

**Statistical Foundations for Microplastic Identification: Efficient Sampling and  
Distribution-Free Uncertainty Quantification**

by

Eduardo Ochoa Rivera

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Statistics)  
in The University of Michigan  
2026

Doctoral Committee:

Professor Ambuj Tewari, Chair  
Professor Yang Chen  
Professor Yuekai Sun  
Professor Paul M. Zimmerman

Eduardo Ochoa Rivera

eochoa@umich.edu

ORCID iD: 0009-0001-8575-8851

© Eduardo Ochoa Rivera 2026

## DEDICATION

To my family.

## ACKNOWLEDGEMENTS

First, I would like to thank Professor Ambuj Tewari. I am deeply grateful to have been part of the community you have cultivated in the CASI group and to have had your guidance throughout these years. I am especially thankful for your encouragement, which helped me navigate moments of uncertainty and believe in myself. I would also like to thank Professor Yang Chen, Professor Yuekai Sun, and Professor Paul M. Zimmerman for their valuable feedback, support, and service as members of my dissertation committee.

I am also grateful to my peers, collaborators, and friends in the Department of Statistics. I especially want to thank the friends from my cohort who became like family: Daniele Bracale, Angelia Estrada, Dr. Felipe Maia Polo, Dr. Alexander Kagan, Dr. Benjamin Osafo Agyare, and Dr. Kevin Wibisono. Daniele, thank you for your sense of humor, your jokes, and the many wonderful conversations that made the Ph.D. experience much more enjoyable. I am also grateful for your emotional support during moments when life felt overwhelming. Angelia, your kindness and friendship are among the most valuable gifts I gained during my time in Ann Arbor. Thank you for your support, our conversations, and for introducing me to so many amazing food spots. Felipe, my brother from another country, it is still funny to remember the times when some faculty members confused us, sometimes asking me whether I was going back to Brazil for the summer. I can only hope that I have become a little more like you, because you are both an outstanding researcher and a wonderful human being. Alex, thank you for your bright personality. The parties at your place, with people playing music, singing, and enjoying delicious barbecue, are among my most treasured memories. Benjamin, thank you for all the lunches we shared. They were not only lunch breaks; they were moments of relief from the pressure and stress of the Ph.D. Our conversations helped reassure me, and in your friendship I found a safe place where I could be truly vulnerable. Kevin, the kindness of your heart is matched only by your ability to optimize trips, from finding the best restaurants to planning the optimal route while taking into account distance, time, and gas costs. Thank you for the trips we have taken together, and I hope there will be many more to come.

I would also like to thank the Measuring, Mapping, and Modeling Microplastics in Michigan (M4AM) student team: Dr. Madeline Clough, Rebecca Parham, Scarlet Aguilar Martinez, Abbygail Ayala, Joseph Pennacchio, Anna Schellin, Randi Libin-Straub, and Paras

Boruah. Maddie, thank you for our weekly meetings. I truly enjoyed our conversations about science, chemistry, statistics, Ph.D. frustrations, and life in general. Those conversations helped me navigate some of the hardest moments of this journey. Rebecca, thank you for all the effort you put into our collaboration. Your dedication and excellence taught me a great deal about how to produce high-quality research. I am also grateful to everyone else who contributed to this amazing project.

I am thankful for the Fulbright community I found in Ann Arbor. Having the opportunity to meet people from around the world and from many different disciplines was one of the experiences that enriched my time here the most. Francisco Renteria and Marco Valdez, I am grateful that we were part of the Fulbright fellows from Mexico. Thank you for your kindness, your generosity, and for the moments we shared, including that memorable lunch at Kanbu Sushi, which I will always remember with great affection. Patricia Torres, thank you for your friendship, your parties, and your cocktails. You are a key part of our lovely community. Cecilia Gavilan, thank you for the many coffee chats, conversations, and moments of encouragement throughout these years. Hye Ri Yang (Mel), thank you for your friendship, all the laughs, study sessions, trips, meals, and your support through both academic and personal challenges. You will always be *mi mejor amiga*.

I would also like to thank my Latin community. Being able to speak Spanish after an exhausting week of thinking twice about everything in English was such a relief. Diana Vergara, thank you for all the *chisme*, the dancing, the parties, and, most of all, your sincere friendship. I will miss our working sessions in your office at the BSB. Gabriel Patron, I was so happy when I learned that someone from Colombia was joining the Statistics Department, and I am so glad we became friends. I admire the passion you put into your work and your dreams. Marc Brooks, although you were born in LA, I think of you as one of my Latin friends. I was always impressed by how genuinely interested you were in Mexican culture, sometimes even teaching me about it.

I also want to thank my close friends from Mexico, who, despite the distance and despite how much our lives have changed, were always there for me. Hugo Aguilera, thank you for the times I visited Mexico and you made time to see me. Even when we did not speak for months, every time we met it felt as if no time had passed. Also, thank you for visiting me in Ann Arbor, that meant a lot to me. Erika Rodriguez, thank you for the constant effort you make to keep us connected and updated, and for always being there whenever I need someone to listen to my problems. That meant a lot to me, especially during the years when distance made it harder to stay close. Karen Perez, thank you for the memories we have created in different places around the world, from a spontaneous visit to Tulum to celebrating my birthday in San Diego. Each of those moments reminded me that true

friendship can remain strong across distance, time, and change.

To my family, thank you for your unconditional love and for believing in me. To my dad, Gonzalo Ochoa, thank you for supporting me and making my education possible. Thank you for your advice and for teaching me the value of hard work. Nadia Casas, you have been my role model for as long as I can remember. You were always my definition of success, intelligence, and maturity, despite being the youngest adult in the house. After knowing and witnessing the challenges you have faced, I admire and respect you even more, and I thank God for giving me such an amazing sister. Daniel Ochoa, thank you for being an amazing brother. I am so proud of the man you have become. Thank you for your kindness, your sobriety, your sincerity, and your wisdom. Although you are my younger brother, I have learned many things from you, and you have inspired me to go further. Luis Ochoa, thank you for all the years you took care of me, for everything you taught me, for challenging me, and for opening the path for me so many times. To my grandparents, Lupita Ontiveros and Javier Rivera, thank you for taking care of me, for sharing your home and all your love. Thank you for passing on your values and helping shape the person I am today. To my mom, Lupita Rivera, I cannot express in words all my gratitude for the sacrifices and effort you made to provide education, well-being, and happiness, even when doing so was beyond your possibilities. Thank you for making my path easier, even when that meant making yours much harder. Above all, thank you for the unconditional love that made all of this possible.

Finally, to my wife, Karen Callejas. I met you at a pivotal moment in my Ph.D. I was going through many doubts about my journey and questioning whether I should continue in the program or leave it behind. Then I met you, and I knew that I wanted to be with you, even though you were so far away. Despite the challenges, you supported me as I finished my Ph.D. and even joined me in Ann Arbor for my last year. Now, more than three years after meeting you, I can hardly believe that I have finished. Thank you, because without you, I am not sure what would have happened to me.

I am thankful to have had the opportunity to study and conduct research alongside brilliant and passionate people. I have learned valuable lessons that will shape not only my future career but also my life more broadly. To everyone who helped me and made this journey more enjoyable, thank you.

# TABLE OF CONTENTS

DEDICATION . . . . .	ii
ACKNOWLEDGEMENTS . . . . .	iii
LIST OF FIGURES . . . . .	ix
LIST OF TABLES . . . . .	xii
LIST OF APPENDICES . . . . .	xiii
ABSTRACT . . . . .	xiv
 CHAPTER	
<b>1 Introduction . . . . .</b>	<b>1</b>
1.1 Spectral Acquisition . . . . .	1
1.2 Pure Exploration for Logistic Bandit Problems . . . . .	3
1.3 Chemical Identification . . . . .	4
1.4 Conformal Prediction . . . . .	5
1.5 Online Conformal Prediction . . . . .	6
<b>2 Near Optimal Pure Exploration in Logistic Bandits . . . . .</b>	<b>8</b>
2.1 Introduction . . . . .	8
2.1.1 Paper Structure . . . . .	9
2.2 Related Work . . . . .	9
2.3 Background and Preliminaries . . . . .	11
2.3.1 Settings . . . . .	11
2.3.2 Maximum Likelihood Estimator . . . . .	12
2.3.3 Concentration . . . . .	13
2.4 Sample Complexity Lower Bound . . . . .	14
2.5 Logistic Track-and-Stop Algorithm . . . . .	16
2.5.1 Stopping rule . . . . .	17
2.5.2 Sampling rule . . . . .	18
2.6 Sample Complexity of Log-TS . . . . .	20
2.7 Experiments . . . . .	21
2.7.1 BAI . . . . .	21
2.7.2 Thresholding Bandit Problem . . . . .	23

2.8	Discussion . . . . .	23
<b>3</b>	<b>Enhancing Confidence in Microplastic Spectral Identification via Conformal Prediction . . . . .</b>	<b>25</b>
3.1	Introduction . . . . .	25
3.2	Methods . . . . .	28
3.2.1	Commercial Software Library Searching . . . . .	28
3.2.2	Conformal Prediction . . . . .	28
3.2.3	Class-Conditional Conformal Prediction . . . . .	30
3.2.4	Similarity Metrics Used to Compute Hit Quality Index . . . . .	30
3.2.5	Application of Conformal Prediction to Real World Environmental Plastics . . . . .	31
3.3	Results and Discussion . . . . .	31
3.3.1	Conformal Prediction . . . . .	31
3.3.2	Comparing the Performance of KnowItAll and the Conformal Prediction Framework . . . . .	33
3.3.3	Takeaways for the Microplastic Community: The Threshold Used for Labeling an Unknown Spectrum Relies on User-Data, Similarity Metric, and Confidence Guarantee . . . . .	35
3.3.4	Using the Open-Access Code Available on GitHub, a Conformal Prediction Workflow Can be Used to Reduce Manual Spectral Comparison and Increase Confidence in Environmental Microplastic Labeling . . . . .	36
<b>4</b>	<b>Conformal Prediction for Ensembles: Improving Efficiency via Score-Based Aggregation . . . . .</b>	<b>37</b>
4.1	Introduction . . . . .	37
4.2	Background . . . . .	39
4.2.1	Conformal Prediction . . . . .	39
4.2.2	Quantile Envelopes . . . . .	40
4.2.3	Related Works . . . . .	41
4.3	Method . . . . .	42
4.3.1	Multivariate Score Quantile . . . . .	42
4.4	Experiments . . . . .	45
4.4.1	Classification Tasks . . . . .	46
4.5	Discussion . . . . .	47
<b>5</b>	<b>Improved Microplastic Identification from Simultaneously Collected Photothermal Infrared and Raman Spectra using Multi-View Conformal Prediction . . . . .</b>	<b>48</b>
5.1	Introduction . . . . .	49
5.1.1	Spectral Preprocessing . . . . .	51
5.1.2	Database Matching with the Normalized Nearest Neighbor ( $\widetilde{NN}$ ) Similarity Metric . . . . .	52
5.1.3	Conformal Prediction (CP) . . . . .	54
5.1.4	Multi-View Conformal Prediction (MVCP) . . . . .	55

5.1.5	Calibration and Test Datasets for Single-View CP and MVCP Experiments . . . . .	57
5.2	Results and Discussion . . . . .	59
5.2.1	Comparing Thresholds between Single-View CP and MVCP Methods . . . . .	59
5.2.2	Performance Evaluation for Single-View CP and MVCP Methods . . . . .	60
5.2.3	Robustness Assessment based on Spectral Identification Capability . . . . .	64
5.2.4	MVCP Method Validation with MPs+Ambient Sample . . . . .	66
5.3	Conclusions . . . . .	70
<b>6</b>	<b>Online Conformal Prediction: Enforcing monotonicity via Online Optimization . . . . .</b>	<b>71</b>
6.1	Introduction . . . . .	71
6.2	Related Work . . . . .	72
6.3	Setting . . . . .	74
6.4	Methods . . . . .	75
6.4.1	Exponentiated Gradient Descent . . . . .	76
6.4.2	Projected Gradient Descent . . . . .	78
6.5	Experiments . . . . .	78
6.5.1	Metrics . . . . .	79
6.5.2	Baselines . . . . .	79
6.5.3	Synthetic Data: Uniform scores with reflected random-walk drift . . . . .	80
6.5.4	Real Data: US Inflation . . . . .	83
6.6	Discussion . . . . .	86
	APPENDICES . . . . .	88
	BIBLIOGRAPHY . . . . .	119

## LIST OF FIGURES

FIGURE	
1.1 Microplastic identification process . . . . .	2
2.1 Logarithm of sample complexity of the benchmark setup for BAI against dimension of the action space $\mathcal{X}$ . . . . .	22
2.2 Logarithm of sample complexity of the benchmark setup for TBP against dimension of the action space $\mathcal{X}$ . . . . .	24
3.1 CP Workflow, Where Exemplative Labels Corresponding to the Red Environmental MP Spectrum are Illustrated on the Bottom Distribution of HQI Scores from True Positive Matches. . . . .	29
3.2 Empirical confidence and mean prediction set size results of conformal prediction using the Pearson correlation coefficient (PCC) and nearest neighbor (NN) and class-conditional conformal prediction. Here, CC-PCC and CC-NN correspond to the use of class-conditional conformal prediction with PCC and NN, respectively. . . . .	32
3.3 Empirical confidence and mean prediction set size results of conformal prediction applied to 28 environmental plastics using PCC and NN similarity metrics. . . . .	33
3.4 Results of comparison between KnowItAll and CP with PCC and NN similarity metrics in terms of empirical confidence and mean prediction set size for a) PE and b) PP. . . . .	34
3.5 Hit quality index threshold and mean prediction set size as a function of the user-defined guaranteed confidence for FLOPP and FLOPP-e data using conformal prediction and PCC with 11 reference labels. . . . .	36
4.1 CSA provides a principled extension to the standard conformal prediction pipeline by leveraging ideas from higher-dimensional quantile regression to define quantile envelopes $\hat{\mathcal{Q}}$ instead of scalar quantiles $\hat{q}$ . It does so by evaluating a collection of score functions (here $s_1$ and $s_2$ ) over the calibration dataset to define $\mathcal{S}$ , finding quantiles $\{\hat{q}_m\}$ over a set of projection directions $\{u_m\}$ , and taking $\hat{\mathcal{Q}}$ to be the intersection of the resulting half-planes $H(u_m, \hat{q}_m)$ . These quantile envelopes result in more informative prediction regions that can be used in downstream tasks. . . . .	38

5.1	Experimental outline for labeling unknown spectra. (a) Database matching for the PTIR and Raman spectra (pink and green traces, respectively) of an unknown particle. The gray traces underneath are the top-matching spectra for each particle type with their respective $\widetilde{NN}$ scores (i.e., HQI scores) next to them. The resulting single-view CP methods for (b) the PTIR $\widetilde{NN}$ scores and (c) the Raman $\widetilde{NN}$ scores are shown with the returned prediction sets based on the methods' theoretical confidence percentages. (d) The results from the MVCP method, utilizing both PTIR and Raman $\widetilde{NN}$ scores, with the returned prediction sets shown based on the theoretical confidence percentage. Note that each numbered point in (b), (c), and (d) corresponds with the particle types denoted in (a). . . . .	53
5.2	Method for generating MVCP envelope and boundary. . . . .	56
5.3	An example calibration with resulting thresholds for the three CP methods. The histograms for the top HQI scores of the correct particle type are shown for the PTIR and Raman CP methods (pink and green, respectively). Their corresponding quantiles at 95% theoretical confidence are denoted with dashed lines that intercept with the PTIR or Raman HQI axes. The spread of these PTIR and Raman scores for each particle are plotted underneath the histograms, with the MVCP envelope at 95% theoretical confidence shown in gold. . . . .	59
5.4	(a) The averaged prediction set sizes compared with the inputted theoretical confidence value for CP (PTIR), CP (Raman), and MVCP. (b) The calculated empirical confidence (top panel) and box-and-whisker plots of average prediction set size (bottom panel) for each split at 95% theoretical confidence ( $n = 100$ ). Error bars shown for the empirical confidence represent the standard deviation. Filled points in the box plots signify outliers. The error bars shown for the empirical confidence represent standard deviation. . . . .	61
5.5	(a) PTIR and Raman spectra for an example HDPE particle and the top-matching reference library HDPE spectra in dark grey below. (b) Top HQI scores for each particle type with the example particle's spectra compared with the CP method's thresholds. A similar structure was used for an example PS particle's (c) spectra and (d) HQI scores. . . . .	63
5.6	Empirical confidence (top panel) and spread of average prediction set sizes over 100 splits (bottom panel) for (a) Spectral Dataset #1 ( $n = 321$ ) and Spectral Dataset #2 ( $n = 209$ ) as well as (b) Spectral Dataset #1: Unchanged Subset ( $n = 161$ ) and Spectral Dataset #1: Noise-Added Subsets ( $n = 160$ ). Error bars shown for the empirical confidence represent standard deviation. Filled points in the box plots signify outliers. . . . .	65
5.7	(a) The false positive rate (FPR), true positive rate (TPR) and $F_1$ score compared over theoretical confidence. (b) Empirical confidence and (c) prediction set size as a function of theoretical confidence. A dotted line denotes the threshold at 73% theoretical confidence, which was selected by the criterion of a 10% maximum FPR. The values of each metric are shown in light gray. The shaded area around traces represents standard deviation. . . . .	67

6.1	Simulated trajectory of the latent process ( $z_t$ ). The process remains confined to the interval while exhibiting gradual drift, inducing a time-varying score distribution in the synthetic experiment. . . . .	81
6.2	Rolling average of the $\ell_1$ distance between $q_t$ and $q_t^*$ . . . . .	82
6.3	Cumulative average of the sum of mis coverage errors $CE_i$ . . . . .	84
6.4	Prediction intervals across multiple coverage levels in the whole window (bottom panels) and zoom after the 2008 crises (top panels): (a) standard quantile tracker method, (b) quantile tracker with the last iteration projected, (c) exponentiated gradient and (d) projected gradient . . . . .	85
B.1	The calibration score evaluations are first split between those used to define the pre-ordering (green) $\mathcal{S}_C^{(1)}$ and those used to define the final multivariate quantile (red) $\mathcal{S}_C^{(2)}$ . . . . .	107
B.2	The pre-ordering points are projected across a number of directions, after which the $\beta$ quantile is used to define a direction quantile. This defines a half-plane of points that are in the region (blue) and those outside (red). . . . .	108
B.3	We use the intersection of hyperplanes to define the quantile envelope, seeking $\beta^*$ that achieves the desired coverage. . . . .	109
B.4	Using the quantile envelope, the family of nested sets $\mathcal{A}_t$ is defined, in turn defining a partial ordering over $\mathbb{R}^K$ . . . . .	110
B.5	Using the nested family of sets, we expand or contract the envelope appropriately using the data of $\mathcal{S}_C^{(2)}$ to find the final adjustment factor. . . . .	110

## LIST OF TABLES

### TABLE

2.1	Average sample complexities, expressed in thousands, of uniform distribution for BAI. The results are shown for various numbers of arms $K$ . Standard deviations and means were computed across 10 randomized trials. . . . .	22
2.2	Average sample complexities, expressed in thousands, of uniform distribution for TBP. The results are shown for various numbers of arms $K$ . Standard deviations and means were computed across 10 randomized trials. . . . .	24
4.1	Average coverages across tasks for $\alpha = 0.10$ are shown in the top rows and average prediction set sizes in the bottom rows. Both were assessed over a batch of i.i.d. test samples (15% of the validation set from ImageNet). Standard deviations and means were computed across 10 randomized draws of the calibration and test sets. . . . .	47

## LIST OF APPENDICES

### APPENDIX

A Near Optimal Pure Exploration in Logistic Bandits . . . . .	88
B Conformal Prediction for Ensembles: Improving Efficiency via Score-Based Aggregation . . . . .	107
C Online Conformal Prediction: Enforcing Monotonicity via Online Optimization . . . . .	113

## ABSTRACT

Microplastics are an emerging pollutant of global concern, with environmental particles documented across the world. Reliable identification of microplastic particles is essential for quantifying their prevalence and assessing environmental exposure, yet current spectroscopic identification pipelines face several statistical challenges. Measurements can be costly and time-consuming, commonly used spectral matching procedures often lack formal guarantees, and environmental samples can vary over time and across locations. This dissertation develops statistical and machine learning methods for adaptive sampling and uncertainty quantification, with a focus on improving the reliability and efficiency of microplastic spectral identification. First, we study adaptive sampling through pure exploration problems in logistic bandits. We introduce Logistic Track-and-Stop, the first track-and-stop algorithm for general pure exploration problems under a logistic bandit model. The method combines adaptive sampling with a stopping rule based on generalized likelihood ratio statistics and asymptotically matches an approximation to the instance-specific lower bound on expected sample complexity. Second, we develop and implement conformal prediction methods for microplastic spectral identification. We first apply conformal prediction to popular database matching pipelines, highlighting the limitations of practitioner-selected similarity thresholds. We then extend the conformal prediction framework to ensemble and multi-view settings by aggregating nonconformity scores across multiple models or data modalities. In particular, we apply multiview conformal prediction to photothermal infrared and Raman spectra, producing more efficient and robust prediction sets than those obtained from single-view methods. The third part studies online conformal prediction across multiple coverage levels. We leverage online optimization algorithms to enforce nestedness of prediction sets across the full risk spectrum while controlling quantile estimation error. Beyond improving interpretability, jointly estimating multiple coverage levels can improve statistical efficiency by enforcing non-crossing constraints and sharing information across quantiles. Together, these contributions provide tools for accelerating microplastic identification and improving the reliability of the resulting scientific conclusions.

# CHAPTER 1

## Introduction

Microplastics (MPs) research has gained attention in the last decades due to increasing concern about their impact on human health [Prata et al., 2021, Zhang et al., 2025a] and their widespread presence in aquatic [Kanhai et al., 2020], terrestrial [Cusworth et al., 2024], and atmospheric systems [Allen et al., 2019]. Understanding the abundance, composition, and sources of MPs is therefore essential for assessing exposure, studying transport mechanisms, and designing mitigation strategies.

However, identifying MPs is a time-intensive task that involves multiple steps to assess their presence and type in the environment (Figure 1.1). A typical workflow begins with sample collection and processing, followed by microscopic image acquisition, particle detection, spectral acquisition, and chemical identification. Automated methods have emerged as a practical solution to accelerate parts of this pipeline. In particular, machine learning and computer vision tools have been used to detect and segment candidate particles from microscopy images. Nevertheless, visual information alone is usually insufficient to confirm whether a particle is plastic or to determine its polymer identity. Chemical characterization, most commonly through vibrational spectroscopy such as Fourier transform infrared (FT-IR), Raman, or photothermal infrared spectroscopy, remains necessary for reliable identification. On the other hand, most spectral matching methods lack uncertainty quantification, which limits the researchers' ability to assess the reliability of MP identification in environmental samples.

This doctoral dissertation aims to address these challenges by developing statistical methods to improve the efficiency and reliability of MP research. In particular, we focus on two key stages of the identification pipeline: spectral acquisition and chemical identification.

### 1.1 Spectral Acquisition

After collecting and processing atmospheric samples, the next steps in the usual pipeline are microscopic image acquisition and particle detection. These steps help identify candidate

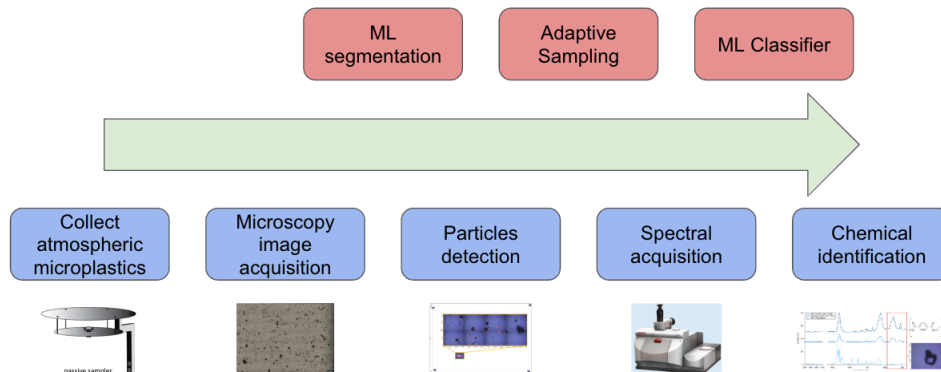


Figure 1.1: Microplastic identification process

particles that may correspond to microplastics. However, visual identification only provides information about suspected microplastics. Chemical characterization is needed to confirm whether a particle is plastic and to determine its polymer identity [Prata et al., 2024]. Therefore, spectral acquisition is a key step for accurate MP identification.

Spectral acquisition typically relies on vibrational spectroscopy methods, especially Fourier transform infrared (FT-IR) and Raman spectroscopy. These techniques are widely used in microplastic analysis because they provide molecular information that can be compared against reference spectra for polymer identification [Araujo et al., 2018, K appler et al., 2016, Xu et al., 2019]. In practice, once candidate particles have been detected, spectra are collected from individual particles or from selected regions of a sample. The resulting spectra are then used in a chemical identification step, often through database matching or related spectral classification methods.

A first challenge in this step is the cost of acquiring reliable spectral information. Spectroscopic measurements can be time-consuming, especially when the number of candidate particles is large. This is particularly relevant in environmental samples, where many particles may be present but only a subset may correspond to MPs. In addition, the measurement time and resolution can depend on the spectroscopic technique, particle size, and desired level of spatial detail. For example, Raman spectroscopy can be useful for smaller particles, but it may require longer acquisition times, while FT-IR methods can be faster in some regimes but may face limitations for very small particles [K appler et al., 2016, Xu et al., 2019, Song et al., 2021].

These practical constraints make spectral acquisition a bottleneck in the MP identification pipeline. Since not every candidate particle can always be measured with the same level of detail, researchers often face a trade-off between throughput and reliability. Measuring more

particles can improve the representativeness of a study, but collecting high-quality spectra for every candidate can be costly. Conversely, reducing the number of spectra can make the analysis more feasible, but may introduce uncertainty about the abundance and composition of MPs in the sample. This motivates the need for principled strategies to decide which particles, regions, or measurements are most informative.

## 1.2 Pure Exploration for Logistic Bandit Problems

The previous section motivates the use of an adaptive sampling strategy. Rather than fixing the sampling plan in advance, an adaptive procedure uses the information collected from previous measurements to decide what should be measured next.

Bandit problems provide a mathematical framework for sequential decision-making under uncertainty. At each round, a learner selects an action, observes a noisy outcome, and then uses this information to decide which action to select next. Classical bandit problems often focus on regret minimization, where the objective is to maximize cumulative reward over time [Thompson, 1933, Robbins, 1952]. However, many scientific problems have a different goal: the objective is not to collect large reward, but to identify the correct answer as efficiently as possible. This is the setting of pure exploration [Bubeck et al., 2009, Garivier and Kaufmann, 2016]. Examples include best-arm identification, where the goal is to identify the action with the largest expected response; thresholding bandits, where the goal is to classify which actions exceed a given threshold [Locatelli et al., 2016]; and top- $m$  identification, where the goal is to recover the best subset of actions.

Pure exploration can be well suited for adaptive sampling in scientific workflows. In these settings, each measurement can be interpreted as an experiment that reduces uncertainty about an underlying scientific question. For spectral acquisition, the actions could correspond to candidate particles, spatial regions of a sample, acquisition settings, or measurement modalities. An adaptive sampling strategy can then prioritize measurements that are most useful for resolving the target question, instead of spending equal effort on all candidates.

Thresholding bandits and related pure exploration tasks are particularly relevant for adaptive scientific measurement. In a thresholding formulation, for instance, one may want to determine which particles or locations have responses above a pre-specified level. In the context of MP analysis, this could correspond to deciding which candidate particles have sufficient evidence of being plastic, which regions contain high concentrations of suspected MPs, or which measurements are informative enough to justify more detailed chemical characterization. The goal is not necessarily to assign a final polymer label at this stage, but to allocate limited measurement effort toward the most informative parts of the sample.

On the other hand, the logistic regression model is a natural choice when observations are binary, providing a model for the log-odds as a linear function of the covariates. For example, a response may indicate whether a candidate particle is a MP or not. In such cases, logistic bandits model the probability of success as a nonlinear function of particle or measurement features. This is useful for adaptive sampling because it allows information to be shared across related actions through their covariates (for example their optical features).

However, pure exploration in logistic and other generalized linear bandits is more technically challenging than in classical multi-armed or linear bandit settings. In linear bandits, the relationship between actions and expected rewards is linear, which leads to relatively tractable confidence sets and sampling rules. In logistic bandits, the curvature of the likelihood depends on the unknown parameter and on the sampled actions, making both estimation and optimal allocation more difficult. Existing work on generalized linear bandits has addressed regret minimization and, more recently, best-arm identification [Filippi et al., 2010, Faury et al., 2020, Kazerouni and Wein, 2021, Jun et al., 2021]. Nevertheless, the broader use of logistic pure exploration for adaptive scientific sampling remains a developing area.

### 1.3 Chemical Identification

After spectra are acquired from candidate particles, the next step is to determine their chemical identity. In microplastic analysis, this is commonly done using vibrational spectroscopy data, such as FT-IR, Raman, or photothermal infrared spectra. These techniques provide information about the molecular structure of the particle, allowing researchers to distinguish plastic particles from non-plastic materials and to identify the corresponding polymer type [Araujo et al., 2018, Kappler et al., 2016, Cowger et al., 2020].

Chemical identification is commonly performed through spectral database matching, where an unknown spectrum is compared against a reference library and assigned one or more candidate labels based on similarity scores. These scores, often reported as hit quality index (HQI) values, measure the likeness between the unknown spectrum and reference spectra [Weisser et al., 2022, Cowger et al., 2020]. Under specific preprocessing choices and similarity metrics, spectral database matching can be viewed as a nearest-neighbor search over reference spectra. In practice, commercial and open-source software tools return ranked lists of candidate matches, and researchers often accept the polymer label with the highest HQI score if it exceeds a chosen threshold.

Although spectral database matching is widely used, the interpretation of HQI scores remains challenging. HQI values do not provide statistical confidence. A high HQI score

indicates that two spectra are similar under a particular similarity metric, but it does not directly quantify the probability that the assigned polymer label is correct. As a result, commonly used thresholding rules may be difficult to interpret and may vary across instruments, reference libraries, preprocessing choices, and similarity metrics.

This issue is especially important because environmental spectra are often noisy and can differ from pristine reference spectra. Weathering, additives, biofilms, surface contamination, and measurement conditions can alter spectral features and make polymer identification more difficult [Renner et al., 2017, Binda et al., 2024]. Moreover, different polymers can have similar spectral signatures, which may lead to ambiguous matches or multiple plausible labels [Jung et al., 2018, Nava et al., 2021]. In these cases, relying on a fixed HQI threshold can lead to misidentification or inconsistent reporting across studies.

Therefore, spectral and chemical identification would benefit from uncertainty quantification methods that describe the reliability of assigned labels and indicate when several polymer identities remain plausible.

## 1.4 Conformal Prediction

Modern machine learning methods can produce highly accurate predictions, but in scientific applications accuracy alone is not enough. Researchers often need to know when a prediction is reliable, when multiple labels are plausible, and how uncertainty should be propagated into downstream conclusions.

Conformal prediction provides a distribution-free framework for constructing prediction sets with finite-sample coverage guarantees [Vovk et al., 2005, Shafer and Vovk, 2008, Angelopoulos and Bates, 2023]. Given a desired confidence level, conformal methods return a set of possible labels or outcomes that contains the true value with a user-specified probability, under appropriate exchangeability assumptions.

In classification problems, conformal prediction returns a set of labels, the size of this set provides a natural measure of uncertainty: a small set indicates that the model is relatively confident, while a large set indicates that several labels remain plausible. In some cases, the conformal set may be empty, suggesting that the observation is not well represented by the calibration data or that none of the candidate labels provides sufficient evidence.

This framework is naturally suited to chemical identification from spectroscopy data. In spectral database matching, an unknown spectrum is compared against reference spectra using similarity scores such as HQI values. These scores rank candidate polymer labels, but they do not directly quantify the probability that a label is correct. Conformal prediction can use calibration spectra to transform these scores into prediction sets with a formal confidence

guarantee. As a result, instead of relying only on a fixed similarity threshold, the analyst can report a set of candidate chemical identities together with an interpretable statistical guarantee.

The validity of conformal prediction depends on the calibration data being representative of future observations. In the context of microplastic identification, this means that the spectra used for calibration should reflect the types of particles, measurement conditions, preprocessing steps, and reference libraries expected in the unknown samples. When this assumption is reasonable, conformal prediction provides a practical way to complement spectral matching with uncertainty quantification, helping researchers distinguish confident identifications from ambiguous cases.

## 1.5 Online Conformal Prediction

Conformal prediction provides finite-sample coverage guarantees under exchangeability assumptions. Although exchangeability is a relatively mild condition compared to parametric modeling assumptions, it may fail in many modern applications. Data may arrive sequentially, the underlying distribution may change over time, or the sequence of observations may exhibit temporal dependence. In some settings, the data can even be viewed as adversarial. These challenges motivate online conformal prediction, which studies how to construct prediction sets when observations are received one at a time and classical exchangeability-based guarantees are no longer directly applicable.

In the online setting, prediction sets must be produced before observing the corresponding outcome. After the outcome is revealed, the method updates its internal parameters, such as a conformal threshold, and then proceeds to the next observation. A common goal is to achieve long-run coverage, meaning that the empirical miscoverage rate over time should approach the target level  $\alpha$ . Thus, instead of requiring finite-sample validity under exchangeability, online conformal methods aim to maintain calibration over a sequential data stream, potentially under distribution shift or adversarial ordering [Gibbs and Candes, 2021, Zaffran et al., 2022, Angelopoulos et al., 2024].

A standard approach updates the conformal threshold based on past coverage errors. If the method undercovers, the threshold is increased to produce larger prediction sets; if it overcovers, the threshold is decreased to produce smaller prediction sets. This feedback mechanism makes online conformal prediction closely related to online learning and online convex optimization. The resulting methods are attractive because they are simple, adaptive, and can be applied to any underlying predictive model through a suitable nonconformity score.

However, in many applications, uncertainty must be reported not at a single confidence level, but across an entire range of risk levels. Different decision-makers may require different levels of conservativeness depending on their tolerance for error. For example, one user may require a narrow prediction set for routine decisions, while another may require a much more conservative set for high-stakes decisions. In such cases, it is natural to construct prediction sets for multiple values of  $\alpha$  simultaneously.

Running standard online conformal methods independently for each coverage level can create an additional difficulty: the resulting prediction sets may fail to be nested. That is, a higher-confidence prediction set may not contain a lower-confidence prediction set. This violates a basic interpretability requirement for uncertainty quantification, since increasing the desired confidence level should lead to a larger or equal prediction set. This issue is analogous to quantile crossing in quantile regression, where independently estimated quantiles can violate their natural ordering [Zou and Yuan, 2008, Liu and Wu, 2011].

Online conformal prediction across multiple coverage levels therefore raises two related problems. The first is calibration: each prediction set should achieve its target long-run coverage level. The second is coherence: prediction sets should be nested across coverage levels. Addressing both problems requires methods that update multiple conformal thresholds jointly, rather than treating each coverage level as an independent task.

## CHAPTER 2

### Near Optimal Pure Exploration in Logistic Bandits

Bandit algorithms have garnered significant attention due to their practical applications in real-world scenarios. However, beyond simple settings such as multi-arm or linear bandits, optimal algorithms remain scarce. Notably, no optimal solution exists for pure exploration problems in the context of generalized linear model (GLM) bandits. In this paper, we narrow this gap and develop the first track-and-stop algorithm for general pure exploration problems under the logistic bandit called logistic track-and-stop (Log-TS). Log-TS is an efficient algorithm that asymptotically matches an approximation for the instance-specific lower bound of the expected sample complexity up to a logarithmic factor.<sup>1</sup>

#### 2.1 Introduction

The multi-arm bandit (MAB) problem is one of the most important and classical problems in sequential decision-making under uncertainty, and it has been studied for nearly a century [Thompson, 1933, Robbins, 1952]. The most common setting involves regret minimization, where the goal is to minimize regret over a finite time horizon. This scenario has been extensively studied under linear reward functions [Abe and Long, 1999, Dani et al., 2008, Rusmevichientong and Tsitsiklis, 2010]. Extensions of the linear case include generalized linear models (GLM) [Filippi et al., 2010, Fauray et al., 2020], Lipschitz bandits [Bubeck et al., 2012] and spectral bandits [Valko et al., 2014].

On the other hand, the pure exploration setting has gained significant attention in recent years, particularly in the context of best arm identification (BAI), which is well understood within the MAB framework [Garivier and Kaufmann, 2016] and in stochastic linear bandits [Soare, 2015]. Other pure exploration problems have also been studied, such as the thresholding bandit problem (TBP) [Locatelli et al., 2016, Kano et al., 2019], and top-m arm identification [Bubeck et al., 2013, Kalyanakrishnan et al., 2012]. Similarly to the regret

---

<sup>1</sup>This chapter is adapted from previous preprint [Ochoa Rivera and Tewari, 2024] work with Ambuj Tewari

minimization problem, there are extensions to the GLM case for BAI [Kazerouni and Wein, 2021, Jun et al., 2021]. However, these algorithms use loose inequalities or require warm-up phases that can be prohibitive in practise and they only can be applied to BAI. In this work we narrow this gap with the following contributions:

- We propose Log-TS, the first track-and-stop type algorithm for general pure exploration problems in the logistic bandit setup.
- We prove both in-expectation and almost sure upper bounds for the sample complexity of Log-TS.
- We provide a lower bound for the expected sample complexity of general pure exploration problems (including BAI, top-m, and thresholding bandits) and a tractable approximation. Log-TS matches this lower bound asymptotically up to a logarithmic factor.
- We confirm the practical performance of Log-TS for the classical hard instance for pure exploration problems and when the number of arms increases.

