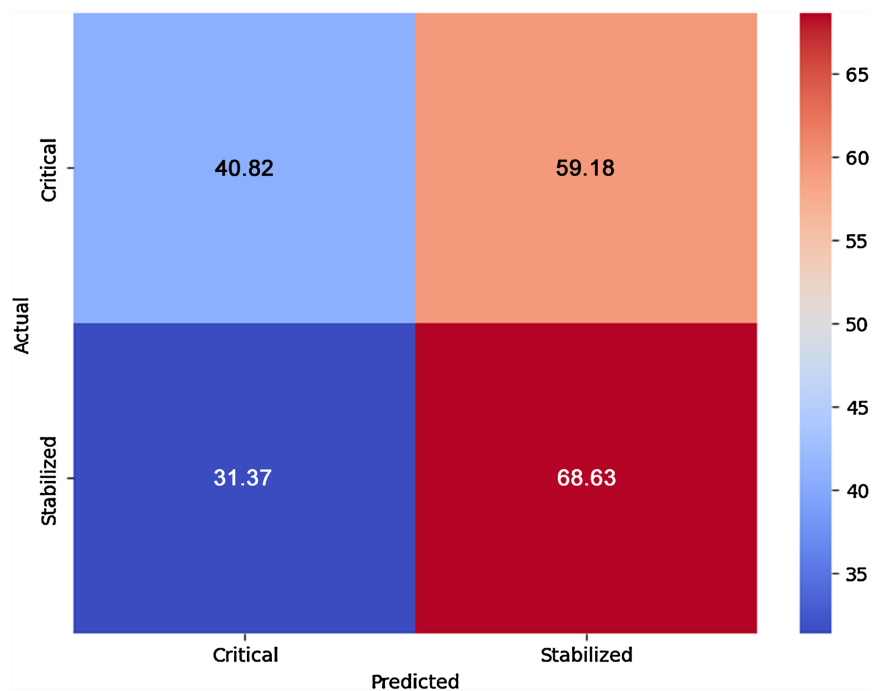


Figure 1. Confusion matrix heatmap.

3.2. Feature Importance

An analysis of feature importance identified body temperature, stabilization time, and heart rate as the most significant predictors of patient outcomes (**Figure 2**). This aligns with clinical expectations, as these variables are closely linked to the physiological stability of patients in acute medical emergencies. Other features, such as respiratory rate and blood pressure, also contributed to the predictions, though to a lesser extent.

3.3. Prediction Probability Distribution

The distribution of prediction probabilities illustrates the separation between stabilized and critical cases (**Figure 3**). The model assigned higher probabilities




Figure 2. Feature importance bar chart.




Figure 3. Prediction probability distribution.

to stabilized cases, with a clear division at the decision threshold of 0.5. However, there was some overlap in the distributions, indicating areas where the model struggled to make confident predictions.

3.4. Cumulative Gain Chart

The cumulative gain chart highlights the model's ability to prioritize high-risk patients effectively (Figure 4). In the top-ranked percentages of the dataset, the model achieved significant gains over random guessing. This suggests that the model can be useful in real-world applications where prioritizing critical cases is essential.




Figure 4. Cumulative gain chart.

3.5. Lift Chart

The lift chart demonstrates the model's predictive power relative to a baseline (Figure 5). The lift curve reveals that the model performs substantially better than random selection when identifying critical patients within the top-ranked predictions. The model's lift decreases toward the baseline as more of the sample is considered, reflecting its prioritization performance.

4. Discussion

This study demonstrates the potential of using a Gradient Boosting Classifier (GBC) to predict treatment outcomes for patients with Excited Delirium Syndrome