### 2.1.1 Paper Structure

The reminder of this paper is organized as follows:

In Section 2.2 we discuss the previous work on GLM bandits and pure exploration problems. In Section 2.3 we formulate pure exploration problems under logistic bandits. We also state the definitions and results needed for the construction of the algorithm. In Section 2.4 we present an instance-specific lower bound for the expected sample complexity of *general* pure exploration problems and a tractable approximation. In Section 2.5 we define the components needed for Log-TS: a stopping rule and a sampling rule. In Section 2.6 we state the asymptotic upper bounds for the sample complexity of Log-TS. In Section 2.7 we describe the numerical experiments for two specific pure exploration problems: BAI and TBP. Finally, in Section 2.8 we discuss the presented results and point out future directions.

## 2.2 Related Work

Extensions of stochastic linear bandits have attracted attention due to the restrictive assumption of linearity in real-world applications. For instance, when observing binary rewards, modeling the mean reward as a linear function can be inaccurate. A natural extension in such cases is generalized linear model (GLM) bandits, particularly logistic bandits. GLM bandits

were studied by Filippi et al. [2010], where they used a tail inequality similar to the one in the linear case [Rusmevichientong and Tsitsiklis, 2010], combined with the worst-case behavior of the non-linearity of the link function  $\mu$ ,  $\kappa = \sup_{x \in \mathcal{X}, \theta \in \Theta} 1/\mu'(x^\top \theta)$ . More precisely,  $|\mu(X_t^\top \theta_t^{(1)}) - \mu(X_t^\top \theta)| \leq \rho(t, \delta)$  with probability  $1 - \delta$  where  $\rho(t, \delta) = \mathcal{O}(\kappa \sqrt{d \log(t) \log(d/\delta)})$ ,  $\theta_t^{(1)}$  is the projection of the maximum likelihood estimator (MLE) estimator and  $X_t$  is any random variable in  $\mathcal{X}$ .

As pointed out in Faury et al. [2020],  $\kappa$  can be restrictively large for certain link functions in real applications, such as in the case of logistic bandits. It can be shown that  $\kappa \geq \exp\left(\max_{x \in \mathcal{X}} |x^\top \theta_*|\right)$  for the logistic model. To address this, Faury et al. [2020] proposed a new tail inequality that takes into account the local curvature of the link function. This inequality is independent of  $\kappa$ , significantly improving the regret upper bound. Another key property exploited in their work is the self-concordance of the logistic loss, which helps to bound the prediction errors  $|\mu(x^\top \theta_*) - \mu(x^\top \hat{\theta}_t)|$  using their tail inequality.

A common assumption in GLM bandits is that  $\|\theta^*\| < S$ . While this assumption helps control errors, it also adds complexity, as some algorithms require projecting the maximum likelihood estimator (MLE) onto the parameter space  $\Theta = \{\theta \in \mathbb{R}^d : \|\theta\| < S\}$ , which can be computationally intensive. In Russac et al. [2021], the authors exploit the self-concordance property more effectively, eliminating the need for this projection step.

Another challenge in GLM bandits is that the MLE and the Fisher information matrix cannot be updated recursively, which increases the number of operations per round. This issue is addressed by Faury et al. [2022], who propose an online procedure with a warm-up phase.

Track and stop (TS) algorithms have been among the most common and successful approaches in pure exploration problems due to their asymptotic optimality [Garivier and Kaufmann, 2016, Jedra and Proutiere, 2020, Degenne et al., 2020, Wang et al., 2021b]. The core of these algorithms is to track the oracle proportions of arm draws defined by the sample complexity lower bound. Then, stopping rules are designed using the generalized likelihood ratio. While this approach often involves complex optimization, Jedra and Proutiere [2020] demonstrated that it is possible to achieve asymptotic optimality even when the optimal weights are not updated at every time, a condition known as the lazy setting.

However, no version of the TS algorithm exists for GLM bandits due to the complexity of the lower bound. The first pure exploration work in GLM bandits that we are aware of is by Kazerouni and Wein [2021]. In their study, they used a loose inequality from Filippi et al. [2010], but as other authors have noted, relying on this worst-case inequality can lead to a dependency on an exponential factor. A more recent approach was proposed by Jun et al. [2021], who used a sharper inequality from Faury et al. [2020], accounting for the curvature

of the logistic function. They designed an algorithm inspired by RAGE [Fiez et al., 2019] which is an algorithm for BAI in the linear bandit case and provided a high-probability upper bound for the sample complexity. Additionally, they derived an instance-specific sample complexity lower bound. However, their algorithm requires a warm-up phase that depends on  $\kappa_0 = \sup_{x \in \mathcal{X}} 1/\mu(x^\top \theta^*)$  and the number of arms, which can be restrictive in real-world applications.

## 2.3 Background and Preliminaries

**Notation** For any vector  $x \in \mathbb{R}^d$  and any positive definite matrix  $\mathbf{M} \in \mathbb{R}^{d \times d}$ , we define  $\|x\|_{\mathbf{M}} := \sqrt{x^\top \mathbf{M} x}$  as the  $\ell^2$ -norm of  $x$  weighted by  $\mathbf{M}$ . When  $\mathbf{M} = \mathbf{I}_{d \times d}$  is the identity matrix, we simply write  $\|x\| := \|x\|_{\mathbf{M}}$ . We define  $\lambda_{\min}(\mathbf{M})$  and  $\lambda_{\max}(\mathbf{M})$  the smallest and largest eigenvalue of  $\mathbf{M}$  respectively. We also denote by  $\text{Tr}(\mathbf{M})$  the trace of the matrix. For two matrices  $\mathbf{A}$  and  $\mathbf{B}$ ,  $\mathbf{A} \succ \mathbf{B}$  means that  $\mathbf{A} - \mathbf{B}$  is positive definite. We define  $\mathcal{B}(d) := \{x \in \mathbb{R}^d : \|x\| \leq 1\}$  the  $d$ -dimensional ball of radius 1 under the norm  $\ell^2$ . For an univariate function  $f$  we define  $\dot{f}$  its derivative. We define  $\Sigma := \{w \in [0, 1]^K : \sum_k w_k = 1\}$  the  $K - 1$  simplex. For any  $w, w' \in \Sigma$ , we define  $d_\infty(w, w') = \max_{k \in [K]} |w_k - w'_k|$ , and for any compact set  $C \subseteq \Sigma$ ,  $d_\infty(w, C) = \min_{w' \in C} d_\infty(w, w')$ . Finally, for  $w \in \Sigma$ , we define  $\text{supp}(w) = \{i \in [K] : w_i > 0\}$ .

### 2.3.1 Settings

Let  $\mathcal{X} \subseteq \mathbb{R}^d$  be a finite set of arms, where  $|\mathcal{X}| = K$  and a unknown parameter  $\theta^* \in \Theta$ . We consider the stochastic logistic bandit, where at each round  $t \geq 1$  the decision maker selects an arm  $x_t \in \mathcal{X}$  according to a sampling rule based on previously observed samples and obtains a reward  $r_t \sim \text{Bernoulli}(\mu(x_t^\top \theta^*))$ , where  $\mu(x) = \frac{1}{1+e^{-x}}$ . It then proceeds to the next round.

We are interested in general pure exploration problems, where the goal is to identify the true answer  $i^*(\theta^*)$  that belongs to a finite set  $\mathcal{I}$  of possible answers (e.g., for best arm identification  $i^*(\theta^*) = \arg \max_{x \in \mathcal{X}} \mu(x^\top \theta^*)$ ). We will assume  $i^*(\theta^*)$  is unique. We consider the fixed confidence setting, where the objective is to accurately identify  $i^*(\theta^*)$  with high probability as soon as possible. Formally, the sampling rule defines for all  $t \geq 1$  a function  $\pi_t$  from  $(\mathcal{X} \times \{0, 1\})^{t-1}$  to the space of probability distributions on  $\mathcal{X}$ , which is measurable with respect to the  $\sigma$ -algebra  $\mathcal{F}_t := \sigma(\{x_s, r_s\}_{s \leq t})$ . We call that  $\sigma$ -algebra history before time  $t$ . At time  $\tau$ , where  $\tau$  is an stopping time with respect to  $\mathcal{F}_t$ , and given an estimator  $\hat{\theta}_\tau$  of  $\theta^*$  the algorithm stops and the decision  $i^*(\hat{\theta}_\tau)$  is made. We say the algorithm is  $\delta$ -correct

if  $\mathbb{P}_\theta[\tau_\delta < \infty, i^*(\hat{\theta}_\tau) \neq i^*(\theta)] < \delta$ . Then, the goal is to design a  $\delta$ -correct algorithm that minimize the expected sample complexity  $\mathbb{E}_\theta[\tau_\delta]$ .

**Assumptions** We will make the usual assumption in the logistic bandit problem

- $\|x\| \in \mathcal{B}(d)$  for all  $x \in \mathcal{X}$
- $\Theta = \{\theta \in \mathbb{R}^d : \|\theta\| \leq S\}$ ,  $S > 0$
- $\mathcal{X}$  spans  $\mathbb{R}^d$
- We have access to  $\kappa_0 := \sup_{x \in \mathcal{X}} 1/\dot{\mu}(x^\top \theta^*)$

### 2.3.2 Maximum Likelihood Estimator

For logistic regression setting, we can estimate the parameter  $\theta^*$  using the MLE. At time  $t$ , the log-likelihood can be expressed as:

$$\begin{aligned} \mathcal{L}_t(\theta) &= \sum_{s=1}^t r_s \log \mu(x_s^\top \theta) \\ &\quad + (1 - r_s) \log(1 - \mu(x_s^\top \theta)) \end{aligned}$$

and the MLE is given by  $\hat{\theta}_t = \arg \max_{\theta \in \mathbb{R}^d} \mathcal{L}_t(\theta)$ . We also define Fisher information or Hessian matrix at  $\theta$  as

$$\mathbf{H}_t(\theta) = \sum_{s=1}^t \dot{\mu}(x_s^\top \theta) x_s x_s^\top, \quad (2.1)$$

and the design matrix as

$$\mathbf{A}_t = \sum_{s=1}^t x_s x_s^\top. \quad (2.2)$$

Similarly, for  $w \in \Sigma$ , we define

$$\mathbf{H}_w(\theta) = \sum_{x \in \mathcal{X}} w_x \dot{\mu}(x^\top \theta) x x^\top.$$

Note that if  $w_x = N_t(x)/t$ , where  $N_t(x)$  is the number of times the arm  $x$  has been selected, then  $t\mathbf{H}_w(\theta) = \mathbf{H}_t(\theta)$ . We also denote  $g_t(\theta) = \sum_{s=1}^t \mu(x_s^\top \theta) x_s$ . This function plays an important role in the concentration inequalities as shown in Faury et al. [2020]. In particular, we have that  $g(\hat{\theta}_t) = \sum_{s=1}^t r_s x_s$  by definition of the MLE.

### 2.3.3 Concentration

We will use the recent concentration tools developed by Faury et al. [2020] for the logistic bandit. In particular, for all  $t \geq 1$ ,  $\|g_t(\hat{\theta}_t) - g_t(\theta_*)\|_{\mathbf{H}_t(\theta_*)^{-1}} \leq \gamma_t(\delta)$  with probability at least  $1 - \delta$  for some function  $\gamma_t(\delta)$ . For completeness, we prove a slightly modified version of this inequality because we need to use the unregularized MLE to guarantee convergence to the true parameter. Additionally, we also use the self-concordant property of the logistic regression so we can apply the inequality to guarantee the algorithm is  $\delta$ -correct as explained in Section 2.5.1.

**Lemma 2.3.1.** *Let  $\delta \in (0, 1]$  and  $\lambda(t) > 0$  for  $t \geq 1$ . If exist  $t_0 \geq 1$  such that for  $t \geq t_0$ ,  $\lambda_{\min}(\mathbf{H}_t(\theta_*)) > \lambda(t)$ , with probability at least  $1 - \delta$ :*

$$\forall t \geq t_0, \quad \|g_t(\hat{\theta}_t) - g_t(\theta_*)\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \gamma_t(\delta), \quad (2.3)$$

$$\text{where } \gamma_t(\delta) := \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log \left( \frac{2^d}{\delta} \left( \frac{Lt}{\lambda(t)d} \right)^{\frac{d}{2}} \right)$$

The main difference compare to the original inequality is the assumption  $\lambda_{\min}(\mathbf{H}_t(\theta_*)) > \lambda(t)$ . In our case this will be guaranteed by the forced exploration component of sampling rule and the knowledge of  $\kappa_0$ . The forced exploration can be thought as an adaptive warm-up phase. Moreover, the amount of forced exploration needed to guarantee  $\delta$ -correctness of the algorithm will depend on  $\lambda(t)$ . As Faury et al. [2020] pointed out, we can use the bound in Abbasi-Yadkori et al. [2011] to derive another high-probability bound

$$\|g_t(\hat{\theta}_t) - g_t(\theta_*)\|_{\mathbf{H}_t^{-1}} = \mathcal{O}(\sqrt{\kappa d \log(t/\delta)}). \quad (2.4)$$

Although the bound in Eq. (2.3) is independent of  $\kappa$ , it has a disadvantage compare to Eq. (2.4). It has an extra factor of  $\sqrt{d \log(t/\delta)}$ . Unfortunately, this impacts the asymptotic behavior of the algorithm, the upper bound will be proportional to  $(\log(1/\delta))^2$  instead of  $\log(1/\delta)$  for small values of  $\delta$ . However, in practice for a fixed  $\delta$ , if we set  $\lambda(t) = \mathcal{O}(\log(t))$  we observe that  $\gamma_t(\delta) = \mathcal{O}(\sqrt{d \log(t)})$ .

**Projection step** We introduce the projection of the MLE estimator  $\hat{\theta}_t$  onto  $\Theta$  as

$$\theta_t^{(1)} = \arg \min_{\theta \in \Theta} \|g_t(\theta) - g_t(\hat{\theta}_t)\|_{\mathbf{H}_t^{-1}(\theta)}. \quad (2.5)$$

Thanks to the fact that  $\theta_t^{(1)} \in \Theta$ , we can establish an upper bound on bound  $\|\theta_t^{(1)} - \theta^*\|_{\mathbf{H}_t(\theta_t^{(1)})}$ , leveraging the self-concordance property of the logistic loss (See Lemma A.1.4). Note that  $\hat{\theta}_t = \theta_t^{(1)}$  when  $\|\hat{\theta}_t\| \leq S$ .

## 2.4 Sample Complexity Lower Bound

In this section, we provide an instance-specific lower bound for the expected sample complexity in *general* pure exploration problems within the logistic bandits setting. Jun et al. [2021] presented a similar lower bound only for BAI under logistic bandits. In contrast, our bound applies to a broader class of pure exploration problems. Moreover, we derive an approximation of the lower bound using the Taylor expansion of the KL divergence, making it tractable for some pure exploration problems such as BAI, TBP and top-m best arm identification (See Section 2.7 and A.3 in the appendix). This approximation provides an optimal proportion of arm draws, which we will track in the sampling rule of Log-TS similarly to previous track-and-stop algorithms [Garivier and Kaufmann, 2016, Jedra and Proutiere, 2020].

**Alternative** For any  $\theta \in \Theta$  we define the alternative to  $i^*(\theta)$ , denoted by  $\text{Alt}(\theta)$ , as the set of parameters where the answer  $i^*(\theta)$  is not correct. Formally,  $\text{Alt}(\theta) := \{\lambda \in \Theta : i^*(\lambda) \neq i^*(\theta)\}$ .

**Theorem 2.4.1.** *For any logistic bandit environment  $(\mathcal{X}, \theta)$  and  $\delta > 0$ , the sample complexity  $\tau_\delta$  of any  $\delta$ -correct strategy satisfies:*

$$\mathbb{E}_\theta[\tau_\delta] \geq \log(1/2.4\delta) \frac{1}{T^*(\theta)^{-1} + C(\theta)},$$

Where  $T^*(\theta)^{-1} := \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2$  and  $C(\theta)$  is an instance-specific constant measuring the precision of the quadratic approximation to the KL divergence.

*Proof.* Let  $\lambda \in \text{Alt}(\theta)$ , we can apply a slightly modified version of Theorem 33.5 from Lattimore and Szepesvári [2020] which has origin in the transportation theorem [Garivier and Kaufmann, 2016] to show.

$$\sum_{x \in \mathcal{X}} \mathbb{E}_\theta [N_x] \text{KL}_x(\theta, \lambda) \geq \log(1/2.4\delta) \tag{2.6}$$

Now, for each  $x \in \mathcal{X}$  we can approximate the KL divergence with the second order Taylor

expansion

$$\begin{aligned}\text{KL}_x(\theta, \lambda) &\approx \text{KL}_x(\theta, \theta) + (\theta - \lambda)^\top \nabla \text{KL}_x(\theta, \cdot)|_\theta \\ &\quad + \frac{1}{2}(\theta - \lambda)^\top \mathbf{H}_{\text{KL}}(\theta, \cdot)|_\theta (\theta - \lambda) + R_x(\lambda) \\ &\approx \frac{1}{2}(\theta - \lambda)^\top \dot{\mu}(x^\top \theta) x x^\top (\theta - \lambda) + R_x(\lambda)\end{aligned}$$

Where  $\text{KL}_x(\theta, \lambda)$  stands for the KL divergence between  $r|x, \theta$  and  $r|x, \lambda$ ,  $\nabla \text{KL}_x(\theta, \cdot)|_\theta$  and  $\mathbf{H}_{\text{KL}}(\theta, \cdot)|_\theta$  are the gradient and Hessian matrix of  $\text{KL}_x(\theta, \lambda)$  with respect to  $\lambda$  evaluated in  $\theta$ . We use the fact that the KL divergence between two distributions from the same exponential family can be expressed as  $\text{KL}(\eta_1, \eta_2) = (\eta_1 - \eta_2)\mu_1 + A(\eta_1) - A(\eta_2)$ . Where their probability distribution is given by  $p(x | \eta) = h(x) \exp\{\eta^\top T(x) - A(\eta)\}$ . In particular, for two GLM models with parameters  $\theta$  and  $\lambda$  and link function  $\mu(\theta)$ , we have  $\text{KL}_x(\theta, \lambda) = x^\top(\theta - \lambda)\mu(x^\top \theta) + A(x^\top \theta) - A(x^\top \lambda)$ ,  $\nabla \text{KL}_x(\theta, \lambda) = x^\top(\mu(x^\top \lambda) - \mu(x^\top \theta))$  and  $\mathbf{H}_{\text{KL}}(\theta, \lambda) = \dot{\mu}(x^\top \lambda) x x^\top$ . Then, after substituting the KL approximation in Eq. (2.6) we obtain

$$\mathbb{E}_\theta[\tau_\delta] \left( \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_{\tau_\delta}(\theta)}^2 + \max_{x \in \mathcal{X}} |R_x(\lambda)| \right) \geq \log(1/2.4\delta)$$

Where  $\mathbf{H}_{\tau_\delta}(\theta) = \left( \sum_{x \in \mathcal{X}} \frac{\mathbb{E}_\theta[N_x]}{\mathbb{E}_\theta[\tau_\delta]} \dot{\mu}(x^\top \theta) x x^\top \right)$ , then

$$\begin{aligned}\mathbb{E}_\theta[\tau_\delta] \sup_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \left( \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2 \right. \\ \left. + \max_{x \in \mathcal{X}} |R_x(\lambda)| \right) \geq \log(1/2.4\delta)\end{aligned}$$

$$\begin{aligned}\mathbb{E}_\theta[\tau_\delta] \left( \sup_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2 \right. \\ \left. + \inf_{\lambda \in \text{Alt}(\theta)} \max_{x \in \mathcal{X}} |R_x(\lambda)| \right) \geq \log(1/2.4\delta)\end{aligned}$$

If we denote  $C(\theta) = \inf_{\lambda \in \text{Alt}(\theta)} \max_{x \in \mathcal{X}} |R_x(\lambda)|$  we conclude the proof. □

As noted previously, Jun et al. [2021] showed an alternative lower bound for BAI under logistic bandits.

$$\mathbb{E}_\theta[\tau_\delta] \geq \log(1/2.4\delta) \frac{1}{\max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \|\theta - \lambda\|_{\mathbf{K}_w(\theta, \lambda)}^2},$$

Where  $\mathbf{K}_w(\theta_1, \theta_2) = \sum_{x \in \mathcal{X}} w_x \beta(\theta_1, \theta_2) x x^\top$  and  $\beta(a, b) = \int_0^1 (1-t) \dot{\mu}(a + t(b-a)) dt$ . The main difference between our bounds is that ours uses a quadratic approximation for the KL divergence, allowing us to establish a direct relationship with the Fisher information

matrix  $\mathbf{H}$ . If the approximation of the KL divergence is accurate for  $(\mathcal{X}, \theta)$ , we expect that  $T^*(\theta)^{-1} + C(\theta) \approx \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \|\theta - \lambda\|_{\mathbf{K}_w(\theta, \lambda)}^2$  and then

$$\max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w}^2 \leq \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \|\theta - \lambda\|_{\mathbf{K}_w}^2$$

where  $\mathbf{H}_w = \mathbf{H}_w(\theta)$  and  $\mathbf{K}_w = \mathbf{K}_w(\theta, \lambda)$ . When the constant  $C(\theta)$  is negligible, our lower bound closely resembles the lower bound for the linear case [Marta Soare, 2014, Degenne et al., 2020], as the constant  $C(\theta) = 0$  in this case. Our goal is to design an algorithm with the following asymptotic sample complexity:

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\theta[\tau]}{(\log(1/\delta))^2} \leq T^*(\theta),$$

**Remark.** We may ask when the assumption  $C(\theta) \approx 0$  is reasonable. From the proof of Theorem 2.4.1 we have  $C(\theta) = \inf_{\lambda \in \text{Alt}(\theta)} \max_{x \in \mathcal{X}} |R_x(\lambda)|$ . From this definition, we can see that if there exists a parameter  $\lambda \in \text{Alt}(\theta)$  that allows a good approximation of the KL divergence, then  $C(\theta)$  will be close to 0. This occurs when there is a  $\lambda \in \text{Alt}(\theta)$  sufficiently close to  $\theta$ , which implies that  $T^*(\theta)^{-1}$  will be small as well.

## 2.5 Logistic Track-and-Stop Algorithm

In this section, we present the first track-and-stop type algorithm for general pure exploration under logistic bandits. First, we propose a modified version of the classical Chernoff stopping rule [Garivier and Kaufmann, 2016], which utilizes the approximation of the generalized likelihood ratio from Theorem 2.4.1. We prove that this stopping rule provides a  $\delta$ -correct algorithm under any sampling rule.

Next, we define the standard tracking rule with forced exploration [Jedra and Proutiere, 2020]. This component of the algorithm tracks the estimated optimal proportion of arm pulls based on the projection of the MLE onto  $\Theta$ . Thanks to forced exploration, we can ensure that the MLE, and consequently its projection, converge almost surely to  $\theta^*$ . As a result, the estimated optimal proportion of arm draws converges to a true optimal proportion due to the continuity of the function defined by the optimization problem.

## 2.5.1 Stopping rule

We will use the approximation of the generalized likelihood ratio given in Theorem 2.4.1. We define the stopping rule as

$$\tau_\delta = \inf \{t \geq 1 : Z(t) > \beta(\delta, t), t \in B\} \quad (2.7)$$

Where  $B = \{t \geq 1 : \lambda_{\min}(\mathbf{A}_s) > \kappa_0 \lambda(s), \forall s \geq t\}$ ,  $Z(t) = \inf_{\lambda \in \text{Alt}(\theta_t^{(1)})} \frac{1}{2} \|\theta_t^{(1)} - \lambda\|_{\mathbf{H}_t(\theta_t^{(1)})}^2$  and  $\beta(\delta, t) = 2((1 + 2S)\gamma_t(\delta))^2$ . Then, by using this stopping rule, we obtain a  $\delta$ -correct algorithm.

**Lemma 2.5.1.** *Under any sampling rule, we have*

$$\mathbb{P}_\theta \left( \tau_\delta < \infty \wedge i^*(\theta_{\tau_\delta}^{(1)}) \neq i^*(\theta) \right) \leq \delta$$

Although Lemma 2.3.1 requires  $\lambda(\mathbf{H}_t(\theta_*)) > \lambda(t)$ , it is not possible to guarantee this directly since we do not have access to the true matrix  $\mathbf{H}_t(\theta_*)$ . Instead, it is enough to ask  $t \in B$ , thanks to the fact that  $\lambda_{\min}(\mathbf{H}_t(\theta_*)) > \frac{1}{\kappa_0} \lambda_{\min}(\mathbf{A}_t)$ . We will see that the condition  $t \in B$  is easily met due to forced exploration. In fact, from Lemma 2.5.2 it is sufficient that  $c_{\mathcal{X}_0} \sqrt{t} > \kappa_0 \lambda(t)$  for some constant  $c_{\mathcal{X}_0}$ . Although this condition may seem easy to satisfy, it will depend on the relationship between  $c_{\mathcal{X}_0}$  and  $\kappa_0$ . If the problem is highly complex or the set  $\mathcal{X}_0$  is flat in some direction, this condition can be restrictive.

In practice, instead of checking  $t \in B$ , we can verify if  $\lambda_{\min}(\mathbf{H}_t(\hat{\theta})) > \lambda(t)$ , which can result in an earlier stopping rule. Another parameter we can adjust is the function  $\lambda(t)$ . It is sufficient that exists  $t^* \geq 1$  such that  $\lambda(t)$  satisfies  $c_{\mathcal{X}_0} \sqrt{t} > \kappa_0 \lambda(t)$  for all  $t \geq t^*$ . This introduces a trade-off between achieving  $Z(t) > \beta(\delta, t)$  and satisfying  $c_{\mathcal{X}_0} \sqrt{t} > \kappa_0 \lambda(t)$ , reflecting the amount of forced exploration required.

One technical challenge in using the concentration inequality from Section 2.3.3 is that  $\mathbf{H}_t(\theta^*)$  requires knowledge of the true parameter. We address this issue by leveraging the generalized self-concordance property of the logistic loss, as noted by Faury et al. [2020]. This property allows us to control the distance between  $\theta_t^{(1)}$  and  $\theta^*$

$$\|\theta_t^{(1)} - \theta^*\|_{\mathbf{H}_t^{(1)}} \leq 2(1 + 2S) \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{(\mathbf{H}_t^*)^{-1}} \quad (2.8)$$

where  $\mathbf{H}_t^{(1)} = \mathbf{H}_t(\theta_t^{(1)})$  and  $\mathbf{H}_t^* = \mathbf{H}_t(\theta^*)$ . In Jun et al. [2021], the authors also used the self-concordance property to control  $\mathbf{H}_t(\hat{\theta}_t)$  using a warm-up phase to control the linear prediction errors  $\max_{s \in [t]} |x_s^\top (\hat{\theta}_t - \theta^*)| \leq 1$  with the advantage that they do not need the assumption  $\|\theta^*\| < S$ .

## 2.5.2 Sampling rule

**Forced exploration** Forced exploration is a crucial component of the Log-TS algorithm, as it allows us to apply Lemma 2.3.1 and guarantees the convergence of the MLE estimator. A sampling rule from the family defined in Lemma 2.5.2 is forced to explore an arm in  $\mathcal{X}_0$  (in a round robin manner) if  $\lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right)$  is too small [Jedra and Proutiere, 2020].

**Lemma 2.5.2.** (Lemma 5 [Jedra and Proutiere, 2020]) Let  $\mathcal{X}_0 = \{x_0(1), \dots, x_0(d)\} : \lambda_{\min}(\sum_{x \in \mathcal{X}_0} x x^\top) > 0$ . Let  $\{b_t\}_{t \geq 0}$  be an arbitrary sequence of arms. Furthermore, define for all  $t \geq 1$ ,  $f(t) = c_{\mathcal{X}_0} \sqrt{t}$  where  $c_{\mathcal{X}_0} = \frac{1}{\sqrt{d}} \lambda_{\min}(\sum_{x \in \mathcal{X}_0} x x^\top)$ . Consider the rule, defined recursively as:  $i_0 = 1$ , and for  $t \geq 0$ ,  $i_{t+1} = (i_t \bmod d) + \mathbb{1}_{\{\lambda_{\min}(\sum_{s=1}^t x_s x_s^\top) < f(t)\}}$  and

$$x_{t+1} = \begin{cases} x_0(i_t) & \text{if } \lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right) < f(t) \\ b_t & \text{otherwise.} \end{cases} \quad (2.9)$$

Then for all  $t \geq \frac{5d}{4} + \frac{1}{4d} + \frac{3}{2}$ , we have

$$\lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right) \geq f(t - d - 1)$$

To ensure the MLE converges almost surely to the true parameter  $\theta^*$ , the ratio between the minimum eigenvalue and the logarithm of the maximum eigenvalue of the matrix  $\sum_{s=1}^t x_s x_s^\top$  must tend to infinity [Chen et al., 1999]. This convergence can be guaranteed by the forced exploration component.

**Lemma 2.5.3.** Under the forced exploration sampling defined in Eq. (2.9), the MLE estimator converge a.s. to the true parameter

$$\lim_{t \rightarrow \infty} \hat{\theta}_t \stackrel{a.s.}{=} \theta^*$$

Another consequence of forced exploration is the following tail inequality. We can upper bound the probability that the distance between the MLE and the true parameter exceeds  $\varepsilon$ . This inequality will play an important role to prove the upper bound for the expected value of the sample complexity.

**Lemma 2.5.4.** Let  $\varepsilon > 0$ , assume that  $\lambda_{\min}(\mathbf{A}_t) \geq ct^{1/2}$  a.s. for all  $t \geq t_0$  for some  $t_0 \geq 1$  and for  $c > \kappa_0 \lambda_0$ . Then

$$\forall t \geq t_0 \quad \mathbb{P} \left( \|\theta_t^{(1)} - \theta^*\| \geq \varepsilon \right) \leq c_2 t^{\beta_2} \exp \left( -c_1 t^{\beta_1} \right)$$

Where  $c_1, c_2, \beta_1, \beta_2$  are positive constants independent of  $\varepsilon$  and  $t$ .

**Tracking** As noted in Jedra and Proutiere [2020], Wang et al. [2021b], we can define the function  $\psi(\theta, w)$  to recover the optimal proportions from the lower bound in Theorem 2.4.1

$$\psi(\theta, w) = \inf_{\lambda \in \text{Alt}(\theta)} \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2 \quad (2.10)$$

This function has a tractable form for BAI, TBP and top-m best arm identification (See Section A.3 in the appendix). The objective is to optimize  $\psi(\theta, w)$  over the simplex  $\Sigma$  so we can estimate the optimal proportions needed in the TS algorithm.

$$\psi^*(\theta) = \max_{w \in \Sigma} \psi(\theta, w) \quad (2.11)$$

$$C^*(\theta) = \arg \max_{w \in \Sigma} \psi(\theta, w) \quad (2.12)$$

Note that  $\psi(\theta^*, w^*) = T^*(\theta^*)^{-1}$ , where  $w^* \in C^*(\theta^*)$ . As pointed out in previous works, the solution to Eq. (2.12) may involve multiple optimal proportions. However, similar to Jedra and Proutiere [2020], we only need to prove that  $\psi^*(\theta)$  is continuous in  $\theta$  and that  $C^*(\theta)$  is convex to guarantee the algorithm converges to an optimal proportion inside  $C^*(\theta^*)$ . We will use the Frank-Wolfe algorithm, as previous works, to solve the optimization problem in Eq. (2.12) [Jedra and Proutiere, 2020, Degenne et al., 2020]. A caveat of using Frank-Wolfe algorithm is that convergence is not guaranteed when the function  $\psi^*(\theta)$  is non-smooth. However, empirical evidence suggests that it can still converge in practice. Moreover, recent work on developing projection-free algorithms for non-smooth functions could be leveraged in the future [Asgari and Neely, 2022].

We use the standard tracking procedure to track the estimated optimal proportions

$$b_t = \arg \min_{x \in \text{supp}(\sum_{s=1}^t w(s))} \left( N_x(t) - \sum_{s=1}^t w_x(s) \right), \quad (2.13)$$

Where  $w(t) \in C^*(\theta_t^{(1)})$ , i.e.

$$w(t) = \arg \max_{w \in \Sigma} \psi(\theta_t^{(1)}, w) \quad (2.14)$$

---

**Algorithm 1** Log Track-and-Stop

---

**Input:** Arms  $\mathcal{X}$ , confidence level  $\delta$ ;  
**Initialize:**  $t = 0, i = 0, \mathbf{A}_0 = 0, Z(0) = 0, N(0) = (N_x(0))_{x \in \mathcal{X}} = 0$ ;  
**while**  $t \notin B$  or  $Z(t) < \beta(\delta, t)$  **do**  
  **if**  $\lambda_{\min}(\mathbf{A}_t) < f(t)$  **then**  
    select  $x$  according to Eq. (2.9)  
  **else**  
    select  $x$  according to Eq. (2.13)  
  **end if**  
   $t \leftarrow t + 1$ ,  
  sample arm  $x$   
  update  $N(t), \hat{\theta}_t, \theta_t^{(1)}, Z(t), \mathbf{A}_t, \mathbf{H}_t$   
   $w(t) = \arg \max_{w \in \Sigma} \psi(\theta_t^{(1)}, w)$   
**end while**  
**Return**  $i^*(\theta_\tau^{(1)})$

---

We can prove that under this sampling rule, the observed proportions of the sampled arms converge to a true optimal proportions.

**Proposition 1.** *Under the sampling rules defined by Eq. (2.9) and Eq. (2.13), the proportions of arm draws approach  $C^*(\theta^*) : \lim_{t \rightarrow \infty} d_\infty \left( (N_x(t)/t)_{x \in \mathcal{X}}, C^*(\theta^*) \right) = 0$ , a.s..*

Although updating  $w(t)$  according to Eq. (2.14) can be computationally expensive, Log-TS can easily adapted to the lazy approach of Jedra and Proutiere [2020], so we do not need to update  $w(t)$  at every step.

## 2.6 Sample Complexity of Log-TS

In this section we state two upper bounds for the sample complexity under Log-TS when we set  $\lambda(t) = c \log(t)$  for some constant  $c > 0$ . The pseudocode is provided in Algorithm 1.

**Theorem 2.6.1.** *Logistic Track-and-stop (Log-TS) satisfies the following sample complexity upper bound*

$$\mathbb{P}_\theta[\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{(\log(\frac{1}{\delta}))^2} \lesssim T^*(\theta)] = 1$$

**Theorem 2.6.2.** *Logistic Track-and-stop (Log-TS) satisfies the following sample complexity upper bound*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\theta[\tau_\delta]}{(\log(\frac{1}{\delta}))^2} \lesssim T^*(\theta)$$

As we mention in Section 2.3.3 these bounds have an extra factor  $\log(\frac{1}{\delta})$  that comes from the threshold  $\gamma_t(\delta)$ .

## 2.7 Experiments

### 2.7.1 BAI

In this section we evaluate the performance of Log-TS for BAI. To the best of our knowledge, there are only two algorithms that address BAI for the logistic bandit [Kazerouni and Wein, 2021, Jun et al., 2021]. We only compare our algorithm against RAGE-GLM-R and random sampling equipped with the same stopping rule as Log-TS, since RAGE-GLM-R is the state of the art algorithm for this task. Lets define  $x^*(\theta) := i^*(\theta) = \arg \max_{x \in \mathcal{X}} \{x^\top \theta\}$ . Then, to implement the algorithm, we use the following lemma:

**Lemma 2.7.1.** *For all  $\theta \in \mathbb{R}^d$ ,*

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \min_{x \neq x^*(\theta)} \frac{(\theta^\top x^*(\theta) - \theta^\top x)^2}{2 \|x^*(\theta) - x\|_{\mathbf{H}_w^{-1}}^2}$$

We evaluated the algorithms under two settings:

**Benchmark for BAI in linear bandits.** The benchmark examples in the linear bandit BAI literature introduced by Marta Soare [2014] is the following. Consider  $\mathcal{X} = \{e_1, \dots, e_d, x'\} \subseteq \mathbb{R}^d$  where  $e_i$  is the  $i$ -standard basis vector,  $x' = \cos(\alpha)e_1 + \sin(\alpha)e_2$  with  $\alpha$  small, and  $\theta$  proportional to  $e_1$  so that  $e_1 = \arg \max_{x \in \mathcal{X}} x^\top \theta$ . This setting is designed to be a hard instance for an algorithm because the rewards  $\mu(\theta^\top e_1)$  and  $\mu(\theta^\top x')$  are close, so distinguishing the best arm becomes a difficult task. In Figure 2.1, we observe comparable results between Log-TS and Rage-GLM-R, suggesting that our algorithm is competitive with the state-of-the-art method.

**Uniform Distribution on a Sphere.** In this example,  $\mathcal{X}$  is sampled from a unit sphere of dimension  $d = 2$  centered at the origin with  $|\mathcal{X}| = K$ . We set  $\theta$  randomly so that  $\|\theta\| = 1$ . To control the variation between experiments, we restricted the complexity  $T^*(\theta)^{-1}$  of the experiments across the different number arms, allowing higher complexity for higher value of  $K$ . In Table 2.1, Log-TS shows a distinctive advantage, achieving almost a 5x reduction in sample complexity. This result highlights the efficiency of Log-TS, particularly in scenarios with many arms. The superiority of Log-TS can be attributed to the independence of the

stopping rule with respect to the number of arms, allowing it to scale better in large action spaces.

In the implementation of the stopping rule, we adopt a heuristic by using  $\lambda_{\min}(\mathbf{H}_t(\hat{\theta})) > \lambda(t)$  instead of  $t \in B$  and select  $\lambda(t) = d \log(t)$ . This choice improves the algorithm’s sample complexity performance avoiding unnecessary exploration. In all instances, all algorithms found the best arm correctly using  $\delta = 0.10$ .

Table 2.1: Average sample complexities, expressed in thousands, of uniform distribution for BAI. The results are shown for various numbers of arms  $K$ . Standard deviations and means were computed across 10 randomized trials.

$K$	Log-TS	Rage-GLM-R	Random
100	<b>4.05 (0.35)</b>	33.22 (19.14)	11.44 (1.59)
200	<b>6.94 (0.39)</b>	49.34 (13.26)	25.10 (5.86)
300	<b>10.66 (1.44)</b>	51.94 (15.96)	44.32 (6.44)
400	<b>14.36 (2.31)</b>	53.39 (18.40)	55.64 (11.86)
500	<b>16.59 (2.32)</b>	68.00 (19.13)	74.10 (14.41)




Figure 2.1: Logarithm of sample complexity of the benchmark setup for BAI against dimension of the action space  $\mathcal{X}$ .

## 2.7.2 Thresholding Bandit Problem

In comparison to BAI, there are not many algorithms that can solve TBP. For the linear case Degenne et al. [2020] developed a general algorithm for pure exploration problems called LinGame. Another example is the work in Mason et al. [2022]. They developed a nearly optimal algorithm for level set estimation, aka TBP. However, to the best of our knowledge, there is no algorithm for TBP under GLM bandit. Therefore, we compare Log-TS with random sampling, both equipped with the same stopping rule. For TBP, given  $\rho \in (0, 1)$ , we have  $i^*(\theta) = \{x \in \mathcal{X} : \mu(x^\top \theta) > \rho\}$ . Similarly to BAI, we have an explicit expression that allows us to implement the algorithm:

**Lemma 2.7.2.** *For all  $\theta \in \mathbb{R}^d$ ,*

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \min_{x \in \mathcal{X}} \frac{(\theta^\top x - \mu^{-1}(\rho))^2}{2 \|x\|_{\mathbf{H}_w^{-1}}^2}$$

**Benchmark for TBP in linear bandits.** We proposed a slightly modified version of the BAI benchmark setting such that it is hard to distinguish if one of the arms is above or below the threshold. Consider a similar setting as before,  $\mathcal{X} = \{e_1, \dots, e_d, x', x''\} \subseteq \mathbb{R}^d$  where  $e_i$  is the  $i$ -standard basis vector,  $x' = p(\cos(\alpha)e_1 + \sin(\alpha)e_2)$ ,  $x'' = (1-p)(\cos(-\alpha)e_1 + \sin(-\alpha)e_2)$  with  $\alpha$  small and  $p$  close to  $1/2$ . Let  $\rho = (\mu(\theta^\top x') + \mu(\theta^\top x''))/2$ ,  $\theta$  proportional to  $e_1$ , such that  $x'^\top \theta, x''^\top \theta \approx \rho$ . In Figure 2.2, we observe a reduction between 5 and 10 times in the sample complexity of Log-TS compare to random sampling.

**Uniform Distribution on a Sphere.** In the same way as in BAI,  $\mathcal{X}$  is sampled from a unit sphere of dimension  $d = 2$  centered at the origin. We set  $\theta$  randomly such that  $\|\theta\| = 1$ , and  $\rho = 0.5$ . We also control the complexity  $T^*(\theta)^{-1}$  of the experiments across the different number arms to avoid high variance. In Table 2.2, it is shown that Log-TS outperforms random sampling by roughly 5 times in sample complexity. This confirms the efficiency of Log-TS in many arms scenarios for different pure exploration problems.

Algorithms were able to identify the arms above and below the threshold  $\rho$  correctly in all instances with  $\delta = 0.10$ .

## 2.8 Discussion

To the best of our knowledge, we propose the first track-and-stop algorithm for general pure exploration problems under logistic bandits, proving that our algorithm is near asymptotically optimal. This means that the sample complexity is upper-bounded in the limit up to

Table 2.2: Average sample complexities, expressed in thousands, of uniform distribution for TBP. The results are shown for various numbers of arms  $K$ . Standard deviations and means were computed across 10 randomized trials.

$K$	Log-TS	Random
20	<b>7.49 (2.36)</b>	15.22 (4.06)
30	<b>9.50 (0.65)</b>	37.64 (10.84)
40	<b>12.86 (0.44)</b>	50.84 (12.21)
50	<b>16.15 (3.64)</b>	94.47 (23.06)




Figure 2.2: Logarithm of sample complexity of the benchmark setup for TBP against dimension of the action space  $\mathcal{X}$ .

logarithmic factor, and the upper bound matches the approximated lower bound. We evaluate our algorithm on BAI and TBP, demonstrating its advantages over existing algorithms and random sampling. The incorporation of forced exploration enables us to apply the tail inequality without a warm-up phase, enhancing the algorithm’s applicability in real-world scenarios. On the other hand, the computation of the MLE and its projection can make the algorithm slow under certain conditions since it can not be computed recursively, so TS-Log could benefit from online approaches like the one presented in Faury et al. [2022]. Additionally, our asymptotic results suffer from an extra logarithmic factor  $\log(1/\delta)$ . This term arises directly from the tail inequality in Lemma 2.3.1; thus, to eliminate it, a tighter tail inequality is necessary, one that depends on  $\gamma_t(\delta) = \mathcal{O}(\sqrt{\log(1/\delta)})$  instead of  $\gamma_t(\delta) = \mathcal{O}(\log(1/\delta))$ .

## CHAPTER 3

# Enhancing Confidence in Microplastic Spectral Identification via Conformal Prediction

Microplastics are an emerging pollutant of concern, with environmental observations recorded across the world. Identifying the type of microplastic is challenging due to spectral similarities among the most common polymers, necessitating methods that can confidently distinguish plastic identities. In practice, a researcher chooses the reference vibrational spectrum that is most like the unknown spectrum, where the likeness between the two spectra is expressed numerically as the hit quality index (HQI). Despite the widespread use of HQI thresholds in the literature, acceptance of a spectral label often lacks any associated confidence. To address this gap, we apply a machine-learning framework called conformal prediction to output a set of possible labels that contain the true identity of the unknown spectrum with a user-defined probability (e.g., 90%). Microplastic reference libraries of environmentally aged and pristine polymeric materials, as well as unknown environmental plastic spectra, were employed to illustrate the benefits of this approach when used with two similarity metrics to compute HQI. We present an adaptable workflow using our open-access code to ensure spectral matching confidence for the microplastic community, reducing manual inspection of spectral matches and enhancing the robustness of quantification in the field. <sup>1</sup>

### 3.1 Introduction

Identifying microplastics (MPs) is essential to address health and environmental concerns caused by plastic production and pollution. Microplastics, defined as small pieces of plastic between 1 and 5000  $\mu\text{m}$  [Frias and Nash, 2019, Hale et al., 2020], have been found in environmental matrices such as seawater [Alfaro-Núñez et al., 2022, Kanhai et al., 2020],

---

<sup>1</sup>This chapter is adapted from previously published work [Clough et al., 2024], with Madeline E. Clough, Eduardo Ochoa Rivera, Rebecca L. Parham, Andrew P. Ault, Paul M. Zimmerman, Anne J. McNeil and Ambuj Tewari. This chapter describes equal contributions of authors Madeline E. Clough and Eduardo Ochoa Rivera

freshwater [Talbot and Chang, 2022], soil [Cusworth et al., 2024], and air [Brahney et al., 2020, Allen et al., 2019]. The extent of microplastic pollution and major MP sources, which are hypothesized based on the polymer identity of the MP [Wang et al., 2021a], must be characterized to develop remediation solutions and regulations. Researchers have employed techniques such as Fourier transform-infrared (FT-IR) and Raman spectroscopy to determine both the concentration and identities of MPs in the environment [Cabernard et al., 2018, Tirkey and Upadhyay, 2021, Araujo et al., 2018, K appler et al., 2016].

Automated spectral matching has prevailed as a practical solution to time-intensive manual spectral matching of unknown environmental species to standard reference polymers [Weisser et al., 2022, Cowger et al., 2020]. Commercial software equipped with library-searching modules, such as Bruker’s OPUS, Thermo Fisher’s OMNIC, and Wiley’s Know-ItAll, have been employed to produce a similarity score between an unknown spectrum and a reference spectrum. This score of spectral likeness, which is computed via similarity metrics, is referred to as hit quality index (HQI) [Weisser et al., 2022, Cowger et al., 2020]. HQI scores are often reported in the range of 0 to 1, where a value of zero indicates the spectra are wholly unlike and a score of one indicates complete similarity [Weisser et al., 2022]. Researchers commonly use a threshold HQI value that must be met or exceeded to identify unknown spectra with a polymer label [Weisser et al., 2022, Cowger et al., 2020]. However, there is disagreement regarding the meaning expressed by the HQI value, occasionally being referenced as a measure of confidence in a polymer label [Abbasi et al., 2023, Clark et al., 2023, Miserli et al., 2023]. This assertion of confidence is misleading, as the HQI does not provide any measure of statistical certainty. Rather, the score simply reflects a measurement of the similarity between two spectra.

Because many polymers have similar structures (and therefore spectra) [Jung et al., 2018, Nava et al., 2021], and because vibrational spectra can change due to environmental aging [Renner et al., 2017, Binda et al., 2024], labeling unknown spectra based on HQI scores alone is insufficient. A better approach for robust MP quantification and source attribution would include a measure of the statistical uncertainty in spectral matching. Acceptance of an identity solely due to an HQI score above an arbitrary threshold may lead to misidentification, as indicated by the findings of L opez-Rosales et al. [2024], wherein a threshold of 0.85 (on a scale of 0 to 1) led to overestimating MP counts [L opez-Rosales et al., 2024]. Indeed, voices in the MP community have called for stricter assurances of polymer identity through manual assessment of likeness between unknown spectra and standard references [Jung et al., 2018, L opez-Rosales et al., 2024].

To answer this call for increased confidence, some researchers have framed uncertainty in spectral labeling via distributions of errors between automated spectral match scores

and expert manual assessments [Primpke et al., 2017, 2018], as well as through automated correlation scores of samples with known composition [Choy et al., 2019, Morgado et al., 2021]. These approaches provide a guaranteed confidence for the whole sample, but for individual spectra, they are limited to accepting or rejecting a single spectral label. Moreover, these methods are asymptotically valid, meaning that the guarantees hold only when the number of observations (here, spectra) nears infinity. In contrast to previous methods, a method known as conformal prediction (CP) presents the advantages of indicating the uncertainty of individual matches and providing a theoretical confidence (or probability) guarantee that the identity of the query spectrum is contained in a set of predicted labels.

Conformal prediction is a framework that provides a principled approach to understanding uncertainty via model-generated predictions [Vovk et al., 1999, 2005]. This method generates a prediction set of reference labels that guarantees a specified level of finite sample confidence [Papadopoulos et al., 2002], meaning the probability that the correct spectral label is included in the prediction set of an unknown spectrum is equivalent to the confidence guarantee (given that the unknown’s identity is reflected in the reference library). CP has been increasingly popular in machine learning, where it is able to provide insight into the uncertainty of a given model without the need for retraining, which is often computationally intensive [Angelopoulos and Bates, 2023, Angelopoulos et al., 2020, Ren et al., 2023]. Our introduction of CP to the MP community provides a tool for enhanced credibility in spectral matching via a statistical assurance that the identity of an unknown spectrum will be returned to the analyst.

Herein, we first evaluate the spectral correlation software, methods, and thresholds reported in the recent microplastic literature to establish a baseline of commonly used automated matching practices in the MP community. We then illustrate that the HQI scores generated by popular commercial software lack the clarity needed to confidently count and label environmental MPs. Further, we demonstrate the advantage of using CP in tandem with two different similarity metrics to provide a confidence guarantee and a measurement of uncertainty at the individual spectrum level. Lastly, we apply this method to authentic, environmental plastics to illustrate CP’s utility with real-world samples. Through this work, we aim to provide a workflow to incorporate statistical assurances in microplastic spectral matching and establish guidance for HQI thresholding and label acceptance for more robust environmental MP identification.

## 3.2 Methods

### 3.2.1 Commercial Software Library Searching

The open-access spectral databases “FT-IR library of plastic particles” (FLOPP, used as a pristine reference library) and “FT-IR library of plastic particles sourced from the environment” (FLOPP-e, used as environmental query spectra) were chosen to evaluate commercial software due to their accessibility and widespread use in the microplastic community [De Frond et al., 2021b]. All individual spectra belonging to the polyethylene (PE) and polypropylene (PP) labels available in the FLOPP-e library were downloaded as .csv files and uploaded to Wiley’s KnowItAll SearchIt function. These polymers were chosen because they had the greatest number of example spectra in the environmentally aged FLOPP-e library (49 and 66 for PE and PP, respectively). For each spectrum, a transmission mode was indicated, and a single component search was performed using the default settings of a correlation method and optimized corrections. A spectral inclusion window was drawn from 4000 to 650  $\text{cm}^{-1}$  to include only regions of the spectra where the data had been collected. Each spectrum was searched for correlation with pristine polymers in the FLOPP library, and the hit list was documented. The highest HQI score and label, as well as the number of matches and number of polymer labels yielding  $\text{HQI} \geq 0.7$  (a common literature threshold of HQI) [Abbasi et al., 2023], were recorded.

### 3.2.2 Conformal Prediction

Conformal prediction is used to generate prediction sets (of polymer labels) for new observations given a trained prediction model (HQI-based spectral matching). CP uses a calibration data set and similarity metric to ensure that the spectral label of a query spectrum will be returned in a prediction set with a user-specified probability (e.g., 90%), affording the confidence guarantee. This confidence guarantee holds when the distribution of new observations (here, HQI scores) is identical to those in the calibration set. Unlike many other uncertainty quantification methods in the literature, CP does not require specific distributional assumptions (e.g., Gaussianity) nor does it require a large data set.

Using all spectra from the 11 labels (acrylonitrile butadiene styrene, cotton [C], polyamide, polycarbonate, PE, polyethylene terephthalate [PET], polyethylene vinyl acetate [PEVA], PP, polystyrene [PS], polyurethane, and polyvinyl chloride [PVC]) that appear in both the FLOPP and FLOPP-e data sets (142 and 189 spectra, respectively), we split the FLOPP-e spectra into calibration (90%) and test (10%) sets for matching to FLOPP reference spectra (Scheme 3.1). We then calculated the HQI scores of the true positive matches (FLOPP-e

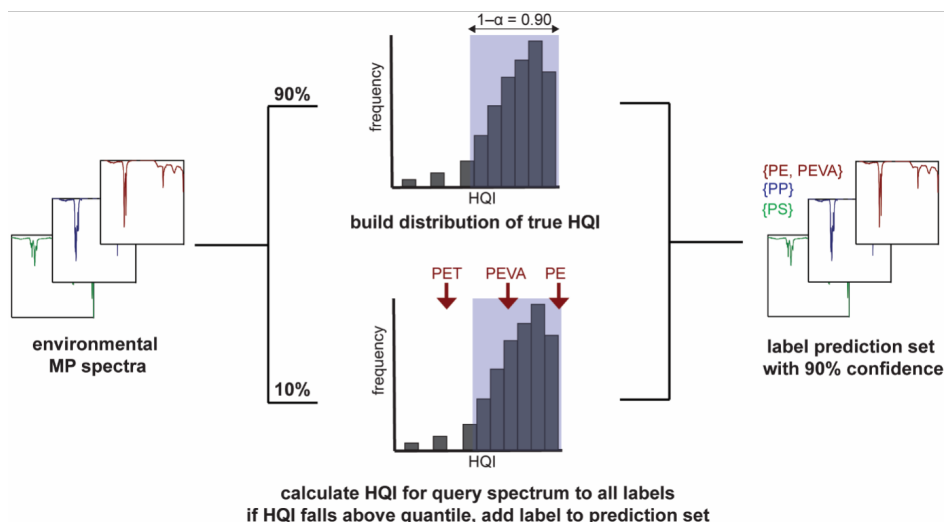


Figure 3.1: CP Workflow, Where Exemplative Labels Corresponding to the Red Environmental MP Spectrum are Illustrated on the Bottom Distribution of HQI Scores from True Positive Matches.

spectrum matching to an FLOPP spectrum of the same polymer label) of the calibration set and tabulated them in frequency distributions relating to the two different similarity metrics. Next, the remaining 10% of FLOPP-e spectra was used as a test set (meaning that they were treated as spectra of unknown identity) and the HQI scores for each spectrum's similarity to all reference polymers in the FLOPP library were calculated. These HQI values were compared to the threshold defined by the corresponding quantile of the calibration set distribution that resembles the guaranteed confidence (here, 90% [ $\alpha = 0.1$ ]) for both similarity metrics. If the HQI of a polymer label was above the threshold, meaning that only 10% of the calibration data's true positive scores were calculated to be less than the HQI, then the label was added to the prediction set with an associated 90% confidence. In Figure 3.1, this workflow is illustrated using example labels and HQI scores for the red FLOPP-e spectrum, which are then compared to the calibration distribution and added to a prediction set of labels (on the far right of the graphic) if the scores are greater than the established quantile.

The FLOPP-e spectra were randomly split into calibration and test data sets 500 times in a cross-validation fashion. We then calculated the empirical confidence (proportion of instances where the correct label was observed in the prediction set) and prediction set size (number of labels included in the prediction sets) for all of the test data.

When using CP, the analyst chooses the confidence guarantee that will determine the HQI threshold. While the confidence is always guaranteed given that the reference library contains spectra of each unknown's identity, the magnitude of the confidence guarantee influences the

number of labels returned in the prediction set [Vovk et al., 2005]. With a high confidence guarantee, the HQI quantile (i.e., threshold) is lower, leading to more polymer labels in the prediction set. On the other hand, a low guaranteed confidence may lead to very few, if any, labels in the prediction set. While this trade-off between confidence and prediction set size should be noted as inherent to the conformal prediction workflow, we provide insight into tailoring the confidence guarantee and leveraging the performance of the chosen similarity metric to produce the most insightful predictions in future sections.

### 3.2.3 Class-Conditional Conformal Prediction

One limitation of the traditional CP framework is that the confidence guarantee is marginal, meaning that some labels may be under or overguaranteed, as only the average theoretical confidence across all labels is assured [Angelopoulos and Bates, 2023]. To address this issue, we also employ the class-conditional conformal prediction (CC-CP) approach [Vovk, 2012]. The basis of the CP methodology framework (i.e., create a calibration set, determine the HQI threshold, compare test set HQI scores to calibration distribution, and add labels with HQI scores represented above the quantile to the prediction set) remains the same. The only difference is that instead of constructing one distribution from the calibration set, we created separate distributions for each label to ensure class-balanced confidence and that the HQI scores per polymer type were compared to label-specific distributions.

Because a unique distribution is built for each label, we used the five labels with 10 or more example spectra in the FLOPP data set (i.e., C, PE, PET, PP, and PS) as possible spectral labels, limiting the pool of FLOPP and FLOPP-e spectra to 76 and 161 examples, respectively. The FLOPP-e calibration and test data were then randomly split 500 times in a cross-validated fashion. Although the traditional CP results using these five labels are presented herein for direct comparison to the class-conditional results, it should be noted that additional labels can be used with a CP framework and that this addition of labels will result in larger average prediction sets.

### 3.2.4 Similarity Metrics Used to Compute Hit Quality Index

In CP, the choice of the similarity metric used to compute the HQI is crucial because while the confidence guarantee holds for any similarity metric, the quality of the similarity metric directly governs the size of the prediction set. The size of, or number of labels in, the prediction set can be used to assess uncertainty in the predictions [Angelopoulos and Bates, 2023], where prediction sets with few labels indicate a low uncertainty and prediction sets with many labels indicate a high uncertainty. Efficient similarity metrics give prediction

sets with fewer labels, therefore providing a more accurate depiction of the inherent spectral uncertainty in the prediction sets [Shafer and Vovk, 2008]. Thus, exploring different similarity metrics, which rely on different measures of spectral likeness [Weisser et al., 2022], is necessary to compare the uncertainty of resultant prediction sets. This aspect of CP led us to investigate nearest neighbor (NN) [Shafer and Vovk, 2008], and the more commonly used Pearson correlation coefficient (PCC), as similarity metrics to calculate HQI.

### 3.2.5 Application of Conformal Prediction to Real World Environmental Plastics

Environmental plastics of unknown identity were sourced from a residential parking lot in Brighton, Michigan and rinsed with tap water to remove debris. After the plastic was dried at room temperature, attenuated total reflection FT-IR measurements of each plastic were collected in percent transmission using an Agilent Cary 630 FT-IR spectrometer. Background spectra were collected prior to sample analysis, and the diamond crystal was cleaned with ethanol between samples. To best mimic the FLOPP-e library (treated as the calibration set), the same spectral collection parameters of 32 coadded scans at  $4\text{ cm}^{-1}$  resolution were used to collect signal between  $4000$  and  $650\text{ cm}^{-1}$ . The unidentified spectra were manually evaluated by one researcher to label each spectrum with a primary polymeric identity based on labeled spectra in the FLOPP and FLOPP-e libraries. Prior to accessing the manually identified labels, the CP method was applied to the environmental spectra (treated as the test set) by a different researcher. The spectral labels determined by manual evaluation were then used to calculate empirical confidence once the CP results were obtained.

## 3.3 Results and Discussion

### 3.3.1 Conformal Prediction

Using the five reference labels with 10 or more spectra in the FLOPP and FLOPP-e libraries, we quantified the mean empirical confidence and prediction set size of using CP with PCC and NN similarity metrics (Figure 3.2). As expected, the empirical confidence closely aligns with the confidence guarantee (90%) for both metrics. However, we observed a lower average prediction set size using the NN metric compared to PCC. The fact that PCC is an absolute metric whereas NN is a relative metric may contribute to this difference in the number of reference labels in the prediction set, as the criteria for generating a high HQI with NN are more restrictive than with PCC. A further explanation for this difference in metric efficiency

relates to the distributions of the positive matches (i.e., correctly identified) and negative matches (i.e., mismatch of query and reference labels) between spectra of the same FLOPP-e and FLOPP identity. In contrast to the NN, the distribution of positive and negative matches for PCC shows higher overlap, which indicates that many FLOPP-e spectra have high HQI scores relative to incorrect labels.




Figure 3.2: Empirical confidence and mean prediction set size results of conformal prediction using the Pearson correlation coefficient (PCC) and nearest neighbor (NN) and class-conditional conformal prediction. Here, CC-PCC and CC-NN correspond to the use of class-conditional conformal prediction with PCC and NN, respectively.

Class-conditional CP builds a unique HQI distribution for each label. CC-CP is particularly useful when the distributions of positive matches between sample and reference spectra differ among the spectral labels. In the case of the five labels used in our analysis, we noted that the HQI distribution of PP is distinct from the other labels when treated with PCC, as it exhibits HQI scores of true positive matches down to 0.4. We hypothesize that the difference in score distributions between labels could be due to the impacts of signal-to-noise ratio, preprocessing steps performed by DeFrond et al. in creating the library, or a greater degree of environmental aging when compared to pristine reference spectra [Cowger et al., 2020]. As mentioned earlier, CC-CP ensures a class-balanced confidence guarantee rather than an average guaranteed confidence, which results in more robust predictions at the expense of slightly larger set sizes (Figure 3.2). Similar to traditional CP, we observe an advantage of NN over PCC with respect to the mean set size, which we relate to the efficiency of the metric.

Despite spectral changes that may have occurred due to weathering, the manually assigned labels of 28 environmental unknowns were frequently returned in the prediction sets

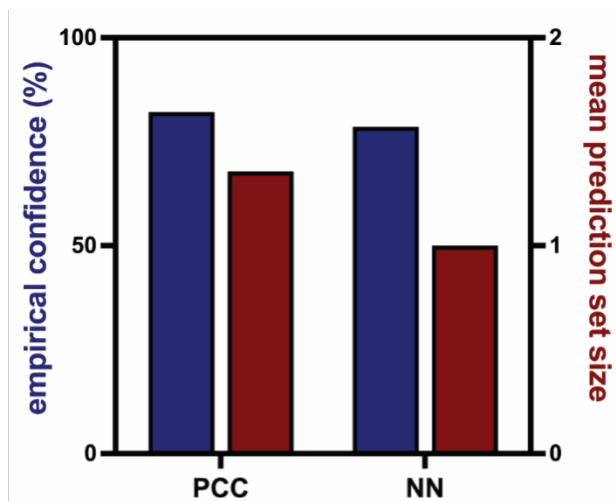


Figure 3.3: Empirical confidence and mean prediction set size results of conformal prediction applied to 28 environmental plastics using PCC and NN similarity metrics.

generated with both PCC (82%) and NN (79%), approaching the 90% confidence guarantee afforded by CP (Figure 3.3). As CP’s guaranteed confidence relies on an assumption that the HQI distributions of the calibration and test sets be similar, we can expect the empirical confidence to fluctuate from the theoretical guarantee when the distributions differ, perhaps due to different instrumentation used between this study and that of De Frond et al. [2021b]. On average, PCC returned more than one label in the prediction set, and NN returned a single correct label. NN, as the more efficient metric, better describes the inherent uncertainty in the environmental spectra. In this vein, the proportion of prediction sets with one or zero labels is an important indicator of metric efficiency and prediction quality. We further used the environmental spectra to compare the traditional CP method’s performance to that of CC-CP and KnowItAll and found that traditional CP outperformed both alternative methods. While the proprietary matching algorithm of KnowItAll limits our understanding of its performance, we attribute the lower performance of CC-CP to discrepancies between the distributions of label-specific test and calibration sets.

### 3.3.2 Comparing the Performance of KnowItAll and the Conformal Prediction Framework

To showcase the advantage of using CP in routine environmental microplastic identification, we compared the matching of FLOPP-e PE and PP spectra to all labels of FLOPP spectra with both KnowItAll’s default correlation metric and CP using PCC and NN metrics. While the analysis pipeline of KnowItAll’s SearchIt function does not guarantee confidence in spec-




Figure 3.4: Results of comparison between KnowItAll and CP with PCC and NN similarity metrics in terms of empirical confidence and mean prediction set size for a) PE and b) PP.

tral identification, we chose to view the results through a conformal prediction lens. To do so, we equated the proportion of instances that KnowItAll returned the correct label above an HQI threshold of 0.7 for a query spectrum with the empirical confidence associated with a prediction set generated with CP. Thus, KnowItAll’s “empirical confidence” was defined as the proportion of query spectra per polymer label (FLOPP-e PE and PP) that were matched to a FLOPP spectrum of the correct label with an  $HQI \geq 0.7$ . The empirical confidence of KnowItAll ( $\sim 94\%$  for PE and  $\sim 82\%$  for PP) was used to inform the confidence guarantee used to compare CP methods ( $94\%$  [ $\alpha = 0.06$ ] for PE and  $82\%$  [ $\alpha = 0.18$ ] for PP), such that the differences in the mean prediction set size were highlighted. As opposed to the mean set size for KnowItAll, here defined as the average number of spectral labels returned with  $HQI \geq 0.7$  per query spectrum, the use of CP with PCC or NN similarity metrics decreases the number of labels returned and provides guaranteed confidence in the match (Figure 3.4).

NN outperforms PCC to give a lower mean prediction set size due to fewer instances of negative matches above the defined HQI quantile, as previously described. Between query spectra identities, the mean set size of predicted labels for PP is lower than that for PE. The HQI scores to reference polymer spectra differ between the FLOPP-e PE and PP spectra, with PE having higher PCC scores to more reference labels than PP. This observation may explain why PE has a larger mean set size than PP, as labels with higher HQI scores are more likely to be included in the prediction set. Additionally, as the confidence guarantee for PP was reduced to reflect KnowItAll’s “empirical confidence”, the resultant HQI threshold is less descriptive of the true positive match score distribution. The higher HQI score associated with a lower guaranteed confidence could limit the number of labels added to the prediction set. Across both polymer identities and magnitudes of confidence guarantees, CP methods provided statistical assurance in the spectral labeling and reduced the number of labels

returned, streamlining manual interpretation of numerous matches.

### **3.3.3 Takeaways for the Microplastic Community: The Threshold Used for Labeling an Unknown Spectrum Relies on User-Data, Similarity Metric, and Confidence Guarantee**

The performance of the CP is dependent on data quality. While we use the FLOPP and FLOPP-e libraries herein, it is important to note that different reference libraries and calibration sets may be used with the CP workflow, but the similarity metric and data analysis routine should remain consistent. As previously noted, the HQI between a reference and query spectrum can be similar for multiple reference labels. Therefore, it is also important that a chosen HQI threshold has a statistical guarantee of differentiation between spectral signatures. As shown in Figure 3.5, when using conformal prediction and PCC with all 11 labels included in both the FLOPP and FLOPP-e libraries, the confidence guarantee that the user selects has implications on both the HQI threshold and the mean prediction set size. If one opts for higher guaranteed confidence in the generated labels, the HQI quantile will be more descriptive of the data and outliers (i.e., lower), and there will be more predicted labels in the set. Therefore, choosing too high of a confidence guarantee could render the prediction sets ineffective unless subjected to tedious manual evaluation of the matches, and it may lead to false positives, causing an overestimation of MP abundance. On the other hand, choosing too low of a confidence guarantee is accompanied by a higher HQI cutoff and fewer labels in the prediction set. Here, false negatives could occur because the environmental spectra are too dissimilar from reference spectra, leading to an underestimation of MP contamination. Given this trade-off, researchers must decide the confidence guarantee with which they are comfortable when considering the balance of HQI threshold, prediction set size, and the possibility of miscounting MP quantities. Our suggestion is to choose a guaranteed confidence that is both descriptive of most of the calibration set’s HQI distribution (e.g.,  $> 70\%$  [ $\alpha < 0.3$ ]) and returns a mean prediction set size for manageable manual evaluation (e.g., 1–3 labels). Using Figure 3.5 as a guide, we chose a guaranteed confidence between 70 and 90%. While the results in Figure 3.5 are specific to the metrics and data discussed herein, it carries that choosing a single HQI threshold without statistical meaning and applying it to diverse environmental campaigns, as is popular in the wider MP literature, may impact the validity of one’s results.

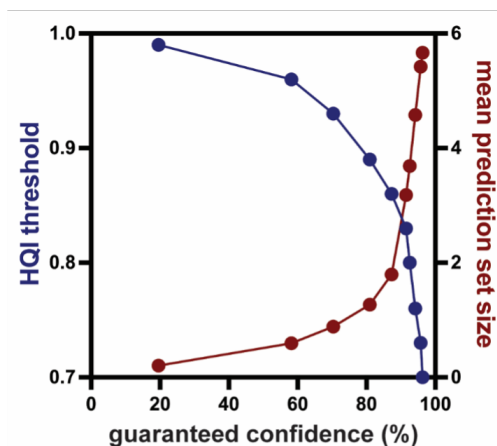


Figure 3.5: Hit quality index threshold and mean prediction set size as a function of the user-defined guaranteed confidence for FLOPP and FLOPP-e data using conformal prediction and PCC with 11 reference labels.

### 3.3.4 Using the Open-Access Code Available on GitHub, a Conformal Prediction Workflow Can be Used to Reduce Manual Spectral Comparison and Increase Confidence in Environmental Microplastic Labeling

Herein, we outlined a flexible framework for statistically guaranteed confidence in environmental microplastic spectral database searching. Based on our findings from a literature review of recent microplastic data analysis routines and evaluating a commercial database matching software, we illustrated the need for guidance on similarity metric comparison, hit quality index thresholding, and confidence in spectral labeling. To fill this gap, we highlighted that conformal prediction streamlines environmental plastic identification with an associated statistical uncertainty measurement and guaranteed confidence. To assist other microplastic researchers in using conformal prediction, we have generated an open-access code, user-guide, and suggested workflow that outlines an adaptable analytical framework for statistically robust spectral labeling. With these tools in hand, we urge the microplastic community to adopt conformal prediction to meet their analytical goals for confidence in hit quality index thresholding and acceptance of environmental microplastic spectral matches.

## CHAPTER 4

# Conformal Prediction for Ensembles: Improving Efficiency via Score-Based Aggregation

Distribution-free uncertainty estimation for ensemble methods is increasingly desirable due to the widening deployment of multi-modal black-box predictive models. Conformal prediction is one approach that avoids such distributional assumptions. Methods for conformal aggregation have in turn been proposed for ensembled prediction, where the prediction regions of individual models are merged as to retain coverage guarantees while minimizing conservatism. Merging the prediction regions directly, however, sacrifices structures present in the conformal scores that can further reduce conservatism. We, therefore, propose a novel framework that extends the standard scalar formulation of a score function to a multivariate score that produces more efficient prediction regions. We then demonstrate that such a framework can be efficiently leveraged in both classification and predict-then-optimize regression settings downstream and empirically show the advantage over alternate conformal aggregation methods.<sup>1</sup>

### 4.1 Introduction

Ensemble methods are an oft-used class of statistical modeling techniques due to their ability to reduce variance or improve predictive accuracy [Schapire et al., 1999, Zhang and Ma, 2012, Dietterich, 2000]. Such methods are increasingly being coupled with complex, black-box models, such as in multi-modal language models [Zhang et al., 2023a, Radford et al., 2021, Sun, 2013, Zhao et al., 2017, Yan et al., 2021]. Couplings of this sort are seeing ever-widening deployment in safety-critical settings, such as medicine [Yuan et al., 2018, Li et al., 2018, Yuan et al., 2017] and robotics [Brena et al., 2020, Blasch et al., 2021, Alatise and Hancke, 2020].

---

<sup>1</sup>This chapter is adapted from previously published work [Ochoa-Rivera et al., 2026], with Eduardo Ochoa Rivera, Yash Patel and Ambuj Tewari. This chapter describes equal contributions of authors Eduardo Ochoa Rivera and Yash Patel. The present version omits some of the sections for which Yash Patel had primary responsibility.




Figure 4.1: CSA provides a principled extension to the standard conformal prediction pipeline by leveraging ideas from higher-dimensional quantile regression to define quantile envelopes  $\widehat{\mathcal{Q}}$  instead of scalar quantiles  $\widehat{q}$ . It does so by evaluating a collection of score functions (here  $s_1$  and  $s_2$ ) over the calibration dataset to define  $\mathcal{S}$ , finding quantiles  $\{\widehat{q}_m\}$  over a set of projection directions  $\{u_m\}$ , and taking  $\widehat{\mathcal{Q}}$  to be the intersection of the resulting half-planes  $H(u_m, \widehat{q}_m)$ . These quantile envelopes result in more informative prediction regions that can be used in downstream tasks.

Increasing interest is, therefore, now being placed on quantifying uncertainty for such models [Subedar et al., 2019, Tian et al., 2020, Denker and LeCun, 1990, Havasi et al., 2020, Malinin and Gales, 2018]. Towards this end, methods of uncertainty quantification have arisen, such as deep ensembles and committee estimation [Rahaman et al., 2021, Abdar et al., 2021, Carrete et al., 2023]. Such methods, however, sacrifice generality with the imposition of distributional assumptions, motivating the need for distribution-free uncertainty quantification for ensemble methods.

One method for performing distribution-free uncertainty quantification is conformal prediction, which provides a principled framework for producing distribution-free prediction regions with marginal frequentist coverage guarantees [Angelopoulos and Bates, 2021, Shafer and Vovk, 2008]. By using conformal prediction on a user-defined score function, prediction regions attain marginal coverage guarantees. While calibration is guaranteed from this procedure, predictive efficiency, quantified as the size of the resulting prediction regions, can be unboundedly large for poorly chosen score functions.

As a result, methods have arisen to perform conformal model aggregation, which both provide uncertainty estimates of the ensembled predictions and do so in ways as to minimize the prediction region size [Gasparin and Ramdas, 2024a, Trunov and V’yugin, 2023, Yang and Kuchibhotla, 2024, V’yugin and Trunov, 2023, Gasparin and Ramdas, 2024b]. While such approaches succeed in reducing the prediction region size over naive aggregation, they all aggregate the *separately conformalized* prediction regions of the predictors in the ensemble. In doing so, they forgo the possibility of automatically leveraging shared structure amongst the scores of the individual predictors, resulting in overly conservative prediction regions.

We instead propose to perform aggregation in *score space* by extending traditional conformal prediction to consider a multivariate score function and defining prediction regions using “quantile envelopes” in place of scalar quantiles. Doing so enables efficient, data-driven, automated conformal model aggregation. We demonstrate that this formulation retains the desired distribution-free coverage guarantees typical of standard conformal prediction and that the resulting prediction regions can be used efficiently in both classification and regression settings. Our contributions, thus, are:

- Providing a multivariate extension to conformal prediction, dubbed “conformal score aggregation” (CSA), that leverages quantile envelopes to enable data-driven, informative uncertainty estimation for model ensembles while retaining coverage guarantees.
- Demonstrating how the prediction regions resulting from CSA can be efficiently leveraged in downstream predict-then-optimize regression tasks.
- Demonstrating the empirical improvement of the CSA framework over alternate conformal aggregation strategies across classification and regression settings.

## 4.2 Background

### 4.2.1 Conformal Prediction

Coverage guarantees of uncertainty quantification methods generally rely on distributional assumptions, often via asymptotics or explicit specification. To alleviate the need for such restrictive assumptions, interest in finite-sample, distribution-free uncertainty quantification methods has risen. Conformal prediction is one such method [Angelopoulos and Bates, 2021, Shafer and Vovk, 2008]. Due to its lack of structural assumptions, conformal prediction pairs well with the increasingly common black-box style of modeling.

Conformal prediction serves as a wrapper around such predictors, producing prediction regions  $\mathcal{C}(x)$  that have formal guarantees of the form  $\mathcal{P}_{X,Y}(Y \notin \mathcal{C}(X)) \leq \alpha$  for some pre-specified level  $\alpha$ . To achieve this, “split conformal” partitions the dataset  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$  into a training set  $\mathcal{D}_T$  and a calibration set  $\mathcal{D}_C$ . The former serves as the data used to fit  $\hat{f}$ .