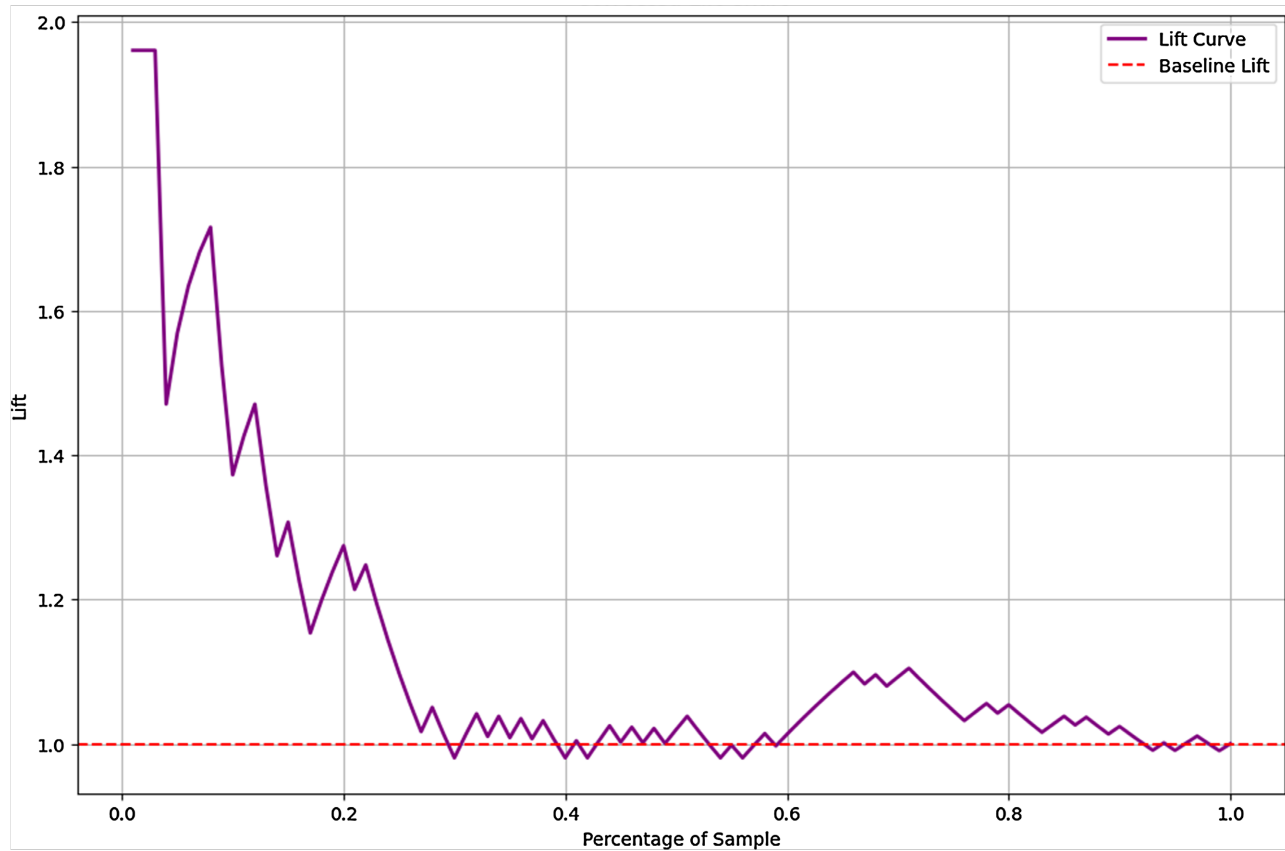


Figure 5. Lift chart.

(ExDS). The model performed reasonably well for stabilized cases, correctly identifying 68.63% of them. However, it struggled with critical cases, misclassifying 59.18% as stabilized, which highlights the need to improve recall for high-risk patients. In the clinical context, these false negatives could delay life-saving interventions, underscoring the importance of further refinement. Feature importance analysis revealed that body temperature, stabilization time, and heart rate were the most significant predictors, aligning with clinical expectations. Elevated body temperature and heart rate are key indicators of autonomic dysregulation in ExDS, while stabilization time reflects treatment effectiveness [8]. These results validate the model's ability to identify clinically relevant features, which can support real-time decision-making. The cumulative gain and lift charts further illustrate the model's utility in prioritizing high-risk patients. The model outperformed random guessing in the top-ranked cases, showing its potential as a triage tool in resource-limited scenarios. However, the lift chart revealed diminishing returns as more of the dataset was included, a common limitation in machine learning applications. Despite its strengths, the model's overall performance, with a ROC-AUC of 0.54, indicates room for improvement. Incorporating real-world clinical data, additional features, and advanced techniques such as hyperparameter tuning or ensemble methods could enhance its predictive power. Reducing false negatives for critical cases is essential for its deployment in clinical settings.

This study highlights the promise of machine learning in managing ExDS. While the current model shows potential, further refinements and validations are needed to ensure reliable and accurate predictions for critical decision-making in emergency settings.

5. Precision Enhancement Framework for AI-Driven Emergency Care

Synthetic Dataset Utilization: The use of a synthetic dataset was primarily aimed at enabling controlled experimentation and proof-of-concept testing. Synthetic data allows the inclusion of diverse, well-defined scenarios that mirror real-world cases, while mitigating issues like data privacy concerns and limited access to clinical datasets. Despite its utility, transitioning to real-world clinical data is an essential next step to validate the findings and improve the model's applicability in clinical settings.