Users of conformal prediction must then design a “score function”  $s(x, y)$ , which should quantify “test error”, often in a domain-specific manner. For instance, a simple score function for a regression setting would be  $s(x, y) = \|\hat{f}(x) - y\|$ . This score function is then evaluated across the calibration set to define  $\mathcal{S}_C = \{s(x, y) \mid (x, y) \in \mathcal{D}_C\}$ . For a desired coverage of  $1 - \alpha$ , we then take  $\hat{q}$  to be the  $\lceil (|\mathcal{D}_C| + 1)(1 - \alpha) \rceil / |\mathcal{D}_C|$  quantile of  $\mathcal{S}_C$ , with which prediction regions for future test queries  $x$  can be defined as  $\mathcal{C}(x) = \{y \mid s(x, y) \leq \hat{q}\}$ .

Under the exchangeability of the score of a test point  $s(X', Y')$  with  $\mathcal{S}_C$ , we have the desired *finite-sample* probabilistic guarantee that  $1 - \alpha \leq \mathcal{P}_{X', Y'}(Y' \in \mathcal{C}(X'))$ .

While this guarantee holds for any  $s(x, y)$ , the informativeness of the resulting prediction regions, quantified as the inverse expected Lebesgue measure across  $X$ , i.e.  $(\mathbb{E}[\mathcal{L}(\mathcal{C}(X))])^{-1}$ , is intimately tied to its specification [Shafer and Vovk, 2008]. Thus, much of the challenge of conformal prediction relates to choosing a score function that retains coverage while minimizing region size.

## 4.2.2 Quantile Envelopes

Generalizations of quantiles have a long history in statistics [Rousseeuw and Struyf, 1998, Serfling, 2002]. Unlike univariate data, multivariate data does not lend itself to an unambiguous definition of a quantile, as there is no canonical ordering in higher dimensional spaces. To account for this, one class of multivariate quantiles, known as “direction quantiles,” prescribes directions along which data points can be compared [Kong and Mizera, 2012, Paindaveine and Šiman, 2011, Hallin et al., 2010]. That is, points  $x \in \mathbb{R}^n$  are projected along some direction  $u \in \mathcal{S}^{n-1}$ , after which the standard notion of a quantile can be leveraged.

The notion of the  $\alpha$  quantile for a random variable  $X \in \mathcal{X}$  clearly then depends on this choice of direction  $u$ , which we denote by  $Q(X, \alpha, u) = \inf\{q \in \mathbb{R} : \mathcal{P}(u^\top X \leq q) \geq \alpha\}$ . When there is no ambiguity, we just denote it as  $Q(\alpha, u)$ . In cases where domain knowledge assigns an unambiguous notion of multivariate ordering, such knowledge can be encoded into the choice of  $u$ . In other cases, however, a more agnostic approach is employed: rather than choosing any particular direction, all directions are considered simultaneously. For any given  $u$ , notice the choice of quantile defines a corresponding halfplane  $H(u, Q(\alpha, u)) = \{x \in \mathcal{X} : u^\top x \leq Q(\alpha, u)\}$ . To consider all directions simultaneously, the quantile envelope is then defined to be

$$D(\alpha) = \bigcap_{u \in \mathcal{S}^{n-1}} H(u, Q(\alpha, u)). \quad (4.1)$$

Notably, while each individual  $H(u, Q(\alpha, u))$  captures  $1 - \alpha$  of the points,  $D(\alpha)$  does *not*, as it is the intersection thereof and hence captures  $< 1 - \alpha$  of the mass. If  $1 - \alpha$  combined coverage is sought, a correction, such as Bonferroni adjustment, is used for the individual planes.

### 4.2.3 Related Works

Ensemble methods consist of  $K$  predictors  $f_k : \mathcal{X}_k \rightarrow \mathcal{Y}$ ; notably, such predictors need not map from the same set of covariates. A naive approach for uncertainty quantification would then be to conformalize the ensembled predictor. That is, for an ensembling algorithm  $\mathcal{F} : \mathcal{Y}^K \rightarrow \mathcal{Y}$ , a score function  $s(\mathcal{F}(f_1(x), \dots, f_K(x)), y)$  would be defined. Denoting the  $[(N_C + 1)(1 - \alpha)]/N_C$  quantile of the score distribution over  $\mathcal{D}_C$  as  $\hat{q}(\alpha)$ ,  $\mathcal{C}(x) = \{y : s(x, y) \leq \hat{q}(\alpha)\}$  would then be marginally calibrated.

Such an approach, however, lacks some desirable properties. In particular, prediction regions  $\mathcal{C}(x)$  should have the quality that, if a particular predictor has less uncertainty in its predictions, as is frequently true of ensemble settings where the predictors span multiple input data modalities, upon routing to that predictor, the corresponding size of the prediction region should be smaller than if it had been routed to a different predictor. While the naive approach does, in principle, support this property, it ultimately relies on defining an *uncertainty-aware* ensembling algorithm  $\mathcal{F}$ . In its typical form, however,  $\mathcal{F}$  simply takes *point predictions*  $f_1(x), \dots, f_K(x)$  in as input, meaning any uncertainty-awareness would need to be baked in a priori into the definition of  $\mathcal{F}$  through domain knowledge of the uncertainties of the predictors  $f_1, \dots, f_K$ , which can seldom be specified precisely, sacrificing the predictive efficiency of  $\mathcal{C}(x)$ .

Conformal model aggregation, thus, seeks to mitigate these deficiencies by aggregating the prediction regions  $\mathcal{C}_1(x), \dots, \mathcal{C}_K(x)$  rather than the individual point predictions [Gasparin and Ramdas, 2024a, Yang and Kuchibhotla, 2024, V'yugin and Trunov, 2023, Gasparin and Ramdas, 2024b]. While there are several methods in this vein, they can be categorized into one of two general approaches.

The first line of work seeks to perform model *selection*, in which a single conformal predictor is selected  $\mathcal{C}_{k^*}$ , typically based on the criterion of minimizing region size  $k^* := \arg \min_k \mathbb{E}[\mathcal{L}(\mathcal{C}_k(X))]$  [Yang and Kuchibhotla, 2024, V'yugin and Trunov, 2023]. Generally, however, methods leveraging the full collection of predictors produce less conservative regions [Gasparin and Ramdas, 2024a,b]. Such works aggregate the individual prediction regions into a final region by defining

$$\mathcal{C}(x) := \left\{ y \mid \sum_{k=1}^K w_k \mathbb{1}[y \in \mathcal{C}_k(x)] \geq \hat{a} \right\} \quad (4.2)$$

for weights  $\{w_k\} \in [0, 1]$  such that  $\sum_{k=1}^K w_k = 1$  and a threshold  $\hat{a}$ . Methods then differ in the procedure by which  $\{w_k\}$  and  $\hat{a}$  are prescribed, several of which were prescribed in [Gasparin and Ramdas, 2024b], whose detailed presentation is deferred to Section B.3 for

space reasons. We note that the methods of [Gasparin and Ramdas, 2024a] are designed for a different setting than that considered herein, namely that in which conformal coverage is sought adaptively over online data streams.

In this vein, Luo and Zhou [2024] is a recent work that has similarly proposed a vector-score extension as that discussed herein, in which candidate weight vectors  $\{w_m\} \in \mathbb{R}^K$  are searched over for score aggregation. That is, a vector  $s(x) := (s_1(x, y), \dots, s_K(x, y)) \in \mathbb{R}^K$  of scores  $s_k(x, y)$  corresponding to each predictor  $f_k(x)$  is predicted and its aggregate prediction region defined on the projection  $\langle w_{m^*}, s \rangle$  for  $w_{m^*}$  the weight resulting in the smallest prediction region. This method, however, has two shortcomings addressed herein. The first is that their method can only be applied in classification settings, whereas our method can be leveraged across both regression and classification problems. The second is that their approach only uses a *single* weighted projection in the end, resulting in suboptimal aggregation and, therefore, conservative prediction regions.

## 4.3 Method

We now introduce a procedure for conformal score aggregation for ensembles. We discuss the approach and its coverage guarantees in Section 4.3.1.

### 4.3.1 Multivariate Score Quantile

We wish to consider the setting typical of conformal model aggregation, as discussed in Section 4.2.3, in which predictors  $f_1(x), \dots, f_K(x)$  and corresponding scores  $s_1(x, y), \dots, s_K(x, y)$  are defined. We further assume a similar premise as that taken by Luo and Zhou [2024], in which the individual scores are stacked into a multivariate score  $s(x, y) := (s_1(x, y), \dots, s_K(x, y))$ . A naive approach would then involve leveraging standard conformal prediction over a pre-defined map  $g : \mathbb{R}^K \rightarrow \mathbb{R}$ , e.g.,  $g(s) = \sum_{k=1}^K s_k$ . Similar to the naive conformalization of an ensembled predictor discussed in Section 4.2.3, using a *fixed*  $g$  fails to adapt to any disparities in uncertainties present across predictors or requires intimate knowledge of such uncertainties. We instead wish to provide a data-adaptive pipeline to automatically produce such a  $g$ .

Importantly, we heretofore assume the score functions are non-negative, i.e.,  $s_k : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ , which is typically the case as the score serves as a generalization of the residual. We highlight that many of the details of the method presented below are geometric in nature and are more easily understood with the supplement of diagrams. We have, thus, provided an accompanying visual walkthrough of the procedure in Section B.1 to clarify its presentation.

### 4.3.1.1 Score Partial Ordering

Intuitively, our method seeks to directly generalize the approach of split conformal, by “ordering” the collection of multivariate calibration scores and taking the  $1 - \alpha$  score under such an ordering to be a threshold  $\hat{\mathcal{Q}}$  with which prediction regions are then implicitly defined. Formally, the multivariate “ordering” is established as a pre-ordering  $\lesssim$  over  $\mathbb{R}^K$ ; a pre-ordering differs from a total ordering in that it need not satisfy the antisymmetric axiom of a total ordering. Roughly speaking, an “acceptance region,” so called as it serves as the criterion used to ultimately decide which  $y$  are accepted into the prediction region, is then defined as  $\hat{\mathcal{Q}} := \{s \mid s \lesssim \hat{q}\}$ , where  $\hat{q}$  is the  $1 - \alpha$  empirical quantile of  $\mathcal{S}_C$  under  $\lesssim$ . Such a  $\hat{\mathcal{Q}}$  naturally generalizes the standard scalar acceptance interval of  $[0, \hat{q}]$  in the case of non-negative score functions. We briefly highlight the distinction between *acceptance regions* and *prediction regions*: the former are subsets of the space of multivariate scores, i.e.  $\subset \mathbb{R}^K$ , that ultimately define the criteria for retaining particular  $y$  values in the prediction region whereas the latter are the subsets of the output space that ultimately have coverage guarantees, i.e.  $\subset \mathcal{Y}$ . The two, however, are directly related; in particular, for a fixed score  $s(x, y)$ , a larger acceptance region will result in a more conservative prediction region.

Crucially, therefore, the problem of choosing this pre-ordering closely parallels that of choosing  $g$ , where a poorly chosen pre-ordering will result in overly large acceptance regions and, hence, conservative prediction regions. For instance, using a lexicographical ordering  $\lesssim_{\text{Lex}}$  will result in axis-aligned hyper-rectangular acceptance regions.

As a result, rather than manually prescribing a pre-ordering, we define  $\lesssim$  in a data-driven fashion by prescribing an indexed family of nested sets  $\{\mathcal{A}_t\}_{t \in \mathbb{R}}$ , such that  $\mathcal{A}_{t_1} \subset \mathcal{A}_{t_2}$  for  $t_1 \leq t_2$  and stating  $s_1 \lesssim s_2$  if  $\forall t, s_2 \in \mathcal{A}_t \implies s_1 \in \mathcal{A}_t$ .

For a family of sets  $\{\mathcal{A}_t\}_{t \in \mathbb{R}}$ , we take each  $\mathcal{A}_t$  to be the region of the positive orthant  $\mathbb{R}_+^K$  bounded by the coordinate axes and an “outer frontier” parameterized by  $t$ .

The shape of this outer frontier remains fixed over the family and is merely scaled outward from the origin with  $t$ . Under this choice, performing a comparison of  $s_1, s_2 \in \mathbb{R}^K$  in the aforementioned fashion, i.e. checking if  $s_1 \lesssim s_2$ , amounts to checking if  $t(s_1) \leq t(s_2)$ , where  $t(s)$  is the smallest  $t$  for which the outer frontier of  $\mathcal{A}_t$  intersects  $s$ . Notably,  $t(s)$  is precisely the aforementioned data-driven score fusion function  $g(s)$  of interest. Defining a data-adaptive  $g(s)$ , therefore, reduces to having a data-driven approach for defining the outer frontier of  $\mathcal{A}_t$ . We restrict this outer frontier to be such that  $\mathcal{A}_t$  is a convex set; if  $\mathcal{A}_t$  were permitted to be nonconvex, computing  $t(s) := \min\{t \in \mathbb{R} : s \in \mathcal{A}_t\}$  would potentially be computationally expensive.

To have tight acceptance regions, we formally wish for the pre-ordering to have the property that the acceptance region given by  $\tilde{\mathcal{Q}}$  has minimal Lebesgue measure and captures  $1 - \alpha$

points of  $\mathcal{S}_C$ . The problem of discovering an optimal pre-ordering can, thus, be equivalently stated as seeking to define the outer frontier of  $\mathcal{A}_t$  to match that of the tightest  $1 - \alpha$  convex cover of  $\mathcal{S}_C$ .

This final reframing, therefore, naturally motivates selecting the outer frontier to be the  $1 - \alpha$  quantile envelope of  $\mathcal{S}_C$ . Using  $\mathcal{S}_C$  to define  $\mathcal{A}_t$  and in turn  $\lesssim$ , however, sacrifices the exchangeability of its points with future test scores  $s'$ , as the very nature of ordering would change in swapping  $s'$  with any  $s \in \mathcal{S}_C$ . The goal, then, follows as seeking to define the outer frontier as the  $1 - \alpha$  quantile envelope of  $\mathcal{S}_C$  without directly using  $\mathcal{S}_C$ . For this reason, we partition  $\mathcal{S}_C = \mathcal{S}_C^{(1)} \cup \mathcal{S}_C^{(2)}$ , where we define  $\lesssim$  using  $\mathcal{S}_C^{(1)}$  and compute the final  $\hat{q}$  over  $\mathcal{S}_C^{(2)}$ . Such a split is predicated on the assumption that the  $1 - \alpha$  quantile envelope defined over  $\mathcal{S}_C^{(1)}$  resembles that of  $\mathcal{S}_C^{(2)}$ , implying the  $|\mathcal{S}_C^{(1)}|$  should be sufficiently large as to capture this structure accurately.

We now focus attention on defining the quantile envelope over  $\mathcal{S}_C^{(1)}$  using a technique paralleling that described in Section 4.2.2. In particular, we start by selecting the projection directions  $\{u_m\}$  of Equation (4.1). The restriction of  $s$  to the positive orthant induces a natural modification of Equation (4.1), namely where  $u \in \mathcal{S}_+^{K-1} := \mathcal{S}^{K-1} \cap \mathbb{R}_+^K$  instead of  $\mathcal{S}^{K-1}$ . To best approximate Equation (4.1), we wish for  $\{u_m\}$  to be uniformly distributed over  $\mathcal{S}_+^{K-1}$ ; however, exactly finding an evenly distributed set of points over hyperspheres in arbitrary  $n$ -dimensional spaces is a classically difficult problem [Schnabel and Janke, 2022]. If  $K = 2$ , we can solve this exactly; for  $K > 2$ , we generate directions stochastically such that  $U \sim \text{Unif}(\mathcal{S}_+^{K-1})$  by drawing  $V_1, \dots, V_M \sim \mathcal{N}(0, I^{K \times K})$  and defining  $U_i := V_i^{|\cdot|} / \sqrt{V_1^2 + \dots + V_M^2}$ , where  $v^{|\cdot|}$  denotes the component-wise absolute values of  $v$ .

We now wish to define the quantile thresholds  $\{\tilde{q}_m\}$  for the selected directions to optimally capture  $1 - \alpha$  of  $\mathcal{S}_C^{(1)}$ . Naively taking the  $1 - \alpha$  quantile per projection direction  $u_m$  results in *joint* coverage by  $\tilde{\mathcal{Q}} := \bigcap_{m=1}^M H(u_m, \tilde{q}_m)$  of  $\mathcal{S}_C^{(1)}$  to be  $< 1 - \alpha$ . A straightforward fix is to replace the  $1 - \alpha$  quantile per direction instead with its Bonferroni-corrected  $1 - \alpha/M$  quantile. While valid, this approach produces overly conservative prediction regions. We, therefore, instead tune a separate  $\beta \in (\alpha/M, \alpha)$  parameter via binary search, finding the maximum  $\beta^*$  such that using the  $\beta^*$  quantile per direction provides the overall desired coverage, i.e.  $|\bigcap_{m=1}^M H(u_m, \tilde{q}_m(1 - \beta^*)) \cap \mathcal{S}_C^{(1)}| / N_{\mathcal{C}_1} \in (1 - \alpha, 1 - \alpha + \epsilon)$  for some fixed, small  $\epsilon > 0$ . With this choice of  $\{(u_m, \tilde{q}_m)\}$ , we have a defined pre-ordering, whose coverage guarantees are formally stated below and proven in Section B.2.

**Theorem 4.3.1.** *Suppose  $\mathcal{D}_C := \{(X_i, Y_i)\}_{i=1}^{N_C}$  and  $(X', Y')$  are exchangeable. Assume further that  $K$  maps  $s_k : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  have been defined and a composite  $s(X, Y) := (s_1(X, Y), \dots, s_K(X, Y))$  is defined. Further denote by  $\mathcal{S}_C$  the evaluation of  $s(X, Y)$  on  $\mathcal{D}_C$ , namely  $\mathcal{S}_C := \{s(X_i, Y_i) \mid (X_i, Y_i) \in \mathcal{D}_C\}$ . For some  $\alpha \in (0, 1)$ , given a pre-order  $\lesssim$  in  $\mathbb{R}^K$*

induced by a collection of nested sets  $\{\mathcal{A}_t\}_{t \geq 0}$  define  $\mathcal{Q}(\alpha) = \{s \in \mathbb{R}^K : s \lesssim s_{\lceil (N_{C_2} + 1)(1 - \alpha) \rceil}\}$ . Then, denoting  $\mathcal{C}(x) := \{y : s(x, y) \in \mathcal{Q}(\alpha)\}$ ,  $\mathcal{P}_{X', Y'}(Y' \in \mathcal{C}(X')) \geq 1 - \alpha$ .

### 4.3.1.2 Score Quantile Threshold

To then compute  $\hat{q}$ , we find  $t^*(s)$  for each  $s \in \mathcal{S}_{\mathcal{C}}^{(2)}$ , defined to be  $\min\{t \in \mathbb{R} : s \in \bigcap_{m=1}^M H(u_m, t\tilde{q}_m)\}$ . This can be efficiently computed as  $t^*(s) = \max_{m=1, \dots, M} (u_m^\top s / \tilde{q}_m)$ . Denoting the  $\lceil (N_{C_2} + 1)(1 - \alpha) \rceil$ -th largest  $t^*(s)$  as  $\hat{t}$ ,  $\hat{q}_m := \hat{t}\tilde{q}_m$  and  $\hat{\mathcal{Q}} := \bigcap_{m=1}^M H(u_m, \hat{q}_m)$ . If the tightest quantile envelope was already discovered over  $\mathcal{S}_{\mathcal{C}}^{(1)}$ , this adjustment factor  $\hat{t} \approx 1$ .

Importantly, while the aforementioned procedure will necessarily result in convex regions  $\hat{\mathcal{Q}}$ , this does **not** mean the downstream prediction regions in  $\mathcal{Y}$  will be convex. Such flexibility is necessary, as many studies have demonstrated the need for nonconvex prediction regions for downstream utility, such as [Patel et al., 2023, Tumu et al., 2023, Feldman et al., 2023]. However, it is unsurprising such flexibility exists, as even a single *scalar* score  $s_1(x, y)$  can produce nonconvex prediction regions. We present the full algorithm in Algorithm 2.

---

**Algorithm 2** CSA: UNIFHYPERSPHERE( $K$ ) is an assumed subroutine that samples  $\sim \text{Unif}(\mathcal{S}^{K-1})$ .

---

- 1: **Inputs:** Score functions  $s_1, \dots, s_K : \mathcal{X} \rightarrow \mathcal{Y}$ , Calibration set  $\mathcal{D}_{\mathcal{C}}$ , Desired coverage  $1 - \alpha$
  - 2:  $[\beta_{\text{lo}}, \beta_{\text{hi}}] \leftarrow [\alpha/M, \alpha]$ ,  $\hat{\mathcal{Q}} \leftarrow \emptyset$
  - 3:  $\mathcal{S}_{\mathcal{C}}^{(1)} \cup \mathcal{S}_{\mathcal{C}}^{(2)} \leftarrow \{(s_k(x_i, y_i))_{k=1}^K\}_{i=1, N_{C_1}+1}^{N_{C_1}, N_{C_2}}$
  - 4:  $\{u_m \leftarrow \text{UNIFHYPERSPHERE}(K)\}_{m=1}^M$
  - 5: **while**  $|\mathcal{S}_{\mathcal{C}}^{(1)} \cap \hat{\mathcal{Q}}|/N_{C_1} \notin 1 - \alpha \pm \epsilon$  **do**
  - 6:    $\beta \leftarrow \frac{\beta_{\text{lo}} + \beta_{\text{hi}}}{2}$
  - 7:    $\left\{ \tilde{q}_m \leftarrow (1 - \beta) \text{ emp. quantile: } \{u_m^\top s_i\}_{s_i \in \mathcal{S}_{\mathcal{C}}^{(1)}} \right\}_{m=1}^M$
  - 8:    $\hat{\mathcal{Q}} \leftarrow \bigcap_{m=1}^M H(u_m, \tilde{q}_m)$
  - 9:   **if**  $|\mathcal{S}_{\mathcal{C}}^{(1)} \cap \hat{\mathcal{Q}}|/N_{C_1} > 1 - \alpha$  **then**  $\beta_{\text{lo}} \leftarrow \beta$
  - 10:   **else**  $\beta_{\text{hi}} \leftarrow \beta$
  - 11: **end while**
  - 12:  $\hat{t} \leftarrow (1 - \alpha) \text{ emp. quantile: } \{\max_{m \in [M]} \frac{u_m^\top s_i}{\tilde{q}_m}\}_{s_i \in \mathcal{S}_{\mathcal{C}}^{(2)}}$
  - 13: **Return**  $\{(u_m, \hat{t}\tilde{q}_m)\}_{m=1}^M$
- 

## 4.4 Experiments

We now study CSA empirically in a classification task, demonstrating its coverage guarantees with reduced conservatism. We demonstrate improvements in an ImageNet classification task

in Section 4.4.1.

To allow for fair comparison, we note that the predictors and calibration and test sets were fixed across choices of calibration procedure for each experiment. This in turn means that some care had to be taken in partitioning  $\mathcal{D}_C = \mathcal{D}_C^{(1)} \cup \mathcal{D}_C^{(2)}$  for CSA, where an insufficiently large  $\mathcal{D}_C^{(1)}$  would result in poor estimation of the  $\alpha$ -quantile envelope and hence require a large adjustment  $\hat{t}$  factor and an insufficiently large  $\mathcal{D}_C^{(2)}$  in the classical reduced predictive efficiency from conformal prediction. For this reason, we fixed the split to be 20%-80% in experiments.

We herein compare against the methods presented in Section 4.2.3, namely the model selection of Yang and Kuchibhotla [2024], the aggregation methods of Gasparin and Ramdas [2024a], and the single weighted score projection (VFCP) of Luo and Zhou [2024]. We specifically consider the following aggregation methods from Gasparin and Ramdas [2024a]: the standard majority-vote  $\mathcal{C}^M$ , partially randomized thresholding  $\mathcal{C}^R$ , and fully randomized thresholding  $\mathcal{C}^U$  approaches, fully described in Section B.3. Notably, these methods do not lend themselves for use in the predict-then-optimize setting, so we eliminate them from consideration therein. Similarly, VFCP can only be applied in classification settings; we, thus, do not compare to it across the regression tasks. Code will be made public upon acceptance.

#### 4.4.1 Classification Tasks

We first study the predictive efficiency of the aforementioned methods on the ImageNet classification task Deng et al. [2009]. In particular, an ensemble was constructed from three separately trained deep learning architectures, namely ResNet-50, VGG-11, and DenseNet-121. Conformalization on the individual models was performed using the standard classification score function across all approaches, namely  $s(x, y) = \sum_{j=1}^l \hat{f}(x)_{\pi_j(x)}$  where  $y = \pi_l(x)$  and  $\pi(x)$  is the permutation of  $\{1, \dots, |\mathcal{Y}|\}$  that sorts  $\hat{f}(x)$  from most to least likely. Calibration was performed using 85% of the ImageNet test set and assessment of the coverage and interval lengths on the remaining 15%, with 10 trials conducted over randomized draws of these calibration and test sets.

The results are presented in Table 4.1. We see that all the approaches exhibit the desired coverage guarantees. However, the CSA score approach consistently produce significantly smaller prediction regions than both the individually conformalized input views and alternate aggregation strategies.

Table 4.1: Average coverages across tasks for  $\alpha = 0.10$  are shown in the top rows and average prediction set sizes in the bottom rows. Both were assessed over a batch of i.i.d. test samples (15% of the validation set from ImageNet). Standard deviations and means were computed across 10 randomized draws of the calibration and test sets.

Dataset	Models	ResNet	VGG	DenseNet	VFCP	$\mathcal{C}^M$	$\mathcal{C}^R$	$\mathcal{C}^U$	CSA
ImageNet	Coverage	0.901 (0.005)	0.902 (0.003)	0.902 (0.003)	0.899 (0.004)	0.938 (0.003)	0.909 (0.004)	0.9 (0.004)	0.9 (0.003)
	Size	137.004 (1.98)	136.116 (2.206)	120.096 (2.427)	46.063 (1.089)	87.337 (1.604)	82.746 (1.692)	131.856 (2.378)	<b>34.006 (0.924)</b>

## 4.5 Discussion

We have presented CSA, a framework for producing informative prediction regions in ensemble predictor pipelines. We additionally demonstrated the generality of the proposed pipeline in settings of both classification and robust predict-then-optimize regression tasks. This work suggests many directions for extension. One point of interest is the integration of CSA with other downstream regression applications. For instance, Chenreddy and Delage [2024] proposed an end-to-end differentiable extension to Patel et al. [2023]: extending CSA for integration to their framework would be of great interest. Additionally, given the prevalence of late-stage fusion in robotics applications, an exciting applied avenue of research would investigate the use of CSA in settings of robust control.

## CHAPTER 5

# Improved Microplastic Identification from Simultaneously Collected Photothermal Infrared and Raman Spectra using Multi-View Conformal Prediction

Microplastics (MPs) have been documented in urban and remote locations across the globe. One of the most popular chemometric methods for identifying microplastics (MPs) is database matching, in which an unknown spectrum is compared with reference library spectra by calculating likeness scores. Threshold minima determine if the score is high enough to consider the reference spectrum as a potential match, yet these thresholds are frequently set arbitrarily. There is a growing consensus that MP identification should involve multiple measurement techniques, but statistically-robust methods to relate multiple database matching scores are lacking. Herein, multi-view conformal prediction (MVCP) incorporates two views (i.e., photothermal infrared (PTIR) and Raman spectra) to calculate multidimensional thresholds for MP identification with statistical confidence. The chemical identities returned for an unknown particle have a statistical assurance that one of the identities is the correct match based on a user-defined uncertainty parameter. The average number of potential matches returned by MVCP was closer to one chemical identity—the ideal number of identities returned—when compared to single-view CP methods that used either the PTIR or Raman spectra. Moreover, MVCP was less affected than its single-view counterparts when one of the two spectra was difficult, or impossible, to identify. To show the utility of MVCP for real-world samples, an ambient particle sample with MPs deposited on it was used to demonstrate that an MVCP threshold at 73% theoretical confidence maximized the fraction of MP particles correctly identified as plastic ( $0.80 \pm 0.07$ ), while limiting the fraction of non-MP environmental particles misidentified as plastic ( $0.10 \pm 0.06$ ). This initial application of MVCP to spectroscopy demonstrates the benefits of utilizing multiple spectral

methods with data analysis routines for identifying MPs with statistical confidence.<sup>1</sup>

## 5.1 Introduction

Microplastics (MPs) have become an increasing concern due to the rising prevalence of mis-managed plastic waste in the environment—the extent of which is largely uncertain [Jambeck et al., 2015, Zhang et al., 2023b]. Pieces of plastic between 1 and 5000  $\mu\text{m}$  in size are considered MPs, and they can be created intentionally (often for commercial products) or by the breakdown of larger pieces of plastic [Hale et al., 2020]. These particles must be identified chemically as visual methods are limited to particles  $> 50 \mu\text{m}$  [Bergmann et al., 2019, Isobe et al., 2019, Kotar et al., 2022, Lenz et al., 2015], with some studies recommending additional chemical validation for any MPs  $< 2 \text{ mm}$  in size [Isobe et al., 2019]. A growing number of spectroscopy and mass spectrometry techniques are being used to identify MPs in field studies as they provide information on the molecular structure of polymers [Cowger et al., 2020, Ivleva, 2021, Zhang et al., 2025b], which aids in distinguishing the polymers from other organic and inorganic species [Käppler et al., 2016, Niu et al., 2024, Parham et al., 2025, Renner et al., 2017, Watteau et al., 2018].

For vibrational spectroscopy techniques specifically, there are many commercial [Bruker, n.d., Thermo Fisher Scientific Inc., n.d., John Wiley & Sons, Inc., n.d.] and in-house [Parham et al., 2025, Cowger et al., 2021, De Frond et al., 2021a, Munno et al., 2020, Dong et al., 2020, Primpke et al., 2018] reference libraries available to help identify MPs from unknown infrared (IR) and Raman spectra using database matching. Hit Quality Index (HQI), a general term for similarity metrics, can be used to calculate a score between the unknown spectrum and a reference library spectrum [Clough et al., 2024, Weisser et al., 2022]. Returned HQI scores are often between 0 and 1 (though they are sometimes reported as 0–1000 or 0-100%) [Weisser et al., 2022, Renner et al., 2019], meaning the spectra are either entirely dissimilar or a perfect match, respectively. Once all of the scores between the unknown spectrum and corresponding reference library spectra are calculated and tabulated, the highest scores are returned to the user as potential matches (also known as labels) for the unknown particle.

The final decision on whether these labels are a correct match can be difficult as HQI scores between the lowest and highest possible values are often arbitrary and carry no inherent significance [Clough et al., 2024, Kozloski et al., 2024]. The scores are highly dependent upon a number of variables such as the metrics used to calculate HQI scores, the particle types included in the reference library, and the MP particle type [Clough et al., 2024, Renner et al.,

---

<sup>1</sup>This chapter is adapted from previously published work [Clough et al., 2024], with Rebecca L. Parham, Eduardo Ochoa Rivera, Abbygail M. Ayala, Madeline E. Clough, Yash Patel, Anne J. McNeil, Ambuj Tewari and Andrew P. Ault. This chapter describes equal contributions of authors Rebecca L. Parham and Eduardo Ochoa Rivera

2019, Kozloski et al., 2024]. Moreover, most commercial database matching software systems do not make their data analysis routines public, making direct comparisons between HQI scores calculated from different software systems challenging [Kozloski et al., 2024]. Many researchers have worked around these issues by selecting a minimum threshold value above which the HQI score is considered high enough for the reference library and unknown spectra to be a potential match [Clough et al., 2024, Weisser et al., 2022]. Methods for selecting a threshold have largely been ambiguous processes based on qualitative observations or the precedent of previous publications, resulting in researchers selecting threshold values ranging anywhere from 0.6 to 0.8—or not reporting a threshold value at all. Given that thresholds are not comparable between different database matching methods yet significantly affect match accuracy [Clough et al., 2024, Weisser et al., 2022, Kozloski et al., 2024], a data-driven approach that determines thresholds with statistical confidence values is crucial to accurately and efficiently determine MP identities.

Conformal prediction (CP) is recently developed statistical approach used to establish thresholds for similarity metrics that carried a theoretical confidence guarantee (i.e., the probability that the correct match was returned) for returned labels for an unknown spectrum [Clough et al., 2024]. Generally, CP is a statistical method that calculates a threshold based on a calibration dataset which can then be used for unknown datasets [Angelopoulos and Bates, 2023, Shafer and Vovk, 2008]. Clough et al. [2024] utilized CP to either return potential matches for an unknown particle’s spectrum in a prediction set or return an empty prediction set if none of the scores surpassed the data-driven threshold. CP not only outperformed a commonly used commercial database matching software for MPs, but the method could also be optimized and employed for environmental plastic identification. However, the use of CP to identify plastics was limited to one spectral input, despite it being established that the combination of IR and Raman spectroscopy can further enhance the certainty of MP identification [Käppler et al., 2016, Parham et al., 2025, Böke et al., 2022].

Optical photothermal infrared (O-PTIR) spectroscopy systems provide an advantage for particle identification as they simultaneously collect photothermal infrared (PTIR) and Raman spectra for individual particles by utilizing a pump-probe mechanism, where the pump is a tunable scanning mid-IR laser and the probe is a visible laser [Olson et al., 2020, Prater et al., 2024, Zhang et al., 2016]. In addition to collecting two spectra for each particle, this technique circumvents the size resolution of traditional Beer’s Law-based IR measurements, which are practically limited to detecting MP particles  $> \sim 20 \mu\text{m}$  in size [Ivleva, 2021, Dong et al., 2022]. The size resolution for O-PTIR depends upon the probe laser (532 nm), not the pump laser ( $\sim 3\text{--}10 \mu\text{m}$ ), and is advantageous for detecting MP particles down to the sub-micron size range [Parham et al., 2025, Böke et al., 2022,

Duswald et al., 2025, Pięta et al., 2025]. Among the few studies combining IR and Raman spectra, Böke et al. [2022] proposed plotting the HQI scores from PTIR and Raman in a two-dimensional score space. However, it was noted the range HQI scores for the correct label was as low as 0.3. Developing a method that can establish reliable thresholds and return labels with statistical confidence values is critical for identifying MPs with multiple spectral inputs.

To achieve the goal of aggregating multiple methods we used a special case of conformal score aggregation (CSA) [Ochoa Rivera et al., 2025], called multiview conformal prediction (MVCP), in which multiple models are obtained from different views of the same object. This adaptation leverages the spectral inputs from O-PTIR+Raman while using preexisting theoretical properties, such as confidence guarantees, that have already been established with CSA. Herein, we apply MVCP to the case of PTIR and Raman spectra collected from one individual particle. This enables the use of PTIR and Raman to identify MP particles with quantified statistical confidence. PTIR and Raman spectra collected from samples containing both MP and non-MP particle types ( $<10 \mu\text{m}$  in size), were used to calibrate and establish an envelope threshold. The efficiency and robustness of this technique was then compared with traditional, single-view CP methods that only considered either the PTIR or Raman spectra independently. Finally, MVCP was applied to an additional data set of aerosolized MP particles impacted on an ambient particle sample (without MPs) to determine the method’s capability to differentiate MP particles from other non-MP particles in an environmental sample.

### 5.1.1 Spectral Preprocessing

For this study, the raw spectra from the samples and reference library were preprocessed in-house using Python. Note that all spectra shown herein have been preprocessed unless otherwise indicated. The PTIR spectra were automatically calibrated in PTIR Studios, and the spectra were interpolated such that all wavenumbers were aligned between the sample and reference library spectra and could be used for database matching. The spacing of the wavenumbers remained  $2 \text{ cm}^{-1}$  before and after interpolation. Although interpolating the PTIR spectra was not strictly necessary in this study, it was included as part of the open-access preprocessing script in case future users included reference library spectra that did not have the same wavenumber values as their unknown particle spectra. The PTIR spectra were then truncated to the wavenumber ranges of  $981\text{--}1799$  and  $2721\text{--}2965 \text{ cm}^{-1}$  to remove regions of high noise due to laser chip transitions. Finally, the spectra were normalized using the standard normal variate (SNV) transformation [Renner et al., 2019, Barnes et al., 1989].

The Raman spectra had to be calibrated manually, so the spectra were first baseline-corrected using Zhao et al. [2007]. This method was shown to improve fluorescence background removal for Raman spectra with low signal-to-noise ratios compared to the commonly used multi-polynomial fitting method. The spectra were then normalized using the SNV transformation, calibrated using a 520  $\text{cm}^{-1}$  peak from internal silicon standard that was collected daily, and interpolated to ensure the wavenumbers aligned between the reference library spectra and sample spectra. The interpolation additionally modified the spacing of wavenumbers from 3  $\text{cm}^{-1}$  to 2  $\text{cm}^{-1}$ . Note that the Raman spectra were collected over a range of 480–3762  $\text{cm}^{-1}$ , though some ranges were slightly extended and changed the wavenumber values. These variances made interpolation necessary for the Raman spectra. Similar to the PTIR spectra, the Raman spectra were truncated to 601–1849 and 2751–3499  $\text{cm}^{-1}$  to exclude the low variance and low specificity regions [Renner et al., 2019], which could influence automated spectral matching.

### 5.1.2 Database Matching with the Normalized Nearest Neighbor ( $\widetilde{NN}$ ) Similarity Metric

Database matching involves calculating likeness between an unknown spectrum and a collection of reference spectra from a library. Figure 5.1 shows example PTIR and Raman spectra for an unknown particle with the reference library spectra below them. Note that these reference library PTIR and Raman spectra are not inherently correlated with each other—meaning that they were not necessarily retrieved from the same particle (see the reference library section above for more information). Additionally, database matching was conducted independently for the particle’s PTIR and Raman spectra.

HQI scores were calculated using the nearest neighbor (NN) similarity metric [Shafer and Vovk, 2008] (Equation (5.1)), which not only takes into account the similarity between an unknown and reference library spectrum, but also whether the unknown spectrum is more similar to other particle types in the reference library. Additionally, it has been shown that the NN similarity metric is more selective than the commonly used Pearson’s Correlation Coefficient (PCC) metric [Clough et al., 2024].

$$NN_y = \frac{\min_{ref} \{ |x - x_{ref}| : y_{ref} = y \}}{\min_{ref} \{ |x - x_{ref}| : y_{ref} \neq y \}} \quad (5.1)$$

$$\widetilde{NN} = 1 - e^{-\frac{NN_y}{\tau}} \quad (5.2)$$

In the numerator for  $NN_y$ , the minimum distance between the unknown spectrum,  $x$ ,

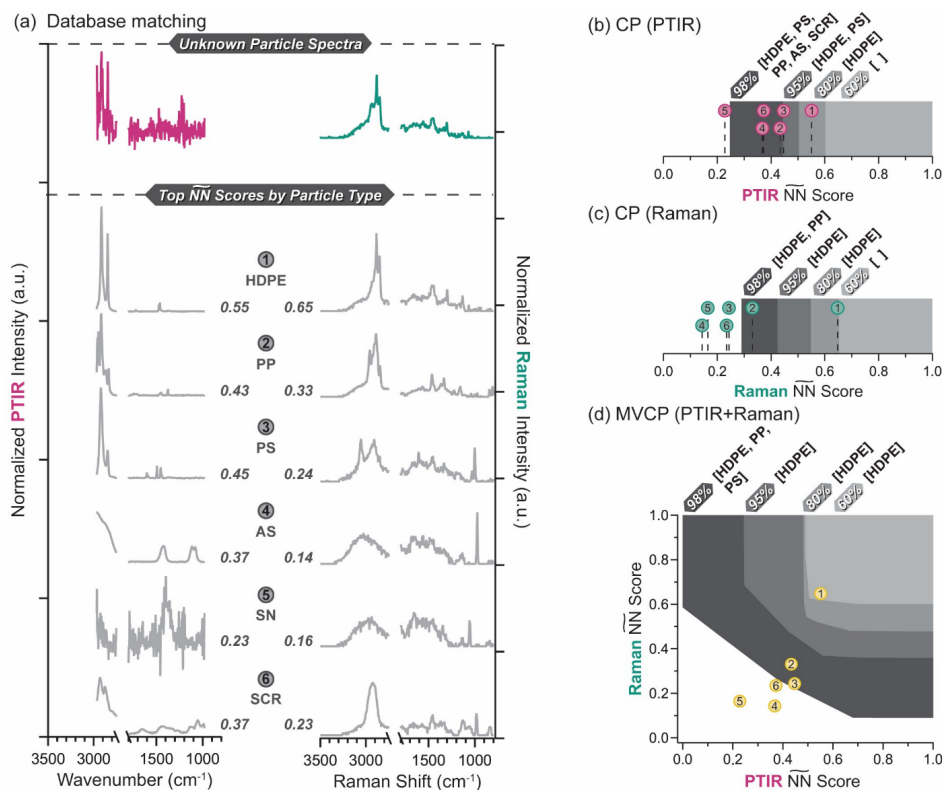


Figure 5.1: Experimental outline for labeling unknown spectra. (a) Database matching for the PTIR and Raman spectra (pink and green traces, respectively) of an unknown particle. The gray traces underneath are the top-matching spectra for each particle type with their respective  $\widetilde{NN}$  scores (i.e., HQI scores) next to them. The resulting single-view CP methods for (b) the PTIR  $\widetilde{NN}$  scores and (c) the Raman  $\widetilde{NN}$  scores are shown with the returned prediction sets based on the methods' theoretical confidence percentages. (d) The results from the MVCP method, utilizing both PTIR and Raman  $\widetilde{NN}$  scores, with the returned prediction sets shown based on the theoretical confidence percentage. Note that each numbered point in (b), (c), and (d) corresponds with the particle types denoted in (a).

and all reference spectra,  $x_{ref}$ , for a specific the label ( $y_{ref} = y$ ) is calculated. This value is then divided by the minimum distance between spectrum  $x$  and every reference spectrum  $x_{ref}$  that does not correspond to the label of interest (i.e.,  $y_{ref} \neq y$ ). For example, if the NN similarity metric was calculated for an unknown spectrum and HDPE was the label of interest, then the Euclidian distance between the unknown spectrum and the closest HDPE reference spectrum would be calculated. This value would then be divided by the Euclidian distance between the unknown spectrum and the closest reference library spectrum of any particle type that is not HDPE. The resulting score,  $NN_{HDPE}$ , would be specific to the HDPE label, and this calculation would be repeated for the other potential labels in the reference library.

Given that the initial NN similarity metric could be any value greater than 0, it can be normalized so that the output value is between 0 and 1 and can be interpreted as an HQI score ((5.2)). This transformation is strictly monotone and is introduced solely to improve interpretability, meaning it will not affect the resulting confidence or prediction sets returned by the CP methods. For this study, an exponential function with a constant  $\tau$  was used to generate this normalized NN similarity metric ( $\widetilde{NN}$ ), which will herein be referred to as an  $\widetilde{NN}$  score. These scores were scaled ( $\tau = 1.4427$ ) such that  $\widetilde{NN} \approx 0.5$  when the first and second nearest neighbors were the same distance from the unknown spectrum.  $\widetilde{NN}$  was more than 0.5 only when the label in question was truly the nearest neighbor to the unknown spectrum, whereas when  $\widetilde{NN}$  was less than 0.5, it signified a different label was the nearest neighbor. Using the example spectra from Figure 5.1, this normalization caused the top  $\widetilde{NN}$  scores for each particle type to be lower than what may be typically expected for an HQI score.

### 5.1.3 Conformal Prediction (CP)

Unlike traditional database matching methods, which seek to assign one label to an unknown spectrum, conformal prediction (CP) returns prediction sets containing however many labels that have scores surpassing a data-determined threshold (i.e., a threshold calculated from a calibration set). The advantage of this CP threshold is that the returned prediction set has a confidence that the correct match is in the prediction set based on the user-defined theoretical confidence percentage. For example, if a user sets a threshold with 80% theoretical confidence, then all prediction sets returned have an 80% guarantee that the correct label (if the correct label is in the reference library) is returned in the prediction set.

An important characteristic of CP is that the confidence guarantees require relatively weak distributional assumptions of the calibration spectra. Whereas other statistical approaches

(such as the bootstrap method [Efron, 1979, Morgado et al., 2021]) give guarantees that are only asymptotically valid—meaning the dataset should be near-infinite in size—CP gives guarantees even when the dataset’s score distribution is not Gaussian in shape and finite. However, it is crucial to note that CP assumes the calibration set used is representative of the sample, meaning the score distributions between the calibration and unknown sample are the same. Thus, if a particle type is not included in the calibration set, then the theoretical confidence guarantees provided by CP do not extend to this type of particle. Figure 5.1 shows the prediction sets returned for the PTIR and Raman spectra’s CP methods based on a range of theoretical confidences. The differences between the calculated CP (PTIR) and CP (Raman) quantiles and  $\widetilde{NN}$  scores mean the prediction sets returned will not always agree.

#### 5.1.4 Multi-View Conformal Prediction (MVCP)

Unlike single-view machine learning methods, multi-view machine learning aims to solve prediction problems when different (often non-overlapping) representations of the same underlying “object” are observed [Sun, 2013]. In this case, our “object” is an unidentified particle and it is being represented via the  $\widetilde{NN}$  scores for its PTIR and Raman spectra. Thus, MVCP is similar to its single-view counterpart in that it utilizes a threshold to determine what particle type labels are returned in the prediction sets based on a user-defined theoretical confidence value [Ochoa Rivera et al., 2025]. However, this boundary is no longer a single number; instead, it is two-dimensional in a score space. The score space encompassed by the boundary, called an envelope, returns labels whose combined scores from multiple views (treated as x and y coordinates) are within it. As shown in Figure 5.1, the example particle’s  $\widetilde{NN}$  scores are now plotted in a two-dimensional score space, and the envelopes cover a wider area of the score space as their theoretical confidence increases. Correspondingly, more potential matches are returned as the envelope increases in size. Although MVCP is used with two views in this work, it is important to note that it could be leveraged for a collection of  $\underline{n}$  views—in which case the threshold would be  $\underline{n}$ -dimensional [Ochoa Rivera et al., 2025].

The challenge with creating an envelope is that there is no canonical ordering of points in higher-dimensional score spaces. For example, if given the  $\widetilde{NN}$  scores for the PTIR spectra of 50 particles, a user could easily order these particles from having the highest to lowest scores. However, there is no clear way to order these particles when also considering their Raman  $\widetilde{NN}$  scores. Thus, the envelope must be constructed using directional quantiles [Kong and Mizera, 2012], meaning the distributional information of the calibration points

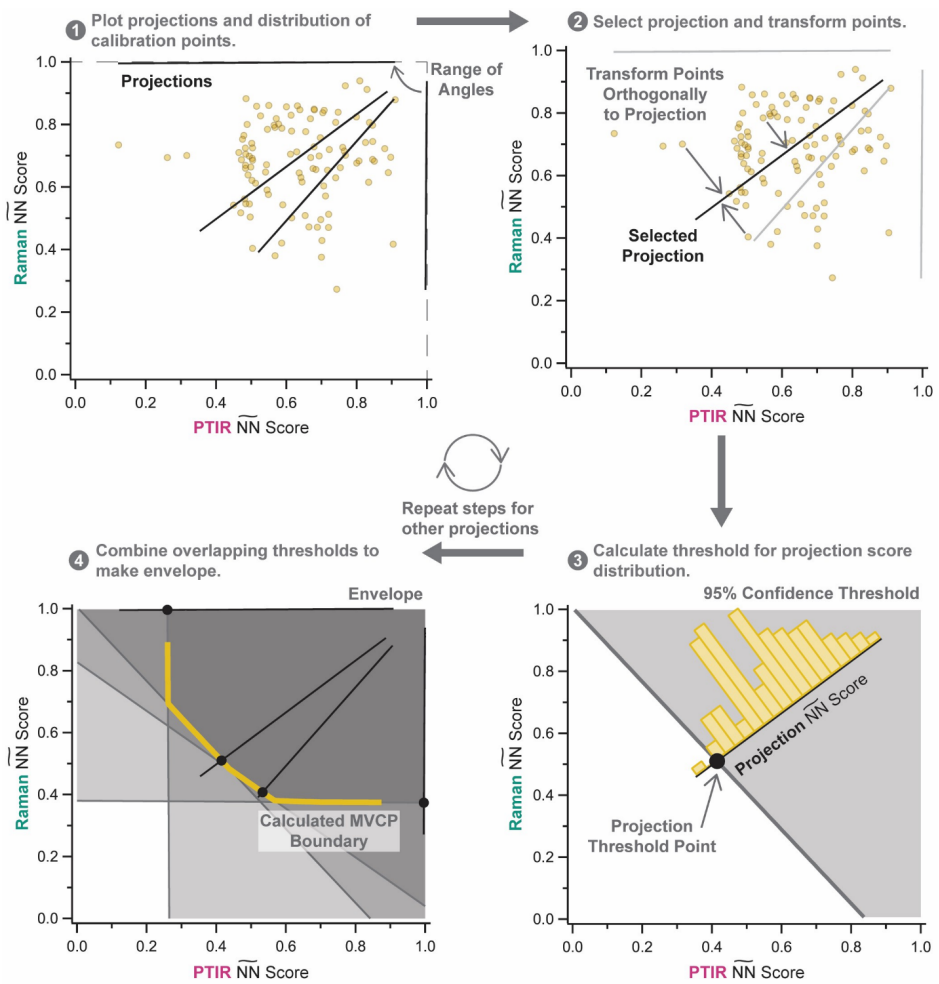


Figure 5.2: Method for generating MVCP envelope and boundary.

are assessed over several directions in the score space, to address the multiple views. An example of this process to generate the shape of an envelope is shown in Figure 5.2. A range of angles originating from the maximum point of the score space (PTIR  $\widetilde{NN} = 1.0$  and Raman  $\widetilde{NN} = 1.0$ ) are evaluated. Each angle is assessed by projecting the calibration data points orthogonally onto a line such that the points are represented by only one dimension, and this projection’s score distribution is used to calculate the threshold at the user’s desired theoretical coverage. By projecting the data points back to a one-dimensional space across different directions (i.e., angles), the weights of the multiple views’ scores on the overall distribution also change. Meaning a projection along the x-axis or y-axis would be 100% weighted towards the PTIR or Raman  $\widetilde{NN}$  scores, respectively, and a projection at a 45° angle between these two axes would be equally weighted between the PTIR and Raman  $\widetilde{NN}$  scores. For each envelope calibration, 400 projection directions were evaluated. The threshold at the user’s selected theoretical confidence is then calculated based on this new distribution of scores. Once the individual threshold windows are calculated at each angle, their overlapping areas are used to create the two-dimensional envelope.

Note that the size of the envelope may need to be adjusted since the data points had to transform from their original positions to make the directional quantile, and the actual coverage of the envelope may be different for the distribution of the calibration points in the two-dimensional score space. Thus, only 24% of the calibration data determines the shape of the envelope while the remaining 76% of the calibration data adjusts the size of the envelope based on the calculated empirical coverage of these data points. A sensitivity analysis for this split for the calibration subsets, along with the number of projection directions used to determine the envelope during calibration, is shown in. Parameters outside of these can be adjusted in the Python script as well, which is available for open access on GitHub<sup>2</sup>. A full list of user-selected parameters for the spectral preprocessing pipeline, database matching, and CP methods are provided in the accompanying user guide.

### 5.1.5 Calibration and Test Datasets for Single-View CP and MVCP Experiments

In this study, calibration and test spectral datasets were generated with the combined manually-labeled particle spectra from the MPs-Only and MPs+Standards samples. To investigate the performance of the single-view CP and MVCP methods, these spectra were filtered based on how challenging the unknown particles were to manually identify based on their respective PTIR and Raman spectra. Spectral Dataset #1 (n = 321) included par-

---

<sup>2</sup><https://github.com/eochoarv/MVCP>

ticles with the most straightforward manual identification: the PTIR and Raman spectra both had an MP or non-MP label and were not mixtures (i.e., neither spectrum was labeled MIX, UNS, or LOW). Additionally, the Raman spectra did not exhibit fluorescence or saturation. Spectral Dataset #2 ( $n = 209$ ) included particles where one spectrum had an MP or non-MP label and the other did not (i.e., this other spectrum was labeled UNS or LOW). Mixtures were not included in this dataset, and Raman spectra that exhibited saturation or fluorescence were included regardless of whether they had an MP or non-MP label.

As a more rigorous test of robustness, a separate dataset was generated by augmenting some of the spectra from Spectral Dataset #1 with noise. For this test, Gaussian noise was added to investigate the effect of spectral noise on the single-view CP and MVCP methods and push their limitations for particle identification. Spectral Dataset #1 was thus split into three subsets: one subset contained half of the particles from Spectral Dataset #1 ( $n = 161$ ) and was left unchanged, one subset contained a quarter of the particles ( $n = 80$ ) with their PTIR spectra augmented with Gaussian noise, and the last subset contained a quarter of the particles ( $n = 80$ ) with their Raman spectra augmented with Gaussian noise. For clarity, the first subset will be referred to herein as Spectral Dataset #1: Unchanged Subset, and the other two subsets will be jointly referred to as Spectral Dataset #1: Noise-Added Subsets.

All datasets or subsets used for an experiment were evaluate using Monte Carlo cross validation [Molinaro et al., 2005, Stone, 1974], which is a method that assesses prediction models by randomly splitting a dataset between calibration and test sets  $n$ -times and using the results for quantitative comparisons. In this study, the dataset used for each experiment was split 84% and 16% between calibration and test data, respectively. These fractions for the split favored the calibration set to produce more accurate estimations of the envelope’s performance during calibration phase. Although the coverage guarantee holds regardless of the calibration dataset’s size, there is an impact on the empirical confidence observed.[Angelopoulos and Bates, 2023] Larger calibration datasets lead to higher precision (i.e., empirical confidence has less error associated with it). In practice, users should include as much data as is reasonable in the calibration set. The calibration data was first run through the database matching algorithm. The top score of the correct polymer type for each PTIR and Raman spectrum was plotted in a distribution for their respective measurements, and quantiles were then calculated based on their respective distributions. The database matching process was then repeated with the test spectra, but this time the scores were plotted along the axis of their corresponding method and used to determine which potential matches were returned in the prediction sets based on the quantiles established. This process was repeated 100 times for each experiment.

## 5.2 Results and Discussion

### 5.2.1 Comparing Thresholds between Single-View CP and MVCP Methods

One of the ways that MVCP utilizes dual information from O-PTIR and Raman is by creating an envelope with both spectral inputs. Figure 5.3 shows the histograms of these top  $\widetilde{NN}$  values for their respective single-view CP methods for Spectral Dataset #1 and Spectral Dataset #2. For this work, the histograms are shown as a visual representation of the one-dimensional spread of data, which is why the thresholds are plotted irrespective of the bin edges. The 95% quantiles for CP (PTIR) and CP (Raman), shown as dashed lines, were nearly identical values of  $\widetilde{NN} = 0.44$  and  $0.42$ , respectively. Although these thresholds are low relative to other thresholds previously used [Clough et al., 2024, Weisser et al., 2022], they are a result of the  $\widetilde{NN}$  scores being scaled so the data points could be interpreted based on which quadrant of the score space they resided in.

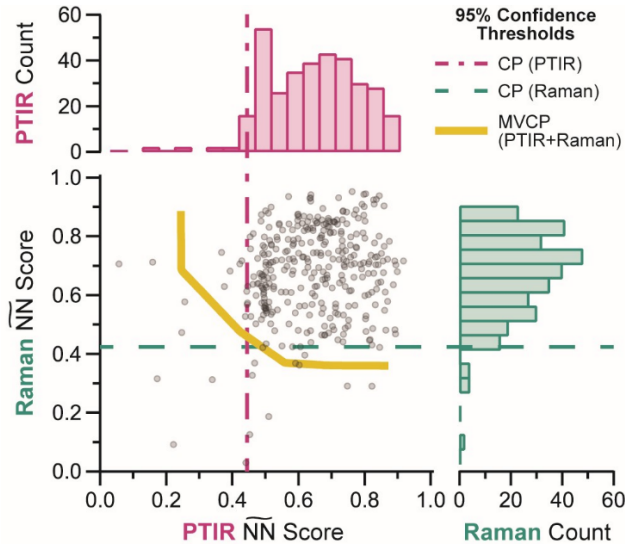


Figure 5.3: An example calibration with resulting thresholds for the three CP methods. The histograms for the top HQI scores of the correct particle type are shown for the PTIR and Raman CP methods (pink and green, respectively). Their corresponding quantiles at 95% theoretical confidence are denoted with dashed lines that intercept with the PTIR or Raman HQI axes. The spread of these PTIR and Raman scores for each particle are plotted underneath the histograms, with the MVCP envelope at 95% theoretical confidence shown in gold.

The MVCP envelope at 95% theoretical confidence extends past the overlapping region from the single-view CP thresholds. However, the convex shape of the envelope covers a

different area of the score space than the rectangular shape of the single-view CP quantiles, which may affect what potential matches are returned for unknown spectra. For example, the left-hand side of the envelope exceeds the CP (PTIR) threshold, but it does not cover a large fraction of the score space within the CP (Raman) threshold. The shared information between the two spectroscopy techniques creates a unique threshold that interacts with labels differently than when considering one of the single-view CP quantiles or the region where the two quantiles overlap.

### 5.2.2 Performance Evaluation for Single-View CP and MVCP Methods

One trade-off of CP methods is the number of potential matches returned in a prediction set increases as the confidence (and thus the probability of the prediction set containing the correct label) increases, which would require manual analysis to decide between the returned labels. Therefore, it is critical to evaluate a CP method’s efficiency over a range of theoretical confidence values to choose a threshold that does not return too many labels, especially when using reference libraries with tens or hundreds of available labels that could be returned. This evaluation is shown in Figure 5.4, where the average prediction set size is used as a metric of a CP method’s efficiency using Spectral Dataset #1 and Spectral Dataset #2. The MVCP and CP (Raman) methods dramatically increase from  $\sim 1$  to  $\sim 4$  labels after 95% theoretical confidence. CP (PTIR)’s average set size is already close to 2 labels at 95% theoretical confidence and reaches as high as 5 labels (out of 6 labels total in the reference library) at 100% theoretical confidence. The 95% theoretical confidence value was thus selected to establish thresholds for all experiments in this work (excluding the assessment of the MPs+Ambient sample) as it maximizes the theoretical confidence guarantee range while keeping the average set size close to one label, thus minimizing the amount of uncertainty (i.e., the need for manual verification of potential matches) associated with the prediction sets returned.

Further comparison shows that MVCP has a higher efficiency than either of the single-view CP methods in most cases (Figure 5.4). MVCP overall has the lowest average prediction set size of  $1.1 \pm 0.1$  labels at 95% theoretical confidence, though its range overlaps with CP (Raman)’s average prediction set size ( $1.3 \pm 0.1$  labels). The average prediction set size for CP (PTIR), which is  $1.8 \pm 0.2$  labels, is significantly higher than MVCP. Despite these differences, the majority of prediction sets on average had a set size of 1 label for each method at 95% theoretical coverage.

These trends between the three CP methods (i.e., CP (PTIR), CP (Raman), and MVCP)

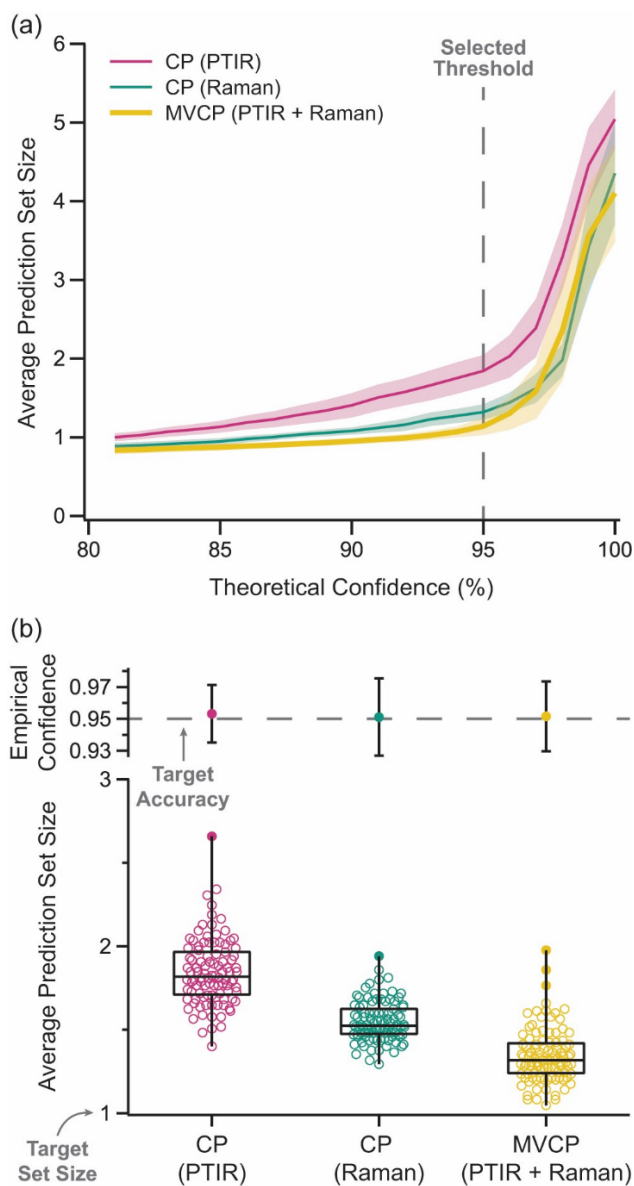


Figure 5.4: (a) The averaged prediction set sizes compared with the inputted theoretical confidence value for CP (PTIR), CP (Raman), and MVCP. (b) The calculated empirical confidence (top panel) and box-and-whisker plots of average prediction set size (bottom panel) for each split at 95% theoretical confidence ( $n = 100$ ). Error bars shown for the empirical confidence represent the standard deviation. Filled points in the box plots signify outliers. The error bars shown for the empirical confidence represent standard deviation.

are even more pronounced in Figure 5.4, where the spread of the average prediction set sizes over 100 splits at 95% theoretical confidence is shown with box-and-whisker plots. MVCP still maintains the lowest median among the three methods (CP (PTIR) median: 1.82 labels; CP (Raman) median: 1.52 labels; MVCP median: 1.32 labels). Although CP (Raman) has a smaller range overall than MVCP (CP (Raman) range: 0.65 labels, MVCP range: 1.93 labels), it is largely due to MVCP’s lowest quantile extending to  $\sim 1$  label—the target prediction set size in this case—whereas the lowest average prediction set size for CP (Raman) is limited to 1.29 labels. CP(PTIR) has the highest median and widest range of labels (median: 1.82 labels, range: 1.26 labels). Although Figure 5.4 shows trends in the average prediction set sizes across 100 splits, the median and range for each prediction set returned from a selected split can also be observed. At 95% theoretical confidence, the largest prediction set returned for MVCP contained 2 labels, whereas CP (PTIR) and CP (Raman) reached set sizes of 6 and 5 labels, respectively. This closer look at returned prediction sets shows that CP (Raman) and CP (PTIR) reach larger set sizes at lower theoretical confidence values, signifying MVCP’s continuous efficiency over a broad dataset.

The three methods are still accurate and precise despite the differences in their prediction set sizes, as shown by the calculated average empirical confidence in Figure 5.4. Empirical confidence is a metric used for CP methods [Angelopoulos and Bates, 2023] to calculate the fraction of prediction sets containing the correct label (which is the manually-assigned label) during testing. It should be noted that empirical confidence is not affected by set size; rather, it only evaluates if the correct label was returned for a test spectrum. All three methods have 95% empirical confidence with a standard deviation of 2–3%, indicating that MVCP’s higher efficiency does not come at the cost of accuracy.

The benefits of having two spectral inputs to distinguish between particle types can also be observed at the single-particle level. Utilizing the complimentary nature of O-PTIR and Raman spectra has been shown to improve manual identification of pristine MP particles, especially in size ranges  $< 10 \mu\text{m}$  where particles often have lower signal [Parham et al., 2025], and this improvement extends to computational analysis of unknown particles as well. Figure 5.5 shows PTIR and Raman spectra from two example particles in Spectral Dataset #2 (i.e., the dataset of particles where one of their spectra was not easily identifiable). In Figure 5.5, an unknown particle has a Raman spectrum with clear peaks that visually agree with a reference library spectrum for HDPE. The Raman spectrum’s  $\widetilde{NN}$  score of 0.56 for HDPE corroborates this likeness, and the score is high enough to indicate that HDPE has closer agreement with the unknown spectrum than neighboring labels. Indeed, Figure 5.5 shows HDPE is the only label to pass the CP (Raman) threshold. The PTIR spectrum in Figure 5.5, on the other hand, is clearly only noise. Its corresponding  $\widetilde{NN}$  score for HDPE

( $\widetilde{NN} = 0.48$ ) is close to 0.5, indicating that it is about the same Euclidean distance from the unknown spectrum as another label in the reference library. In fact, all PTIR  $\widetilde{NN}$  scores for the six labels are between 0.47 and 0.51, and the resulting prediction set contains all six labels. In this case, the MVCP envelope only passes HDPE and excludes all other labels due to the added information from the Raman spectra.

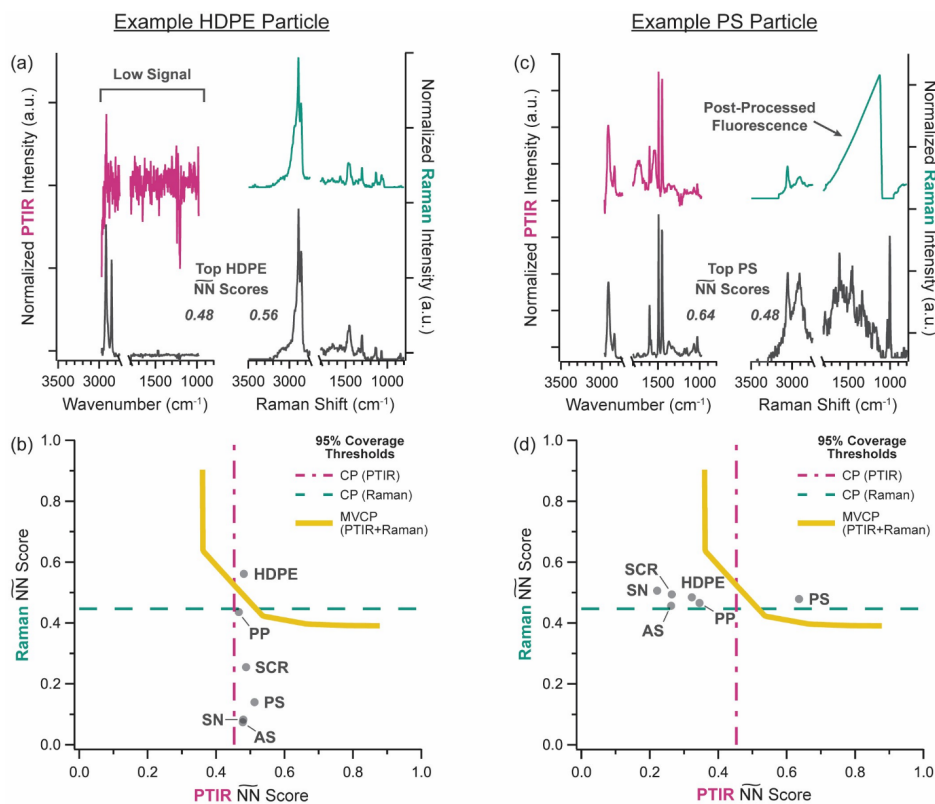


Figure 5.5: (a) PTIR and Raman spectra for an example HDPE particle and the top-matching reference library HDPE spectra in dark grey below. (b) Top HQI scores for each particle type with the example particle’s spectra compared with the CP method’s thresholds. A similar structure was used for an example PS particle’s (c) spectra and (d) HQI scores.

Raman spectra can also suffer from spectral quality issues, especially when peaks are overwhelmed by fluorescence or saturation. Figure 5.5 shows an example PS particle with part of its Raman spectrum saturated. The signal in the fingerprint region has a sharp increase and a gradual decline, which was observed after the saturated signal underwent baseline correction. Similar to the previous example, all six labels pass the CP (Raman) threshold in Figure 5.5. The PTIR spectrum for the PS particle has clear agreement with the PS reference spectrum, and this agreement is further supported by the PS label’s relatively high  $\widetilde{NN}$  score of 0.64. Now the CP (PTIR) and MVCP methods only return the PS label. These example cases clearly exhibit MVCP’s use of a two-dimensional score space

to return correct matches for unknown particles even when one of the spectra cannot be identified based on a similarity metric—thus reducing the uncertainty of the method by returning smaller prediction sets. Moreover, utilizing two complementary spectral inputs further ensures the identification of MP particles that could otherwise be overlooked if only one of the spectroscopy methods were used.

### 5.2.3 Robustness Assessment based on Spectral Identification Capability

The robustness of MVCP is further demonstrated by the quantitative comparisons between spectra based on how readily they can be identified. The earlier results from the combined Spectral Dataset #1 and Spectral Dataset #2 (see **Figure 4b**) are now shown separated by dataset in **Figure 6a**. For the particles in Spectral Dataset #1, the spread of average prediction set sizes for the single-view CP and MVCP methods have a median near 1 label with narrow ranges. On the other hand, the particles that were harder to manually identify using both the PTIR and Raman spectra (Spectral Dataset #2) subsequently have much higher median set sizes—especially for CP (PTIR) (CP (PTIR) median: 3.0 labels, CP (Raman) median: 1.5 labels, MVCP median: 1.2 labels). The empirical confidence shows that all methods for both types of spectra are statistically within bounds for 95% accuracy, which is the target as 95% theoretical coverage was still used to establish the thresholds.

The stark difference between CP (PTIR) and the other two methods in the results for Spectral Database #2 is likely due to the number of particles with PTIR spectra labeled as LOW and Raman spectra labeled as an MP or non-MP ( $n = 124$ ). These particles could lead to an increase in the uncertainty of the CP (PTIR) method during calibration and testing, causing a slightly lower threshold (**Figure 3**) and higher average prediction set sizes (Figure 5.4 and Figure 5.6). The influence of these unidentifiable spectra on the calibration and performance of the CP (PTIR) method is thus observed through the increase in uncertainty of the prediction sets returned via the higher number of potential matches included in the set.

The limits of the single-view CP and MVCP methods were investigated further by augmenting subsets of spectra in Spectral Dataset #1 with Gaussian noise. The results for the CP methods when calibrated and tested with Spectral Dataset #1: Unchanged Subset and Spectral Dataset #1: Noise-Added Subsets are shown in Figure 5.6. As expected, the addition of Gaussian noise to the augmented subsets increased the overall root-mean-square (rms) noise of the spectra. The unchanged subset has similar metrics to Spectral Dataset #1 shown in Figure 5.6, though this unchanged subset has an even narrower range of aver-

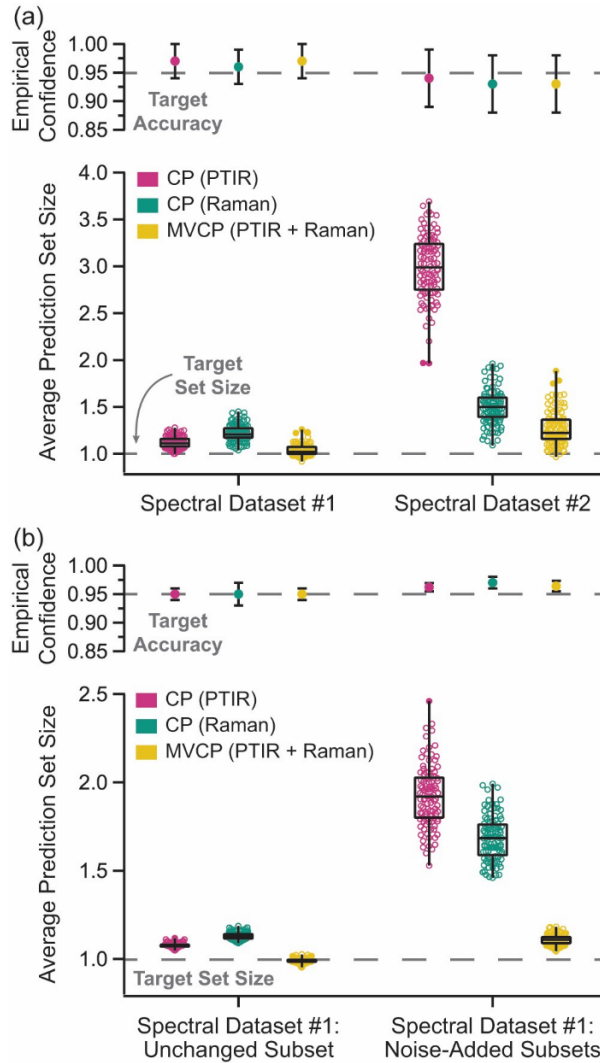


Figure 5.6: Empirical confidence (top panel) and spread of average prediction set sizes over 100 splits (bottom panel) for (a) Spectral Dataset #1 ( $n = 321$ ) and Spectral Dataset #2 ( $n = 209$ ) as well as (b) Spectral Dataset #1: Unchanged Subset ( $n = 161$ ) and Spectral Dataset #1: Noise-Added Subsets ( $n = 160$ ). Error bars shown for the empirical confidence represent standard deviation. Filled points in the box plots signify outliers.

age prediction set sizes with the widest range being 0.1 labels. Meanwhile, the noise-added subset can be seen to greatly affect the single-view CP methods. Indeed, the median for CP (PTIR) increased to 1.92 labels (range: 1.53–2.46 labels), and CP (Raman) also increased to a median of 1.68 labels (range: 1.46–1.99 labels). Only MVCP remained relatively unaffected as its median stayed near 1 and its range was still relatively small (1.04–1.18 labels)—thus maintaining a lower uncertainty in its returned prediction sets. All three methods surprisingly show a slight increase in empirical confidence (CP (PTIR):  $0.962 \pm 0.007$ , CP (Raman):  $0.97 \pm 0.01$ , MVCP:  $0.964 \pm 0.009$ ), which may be due to the smaller number of data points in the subsets when compared to Spectral Dataset #1 and Spectral Dataset #2. The adaptability of MVCP to these noisier spectra demonstrates the advantage of multiple spectra when working with a variety of particles and pushing the limits of MP detection and identification.

#### 5.2.4 MVCP Method Validation with MPs+Ambient Sample

To probe the capabilities of MVCP for environmental MP studies, the method was applied to a controlled experiment where HDPE, PP, and PS particles were impacted on top of an ambient aerosol particle sample on a gridded silicon substrate. A selected grid cell was analyzed after the environmental sample was first collected (no MP particles were identified from the ambient particle spectra), and the same grid cell was reanalyzed after MP impaction [Parham et al., 2025]. The particles from the MPs+Ambient sample were then used as a test dataset to determine the most appropriate threshold for MVCP.