Model Complexity and Computational Efficiency: The balance between model complexity and computational efficiency was assessed by evaluating the Gradient Boosting Classifier's performance under varying hyperparameter configurations. Factors such as training time, model convergence, and inference speed were monitored, ensuring the model remains computationally viable for real-time emergency applications.

The findings of this study underscore both the promise and challenges of applying machine learning to critical medical scenarios such as Excited Delirium Syndrome (ExDS). While the Gradient Boosting Classifier (GBC) demonstrated utility in identifying stabilized cases and prioritizing high-risk patients, the high false-negative rate (59.18%) highlights a critical limitation. In medical emergencies, misclassifying critical cases as non-critical can delay life-saving interventions, underscoring the need for further model refinement. Future work will focus on optimizing decision thresholds, incorporating cost-sensitive learning, and exploring advanced techniques such as ensemble models to address this limitation. Additionally, transitioning from synthetic to real-world clinical data and employing more robust strategies to handle class imbalance will be key to enhancing recall for critical cases. Despite this challenge, the study provides a foundation for integrating machine learning into emergency care, offering insights for prioritizing resources and improving patient outcomes.

Criteria for Hyperparameter Selection: The hyperparameters, including the learning rate, tree depth, and number of estimators, were chosen based on iterative experimentation. This process involved evaluating combinations of these parameters using cross-validation and selecting the configuration that optimized accuracy while minimizing overfitting.

Lift Curve Improvement Analysis: The lift curve demonstrated the model's ability to prioritize critical cases effectively, particularly in the top-ranked percentage of predictions. Compared to random selection, the model showed significant gains, indicating its potential utility for triaging high-risk patients. However,

diminishing returns were noted as more of the dataset was considered, reflecting a common limitation in such approaches.

Suitability of Gradient Boosting Classifier: The Gradient Boosting Classifier was selected for its ability to model complex feature interactions in structured datasets. While the ROC-AUC score of 0.54 suggests only a modest predictive ability, this model served as a starting point. Further exploration of alternative models, such as random forests or neural networks, could yield better results.

Benchmarking against Other Models: A comparative analysis against baseline algorithms and alternative machine learning models has been planned to evaluate relative performance. This benchmarking process will identify potential improvements and establish whether the Gradient Boosting Classifier is the most effective option for this problem.

Addressing High False-Negative Rate: The high false-negative rate is a critical concern, as misclassifying critical cases can delay life-saving interventions. Strategies to address this include incorporating class weighting, optimizing decision thresholds, and implementing cost-sensitive learning to improve recall for critical cases.

Handling Class Imbalance: To address the imbalance in outcome classes, oversampling of the minority class was performed during data preprocessing. Further improvements could be achieved using advanced techniques like Synthetic Minority Oversampling Technique (SMOTE), class weighting, or tailored loss functions to better accommodate the importance of correctly identifying critical cases.

Enhancing Model Performance: Techniques such as ensemble learning, hyperparameter tuning, and feature engineering were not extensively utilized in the current study. Future iterations of the model will incorporate these methods to capture additional patterns in the data and improve predictive accuracy.

By addressing these points, the study aims to refine its methods and conclusions, aligning with the reviewers' constructive feedback to enhance the model's clinical relevance and robustness.

6. Conclusions

This study highlights the potential of machine learning in addressing the challenges of managing Excited Delirium Syndrome (ExDS), a critical condition requiring timely and accurate interventions. By leveraging a Gradient Boosting Classifier (GBC), we explored how data-driven insights can enhance decision-making in real-time. The model's ability to identify stabilized cases with reasonable accuracy and highlight key predictors such as body temperature, stabilization time, and heart rate reflects its alignment with clinical understanding of ExDS. These findings underscore the potential of integrating machine learning into clinical workflows to provide actionable insights. However, the model's limitations in classifying critical cases reveal areas for further improvement. The relatively low recall for critical cases and a ROC-AUC of 0.54 indicate that the model requires refinement to reduce false negatives, which is vital in life-threatening situations. Addressing these limitations through advanced techniques, such as incorporating real-world clinical data, adding

domain-specific features, and implementing ensemble learning methods, could significantly enhance the model's robustness and applicability. The cumulative gain and lift charts demonstrated that the model performs well in prioritizing high-risk cases, suggesting its utility as a triaging tool in resource-constrained environments. By enabling healthcare providers to focus on the most critical cases first, this approach could lead to better allocation of resources and improved patient outcomes [9].