The results from this experiment were treated as a binary classification so the data analysis routine could be assessed using common metrics for models. In particular, only the three MP types were used as potential labels for database matching, and only MP particles from Spectral Dataset #1 and Spectral Dataset #2 were used for the envelope calibration. The calibration data was randomly split 100 times between the two calibration subsets, to account for differences in the envelope’s shape and size, and the averaged metrics were calculated. Test particles from the MPs+Ambient sample were considered correctly identified if 1) one or more labels were returned for MP particles and 2) no labels were returned for ambient or unidentified particles (i.e., ambient, no identifiable peaks, or unsorted). Figure 5.7 shows performance metrics for this binary classification as a function of theoretical confidence. True positive rate (TPR) was defined as the fraction of prediction sets that contained a label for an MP particle. Conversely, the fraction of prediction sets returned with an MP label for an ambient, no identifiable peaks, or unsorted particle was defined as the false positive rate (FPR). For simplicity, particle mixtures were not considered for the TPR and FPR. These

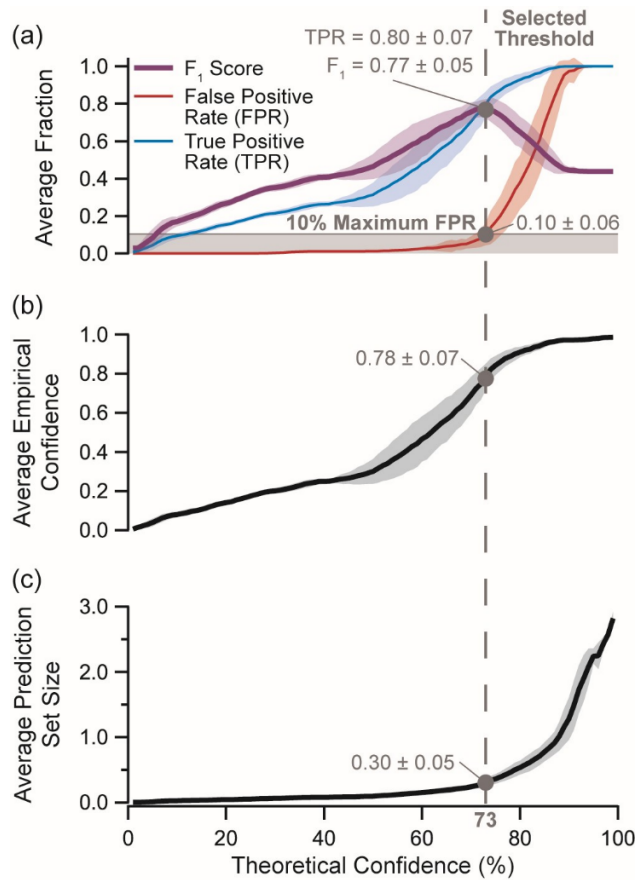


Figure 5.7: (a) The false positive rate (FPR), true positive rate (TPR) and  $F_1$  score compared over theoretical confidence. (b) Empirical confidence and (c) prediction set size as a function of theoretical confidence. A dotted line denotes the threshold at 73% theoretical confidence, which was selected by the criterion of a 10% maximum FPR. The values of each metric are shown in light gray. The shaded area around traces represents standard deviation.

two rates were then used to calculate an  $F_1$  score, which is frequently used as an accuracy metric for a model.

The priority for this study was to minimize the number of false positives to reduce prediction set size, so the threshold was established at 73% theoretical confidence to ensure the average FPR did not exceed 0.10 ( $FPR = 0.10 \pm 0.06$ ). This threshold also coincides with the peak in the  $F_1$  score ( $0.76 \pm 0.05$ )—which indicates the number of true positives is greater than the sum of false positives and false negatives [Liemohn, 2024]. The TPR is also fairly high ( $0.80 \pm 0.07$ ), which is further corroborated by the selectivity of the PP and PS particles (**Figure S7**). The HDPE test particles had more HDPE labels outside of the envelope while other MP labels were within the envelope, which contributes to lowering the TPR fraction. These metrics could be considered satisfactory by general standards for models, but they indicate lower performance compared to other data analysis routines employed for MP identification [Morgado et al., 2021, Smolen et al., 2025].

Although these results show the potential of using MVCP in environmental analysis, more can be done to improve the overall performance of the data analysis routine. Fortunately, this routine can be assessed with tools from the MVCP method. The performance of the selected envelope itself should be evaluated first using the average empirical confidence and prediction set size in **Figure 7b** and **Figure 7c**, respectively. The empirical confidence, which only considers if the correct MP label is returned, is higher on average than the target confidence for this threshold ( $0.78 \pm 0.07$ ). In terms of efficiency, the prediction sets for the MP particles should ideally return 1 label and all other particle types should not have a label returned, so the target prediction set size should reflect the fraction of MP particles ( $n = 71$ ) within all particles analyzed ( $n = 253$ ), which was  $\sim 0.28$ . The average prediction set size at 73% theoretical confidence is  $0.30 \pm 0.05$ , which is in good agreement with the target set size. It should be noted that this fraction of MPs is much higher than what would be expected in an ambient atmospheric sample. A recent modeling study from Evangelou et al. [2026] compiled results from 76 studies and computed an average global MP concentration of 31 particles  $m^{-3}$ , which is several orders of magnitude lower than typical ambient aerosol number concentrations [Seinfeld and Pandis, 2006]. These metrics combined signify the capability of MVCP to not only identify MPs, but also its capacity to differentiate MPs from ambient particles.

Given that the MVCP envelope is performing well, the score space and envelopes used for MVCP can be employed to evaluate database matching with the ambient and unidentified particles. The  $\widetilde{NN}$  scores of the MP labels for ambient particles (AMB) or particles with no peaks in their spectra (LOW) are clustered in the center of the score space, indicating the database matching method employed did not have the correct label to identify the particles.

In fact, one of the primary reasons why an envelope with higher theoretical confidence (e.g., 95% theoretical confidence) was not selected was due to MP labels being returned for a majority of the AMB and LOW particles. More MP particles, and even some particle mixtures, could be captured, but it would also result in more misidentifications of ambient or unidentified particles and thus increasing analysis time during manual verification of the returned prediction sets. Utilizing the MVCP envelope to contextualize results from database matching with multiple spectral inputs is advantageous for determining areas of improvement.

Although non-MP particle types were deliberately excluded from the reference library to employ a binary classification, restricting the reference library and calibration set to only pristine MPs severely limits the applicability of the envelope and confidence guarantee to other samples beyond the one evaluated in this work. A confidence guarantee assumes the unknown particles will have the same distribution of scores as those used in the calibration set (i.e., the same particle types are in the unknown and calibration sets), meaning the envelope used for the MPs+Ambient sample would only be valid if a future user intended to identify pristine HDPE, PS, or PP particles in their sample. The reference library would also need to include any particle types that are expected in an unknown sample. Studies have shown that expanding reference libraries to contain more particle types encountered in an environmental sample, as well as natural polymeric species and aged plastic particles, improves database matching and HQI scores [De Frond et al., 2021a, Munno et al., 2020, Primpke et al., 2018, Kozloski et al., 2024]. In other words, the reference libraries and calibration sets employed in this study would be deficient on their own for calibrating a CP threshold to identify MPs in environmental samples, but they could be used to augment existing reference libraries and calibration sets. Guidelines are available for determining calibration set size [Angelopoulos and Bates, 2023] and it is recommended to include particle types as expected in a sample.

Methods beyond traditional database matching could further improve the identification of these chemically complex particles. Statistical [Renner et al., 2017] and deep-learning [Smolen et al., 2025, Lim et al., 2025] methods have shown potential for identifying MPs non-pristine conditions (e.g., weathering and particle mixtures). MVCP can adapt to these improvements to data analysis routines, making the amenable nature of MVCP a powerful tool for quantitatively evaluating the overall performance of MP identification methods and enhancing them.

### 5.3 Conclusions

This study demonstrated the potential for MVCP to identify MPs by assessing HQI scores from PTIR and Raman spectra with statistical confidence. By treating the  $\widetilde{NN}$  scores from individual particles as coordinates in a score space, the information between these two inputs was shared to calculate multidimensional regions that generalize the thresholds used in the one-dimensional case. This envelope allowed MVCP to maintain accuracy while returning fewer labels than its single-view counterparts (CP), as well as be less susceptible to particle misidentification when one of the spectra was difficult to identify or could not be identified. Although this method is designed to work with any combination of analysis methods employed with database matching, it is implied that these chemical inputs are collected from the same particle. The mIRage instrument that provides O-PTIR+Raman has a distinct advantage as it simultaneously collects PTIR and Raman spectra for individual particles, thus enhancing the simplicity and throughput of multi-technique analysis. Directly comparing the two spectral inputs with these CP methods quantitatively demonstrates the benefits of using both types of spectra for MP identification.

MVCP was also shown to differentiate between MP and environmental non-MP particles, and metrics for binary classification models combined with CP performance metrics to determine the best-performing threshold. Minimizing misidentification of ambient or unidentified particles as MPs was emphasized in this study, but this did result in a few MP particles not being identified. Although the overall performance of MVCP was positive, we recommend incorporating additional particle types encountered in environmental samples (e.g., natural inorganic and organic materials, natural polymers, oxidized MPs, particle mixtures, etc.) before implementing it broadly across field datasets. Improvements to data analysis routines, such as the spectral preprocessing pipeline and the method used for determining a scores of likeness, could be further assessed using sensitivity studies as well as the statistical confidence provided by the MVCP or single-view CP methods. Such advancements paired with reporting statistical confidence for targeted analysis of MPs will ensure confidence in MP identification and contextualize results between MP studies.

## CHAPTER 6

# Online Conformal Prediction: Enforcing monotonicity via Online Optimization

### 6.1 Introduction

Conformal prediction has emerged as a powerful tool for constructing prediction sets with distribution-free coverage guarantees [Papadopoulos et al., 2002, Vovk et al., 2005, Shafer and Vovk, 2008]. Given a desired coverage level  $1 - \alpha$ , conformal methods output a prediction set that contains the true outcome with probability at least  $1 - \alpha$  under minimal assumptions. These guarantees have made conformal prediction attractive in safety-critical and risk-sensitive applications [Lu et al., 2022, 2023, Kuchibhotla and Berk, 2023, Lindemann et al., 2023, Clough et al., 2024].

Many modern applications, however, operate in an online or streaming regime, where data arrives sequentially and predictions must be produced in real time [Chernozhukov et al., 2018, Gibbs and Candes, 2021, Zaffran et al., 2022]. This has motivated the development of online conformal prediction methods, which aim to maintain validity while update predictions as new data is observed. In contrast to the batch setting, where nestedness of prediction sets naturally arises for standard nonconformity scores, this property is generally lost in the online setting. To the best of our knowledge, existing online conformal prediction methods operate at a single coverage level and must be rerun independently for different values of  $\alpha$ , resulting in prediction sets that are not necessarily nested. This limitation restricts their applicability in settings where uncertainty must be assessed across multiple risk levels simultaneously.

In practice, decision-makers often require prediction sets at multiple confidence levels simultaneously, reflecting heterogeneous risk preferences across users or decision contexts. For example, in hurricane forecasting, emergency planners may require different levels of conservativeness depending on location and evacuation costs [Regnier, 2008]. Similarly, in macroeconomic forecasting, policymakers rely on predictive distributions that reflect varying degrees of uncertainty, corresponding to different implicit risk levels [Dowd, 2007]. In such

cases, it is natural to require that prediction sets are nested, so that higher coverage levels correspond to larger sets. However, direct application of current online conformal prediction to multiple  $\alpha$  values do not guarantee this property.

Beyond interpretability, there is also a strong statistical motivation for jointly estimating multiple coverage levels. In the quantile regression literature, estimating quantiles independently can lead to suboptimal performance and violations of monotonicity (quantile crossing). Joint estimation methods that enforce non-crossing constraints produce coherent quantile functions and improve statistical efficiency by sharing information across quantiles [Zou and Yuan, 2008, Liu and Wu, 2011]. This suggests that, in the online conformal setting, jointly update multiple quantiles may similarly improve stability and efficiency, compared to running independently for each coverage level.

This paper addresses these challenges by proposing two novel online conformal prediction methods that produce nested prediction sets across a range of coverage levels while preserving long-run coverage guarantees.

**Contributions.** The main contributions of this paper are as follows:

- We introduce a novel formulation of online conformal prediction that produces nested prediction sets across multiple coverage levels, enabling calibrated uncertainty quantification over the entire risk spectrum.
- We propose two online algorithms based on convex optimization algorithms that jointly estimate multiple conformal thresholds while guaranteeing monotonicity: exponentiated gradient and projected gradient.
- We establish theoretical guarantees for long-run coverage across all coverage levels, together with guarantees ensuring nested prediction sets at every time step.
- We demonstrate empirically, on both synthetic and real-world datasets, that our methods achieve improved calibration, efficiency, and stability compared to independent single-level conformal approaches.

## 6.2 Related Work

Conformal prediction has traditionally relied on exchangeability assumptions to obtain finite-sample coverage guarantees [Vovk et al., 2005, Shafer and Vovk, 2008]. Under exchangeability, prediction sets can be constructed using empirical quantiles of past nonconformity scores,

and nestedness across coverage levels follows naturally when using an appropriate nonconformity score function. However, these guarantees break down in online settings where the data stream may exhibit temporal dependence, distribution shift, or adversarial behavior.

To address this limitation, recent work has extended conformal prediction to the online setting without relying on exchangeability. A key line of work, initiated by Gibbs and Candès [2021], proposes update the conformal threshold  $q_t$  via online gradient descent with a fixed step size known as Adaptive Conformal Inference (ACI). This approach provides long-run coverage guarantees under arbitrary (potentially adversarial) data sequences. Subsequent works have proposed refinements to improve practical performance, particularly in terms of prediction set size and adaptivity to distributional changes [Bhatnagar et al., 2023, Angelopoulos and Bates, 2023, Gibbs and Candès, 2024, Areces et al., 2025, Sun and Yu, 2025].

More recent developments have explored the use of time-varying or adaptive step sizes to better handle non-stationarity and to bridge adversarial and stochastic regimes. These methods often draw on tools from online learning, such as adaptive regret minimization and expert aggregation, and aim to provide stronger notions of validity beyond long-run coverage [Zaffran et al., 2022, Bastani et al., 2022]. Notably, Angelopoulos et al. [2024] and related works establish guarantees that hold simultaneously in adversarial and i.i.d. settings, achieving a “best-of-both-worlds” behavior.

Parallel to this line of work, several approaches have studied conformal prediction under relaxed assumptions on the data-generating process, such as weak dependence or time-series structure [Chernozhukov et al., 2018, Barber et al., 2023]. While these methods can provide valid coverage under specific assumptions, they generally do not offer distribution-free guarantees in fully adversarial settings.

Despite this progress, existing online conformal prediction methods are designed to operate at a single coverage level. When multiple coverage levels are required, standard approaches must be run independently for each  $\alpha$ , which does not guarantee nested prediction sets and can lead to inefficiencies. In addition, the quantile regression literature has shown that joint estimation of multiple quantiles, with constraints to prevent crossing, improves both coherence and statistical efficiency [Zou and Yuan, 2008, Liu and Wu, 2011].

Our work builds on the online optimization perspective of conformal prediction and extends it to the joint estimation of multiple coverage levels. By enforcing structural constraints across quantiles, we obtain nested prediction sets and calibrated uncertainty estimates over the full risk spectrum, a setting that has not been addressed in prior online conformal prediction work.

### 6.3 Setting

Consider a bounded score function  $s_t : \mathcal{X} \times \mathcal{Y} \rightarrow [0, B]$  at each time  $t$  and a stream of data  $\{(X_t, Y_t)\}_{t=1}^T$ . The score functions come from a pre-trained or trained online model (e.g.  $\hat{f}_t$ ) that predicts  $y$  using  $x$ . The function  $s_t(x, y)$ , often times called non-conformity score, measures how well the value  $y$  conforms with the predictions of our fitted model. A common example in the regression setting is  $s_t(x, y) = |y - \hat{f}_t(x)|$ . Given a coverage level  $1 - \alpha$  and a threshold  $q_t(1 - \alpha)$ , we can define the prediction set as

$$\mathcal{C}_t^{1-\alpha}(x) = \{y \in \mathcal{Y} : s_t(x, y) \leq q_t(1 - \alpha)\} \quad (6.1)$$

Under the exchangeability of the data, and symmetry of the given model fitting algorithm as a function of the data, this definition of prediction sets, with  $q_t(1 - \alpha)$  the  $\frac{[(n+1)(1-\alpha)]}{n}$  quantile of the calibration scores  $\{s(X_t, Y_t)\}_{t=1}^T$ , the following holds:

$$\mathbb{P}[Y_t \in \mathcal{C}_t^{1-\alpha}(X_t)] \geq 1 - \alpha$$

However, when these assumptions do not hold, we lose this appealing guarantee. Gibbs and Candès [2021] introduce a new method to set the threshold  $q_t$ , namely ACI. Currently, most of the online conformal prediction algorithms update the threshold  $q_t(1 - \alpha)$  in an online fashion:

$$q_{t+1}(1 - \alpha) = q_t(1 - \alpha) - \eta_t(\alpha - \text{err}_t(\alpha)), \quad \text{err}_t(\alpha) = \mathbb{1}_{Y_t \notin \mathcal{C}_t^{1-\alpha}(X_t)} \quad (6.2)$$

Note that  $\text{err}_t(\alpha)$ , the  $\alpha$  miscoverage error, is equivalent to  $\mathbb{1}_{s_t > q_t(1-\alpha)}$ . As pointed out in previous works, the update rule in Eq. (6.2) does something intuitive, adjusts the threshold  $q_t(1 - \alpha)$  based on the discrepancy between observed and target miscoverage  $\alpha$ . Some authors have referred to this update rule as the quantile tracker [Angelopoulos and Bates, 2023, Angelopoulos et al., 2024]. It has been shown that, although simple, it has long run coverage guarantees, meaning

$$\frac{1}{T} \sum_{t=1}^T \mathbb{1}_{Y_t \in \mathcal{C}_t(X_t)} \rightarrow 1 - \alpha$$

However, this definition has a subtle caveat. If we run it for different levels of coverage, the prediction sets are not necessarily nested, which is undesirable for applications that multiple coverage levels are needed at the same time. Formally, consider an increasing sequence of miscoverage levels  $0 = \alpha_0 < \alpha_1 < \alpha_2 < \dots < \alpha_K < \alpha_{K+1} = 1$ . To simplify notation, we denote  $\mathcal{C}_t^i := \mathcal{C}_t^{1-\alpha_i}(X_t)$ ,  $q_{t,i} := q_t(1 - \alpha_i)$  and  $\text{err}_{t,i} := \text{err}_t(\alpha_i)$ . For every

$i \in [K]$ , our goal is to control miscoverage while preserving monotonicity in the thresholds  $q_{t,i} > q_{t,i+1}$  for all  $t \geq 1$ . In other words, nestedness in the prediction sets  $\mathcal{C}_t^{i+1} \subset \mathcal{C}_t^i$ .

## 6.4 Methods

We explore alternative approaches to enforce nestedness of prediction sets across increasing miscoverage levels. Our starting point is the standard online update of quantile thresholds, which it can also be interpreted through an online optimization lens. In particular, the update in Equation (6.2) can be viewed as an instance of online (sub)gradient descent applied to the pinball (quantile) loss  $\rho_t(q, 1 - \alpha) = (s_t - q) (\mathbb{1}_{\{s_t > q\}} - \alpha)$ ,

$$q_{t+1} = q_t - \eta_t \nabla \rho_t(q_t, 1 - \alpha) \quad (6.3)$$

We can further extend this definition to estimate quantiles for multiple miscoverage levels. Let's define the joint loss function as  $f_t(\bar{q}) = \sum_{i=1}^K \rho_t(q_i, 1 - \alpha_i)$  where  $\bar{q} = (q_1, \dots, q_K)$ . We can then express the joint update rule as

$$\bar{q}_{t+1} = \bar{q}_t - \eta_t \nabla f_t(\bar{q}_t) \quad (6.4)$$

Note that this update rule is just individually estimating different quantile at the same time, but the estimations are not sharing information and they won't necessarily be monotonic. Using the common online optimization perspective, we can impose constraints that will modify the update rule to guarantee monotonicity. Let's revisit the definition of the online mirror descent algorithm

$$x_{t+1} = \arg \min_{x \in \mathcal{X}} \eta \langle \nabla f(x_t), x \rangle + D_\phi(x, x_t)$$

Where  $D_\phi$  is the Bregman divergence generated by a convex function  $\phi$ . When  $\phi(x) = \frac{1}{2} \|x\|_2^2$  and  $\mathcal{X} = \mathbb{R}^n$ , we recover the usual gradient descent algorithm and then the update rule in Eq. (6.4). If we constrain  $\mathcal{X}$  to be a convex set (for example, the set of decreasing sequences in  $\mathbb{R}^n$ ) we are under the projected gradient descent (PG).

$$x_{t+1} = \Pi_{\mathcal{X}}(x_t - \eta \nabla f(x_t))$$

Where  $\Pi_{\mathcal{X}}$  denotes the projection on the the convex set  $\mathcal{X}$ . Finally, if we set  $\phi(x) = \sum_{i=1}^n x_i \log(x_i)$  and  $\mathcal{X} = \{x \in \mathbb{R}^n \mid x_i \geq 0, \sum_{i=1}^n x_i = 1\}$  we end up with exponentiated

gradient descent (EG), i.e.

$$(x_{t+1})_i = \frac{(x_t)_i \exp(-\eta \nabla f(x_t)_i)}{\sum_{j=1}^K (x_t)_j \exp(-\eta \nabla f(x_t)_j)}$$

In the next section we explain how we can leverage these variations of the online mirror descent algorithm to achieve our goal of enforce monotonicity in the quantiles and how we can leverage classic dynamic regret upper bounds to control the error between the true quantiles and the estimations.

### 6.4.1 Exponentiated Gradient Descent

We first consider the EG update restricted to the truncated simplex

$$\Delta_{K+1}^\mu = \left\{ w \in \mathbb{R}^{K+1} \mid w_i \geq \mu, \sum_{i=0}^K w_i = 1 \right\}$$

for some  $0 < \mu < \frac{1}{K+1}$ . Rather than updating the quantiles directly, we work with the normalized gaps between consecutive quantiles. Specifically, define

$$w_{t,i} = \frac{d_{t,i}}{B} \in [0, 1], \quad \text{where } d_{t,i} = q_{t,i} - q_{t,i+1},$$

with boundary conditions  $q_{t,K+1} = 0$  and  $q_{t,0} = B$  for all  $t \geq 1$ . We initialize  $w_{1,i} = \frac{1}{K+1}$  for all  $i \geq 0$ . This parametrization is equivalent to the change of variables

$$q_{t,i} = B \sum_{j=i}^K w_{t,j},$$

which can be written as  $\bar{q}_t = J\bar{w}_t$ , where  $J_{i,j} = B \mathbb{1}_{\{i \leq j\}}$ . Under this transformation, we can express the loss function in terms of  $\bar{w}_t$  as follows:

$$\begin{aligned} g_t(\bar{w}) &:= f_t(J\bar{w}) \\ &= \sum_{i=1}^K \rho_t \left( B \sum_{j=i}^K w_j, 1 - \alpha_i \right) \end{aligned}$$

This formulation makes explicit how each coordinate  $w_i$  influences all quantiles  $\{q_j\}_{j \leq i}$ , thereby introducing a structured coupling across coverage levels. We now apply the EG algorithm to minimize  $g_t$  over the truncated simplex  $\Delta_{K+1}^\mu$ :

$$\tilde{w}_{t+1,i} = w_{t,i} \exp(-\eta (\nabla g_t)_i), \quad \bar{w}_{t+1} = \arg \min_{\bar{w} \in \Delta_{K+1}^\mu} D_{\text{KL}}(\bar{w}, \tilde{w}_{t+1}), \quad (6.5)$$

The projection step admits a simple closed form.

**Proposition 2.** *The KL projection onto the truncated simplex  $\Delta_{K+1}^\mu$  is given by*

$$w_{t+1,i} = \max \{ \mu, c_t \tilde{w}_{t+1,i} \}$$

where

$$c_t = \frac{1 - |S_t^c| \mu}{\sum_{i \in S_t} \tilde{w}_{t+1,i}}, \quad S_t := \{i : \tilde{w}_{t+1,i} \geq \mu\}$$

Furthermore, we can calculate the gradient using the chain rule as  $\nabla g_t = J^\top \nabla f_t$  which yields the coordinate-wise expression

$$(\nabla g_t)_i = \sum_{j=1}^K \frac{\partial \rho_t(q_{t,j}, 1 - \alpha_j)}{\partial q_{t,j}} \frac{\partial q_{t,j}}{\partial w_{t,i}} = B \sum_{j=i}^K (\alpha_j - \text{err}_{t,j}),$$

since  $\partial_{q_{t,j}} \rho_t(q_{t,j}, 1 - \alpha_j) = \alpha_j - \text{err}_{t,j}$  and  $\frac{\partial q_{t,j}}{\partial w_{t,i}} = B \mathbb{1}_{\{i \leq j\}}$ .

This reveals a key feature of the EG update: it induces a global coupling across quantile levels. Each coordinate  $w_{t,i}$  controls the gap between consecutive quantiles and influences all prediction sets up to level  $i$ . Correspondingly, the gradient  $(\nabla g_t)_i$  aggregates the miscoverage discrepancies across these levels. When the method is overly conservative (i.e.,  $\text{err}_{t,j} < \alpha_j$  on average), the cumulative term is positive and the corresponding gap is reduced; when it undercovers, the gap is increased. As a result, the EG update reallocates mass across quantile gaps in a coordinated manner, adjusting the entire quantile curve simultaneously while preserving nestedness by construction. This global interaction contrasts with projection-based approaches, where corrections are applied locally.

Standard regret guarantees for EG, combined with mild assumptions on the distribution of  $s_t$ , allow us to control the deviation of the estimated quantiles from their targets.

**Theorem 6.4.1.** *Let  $s_1, \dots, s_T$  be a sequence of random variables with densities  $p_t(s) > p > 0$ . Then, the quantiles  $\bar{q}_t = J\bar{w}_t$  using the updating rule in Eq. (6.5) satisfy*

$$\frac{1}{T} \sum_{t=1}^T p \frac{\|\bar{q}_t - \bar{q}_t^*\|_2^2}{2} \leq \frac{(1 + \log(1/\mu))(1 + V_T^w)}{\eta T} + (BK)^2 \frac{\eta}{2} \quad (6.6)$$

Where  $q_{t,i}^*$  is such that  $\mathbb{P}(s_t < q_{t,i}^*) = 1 - \alpha_i$  and  $V_T^w = \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1$ .

Similar to Gibbs and Candès [2024], if we additionally assume the map

$$q_i \rightarrow \mathbb{P}(s_t \leq q_i | \{s_r\}_{r < t})$$

is  $L$ -Lipschitz, then  $|\mathbb{P}(s_t \leq q_i | \{s_r\}_{r < t}) - (1 - \alpha_i)| < L|q_{t,i} - q_{t,i}^*|$ , then Theorem 6.4.1 translates to coverage error bound.

## 6.4.2 Projected Gradient Descent

For our second method, we will use a modified version of projected gradient descent such that we keep a minimum difference between quantiles to avoid collapsing to single points. We first update the whole quantile vector according to the gradient. Finally, we project it onto the set with decreasing elements.

$$\tilde{q}_{t+1} = \bar{q}_t - \eta(\bar{\alpha} - \overline{\text{err}}_t) \quad \bar{q}_{t+1} = \Pi_{\mathcal{Q}}(\tilde{q}_{t+1}) \quad (6.7)$$

Where  $\overline{\text{err}}_t = (\text{err}_{t,1}, \dots, \text{err}_{t,K})$  and  $\mathcal{Q} = \{\bar{q} : B \geq q_1 \geq \dots \geq q_K \geq 0\}$ . The projection  $\Pi_{\mathcal{Q}}$  reduces to the PAVA algorithm [Ayer et al., 1955].

In contrast to EG, the PG method operates via local correction: it first performs independent gradient updates on each quantile and then enforces monotonicity through a projection step. This projection only adjusts neighboring quantiles as needed to restore the ordering constraint, without explicitly redistributing information across all levels. Consequently, PG ensures feasibility (nestedness) but lacks the intrinsic global coordination present in EG, which can limit its ability to exploit shared structure across coverage levels.

**Theorem 6.4.2.** *Let  $s_1, \dots, s_T$  be a sequence of random variables with densities  $p_t(s) > p > 0$ . Then, the quantiles  $\bar{q}_t$  from the update rule in Eq. (6.7) satisfy*

$$\frac{1}{T} \sum_{t=1}^T \frac{p \|\bar{q}_t - \bar{q}_t^*\|_2^2}{2} \leq \frac{3B^2K(1 + V_T^q)}{\eta T} + \eta K \quad (6.8)$$

Where  $q_{t,i}^*$  is such that  $\mathbb{P}(s_t < q_{t,i}^*) = 1 - \alpha_i$  and  $V_T^q = \sum_{t=1}^{T-1} \|q_{t+1}^* - q_t^*\|_1$

## 6.5 Experiments

In this section we evaluate our methods as well as the previously proposed quantile tracker and naive adaptations to enforce nestedness. We measure three key properties: (i) calibration across the full risk spectrum, (ii) nestedness of prediction sets, and (iii) statistical efficiency of joint quantile estimation. We evaluate performance across multiple coverage levels  $\{\alpha_i\}_{i=1}^K$  using the following metrics:

### 6.5.1 Metrics

**Full risk spectrum calibration error.** This metric evaluates the deviation between the empirical miscoverage rate and the target level  $\alpha_i$  across all coverage levels. While the quantile tracker is expected to perform well under this metric, it is known to exhibit over-correction behavior [Angelopoulos et al., 2024]. In particular, the algorithm may alternate between periods of undercoverage and overly conservative predictions, producing wider intervals that compensate for past errors. As a result, a low average calibration error can be misleading, as it may reflect oscillatory behavior rather than stable and well-calibrated uncertainty estimates.

$$CE_i = \left| \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{Y_t \notin \mathcal{C}_t^i(X_t)\} - \alpha_i \right|$$

**Quantile estimation error.** In the simulated setting, the true quantiles are known, allowing us to directly compare them with the online estimates. This metric evaluates how accurately each method tracks the underlying quantile dynamics over time. In contrast to calibration error, which reflects long-run coverage, this measure provides a more direct assessment of statistical efficiency and the ability to adapt to stochastic variability in the data-generating process.

$$\|q_t - q_t^*\|_1 = \sum_{i=1}^K |q_{t,i} - q_{t,i}^*|$$

**Prediction set size.** We calculate the rolling average width of  $\mathcal{C}_t^i$  across  $t$ :, this jointly

$$\sum_{s=t-w}^t |2q_s|$$

**Nestedness violations.** Gap between consecutive quantiles  $q_{t,i} - q_{t,i+1}$ . This will be positive for all  $t, i$  and all methods except the standard quantile tracking which does not guarantee nestedness.

### 6.5.2 Baselines

We compare our methods against two baselines: the classical quantile tracker and a projected quantile tracker. The first baseline applies the quantile tracker update in Eq. (6.4) independently across coverage levels. As discussed earlier, this update rule does not enforce monotonicity across quantiles and therefore does not guarantee nested prediction sets.

The second baseline is a simple extension of the quantile tracker that enforces mono-

tonicity via projection. Specifically, after performing the standard update in Eq. (6.4), the resulting quantiles are projected onto the set of monotone sequences. While this approach guarantees nestedness, it does not incorporate any coupling between quantile levels during the update step, and therefore does not fully exploit the benefits of joint estimation.

$$\tilde{q}_{t+1} = \tilde{q}_t + \eta (\overline{\text{err}}_t - \bar{\alpha}), \quad \bar{q}_{t+1} = \Pi_{\mathcal{Q}}(\tilde{q}_{t+1}) \quad (6.9)$$

Note that this approach differs from our projected gradient method. Here, the quantiles are updated independently using the gradient step, and only the current iterate is projected to enforce monotonicity. While this guarantees nested prediction sets, the projection acts as a post-processing step and does not introduce coupling across quantile levels during the update. As a result, it does not benefit from the joint estimation structure that our method exploits.

### 6.5.3 Synthetic Data: Uniform scores with reflected random-walk drift

To evaluate the ability of our online conformal procedures to track time-varying quantiles, we consider a synthetic setting in which the score distribution evolves according to a latent random walk with reflecting boundaries.

Let  $(z_t)_{t \geq 1}$  be a latent process taking values in a bounded interval  $[a, b]$ , initialized at  $z_1 \in [a, b]$ , and evolving as

$$\tilde{z}_{t+1} = z_t + \sigma \varepsilon_t, \quad \varepsilon_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1), \quad (6.10)$$

followed by reflection onto  $[a, b]$ :

$$z_{t+1} = \begin{cases} 2a - \tilde{z}_{t+1}, & \tilde{z}_{t+1} < a, \\ 2b - \tilde{z}_{t+1}, & \tilde{z}_{t+1} > b, \\ \tilde{z}_{t+1}, & \text{otherwise.} \end{cases} \quad (6.11)$$

If multiple reflections are required (i.e., when  $|\tilde{z}_{t+1} - z_t|$  is large), the reflection step is applied iteratively until  $z_{t+1} \in [a, b]$ . This construction yields a random walk that behaves locally like a Gaussian walk but is confined to a compact domain. Example simulated trajectories of  $(z_t)$  are shown in Figure 6.1. Conditional on  $z_t$ , the conformity score is generated as




Figure 6.1: Simulated trajectory of the latent process ( $z_t$ ). The process remains confined to the interval while exhibiting gradual drift, inducing a time-varying score distribution in the synthetic experiment.

$$s_t \mid z_t \sim \text{Unif}\left[z_t - \frac{w}{2}, z_t + \frac{w}{2}\right], \quad (6.12)$$

where  $w > 0$  controls the width of the score distribution. Thus, the score distribution has fixed spread but a time-varying center governed by the reflected random walk which induces a non-stationary environment. Hence, for any target miscoverage level  $\alpha \in (0, 1)$ , the ground truth  $(1 - \alpha)$ -quantile is available in closed form:

$$q_t^*(1 - \alpha) = z_t + \frac{w}{2} - \alpha w. \quad (6.13)$$

For a collection of levels  $\alpha_1 < \dots < \alpha_K$ , the true quantiles are

$$q_{t,i}^* = z_t + \frac{w}{2} - \alpha_i w, \quad i = 1, \dots, K. \quad (6.14)$$

This setting provides a controlled form of nonstationarity. The parameter  $\sigma$  governs the local variability of the latent process, while the reflecting boundaries ensure that  $z_t$  remains in  $[a, b]$ , preventing unbounded drift. As a result, the score distribution continuously evolves over time but within a fixed range, yielding a stable yet nonstationary benchmark.

A key advantage of a synthetic data experiment is that the true quantiles are known exactly, allowing us to directly assess the tracking performance of the online estimates. In particular, we evaluate the rolling average tracking error

$$\frac{1}{t - dt} \sum_{s=t-dt}^t \|\bar{q}_s - \bar{q}_s^*\|_1, \quad (6.15)$$

Calibration is assessed through empirical coverage at each level using the indicators  $\mathbb{1}\{s_t \leq q_{t,i}\}$ .

We fix the reflecting boundaries to  $a = 0.5$  and  $b = 9.5$ , and initialize the latent process at  $z_1 = 5$ . The random-walk scale is set to  $\sigma = 0.025$ . The score distribution has fixed width  $w = 1$ . We consider  $K = 9$  coverage levels given by

$$\alpha_i = 0.1 i, \quad i = 1, \dots, 9,$$

corresponding to target coverages  $1 - \alpha_i \in \{0.9, 0.8, \dots, 0.1\}$ . The simulation is run over  $t \in \{1, \dots, 50,000\}$ , and performance metrics are computed using evaluation windows of size  $dt = 10,000$ . These choices produce a moderate level of drift while keeping the score distribution well within the interval  $[0, 10]$ , ensuring stable behavior throughout the simulation.

Figure 6.2 shows that the EG update achieves a lower  $\ell_1$  tracking error, supporting the hypothesis that sharing information across quantiles improves statistical efficiency. In contrast, the projected gradient method closely mirrors the behavior of the independent quantile tracker, indicating limited gains from joint optimization in this setting.

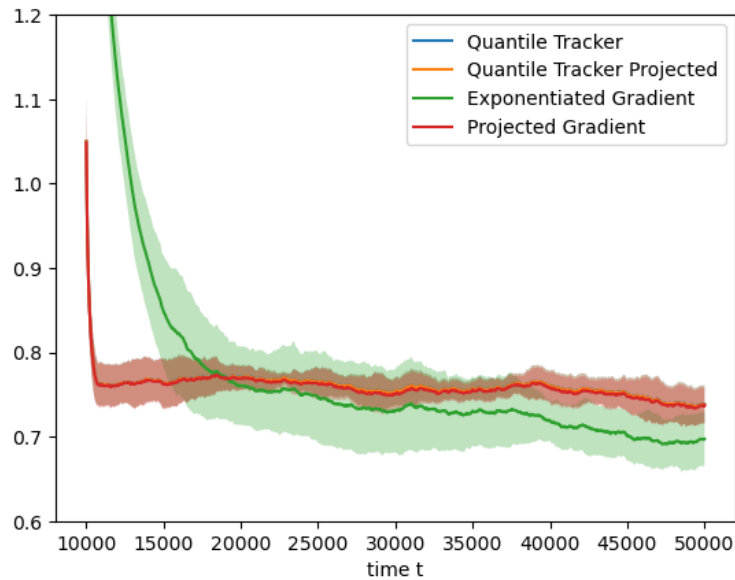


Figure 6.2: Rolling average of the  $\ell_1$  distance between  $q_t$  and  $q_t^*$

### 6.5.4 Real Data: US Inflation

We evaluate our methods on the US Consumer Price Index (CPI) for All Urban Consumers, published by the Bureau of Labor Statistics and distributed via FRED [U.S. Bureau of Labor Statistics, n.d.]. Inflation forecasting is a natural application for our framework, as policymakers and market participants often require uncertainty estimates at multiple confidence levels to support decisions under heterogeneous risk preferences.

We use monthly CPI data from 1950 to 2025 and construct the yearly inflation series as

$$y_t = \frac{\text{CPI}_t - \text{CPI}_{t-12}}{\text{CPI}_{t-12}}.$$

To generate predictions, we employ an autoregressive model of order three, AR(3), which captures short-term temporal dependencies in the inflation series. At each time  $t$ , we fit the model

$$y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \beta_3 y_{t-3} + \varepsilon_t,$$

using a rolling window of the previous five years of monthly observations (i.e., 60 data points). The model is retrained sequentially at each time step, and the fitted model is used to produce a one-step-ahead prediction  $\hat{y}_{t+1}$ . This rolling-window procedure allows the predictive model to adapt to evolving dynamics in the inflation process, reflecting its non-stationary nature.

We use as the nonconformity score the absolute prediction error,

$$s_t = |y_t - \hat{y}_t|,$$

which measures the deviation between observed and predicted inflation. These scores are then used to construct prediction sets of the form

$$\mathcal{C}_t^i = \{y \in \mathbb{R} : |y - \hat{y}_t| \leq q_{t,i}\},$$

where the thresholds  $q_{t,i}$  are updated online using the proposed methods corresponding to the coverage level  $1 - \alpha_i$  for  $i \in [K]$ . We consider  $K = 99$  coverage levels given by

$$\alpha_i = 0.01 i, \quad i = 1, \dots, 99,$$

corresponding to target coverages  $1 - \alpha_i \in \{0.99, 0.8, \dots, 0.01\}$ .

Figure 6.3 reports the cumulative average calibration error aggregated across all miscoverage levels  $\{\alpha_i\}_{i=1}^K$ .

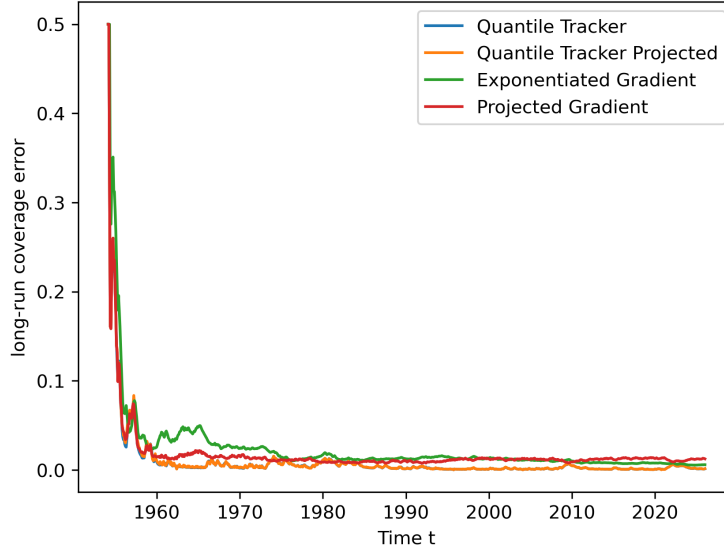


Figure 6.3: Cumulative average of the sum of mis coverage errors  $CE_i$

$$\sum_{i=1}^K \left| \frac{1}{T} \sum_{t=1}^T \mathbb{1}\{Y_t \notin \mathcal{C}_i^i(X_t)\} - \alpha_i \right|$$

Specifically, at each time  $t$  we compute the average deviation between the empirical mis-coverage and the target levels, and report its cumulative average over time. This metric captures how well each method calibrates the entire risk spectrum, rather than a single coverage level. We observe that all methods eventually achieve small calibration error, indicating convergence to the desired long-run coverage. However, their transient behaviors differ. The unconstrained quantile tracker converges more rapidly toward zero, which is consistent with its tendency to overcorrect by alternating between undercoverage and overly conservative intervals. In contrast, the proposed methods exhibit slightly slower but more stable trajectories. This suggests that enforcing structural constraints to preserve nestedness introduces a controlled trade-off, leading to more stable calibration dynamics while maintaining accurate long-run coverage.

Figure 6.4 visualizes the resulting prediction sets across multiple coverage levels. Each color corresponds to a different miscoverage level  $\alpha$ , forming a fan chart representation of the estimated quantile function over time. All methods adapt to changes in uncertainty, with prediction sets widening during periods of high volatility (e.g., around 1980, 2010, and 2020) and contracting during more stable regimes.

However, important qualitative differences emerge across methods. Baseline approaches (top panels) exhibit higher variability and occasional distortions in the quantile structure,

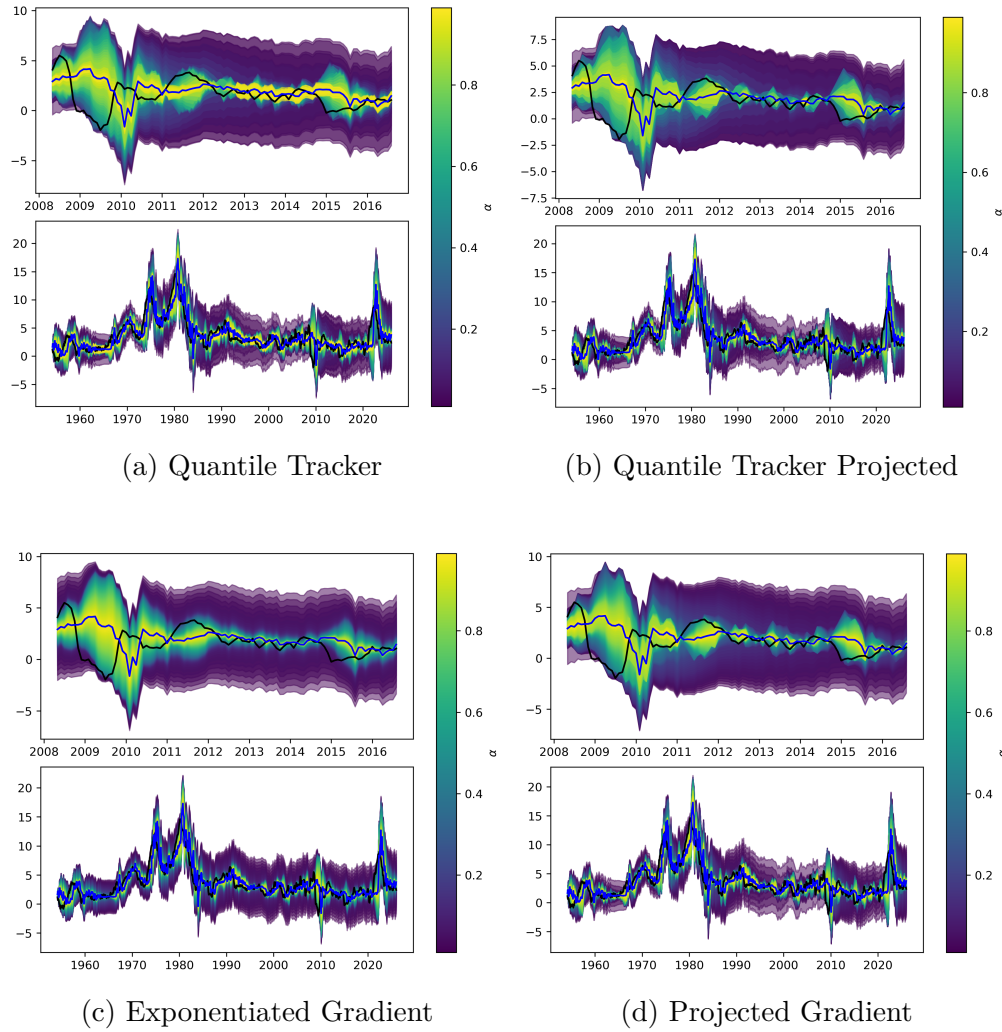


Figure 6.4: Prediction intervals across multiple coverage levels in the whole window (bottom panels) and zoom after the 2008 crises (top panels): (a) standard quantile tracker method, (b) quantile tracker with the last iteration projected, (c) exponentiated gradient and (d) projected gradient

leading to irregular and sometimes non-nested prediction bands. In contrast, the proposed constrained methods (bottom panels) produce smoother and more coherent prediction sets, consistently preserving the nested structure across all coverage levels. This behavior highlights the benefit of enforcing monotonicity jointly with online updates: beyond guaranteeing structural consistency, it stabilizes the learning dynamics and yields more interpretable and reliable uncertainty estimates in practice. In addition, The EG update yields smoother prediction sets by coordinating adjacent quantiles through cumulative discrepancies, reducing abrupt changes across levels.

## 6.6 Discussion

This work is motivated by practical settings in which uncertainty must be assessed across multiple risk levels simultaneously. In many applications—including forecasting, risk management, and scientific decision-making—different stakeholders operate under heterogeneous risk tolerances and require prediction sets at several coverage levels. Standard online conformal methods address each level independently, which can lead to incoherent and non-nested prediction sets. By jointly estimating multiple quantiles under structural constraints, our approach provides a coherent risk spectrum with valid coverage guarantees and strictly nested prediction sets, improving both interpretability and usability in downstream decision-making.

A key insight of our methods lies in how information is shared across coverage levels. The exponentiated gradient (EG) update induces a form of global coupling, where each update aggregates miscoverage discrepancies across multiple levels. This leads to coordinated adjustments of the entire quantile curve, allowing the method to borrow strength across levels and improving statistical efficiency. In contrast, the projected gradient (PG) update operates through local corrections: quantiles are first updated independently and then minimally adjusted via projection to restore monotonicity. While this guarantees feasibility, it lacks the intrinsic global coordination of EG. Empirically, this distinction manifests in improved quantile estimation accuracy for EG, particularly in non-stationary environments where sharing information across levels is beneficial.

Beyond accuracy, the global structure of the EG update also impacts the smoothness of the estimated quantile function across coverage levels. Because updates are driven by cumulative discrepancies, adjacent quantiles evolve in a coordinated manner, reducing irregularities and abrupt changes across levels. This leads to smoother and more stable prediction bands, which is particularly desirable in applications such as forecasting, where interpretability of uncertainty is critical. In contrast, methods based on independent updates or post-hoc

projection may exhibit more variability and local distortions in the quantile structure.

An important direction for future work is the incorporation of adaptive or data-driven learning rates. Recent advances in online conformal prediction have shown that carefully designed step-size sequences can improve adaptivity to distributional shifts and provide stronger guarantees across stochastic and adversarial regimes. Integrating such adaptive mechanisms into the joint quantile estimation framework may further enhance performance, allowing the methods to respond more effectively to non-stationarity while preserving the structural constraints imposed here.

Finally, our work contributes to a broader understanding of online conformal prediction beyond worst-case guarantees. While many existing methods focus on adversarial settings, practical applications often exhibit stochastic structure with evolving distributions. By framing the problem as one of tracking time-varying quantiles and analyzing performance through dynamic regret, our approach helps bridge the gap between adversarial robustness and statistical efficiency. This perspective opens the door to a richer analysis of online conformal methods in stochastic environments, where the goal is not only to guarantee coverage, but also to accurately and efficiently track the underlying uncertainty over time.

# APPENDIX A

## Near Optimal Pure Exploration in Logistic Bandits

### A.1 Concentration

The proof of Lemma 2.3.1 follows tightly Faury et al. [2020], with the only difference that there is no regularization. Instead, we assume we can control the minimum eigenvalue of the fisher information matrix. First we will prove the general concentration bound

**Lemma A.1.1.** *Let  $\{\mathcal{F}_t\}_{t=1}^\infty$  be a filtration. Let  $\{x_t\}_{t=1}^\infty$  be a stochastic process in  $\mathcal{B}_2(d)$  such that  $x_t$  is  $\mathcal{F}_t$  measurable. Let  $\{\varepsilon_t\}_{t=1}^\infty$  be a martingale difference sequence such that  $\varepsilon_t$  is  $\mathcal{F}_t$  measurable. Furthermore, assume that conditionally on  $\mathcal{F}_t$  we have  $|\varepsilon_t| \leq 1$  almost surely, and note  $\sigma_t^2 := \mathbb{E}[\varepsilon_t^2 \mid \mathcal{F}_{t-1}]$ . Let  $\lambda(t) > 0$  and for any  $t \geq 1$  define:*

$$\mathbf{H}_t := \sum_{s=1}^t \sigma_s^2 x_s x_s^T, \quad S_t := \sum_{s=1}^t \varepsilon_s x_s.$$

*Then, if exist  $t_0 \geq 1$  such that for  $t \geq t_0$ ,  $\lambda_{\min}(\mathbf{H}_t) > \lambda(t)$ , for any  $\delta \in (0, 1]$  :*

$$\mathbb{P} \left( \exists t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}} \geq 2\sqrt{\lambda(t)} + \frac{2}{\sqrt{\lambda(t)}} \log \left( \frac{\det(\mathbf{H}_t)^{\frac{1}{2}} \lambda(t)^{-\frac{d}{2}}}{\delta} \right) + \frac{2}{\sqrt{\lambda(t)}} d \log(2) \right) \leq \delta$$

*Proof.* Let  $\bar{\mathbf{H}}_t := \mathbf{H}_t - \lambda(t)\mathbf{I}_d$ , for  $\xi \in \mathbb{R}^d$  let  $M_0(\xi) = 1$  and for  $t \geq t_0$  define:

$$M_t(\xi) := \exp \left( \xi^T S_t - \|\xi\|_{\bar{\mathbf{H}}_t}^2 \right)$$

Using lemma 5 from Faury et al. [2020], for all  $\xi \in \mathcal{B}_2(d)$ ,  $\{M_t(\xi)\}_{t=t_0}^\infty$  is a non-negative super-martingale.

Let  $h(\xi)$  be a probability density function with support on  $\mathcal{B}_2(d)$  (to be defined later). For  $t \geq t_0$  let:

$$\bar{M}_t := \int_{\xi} M_t(\xi) dh(\xi)$$

By Lemma 20.3 of Lattimore and Szepesvári [2020]  $\bar{M}_t$  is also a non-negative supermartingale, and  $\mathbb{E}[\bar{M}_0] = 1$ . Let  $\tau$  be a stopping time with respect to the filtration  $\{F_t\}_{t=0}^{\infty}$ . We can follow the proof of Lemma 8 in Abbasi-Yadkori et al. [2011] to justify that  $\bar{M}_\tau$  is well-defined (independently of whether  $\tau < \infty$  holds or not) and that  $\mathbb{E}[\bar{M}_\tau] \leq 1$ . Therefore, with  $\delta \in (0, 1)$  and thanks to the maximal inequality:

$$\mathbb{P}\left(\log(\bar{M}_\tau) \geq \log\left(\frac{1}{\delta}\right)\right) = \mathbb{P}\left(\bar{M}_\tau \geq \frac{1}{\delta}\right) \leq \delta \quad (\text{A.1})$$

We now proceed to compute  $\bar{M}_t$  (more precisely a lower bound on  $\bar{M}_t$ ). Let  $\lambda(t)$  be a strictly positive scalar, and set  $h$  to be the density of an isotropic normal distribution of precision  $2\lambda(t)$  truncated on  $\mathcal{B}_2(d)$ . We will denote  $N(h)$  its normalization constant. Simple computations show that:

$$\begin{aligned} \bar{M}_t &= \frac{1}{N(h)} \int_{\mathcal{B}_2(d)} \exp\left(\xi^T S_t - \|\xi\|_{\mathbf{H}_t}^2 - \lambda(t)\|\xi\|^2\right) d\xi \\ &\geq \frac{1}{N(h)} \int_{\mathcal{B}_2(d)} \exp\left(\xi^T S_t - 2\|\xi\|_{\mathbf{H}_t}^2\right) d\xi \end{aligned}$$

To ease notations, let  $f(\xi) := \xi^T S_t - 2\|\xi\|_{\mathbf{H}_t}^2$  and  $\xi_* = \arg \max_{\|\xi\|_2 \leq 1/2} f(\xi)$ . Because:

$$f(\xi) = f(\xi_*) + (\xi - \xi_*)^T \nabla f(\xi_*) - 2(\xi - \xi_*)^T \mathbf{H}_t (\xi - \xi_*)$$

we obtain that:

$$\begin{aligned} \bar{M}_t &= \frac{e^{f(\xi_*)}}{N(h)} \int_{\mathbb{R}^d} \mathbf{1}_{\|\xi\|_2 \leq 1} \exp\left((\xi - \xi_*)^T \nabla f(\xi_*) - 2(\xi - \xi_*)^T \mathbf{H}_t (\xi - \xi_*)\right) d\xi \\ &= \frac{e^{f(\xi_*)}}{N(h)} \int_{\mathbb{R}^d} \mathbf{1}_{\|\xi + \xi_*\|_2 \leq 1} \exp\left(\xi^T \nabla f(\xi_*) - 2\xi^T \mathbf{H}_t \xi\right) d\xi \quad (\text{change of variable } \xi + \xi_*) \\ &\geq \frac{e^{f(\xi_*)}}{N(h)} \int_{\mathbb{R}^d} \mathbf{1}_{\|\xi\|_2 \leq 1/2} \exp\left(\xi^T \nabla f(\xi_*) - 2\xi^T \mathbf{H}_t \xi\right) d\xi \quad (\text{as } \|\xi_*\|_2 \leq 1/2) \\ &= \frac{e^{f(\xi_*)}}{N(h)} \int_{\mathbb{R}^d} \mathbf{1}_{\|\xi\|_2 \leq 1/2} \exp\left(\xi^T \nabla f(\xi_*)\right) \exp\left(-\frac{1}{2}\xi^T (4\mathbf{H}_t) \xi\right) d\xi \end{aligned}$$

By defining  $g(\xi)$  the density of the normal distribution of precision  $4\mathbf{H}_t$  truncated on the

ball  $\{\xi \in \mathbb{R}^d, \|\xi\|_2 \leq 1/2\}$  and noting  $N(g)$  its normalizing constant, we can rewrite:

$$\begin{aligned}
\bar{M}_t &\geq \exp(f(\xi_*)) \frac{N(g)}{N(h)} \mathbb{E}_g \left[ \exp(\xi^T \nabla f(\xi_*)) \right] \\
&\geq \exp(f(\xi_*)) \frac{N(g)}{N(h)} \exp(\mathbb{E}_g [\xi^T \nabla f(\xi_*)]) && \text{(Jensen's inequality)} \\
&\geq \exp(f(\xi_*)) \frac{N(g)}{N(h)} && \text{(as } \mathbb{E}_g[\xi] = 0 \text{)}
\end{aligned}$$

Unpacking this results and assembling (10) and (11), we obtain that for any  $\xi_0$  such that  $\|\xi_0\|_2 \leq 1/2$  :

$$\begin{aligned}
\mathbb{P}\left(\bar{M}_t \geq \frac{1}{\delta}\right) &\geq \mathbb{P}\left(\exp(f(\xi_*)) \frac{N(g)}{N(h)} \geq 1/\delta\right) \\
&= \mathbb{P}\left(\log\left(\exp(f(\xi_*)) \frac{N(g)}{N(h)}\right) \geq \log(1/\delta)\right) \\
&= \mathbb{P}\left(f(\xi_*) \geq \log(1/\delta) + \log\left(\frac{N(h)}{N(g)}\right)\right) \\
&= \mathbb{P}\left(\max_{\|\xi\|_2 \leq 1/2} \xi^T S_t - 2\|\xi\|_{\mathbf{H}_t}^2 \geq \log(1/\delta) + \log\left(\frac{N(h)}{N(g)}\right)\right) \\
&\geq \mathbb{P}\left(\xi_0^T S_t - 2\|\xi_0\|_{\mathbf{H}_t}^2 \geq \log(1/\delta) + \log\left(\frac{N(h)}{N(g)}\right)\right) && \text{(A.2)}
\end{aligned}$$

In particular, we can use:

$$\xi_0 := \frac{\mathbf{H}_t^{-1} S_t}{\|S_t\|_{\mathbf{H}_t^{-1}}} \frac{\lambda(t)^{1/2}}{4}$$

since

$$\|\xi_0\|_2 \leq \frac{\lambda(t)^{1/2}}{4} (\lambda_{\min}(\mathbf{H}_t))^{-1/2} \leq 1/2$$

Using this value of  $\xi_0$  in Eq. (A.2) yields:

$$\mathbb{P}\left(\|S_t\|_{\mathbf{H}_t^{-1}} \geq \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log(1/\delta) + \frac{4}{\sqrt{\lambda(t)}} \log\left(\frac{N(h)}{N(g)}\right)\right) \leq \mathbb{P}\left(\bar{M}_t \geq \frac{1}{\delta}\right)$$

We can use the upper-bound from lemma 6 of Faury et al. [2020] for the log of their ratio  $\log\left(\frac{N(h)}{N(g)}\right)$ . Therefore with probability at least  $1 - \delta$  and by using Eq. (A.1):

$$\|S_\tau\|_{\mathbf{H}_\tau^{-1}} \leq \frac{\sqrt{\lambda(\tau)}}{2} + \frac{4}{\sqrt{\lambda(\tau)}} \log(1/\delta) + \frac{4}{\sqrt{\lambda(\tau)}} \log\left(\frac{2^d \det(\mathbf{H}_\tau)^{1/2}}{\lambda(\tau)^{d/2} \delta}\right) + \frac{4}{\sqrt{\lambda(\tau)}} d \log(2)$$

Directly following the stopping time construction argument in the proof of Theorem 1 of (Abbasi-Yadkori et al., 2011) we obtain that with probability at least  $1 - \delta$ , for all  $t \in \mathbb{N}$  :

$$\|S_t\|_{\mathbf{H}_t^{-1}} \leq \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log \left( \frac{2^d \det(\mathbf{H}_t)^{1/2}}{\lambda(t)^{d/2} \delta} \right) + \frac{4}{\sqrt{\lambda(t)}} d \log(2)$$

□

**Lemma A.1.2.** (lemma 6 [Fauray et al., 2020]) *The following inequality holds:*

$$\log \left( \frac{N(h)}{N(g)} \right) \leq \log \left( 2^d \left( \frac{\det(\mathbf{H}_t)}{\lambda(t)^d} \right)^{1/2} \right) + d \log(2)$$

**Lemma 2.3.1** Let  $\delta \in (0, 1]$  and  $\lambda(t) > 0$  for  $t \geq 1$ . If exist  $t_0 \geq 1$  such that for  $t \geq t_0$ ,  $\lambda_{\min}(H_t(\theta_*)) > \lambda(t)$ , with probability at least  $1 - \delta$ :

$$\forall t \geq t_0, \quad \left\| g_t(\hat{\theta}_t) - g_t(\theta_*) \right\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \gamma_t(\delta)$$

$$\text{Where } \gamma_t(\delta) := \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log \left( \frac{2^d}{\delta} \left( \frac{Lt}{d\lambda(t)} \right)^{\frac{d}{2}} \right)$$

*Proof.* Recall that  $\hat{\theta}_t$  is the unique maximizer of the log-likelihood:

$$\mathcal{L}_t^\lambda(\theta) := \sum_{s=1}^t \left[ r_s \log \mu(x_s^\top \theta) + (1 - r_s) \log(1 - \mu(x_s^\top \theta)) \right]$$

and therefore  $\hat{\theta}_t$  is a critical point of  $\mathcal{L}_t^\lambda(\theta)$ . Solving for  $\nabla_\theta \mathcal{L}_t^\lambda = 0$  and using the fact that  $\dot{\mu} = \mu(1 - \mu)$  we obtain:

$$\sum_{s=1}^t \mu(x_s^\top \hat{\theta}_t) x_s = \sum_{s=1}^t r_s x_s$$

This result, combined with the definition of  $g_t(\theta_*) = \sum_{s=1}^{t-1} \mu(x_s^\top \theta_*) x_s$  yields:

$$\begin{aligned} g_t(\hat{\theta}_t) - g_t(\theta_*) &= \sum_{s=1}^t \varepsilon_s x_s \\ &= S_t \end{aligned}$$

where we denoted  $\varepsilon_s := r_s - \mu(x_s^\top \theta_*)$  for all  $s \geq 1$  and  $S_t := \sum_{s=1}^t \varepsilon_s x_s$  for all  $t \geq 1$ . Then:

$$\left\| g_t(\hat{\theta}_t) - g_t(\theta_*) \right\|_{\mathbf{H}_t^{-1}(\theta_*)} = \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \tag{A.3}$$

Note that  $\{\varepsilon_t\}_{t=1}^\infty$  is a martingale difference sequence adapted to  $\mathcal{F}$  and almost surely bounded by 1. Also, note that for all  $s \geq 1$ :

$$\dot{\mu}(x_s^\top \theta_*) = \mu(x_s^\top \theta_*) (1 - \mu(x_s^\top \theta_*)) = \mathbb{E}[\varepsilon_s^2 | \mathcal{F}_{s-1}] =: \sigma_s^2$$

and thus  $\mathbf{H}_t(\theta_*) = \sum_{s=1}^t \sigma_s^2 x_s x_s^\top$ . All the conditions of Lemma A.1.1 are checked and therefore:

$$\begin{aligned} 1 - \delta &\leq \mathbb{P}\left(\forall t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log\left(\frac{\det(\mathbf{H}_t(\theta_*))^{1/2}}{\lambda(t)^{d/2}\delta}\right) + \frac{4d}{\sqrt{\lambda(t)}} \log(2)\right) \\ &\leq \mathbb{P}\left(\forall t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log\left(\frac{(Lt/d)^{d/2}}{\lambda(t)^{d/2}\delta}\right) + \frac{4d}{\sqrt{\lambda(t)}} \log(2)\right) \\ &\leq \mathbb{P}\left(\forall t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \frac{\sqrt{\lambda(t)}}{2} + \frac{4}{\sqrt{\lambda(t)}} \log\left(\frac{1}{\delta} \left(\frac{Lt}{d\lambda(t)}\right)^{d/2}\right) + \frac{4d}{\sqrt{\lambda(t)}} \log(2)\right) \\ &= \mathbb{P}\left(\forall t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \gamma_t(\delta)\right) \end{aligned} \tag{A.4}$$

where we used that:

$$\det(\mathbf{H}_t(\theta_*)) \leq L^d \det\left(\sum_{s=1}^t x_s x_s^\top\right) \leq L^d \left(\frac{t}{d}\right)^d \leq \left(\frac{Lt}{d}\right)^d$$

thanks to Lemma A.4.1. Assembling Eq. (A.3) with Eq. (A.4) yields:

$$\begin{aligned} \mathbb{P}\left(\forall t \geq 1, \|g_t(\hat{\theta}_t) - g_t(\theta_*)\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \gamma_t(\delta)\right) &= \mathbb{P}\left(\forall t \geq 1, \|S_t\|_{\mathbf{H}_t^{-1}(\theta_*)} \leq \gamma_t(\delta)\right) \\ &\geq 1 - \delta \end{aligned}$$

hence the announced result. □

For next results, we will use the following notations:

$$\begin{aligned} \alpha(x, \theta_1, \theta_2) &:= \int_{v=0}^1 \dot{\mu}(vx^\top \theta_2 + (1-v)x^\top \theta_1) dv > 0 \\ \mathbf{G}_t(\theta_1, \theta_2) &:= \sum_{s=1}^{t-1} \alpha(x, \theta_1, \theta_2) x_s x_s^\top \mathbf{I}_d \end{aligned}$$

where  $\theta_1, \theta_2$  and  $x$  are vectors in  $\mathbb{R}^d$ . The quantities  $\alpha(x, \theta_1, \theta_2)$  and  $\mathbf{G}_t(\theta_1, \theta_2)$  naturally arise when studying GLMs. Indeed, note that for all  $x \in \mathbb{R}^d$  and  $\theta \in \mathbb{R}^d$ , the following

equality holds:

$$\mu \left( x^\top \theta_1 \right) - \mu \left( x^\top \theta_2 \right) = \alpha \left( x, \theta_2, \theta_1 \right) x^\top \left( \theta_1 - \theta_2 \right)$$

This result is classical (see Filippi et al. [2010]) and can be obtained by a straight-forward application of the mean-value theorem. It notably allows us to link  $\theta_1 - \theta_2$  with  $g_t(\theta_1) - g_t(\theta_2)$ . Namely, it is straightforward that:

$$\begin{aligned} g_t(\theta_1) - g_t(\theta_2) &= \sum_{s=1}^{t-1} \alpha(x_s, \theta_2, \theta_1) x_s x_s^\top (\theta_1 - \theta_2) \\ &= \mathbf{G}_t(\theta_2, \theta_1) (\theta_1 - \theta_2) \end{aligned}$$

Because  $\mathbf{G}_t(\theta_1, \theta_2) \succ \mathbf{0}_{d \times d}$  this yields:

$$\|\theta_1 - \theta_2\|_{\mathbf{G}_t(\theta_2, \theta_1)} = \|g_t(\theta_1) - g_t(\theta_2)\|_{\mathbf{G}_t^{-1}(\theta_2, \theta_1)}$$

**Lemma A.1.3.** (Lemma 10 [Fauray et al., 2020]) For all  $\theta_1, \theta_2 \in \Theta$  the following inequalities hold:

$$\begin{aligned} \mathbf{G}_t(\theta_1, \theta_2) &\geq (1 + 2S)^{-1} \mathbf{H}_t(\theta_1) \\ \mathbf{G}_t(\theta_1, \theta_2) &\geq (1 + 2S)^{-1} \mathbf{H}_t(\theta_2) \end{aligned}$$

**Lemma A.1.4.**

$$\|\theta_t^{(1)} - \theta^*\|_{\mathbf{H}_t(\theta_t^{(1)})} \leq 2(1 + 2S) \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)}$$

*Proof.*

$$\begin{aligned} \|\theta_t^{(1)} - \theta^*\|_{\mathbf{H}_t(\theta_t^{(1)})} &\leq \sqrt{1 + 2S} \|\theta_t^{(1)} - \theta^*\|_{\mathbf{G}_t(\theta^*, \theta_t^{(1)})} \\ &= \sqrt{1 + 2S} \|g_t(\theta_t^{(1)}) - g_t(\theta^*)\|_{\mathbf{G}_t^{-1}(\theta^*, \theta_t^{(1)})} \\ &\leq \sqrt{1 + 2S} \left( \|g_t(\theta_t^{(1)}) - g_t(\hat{\theta}_t)\|_{\mathbf{G}_t^{-1}(\theta^*, \theta_t^{(1)})} + \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{G}_t^{-1}(\theta^*, \theta_t^{(1)})} \right) \\ &\leq (1 + 2S) \left( \|g_t(\theta_t^{(1)}) - g_t(\hat{\theta}_t)\|_{\mathbf{H}_t^{-1}(\theta_t^{(1)})} + \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)} \right) \\ &\leq 2(1 + 2S) \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)} \end{aligned}$$

□

**Lemma 2.5.4**

Let  $\epsilon > 0$ , assume that  $\lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right) \geq ct^\alpha \geq \kappa_0 \lambda_0$  a.s. for all  $t \geq t_0$ , some  $t_0 \geq 1$

and some constants  $\alpha, c > 0$ . Then

$$\forall t \geq t_0 \quad \mathbb{P} \left( \|\hat{\theta}_t - \theta^*\| \geq \varepsilon \right) \leq c_2 t^{d/2} \exp \left( -c_1 \varepsilon t^{\alpha/2} \right)$$

Where  $c_1, c_2$  are positive constants independent of  $\varepsilon$  and  $t$ .

*Proof.* Similar to Lemma A.1.4, we have

$$\begin{aligned} \|\hat{\theta}_t - \theta^*\| &\leq \frac{\|\hat{\theta}_t - \theta^*\|_{\mathbf{H}_t(\theta^*)}}{\sqrt{\lambda_{\min}(\mathbf{H}_t(\theta^*))}} \\ &\leq \frac{2(1+2S)}{\sqrt{\lambda_{\min}(\mathbf{H}_t(\theta^*))}} \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)} \\ &\leq \frac{2(1+2S)}{\sqrt{c/\kappa_0}} t^{-\alpha/2} \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)} \end{aligned}$$

And, with probability at least  $1 - \delta$

$$\|\hat{\theta}_t - \theta^*\| \leq \frac{2(1+2S)}{\sqrt{c/\kappa_0}} t^{-\alpha/2} \gamma_t(\delta)$$

which we may rewrite after substitution as

$$\mathbb{P} \left( \|\hat{\theta}_t - \theta^*\| > \varepsilon \right) < c_2 t^{d/2} \exp \left( -c_1 \varepsilon t^{\alpha/2} \right)$$

□

**Lemma 2.5.1.**

Under any sampling rule, we have

$$\mathbb{P}_\theta \left( \tau_\delta < \infty \wedge i^*(\theta_{\tau_\delta}^{(1)}) \neq i^*(\theta) \right) \leq \delta$$

*Proof.* Lets consider

$$\mathcal{E}_1 = \{\tau_\delta < \infty\} = \{\exists t \geq 1 : Z(t) > \beta(\delta, t), t \in B\}$$

$$\mathcal{E}_2 = \{i^*(\theta_{\tau_\delta}^{(1)}) \neq i^*(\theta)\} = \{\theta \in \text{Alt}(\theta_{\tau_\delta}^{(1)})\}$$

$$\begin{aligned}
\mathcal{E}_1 \cap \mathcal{E}_2 &= \left\{ \exists t \geq 1 : Z(t) > \beta(\delta, t), t \in B, \theta \in \text{Alt}(\theta_t^{(1)}) \right\} \\
&= \left\{ \exists t \geq 1 : \inf_{\lambda \in \text{Alt}(\theta_t^{(1)})} \frac{1}{2} \|\theta_t^{(1)} - \lambda\|_{\mathbf{H}_t(\theta_t^{(1)})}^2 > \beta(\delta, t), t \in B, \theta \in \text{Alt}(\theta_t^{(1)}) \right\} \\
&\subseteq \left\{ \exists t \geq 1 : \frac{1}{2} \|\theta_t^{(1)} - \theta\|_{\mathbf{H}_t(\theta_t^{(1)})}^2 > \beta(\delta, t), t \in B \right\} \\
&\subseteq \left\{ \exists t \geq 1 : \sqrt{2}(1+2S) \|g_t(\hat{\theta}_t) - g_t(\theta^*)\|_{\mathbf{H}_t^{-1}(\theta^*)} > \sqrt{2}(1+2S)\gamma_t(\delta), t \in B \right\}.
\end{aligned}$$

Where the first containment is because of Lemma A.1.4. Using Lemma 2.3.1 we have

$$\mathbb{P}_\theta[\mathcal{E}_1 \cap \mathcal{E}_2] < \delta$$

Thus

$$\mathbb{P}_\theta \left( \tau_\delta < \infty \wedge i^*(\theta_{\tau_\delta}^{(1)}) \neq i^*(\theta) \right) = \mathbb{P}_\theta[\mathcal{E}_1 \cap \mathcal{E}_2] < \delta$$

□

### Lemma 2.5.3

Under forced exploration, the MLE estimator convergence a.s. to the true parameter

$$\lim_{t \rightarrow \infty} \hat{\theta}_t \stackrel{a.s.}{=} \theta^*$$

*Proof.* Using Lemma 2.5.2 we have that exists  $t_0$  such that  $\lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right) \geq c\chi_0 \sqrt{t}$  if  $t \geq t_0$ . On the other hand,

$$\begin{aligned}
\lambda_{\max} \left( \sum_{s=1}^t x_s x_s^\top \right) &\leq \text{Tr} \left( \sum_{s=1}^t x_s x_s^\top \right) \\
&= \sum_{s=1}^t \text{Tr}(x_s x_s^\top) \\
&= \sum_{s=1}^t \|x_s\|^2 \\
&\leq t
\end{aligned}$$

then

$$\lim_{t \rightarrow \infty} \frac{\lambda_{\min} \left( \sum_{s=1}^t x_s x_s^\top \right)}{\log \left( \lambda_{\max} \left( \sum_{s=1}^t x_s x_s^\top \right) \right)} = \infty$$

Using Theorem 2 of Chen et al. [1999] we have  $\lim_{t \rightarrow \infty} \hat{\theta}_t \stackrel{a.s.}{=} \theta^*$ .