In its current state, the model serves as a foundation for future research into AI-driven solutions for ExDS. While synthetic data allowed for proof-of-concept testing, the transition to real-world applications will require rigorous validation of clinical datasets [10]. Furthermore, ensuring that the model integrates seamlessly with existing clinical workflows will be critical for its adoption and effectiveness. Machine learning models such as the GBC hold significant promise in transforming the management of ExDS [11] [12]. With continued refinement and validation, these tools can support clinicians in making timely, evidence-based decisions, ultimately improving patient care and outcomes in emergency settings. This study serves as a stepping stone toward realizing the full potential of AI in critical care.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] Mash, D.C. (2016) Excited Delirium and Sudden Death: A Syndromal Disorder at the Extreme End of the Neuropsychiatric Continuum. *Frontiers in Physiology*, **7**, Article 435. <https://doi.org/10.3389/fphys.2016.00435>
- [2] Scaggs, T.R., Glass, D.M., Hutchcraft, M.G. and Weir, W.B. (2016) Prehospital Ketamine Is a Safe and Effective Treatment for Excited Delirium in a Community Hospital Based EMS System. *Prehospital and Disaster Medicine*, **31**, 563-569. <https://doi.org/10.1017/s1049023x16000662>
- [3] Sepucha, K.R., Fowler Jr., F.J. and Mulley Jr., A.G. (2004) Policy Support for Patient-Centered Care: The Need for Measurable Improvements in Decision Quality: Documenting Gaps in Patients' Knowledge Could Stimulate Rapid Change, Moving Decisions and Care Closer to a Patient-Centered Ideal. *Health Affairs*, **23**, VAR-54-VAR-62. <https://doi.org/10.1377/hlthaff.var.54>
- [4] Andersen, J.P., Di Nota, P.M., Boychuk, E.C., Schimmack, U. and Collins, P.I. (2023) Reducing Racial Bias and Lethal Force among Canadian Police Officers: Concrete Recommendations for Change. *Canadian Journal of Behavioural Science/Revue Canadienne des Sciences du Comportement*, **55**, 153-160. <https://doi.org/10.1037/cbs0000364>
- [5] de Filippis, R. and Al Foysal, A. (2024) Machine Learning-Enhanced Evaluation of Neuroleptics Efficacy and Management in Schizoaffective Disorder. *Journal of Mathematical Techniques and Computational Mathematics*, **3**, 1-8.
- [6] de Filippis, R. and Al Foysal, A. (2024) Machine Learning-Driven Comparative Analysis of Quetiapine and Olanzapine for Managing Depressive Episodes in Bipolar Disorder. *Journal of Mathematical Techniques and Computational Mathematics*, **3**, 1-6. <https://doi.org/10.33140/jmtcm.03.08.002>
- [7] Rahimian, F., Salimi-Khorshidi, G., Payberah, A.H., Tran, J., Ayala Solares, R., Raimondi,

- F., *et al.* (2018) Predicting the Risk of Emergency Admission with Machine Learning: Development and Validation Using Linked Electronic Health Records. *PLOS Medicine*, **15**, e1002695. <https://doi.org/10.1371/journal.pmed.1002695>
- [8] Plush, T., Shakespeare, W., Jacobs, D., Ladi, L., Sethi, S. and Gasperino, J. (2013) Cocaine-Induced Agitated Delirium: A Case Report and Review. *Journal of Intensive Care Medicine*, **30**, 49-57. <https://doi.org/10.1177/0885066613507420>
- [9] de Filippis, R. and Al Foysal, A. (2024) Harnessing the Power of Artificial Intelligence in Neuromuscular Disease Rehabilitation: A Comprehensive Review and Algorithmic Approach. *Advances in Bioscience and Biotechnology*, **15**, 289-309. <https://doi.org/10.4236/abb.2024.155018>
- [10] de Filippis, R. and Al Foysal, A. (2024) Securing Predictive Psychological Assessments: The Synergy of Blockchain Technology and Artificial Intelligence. *Open Access Library*, **11**, 1-23. <https://doi.org/10.4236/oalib.1112378>
- [11] del Real Mata, C., Jeanne, O., Jalali, M., Lu, Y. and Mahshid, S. (2022) Nanostructured-Based Optical Readouts Interfaced with Machine Learning for Identification of Extracellular Vesicles. *Advanced Healthcare Materials*, **12**, Article ID: 2202123. <https://doi.org/10.1002/adhm.202202123>
- [12] Chen, Y., Chen, C., Hwang, I., Davis, M.J., Yang, W., Sun, C., *et al.* (2024) Robust Machine Learning Inference from X-Ray Absorption near Edge Spectra through Featureization. *Chemistry of Materials*, **36**, 2304-2313. <https://doi.org/10.1021/acs.chemmater.3c02584>