□

## A.2 Tracking

$$\psi(\theta, w) = \inf_{\lambda \in \text{Alt}(\theta)} \frac{1}{2} \|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2$$

**Lemma A.2.1.** *If exists  $\epsilon_0 > 0$  such that  $i^*(\theta_t) = i^*(\theta)$  for all  $\theta_t$  such that  $\|\theta_t - \theta\| < \epsilon_0$ , then  $\psi$  is continuous in both  $\theta$  and  $w$ , and  $w \mapsto \psi(\theta, w)$  attains its maximum in  $\Sigma$  at a point  $w_\delta^*$  such that  $\sum_{x \in \mathcal{X}} (w_\delta^*)_x x x^\top$  is invertible.*

**Remark.** It is easy to check that BAI and TBP meet the assumption that exists  $\epsilon_0 > 0$  such that  $i^*(\theta_t) = i^*(\theta)$  for all  $\theta_t$  that  $\|\theta_t - \theta\| < \epsilon_0$

*Proof.* Let  $\theta \in \mathbb{R}^d$  such that  $i^*(\theta)$  is unique. Consider the alternative set  $\text{Alt}(\theta)$  and denote

$$f(\theta, \lambda, w) = \frac{1}{2} (\theta - \lambda)^\top \left( \sum_{x \in \mathcal{X}} w_x \dot{\mu}(x^\top \theta) x x^\top \right) (\theta - \lambda).$$

Let  $(\theta_t, w_t)_{t \geq 1}$  be a sequence taking values in  $\Theta \times \Sigma$  and converging to  $(\theta, w)$ . Let  $\epsilon < \epsilon_0$  and  $t_1 \geq 1$  such that for all  $t \geq t_1$  we have  $\|(\theta_t, w_t) - (\theta, w)\| < \epsilon$ . Now,  $\text{Alt}(\theta_t) = \text{Alt}(\theta)$  because  $i^*(\theta)$  is unique and  $i^*(\theta_t) = i^*(\theta)$ . Furthermore, note that  $f(\theta, \lambda, w)$  is a continuous function in  $\theta, \lambda, w$ , thus it is in particular continuous in  $\theta, w$ , and there exists  $t_2 \geq 1$  such that for all  $t \geq t_2$  and for all  $\lambda \in \mathbb{R}^d$ , it holds that  $|f(\theta, \lambda, w) - f(\theta_t, \lambda, w_t)| \leq \epsilon f(\theta, \lambda, w)$ . Hence, with our choice of  $\epsilon$ , we have for all  $t \geq t_1 \vee t_2$

$$\begin{aligned} |\psi(\theta, w) - \psi(\theta_t, w_t)| &= \left| \min_{\lambda \in \text{Alt}(\theta)} f(\theta, \lambda, w) - \min_{\lambda \in \text{Alt}(\theta_t)} f(\theta_t, \lambda, w_t) \right| \\ &\leq \epsilon \left| \min_{\lambda \in \text{Alt}(\theta)} f(\theta, \lambda, w) \right| \\ &\leq \epsilon |\psi(\theta, w)|. \end{aligned}$$

Thus  $\psi$  is continuous in  $\theta, w$ . Now, we know that  $w \mapsto \psi(\theta, w)$  is continuous on  $\Sigma$ , and by compactness of the simplex, the maximum is attained at some  $w_\delta^* \in \Sigma$ . Furthermore, since  $\mathcal{X}$  spans  $\mathbb{R}^d$ , we may construct an allocation  $\tilde{w}$  such that  $\sum_{a \in \mathcal{A}} \tilde{w}_a x x^\top$  is a positive definite matrix. In addition, by construction of  $\text{Alt}(\theta)$ , there exists some  $M > 0$  such that for all  $\lambda \in \text{Alt}(\theta)$  we have  $\|\theta - \lambda\| > M$ , which implies that  $\psi(\theta, \tilde{w}) \geq M^2 \lambda_{\min} \left( \sum_{x \in \mathcal{X}} \tilde{w}_x x x^\top \right) > 0$ . On the other hand, for any allocation  $w \in \Sigma$  such that  $\sum_{x \in \mathcal{X}} w_x x x^\top$  is rank deficient, we may

find a  $\lambda \in \text{Alt}(\theta)$  where  $\lambda - \theta$  is in the null space of  $\sum_{x \in \mathcal{X}} w_x x x^\top$ . Therefore,  $\sum_{x \in \mathcal{X}} (w_\theta^*)_x x x^\top$  is invertible. □

**Lemma A.2.2.** (*Maximum theorem*) Let  $\theta \in \mathbb{R}^d$ . Define

$$\psi^*(\theta) = \max_{w \in \Sigma} \psi(\theta, w)$$

and  $C^*(\theta) = \arg \max_{w \in \Sigma} \psi(\theta, w)$ . Then  $\psi^*$  is continuous at  $\theta$ , and  $C^*(\theta)$  is convex, compact and non-empty. Furthermore, we have for any open neighborhood  $\mathcal{V}$  of  $C^*(\theta)$ , there exists an open neighborhood  $\mathcal{U}$  of  $\theta$ , such that for all  $\theta' \in \mathcal{U}$ , we have  $C^*(\theta') \subseteq \mathcal{V}$ .

**Lemma A.2.3.** (*Lemma 6 [Jedra and Proutiere, 2020]*) Let  $(w(t))_{t \geq 1}$  be a sequence taking values in  $\Sigma$ , such that there exists a compact, convex and non empty subset  $C \subseteq \Sigma$ , there exists  $\varepsilon > 0$  and  $t_0(\varepsilon) \geq 1$  such that  $\forall t \geq t_0, d_\infty(w(t), C) \leq \varepsilon$ . Consider a sampling rule defined by Eq. (2.9) and

$$b_t = \arg \min_{x \in \text{supp}(\sum_{s=1}^t w(s))} \left( N_x(t) - \sum_{s=1}^t w_x(s) \right),$$

where  $N_x(0) = 0$  and for  $t \geq 0, N_x(t+1) = N_x(t) + \mathbb{1}_{\{x_t=x\}}$ . Then there exists  $t_1(\varepsilon) \geq t_0(\varepsilon)$  such that  $\forall t \geq t_1(\varepsilon), d_\infty\left(\left(N_x(t)/t\right)_{x \in \mathcal{X}}, C\right) \leq (p_t + d - 1)\varepsilon$  where  $p_t = |\text{supp}\left(\sum_{s=1}^t w(s)\right) \setminus \mathcal{X}_0| \leq K - d$ .

**Proposition 1** Under the sampling rules Eq. (2.9) and Eq. (2.13), the proportions of arm draws approach  $C^*(\theta^*) : \lim_{t \rightarrow \infty} d_\infty\left(\left(N_x(t)/t\right)_{x \in \mathcal{X}}, C^*(\theta^*)\right) = 0$ , a.s..

*Proof.* Let  $\varepsilon > 0$ . First, by Lemma A.2.2, there exists  $\xi(\varepsilon) > 0$  such that for all  $\theta' \in \mathbb{R}^d$  such that  $\|\theta - \theta'\| < \xi(\varepsilon)$ , it holds that  $\max_{w \in C^*(\theta')} d_\infty(w, C^*(\theta)) < \varepsilon/2$ . By Lemma 2.5.2, we have a sufficient exploration. That is  $\liminf_{t \rightarrow \infty} t^{-1/2} \lambda_{\min}\left(\sum_{s=1}^t x_s x_s^\top\right) > 0$ . Thus, by Lemma 2.5.3,  $\hat{\theta}_t$  converges almost surely to  $\theta^*$  with a rate of order  $o\left(t^{1/4}\right)$ . Consequently, there exists  $t_0 \geq 0$  such that for all  $t \geq t_0$ , we have  $\|\theta - \hat{\theta}_t\| \leq \xi(\varepsilon)$ . Then, we have

$$d_\infty(w(t), C^*(\theta^*)) \leq \max_{w \in C^*(\hat{\theta}_t)} d_\infty(w, C^*(\mu)) < \varepsilon.$$

We have shown that  $d_\infty(w(t), C^*(\theta^*)) \xrightarrow{t \rightarrow \infty} 0$  a.s. Next, we recall that by Lemma A.2.2,  $C^*(\theta)$  is non empty, compact and convex. Thus, applying Lemma A.2.3 yields immediately that  $d_\infty\left(\left(N_x(t)/t\right)_{x \in \mathcal{X}}, C^*(\theta^*)\right) \xrightarrow{t \rightarrow \infty} 0$  a.s.. □

**Theorem 2.6.1**

Log Track-and-Stop satisfies the same sample complexity upper bound

$$\mathbb{P}_\theta[\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{(\log(\frac{1}{\delta}))^2} \lesssim T^*(\theta)] = 1$$

*Proof.* From Lemma A.2.1 and Lemma A.2.2 we know  $\psi(\theta, w)$  is continuous in both  $\theta$  and  $w$  and  $C^*(\theta)$  is continuous in  $\theta$ . Note that

$$\mathcal{E} = \left\{ d_\infty((N_x(t)/t)_{x \in \mathcal{X}}, C^*(\theta)) \rightarrow 0 \wedge \hat{\theta}_t \rightarrow \theta \right\}$$

holds with probability 1 (Lemma 3, 5 and Proposition 1 in [Jedra and Proutiere, 2020]). Let  $\xi > 0$ , By continuity of  $\psi$ , there exists an open neighborhood  $\mathcal{V}(\xi)$  of  $\{\theta\} \times C^*(\theta)$  such that for all  $(\theta', w') \in \mathcal{V}(\xi)$ , it holds that

$$\psi(\theta', w') \geq (1 - \xi)\psi(\theta, w^*)$$

for any  $w^* \in C^*(\theta)$ . Under  $\mathcal{E}$ , there exists  $t_0 \geq 1$  such that for all  $t \geq t_0$  it holds that  $(\hat{\theta}_t, (N_x(t)/t)_{x \in \mathcal{X}}) \in \mathcal{V}(\xi)$ , this for all  $t \geq t_0$ , it follows that

$$\psi(\hat{\theta}_t, (N_x(t)/t)_{x \in \mathcal{X}}) \geq (1 - \xi)\psi(\theta, w^*)$$

By Lemma 2.5.2, there exists  $t_1 \geq 1$  such that for all  $t \geq t_1$  we have  $\lambda_{\min}(\mathbf{A}_t) > c_{\mathcal{X}_0} \sqrt{t - d - t} > \kappa_0 \log(t)$ , then  $\lambda_{\min}(\mathbf{A}_t) \geq \kappa_0 \log(t)$  which implies  $t \in B$ .

We also have that under  $\mathcal{E}$ , there exists  $t_2 \geq 1$  such that for all  $t \geq t_0$  it holds that  $\theta_t^{(1)} = \hat{\theta}_t$ . Then, we can write

$$\begin{aligned} Z(t) &= \inf_{\lambda \in \text{Alt}(\theta_t^{(1)})} \frac{1}{2} \|\theta_t^{(1)} - \lambda\|_{\mathbf{H}_t(\theta_t^{(1)})}^2 \\ &= t \inf_{\lambda \in \text{Alt}(\theta_t^{(1)})} \frac{1}{2} \|\theta_t^{(1)} - \lambda\|_{\mathbf{H}_{w_t}(\theta_t^{(1)})}^2 \\ &= t\psi(\hat{\theta}_t, (N_x(t)/t)_{x \in \mathcal{X}}) \end{aligned}$$

Hence, under  $\mathcal{E}$  and for  $t \geq \max\{t_0, t_1, t_2\}$

$$Z(t) \geq t(1 - \xi)\psi(\theta, w^*)$$

Then

$$\begin{aligned}
\tau_\delta &= \inf \{t \geq 1 : Z(t) > \beta(\delta, t), t \in B\} \\
&\leq \max\{t_0, t_1, t_2\} \vee \inf \{t \geq 1 : t(1 - \xi)\psi(\theta, w^*) > \beta(\delta, t)\} \\
&\lesssim \max \left\{ t_0, t_1, t_2, \frac{1}{1 - \xi} T^*(\theta) \left( \log\left(\frac{1}{\delta}\right) \right)^2 \right\}
\end{aligned}$$

Where the last inequality uses Lemma 8 in [Jedra and Proutiere, 2020]. Thus

$$\mathbb{P}_\theta \left( \limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\left(\log\left(\frac{1}{\delta}\right)\right)^2} \leq T^*(\theta) \right) = 1$$

□

### Theorem 2.6.2

Log Track-and-Stop satisfies the same sample complexity upper bound

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\theta[\tau]}{\left(\log\left(\frac{1}{\delta}\right)\right)^2} \lesssim T^*(\theta)$$

*Proof.* Let  $\varepsilon > 0$

**Step 1.** By continuity of  $\psi$  (see Lemma A.2.1), there exists  $\xi_1(\varepsilon) > 0$  such that for all  $\theta' \in \mathbb{R}^d$  and  $w' \in \Sigma$

$$\begin{cases} \|\theta' - \theta\| &\leq \xi_1(\varepsilon) \\ d_\infty(w', C^*(\theta)) &\leq \xi_1(\varepsilon) \end{cases} \implies |\psi(\theta, w^*) - \psi(\theta', w')| \leq \varepsilon \psi(\theta, w^*) = \varepsilon (T^*(\theta))^{-1} \quad (\text{A.5})$$

for any  $w^* \in \arg \min_{w \in C^*(\theta)} d_\infty(w', w)$  (we have  $w^* \in C^*(\theta)$ ). Furthermore, by the continuity properties of the correspondence  $C^*$  (see Lemma A.2.2), there exists  $\xi_2(\varepsilon) > 0$  such that for all  $\theta' \in \mathbb{R}^d$

$$\|\theta - \theta'\| \leq \xi_2(\varepsilon) \implies \max_{w'' \in C^*(\theta')} d_\infty(w'', C^*(\theta)) < \frac{\xi_1(\varepsilon)}{K - 1}$$

Additionally, let  $\xi_3(\varepsilon) \leq (S - \|\theta\|)/2$ , then

$$\|\theta - \hat{\theta}_t\| \leq \xi_3(\varepsilon) \implies \|\hat{\theta}_t\| \leq S \implies \hat{\theta}_t = \theta_t^{(1)}$$

Let  $\xi(\varepsilon) = \min(\xi_1(\varepsilon), \xi_2(\varepsilon), \xi_3(\varepsilon))$ . In the following, we construct  $T_0$ , and for each  $T \geq T_0$  an event  $\mathcal{E}_T$ , under which for all  $t \geq T$ , it holds

$$\|\theta - \hat{\theta}_t\| \leq \xi(\varepsilon) \implies d_\infty\left(\left(N_x(t)/t\right)_{x \in \mathcal{X}}, C^*(\theta)\right) \leq \xi_1(\varepsilon)$$

Let  $T \geq 1$ , and define the following event

$$\mathcal{E}_{1,T} = \bigcap_{t=T}^{\infty} \left\{ \|\theta - \hat{\theta}_t\| \leq \xi(\varepsilon) \right\}$$

Note that, under the event  $\mathcal{E}_{1,T}$ , we have for all  $t \geq T$

$$\begin{aligned} d_{\infty}(w(t), C^*(\theta)) &\leq \max_{w' \in C^*(\hat{\theta}_t)} d_{\infty}(w', C^*(\theta)) \\ &< \frac{\xi_1(\varepsilon)}{K-1} \end{aligned}$$

Define  $\varepsilon_1 = \xi_1(\varepsilon)/(K-1)$ . By Lemma A.2.3 (6 of Jedra and Proutiere [2020]), there exists  $t_1(\varepsilon_1) \geq T$  such that

$$d_{\infty}\left(\left(N_a(t)/t\right)_{a \in \mathcal{A}}, C^*(\mu)\right) \leq (p_t + d - 1) \frac{\xi_1(\varepsilon)}{K-1} \leq \xi_1(\varepsilon)$$

where  $p_t = \left| \text{supp} \left( \sum_{s=1}^t w(s) \right) \setminus \mathcal{X}_0 \right|$  and more precisely

$$t_1(\varepsilon_1) = \max \left\{ 1/\varepsilon_1^3, 1/(\varepsilon_1^2 d), T/\varepsilon_1^3, 10/\varepsilon_1 \right\}$$

(see the proof of Lemma 6 of Jedra and Proutiere [2020]). Thus for  $T \geq \max \{10\varepsilon_1^2, \varepsilon_1/d, 1\}$ , we have  $t_1(\varepsilon_1) = \lceil T/\varepsilon_1^3 \rceil$ . Hence, defining for all  $T \geq \varepsilon_1^{-3}$ , the event

$$\mathcal{E}_T = \mathcal{E}_{1, \lceil \varepsilon_1^3 T \rceil}$$

we have shown that for all  $T \geq T_0 = \max(10\varepsilon_1^5, \varepsilon_1^4/d, \varepsilon_1^3, 1/\varepsilon_1^3)$ , the following holds

$$\forall t \geq T, \quad \|\theta - \theta_t\| \leq \xi(\varepsilon) \implies d_{\infty}\left(\left(N_x(t)/t\right)_{x \in \mathcal{X}}, C^*(\theta)\right) \leq \xi_1(\varepsilon). \quad (\text{A.6})$$

Finally, combining the implication Eq. (A.6) with the fact that Eq. (A.5) holds under  $\mathcal{E}_T$  we conclude that for all  $T \geq T_0$ , under  $\varepsilon_T$  we have

$$\psi\left(\hat{\theta}_t, \left(N_x(t)/t\right)_{x \in \mathcal{X}}\right) \geq (1 - \varepsilon)\psi^*(\theta) \quad (\text{A.7})$$

**Step 2.:** Let  $T \geq T_0 \vee T_1 \vee T_2$  where  $T_1$  is defined as

$$T_1 = \inf \{t \geq 1 : t \in B\}.$$

and  $T_2$  is defined as

Under the event  $\mathcal{E}_T$ , for all  $t \geq T$  we have

$$Z(t) = t\psi\left(\hat{\theta}_t, (N_x(t)/t)_{x \in \mathcal{X}}\right),$$

Thus under the event  $\mathcal{E}_T$ , the inequality Eq. (A.7) holds, and for all  $t \geq T$  we have

$$Z(t) > t(1 - \varepsilon) (T^*(\theta))^{-1}$$

Under the event  $\mathcal{E}_T$ , we have

$$\begin{aligned} \tau &= \inf \{t \geq 1 : Z(t) > \beta(\delta, t), t \in B\} \\ &\leq \inf \{t \geq T : Z(t) > \beta(\delta, t)\} \\ &\leq \inf \left\{ t \geq T : t(1 - \varepsilon) (T^*(\theta))^{-1} > \sqrt{2}(1 + 2S) \left( \frac{\sqrt{\log(t)}}{2} + \frac{4}{\sqrt{\log(t)}} \log \left( \frac{2^d}{\delta} \left( \frac{Lt}{d} \right)^{\frac{d}{2}} \right) \right) \right\} \end{aligned}$$

Applying Lemma 8 in [Jedra and Proutiere, 2020] yields

$$\tau \leq T_2^*(\delta),$$

where  $T_2^*(\delta) = \frac{c_1}{1-\varepsilon} T^*(\theta) (\log(1/\delta))^2 + o(\log(1/\delta))$  for some  $0 < c_1$  independent of  $\delta$ . This means for  $T \geq \max \{T_0, T_1, T_2^*(\delta)\}$ , we have shown that

$$\mathcal{E}_T \subseteq \{\tau \leq T\} \tag{A.8}$$

Define  $T_3^*(\delta) = \max \{T_0, T_1, T_2^*(\delta)\}$ . We may then write for all  $T \geq T_3^*(\delta)$

$$\tau \leq \tau \wedge T_3^*(\delta) + \tau \vee T_3^*(\delta) \leq T_3^*(\delta) + \tau \vee T_3^*(\delta).$$

Taking the expectation of the above inequality, and using the set inclusion Eq. (A.8), we obtain that

$$\mathbb{E}[\tau] \leq T_3^*(\delta) + \mathbb{E}[\tau \vee T_3^*(\delta)]$$

Now we observe that

$$\begin{aligned}
\mathbb{E}[\tau \vee T_3^*(\delta)] &= \sum_{T=0}^{\infty} \mathbb{P}(\tau \vee T_3^*(\delta) > T) \\
&= \sum_{T=T_3^*(\delta)+1}^{\infty} \mathbb{P}(\tau \vee T_3^*(\delta) > T) \\
&= \sum_{T=T_3^*(\delta)+1}^{\infty} \mathbb{P}(\tau > T) \\
&\leq \sum_{T=T_3^*(\delta)+1}^{\infty} \mathbb{P}(\mathcal{E}_T^c) \\
&\leq \sum_{T=T_0 \vee T_1}^{\infty} \mathbb{P}(\mathcal{E}_T^c)
\end{aligned}$$

We have thus shown that

$$\mathbb{E}[\tau] \leq \frac{c_1}{1-\varepsilon} T^*(\theta) (\log(1/\delta))^2 + \mathcal{O}(\log(1/\delta)) + T_0 \vee T_1 + \sum_{T=T_0 \vee T_1}^{\infty} \mathbb{P}(\mathcal{E}_T^c). \quad (\text{A.9})$$

**Step 3:** We now show that  $\sum_{T=T_0 \vee T_1+1}^{\infty} \mathbb{P}(\mathcal{E}_T^c) < \infty$  and that it can be upper bounded by a constant independent of  $\delta$ . Let  $T \geq T_0 \vee T_1$ , we have

$$\mathbb{P}(\mathcal{E}_T^c) \leq \mathbb{P}\left(\mathcal{E}_{1, \lceil \varepsilon_1^3 T \rceil}^c\right).$$

We observe, using a union bound, Lemma 5 from Jedra and Proutiere [2020] and Lemma 2.5.4, that there exists strictly positive constants  $c_3, c_4$  that are independent of  $\varepsilon$  and  $T$ , and such that

$$\begin{aligned}
\mathbb{P}\left(\mathcal{E}_{1, \lceil \varepsilon_1^3 T \rceil}^c\right) &\leq \sum_{t=\ell(\lceil \varepsilon_1^3 T \rceil)}^{\infty} \mathbb{P}\left(\|\theta - \hat{\theta}_t\| > \xi(\varepsilon)\right) \\
&\leq \sum_{t=\ell(\lceil \varepsilon_1^3 T \rceil)}^{\infty} c_2 t^{d/2} \exp\left(-c_1 \xi(\varepsilon) t^{1/2}\right)
\end{aligned}$$

For  $t$  large enough, the function  $t \mapsto t^{d/2} \exp\left(-c_1 \xi(\varepsilon) \sqrt{t}\right)$  becomes decreasing. Hence, for  $T \geq T_2$ , we have

$$\mathbb{P}\left(\mathcal{E}_{2, \lceil \varepsilon_1^3 T \rceil}^c\right) \leq c_3 \int_{\lceil \varepsilon_1^3 T \rceil - 1}^{\infty} t^{d/2} \exp\left(-c_1 \xi(\varepsilon) \sqrt{t}\right) dt.$$

Furthermore, for some  $T_3 \geq T_2$  large enough, we may bound the integral for all  $T \geq T_3$

as follows

$$\int_{\ell([\varepsilon_1^3 T])^{-1}}^{\infty} t^{d/4} \exp(-c_1 \xi(\varepsilon)^2 \sqrt{t}) dt \lesssim \frac{([\varepsilon_1^3 T] - 1)^{d/2+1}}{\xi(\varepsilon)^4 \exp(c_4 \xi(\varepsilon)^2 \sqrt{[\varepsilon_1^3 T] - 1})}.$$

We spare the details of this derivation as the constants are irrelevant in our analysis. Essentially, the integral can be expressed through the upper incomplete Gamma function which can be upper bounded using some classical inequalities [23, 24]. We then obtain that for  $T \geq T_3$ ,

$$\mathbb{P}\left(\mathcal{E}_{1, [\varepsilon_1^3 T]}^c\right) \lesssim \frac{([\varepsilon_1^3 T] - 1)^{d/2+1}}{\xi(\varepsilon)^4 \exp(c_4 \xi(\varepsilon)^2 \sqrt{[\varepsilon_1^3 T] - 1})}.$$

Thus there exists  $T_4 \geq T_3$  such that for all  $T \geq T_4$ ,

$$\mathbb{P}\left(\mathcal{E}_{2, [\varepsilon_1^3 T]}^c\right) \lesssim \frac{\ell([\varepsilon_1^3 T] - 1)^{d/2+1}}{\xi(\varepsilon)^4 \exp(c_4 \xi(\varepsilon)^2 \sqrt{\ell([\varepsilon_1^3 T] - 1)})} \lesssim \frac{T^{d/2+1}}{\exp(c_5(\varepsilon) T^{\gamma/2})}.$$

This shows that

$$\begin{aligned} \sum_{T=T_0 \vee T_1}^{\infty} \mathbb{P}\left(\mathcal{E}_{1, [\varepsilon_1^3 T]}^c\right) &= \sum_{T=T_0 \vee T_1}^{T_4} \mathbb{P}\left(\mathcal{E}_{1, [\varepsilon_1^3 T]}^c\right) + \sum_{T=T_4+1}^{\infty} \mathbb{P}\left(\mathcal{E}_{1, [\varepsilon_1^3 T]}^c\right) \\ &\lesssim \sum_{T=T_0 \vee T_1}^{T_4} \mathbb{P}\left(\mathcal{E}_{1, [\varepsilon_1^3 T]}^c\right) + \sum_{T=T_4+1}^{\infty} \frac{T^{d/2+1}}{\exp(c_5(\varepsilon) T^{\gamma/2})} \\ &< \infty \end{aligned}$$

where the last inequality follows from the fact that we can upper bound the infinite sum by a Gamma function, which is convergent as long as  $\gamma > 0$ .

Finally, we have thus shown that

$$\sum_{T=T_0 \vee T_1+1}^{\infty} \mathbb{P}(\mathcal{E}_T^c) < \infty$$

We note that this infinite sum depends  $\varepsilon$  only.

**Last step:** Finally, we have shown that for all  $\varepsilon > 0$

$$\mathbb{E}[\tau] \leq \frac{c_1}{1-\varepsilon} T^*(\theta) (\log(1/\delta))^2 + \mathcal{O}(\log(1/\delta)) + T_0 \vee T_1 + \sum_{T=T_0 \vee T_1}^{\infty} \mathbb{P}(\mathcal{E}_T^c)$$

where  $\sum_{T=T_0 \vee T_1}^{\infty} \mathbb{P}(\mathcal{E}_T^c) < \infty$  and is independent of  $\delta$ . Hence,

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_{\theta}[\tau_{\delta}]}{(\log(1/\delta))^2} \lesssim \frac{c_1}{1-\varepsilon} T^*(\theta)$$

Letting  $\varepsilon$  tend to 0, we conclude that

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_{\theta}[\tau_{\delta}]}{(\log(1/\delta))^2} \lesssim T^*(\theta)$$

□

### A.3 Examples

We present three concrete examples of pure exploration problems under logistic bandits. First we need an auxiliary Lemma.

**Lemma A.3.1.** (Lemma 5 [Degenne et al., 2020]) For  $\theta, \lambda \in \mathbb{R}^d$ ,  $w$  in the interior of the probability simplex  $\Sigma$ ,  $y \in \mathbb{R}^d$ ,  $x \in \mathbb{R}$ , we have

$$\inf_{\lambda: \lambda^{\top} y \geq x} \frac{\|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2}{2} = \begin{cases} \frac{(x - \theta^{\top} y)^2}{2\|y\|_{\mathbf{H}_w(\theta)^{-1}}^2} & \text{if } x \geq \theta^{\top} y \\ 0 & \text{otherwise} \end{cases}.$$

And

$$\inf_{\lambda: \lambda^{\top} y \leq x} \frac{\|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2}{2} = \begin{cases} \frac{(x - \theta^{\top} y)^2}{2\|y\|_{\mathbf{H}_w(\theta)^{-1}}^2} & \text{if } x \leq \theta^{\top} y \\ 0 & \text{otherwise} \end{cases}.$$

*Proof.* We consider the Lagrangian of the problem, and we obtain

$$\begin{aligned} \inf_{\lambda: \lambda^{\top} y \geq x} \frac{\|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2}{2} &= \sup_{\alpha \geq 0} \inf_{\lambda \in \mathbb{R}^d} \frac{\|\theta - \lambda\|_{\mathbf{H}_w(\theta)}^2}{2} + \alpha(x - \lambda^{\top} y) \\ &= \sup_{\alpha \geq 0} \alpha(x - \theta^{\top} y) - \alpha^2 \frac{\|y\|_{\mathbf{H}_w(\theta)}^2}{2} \\ &= \begin{cases} \frac{(x - \theta^{\top} y)^2}{2\|y\|_{\mathbf{H}_w(\theta)^{-1}}^2} & \text{if } x \geq \theta^{\top} y \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where the infimum in the first equality is reached at  $\lambda = \theta + \alpha \mathbf{H}_w^{-1}(\theta)y$  and the supremum in the last equality is reached at  $\alpha = (x - \langle \theta, y \rangle) / \|y\|_{\mathbf{H}_w^{-1}(\theta)}^2$  if  $x \geq \langle \theta, y \rangle$  and at  $\alpha = 0$  else. The second equality can be solved with the same steps. □

### A.3.1 Best arm identification

In this example  $i^*(\theta) = \arg \max_{x \in \mathcal{X}} \{\mu(x^\top \theta)\} = \arg \max_{x \in \mathcal{X}} \{x^\top \theta\}$ . Lets define  $x^*(\theta) := i^*(\theta)$ .

**Lemma 2.7.1** For all  $\theta \in \mathbb{R}^d$  such that  $i^*(\theta)$  is unique,

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \min_{x \neq x^*(\theta)} \frac{(\theta^\top x^*(\theta) - \theta^\top x)^2}{2 \|x^*(\theta) - x\|_{\mathbf{H}_w^{-1}}^2}$$

*Proof.* Recall that the characteristic time is given by

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2}$$

$$\begin{aligned} T^*(\theta)^{-1} &= \max_{w \in \Sigma} \min_{x \neq x^*(\theta)} \inf_{\lambda: \lambda^\top x > \lambda^\top x^*(\theta)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \\ &= \max_{w \in \Sigma} \min_{x \neq x^*(\theta)} \frac{(\theta^\top x^*(\theta) - \theta^\top x)^2}{2 \|x^*(\theta) - x\|_{\mathbf{H}_w^{-1}}^2} \end{aligned}$$

□

### A.3.2 Thresholding bandit problem

In this example, given  $\rho \in (0, 1)$ ,  $i^*(\theta) = \{x \in \mathcal{X} : \mu(x^\top \theta) > \rho\} = \{x \in \mathcal{X} : x^\top \theta > \mu^{-1}(\rho)\}$ .

**Lemma 2.7.2** For all  $\theta \in \mathbb{R}^d$  such that  $i^*(\theta)$  is unique,

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \min_{x \in \mathcal{X}} \frac{(\theta^\top x - \mu^{-1}(\rho))^2}{2 \|x\|_{\mathbf{H}_w^{-1}}^2}$$

*Proof.*

$$\begin{aligned} T^*(\theta)^{-1} &= \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \\ &= \max_{w \in \Sigma} \min \left( \min_{x: \mu(\theta^\top x) < \rho} \inf_{\lambda: \mu(x^\top \lambda) > \rho} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2}, \min_{x: \mu(\theta^\top x) > \rho} \inf_{\lambda: \mu(x^\top \lambda) < \rho} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \right) \\ &= \max_{w \in \Sigma} \min \left( \min_{x: \mu(\theta^\top x) < \rho} \inf_{\lambda: x^\top \lambda > \mu^{-1}(\rho)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2}, \min_{x: \mu(\theta^\top x) > \rho} \inf_{\lambda: x^\top \lambda < \mu^{-1}(\rho)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \right) \\ &= \max_{w \in \Sigma} \min_{x \in \mathcal{X}} \frac{(\theta^\top x - \mu^{-1}(\rho))^2}{2 \|x\|_{\mathbf{H}_w^{-1}}^2} \end{aligned}$$

□

### A.3.3 Best-M arm identification

In this example, given  $M \in [K - 1]$ ,  $i^*(\theta) = \{x_{(i)}(\theta) : i \in [M]\}$  where  $x_{(1)}(\theta), \dots, x_{(K)}(\theta)$  are the ordered arms with respect to their expected rewards.

**Lemma A.3.2.** *For all  $\theta \in \mathbb{R}^d$  such that  $x_{(1)}(\theta)^\top \theta < \dots < x_{(K)}(\theta)^\top \theta$ ,*

$$T^*(\theta)^{-1} = \max_{w \in \Sigma} \min_{x \neq x_{(M)}(\theta)} \frac{(\theta^\top x - \theta^\top x_{(M)}(\theta))^2}{2 \|x\|_{\mathbf{H}_w^{-1}}^2}$$

*Proof.*

$$\begin{aligned} T^*(\theta)^{-1} &= \max_{w \in \Sigma} \inf_{\lambda \in \text{Alt}(\theta)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \\ &= \max_{w \in \Sigma} \min \left( \min_{x: x \notin i^*(\theta)} \inf_{\lambda: \mu(x^\top \lambda) \geq \mu(x_{(M)}(\theta)^\top \lambda)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2}, \right. \\ &\quad \left. \min_{x: x \in i^*(\theta)} \inf_{\lambda: \mu(x^\top \lambda) < \mu(x_{(M)}(\theta)^\top \lambda)} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \right) \\ &= \max_{w \in \Sigma} \min \left( \min_{x: x \notin i^*(\theta)} \inf_{\lambda: x^\top \lambda \geq x_{(M)}(\theta)^\top \lambda} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2}, \min_{x: x \in i^*(\theta)} \inf_{\lambda: x^\top \lambda < x_{(M)}(\theta)^\top \lambda} \frac{\|\theta - \lambda\|_{\mathbf{H}_w}^2}{2} \right) \\ &= \max_{w \in \Sigma} \min_{x \neq x_{(M)}(\theta)} \frac{(\theta^\top x - \theta^\top x_{(M)}(\theta))^2}{2 \|x\|_{\mathbf{H}_w^{-1}}^2} \end{aligned}$$

□

## A.4 Useful lemmas

**Lemma A.4.1.** *(Determinant-Trace inequality). Let  $\{x_s\}_{s=1}^\infty$  a sequence in  $\mathbb{R}^d$  such that  $\|x_s\|_2 \leq X$  for all  $s \in \mathbb{N}$ , and let  $\lambda$  be a non-negative scalar. For  $t \geq 1$  define  $\mathbf{A}_t := \sum_{s=1}^t x_s x_s^\top$ . The following inequality holds:*

$$\det(\mathbf{A}_t) \leq (tX^2/d)^d$$

## APPENDIX B

### Conformal Prediction for Ensembles: Improving Efficiency via Score-Based Aggregation

#### B.1 CSA Visual Walkthrough

We walk through a visual presentation of the approach below to supplement the textual description in the main text. We start with a collection of multivariate calibration scores  $\mathcal{S}_C$  with  $s \in \mathcal{S}$  being  $\in \mathbb{R}^K$ . For the purposes of visualization in this section, we have  $K = 2$ . We first partition the score evaluations  $\mathcal{S}_C = \mathcal{S}_C^{(1)} \cup \mathcal{S}_C^{(2)}$ , with a subset  $\mathcal{S}_C^{(1)}$  used to define the pre-ordering and the remainder  $\mathcal{S}_C^{(2)}$  to define the multivariate quantile.




Figure B.1: The calibration score evaluations are first split between those used to define the pre-ordering (green)  $\mathcal{S}_C^{(1)}$  and those used to define the final multivariate quantile (red)  $\mathcal{S}_C^{(2)}$ .

We first wish to define the pre-ordering over  $\mathcal{S}_C^{(1)}$ . As described in the main text, the goal is to define this using an indexed family of sets  $\mathcal{A}_t$  with index  $t \in \mathbb{R}$ , after which the multivariate quantile approach reduces to the univariate quantile formulation. To ensure the final envelope over  $\mathcal{S}_C^{(2)}$  remains as tight as possible, we wish to define this family in a data-driven fashion. Critically, the *shape* of this tightest envelope around  $\mathcal{S}_C^{(2)}$  will vary across  $\alpha$ , meaning we must define the family *separately* for each choice of  $\alpha$ . We expect the

contour of the tightest  $\alpha$  envelope for  $\mathcal{S}_C^{(1)}$  will be similar to that over  $\mathcal{S}_C^{(2)}$ , motivating such a choice to define the indexing family. To do this, we project  $\mathcal{S}_C^{(1)}$  along a number of directions, finding the  $\beta$  quantile along each, in turn defining a half-plane, where  $\beta$  is as described in 4.3.1.




Figure B.2: The pre-ordering points are projected across a number of directions, after which the  $\beta$  quantile is used to define a direction quantile. This defines a half-plane of points that are in the region (blue) and those outside (red).

We then iteratively update  $\beta$  in the manner described in 2 to obtain  $\beta^*$ , namely the minimum value for which the region given by the intersection of the corresponding half-planes covers roughly  $1 - \alpha$  of  $\mathcal{S}_C^{(1)}$ .




Figure B.3: We use the intersection of hyperplanes to define the quantile envelope, seeking  $\beta^*$  that achieves the desired coverage.

Once this  $1 - \alpha$  quantile envelope of  $\mathcal{S}_C^{(1)}$  is found, we define  $\mathcal{A}_1$  to be such an envelope, with which future points can now be partially ordered. That is, for any point  $s \in \mathbb{R}^K$  notice that we can unambiguously associate it with  $t(s) := \min\{t \in \mathbb{R} : s \in \mathcal{A}_t\}$ . Intuitively, this is the  $t$  where the contour “intersects”  $s$ . Notably, now that the partial ordering has been defined, the points of  $\mathcal{S}_C^{(1)}$  are no longer used. It would be of interest to investigate whether a concurrent definition of the partial ordering and final calibration is possible without such data splitting in future work.




Figure B.4: Using the quantile envelope, the family of nested sets  $\mathcal{A}_t$  is defined, in turn defining a partial ordering over  $\mathbb{R}^K$ .

With this  $\mathcal{A}_t$ , we find the final  $\hat{q}$  simply by mapping the points of  $\mathcal{S}_C^{(2)}$  to their corresponding  $t(s)$  values in the aforementioned fashion and performing standard conformal prediction. As discussed, if the envelope has a similar structure to that found over  $\mathcal{S}_C^{(1)}$ , the envelope should be adjusted by only a minor amount.




Figure B.5: Using the nested family of sets, we expand or contract the envelope appropriately using the data of  $\mathcal{S}_C^{(2)}$  to find the final adjustment factor.

## B.2 Multivariate Score Coverage

The proof of the multivariate extension of conformal prediction follows in precisely the same manner as that of standard conformal prediction with the pre-order  $\lesssim$  replacing the complete ordering used in traditional conformal prediction.

**Theorem B.2.1.** *Suppose  $\mathcal{D}_C := \{(X_i, Y_i)\}_{i=1}^{N_C}$  and  $(X', Y')$  are exchangeable. Assume further that  $K$  maps  $s_k : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  have been defined and a composite  $s(X, Y) := (s_1(X, Y), \dots, s_K(X, Y))$  is defined. Further denote by  $\mathcal{S}_C$  the evaluation of  $s(X, Y)$  on  $\mathcal{D}_C$ , namely  $\mathcal{S}_C := \{s(X_i, Y_i) \mid (X_i, Y_i) \in \mathcal{D}_C\}$ . For some  $\alpha \in (0, 1)$ , given a pre-order  $\lesssim$  in  $\mathbb{R}^K$  induced by a collection of nested sets  $\{\mathcal{A}_t\}_{t \geq 0}$  define  $\mathcal{Q}(\alpha) = \{s \in \mathbb{R}^K : s \lesssim s_{\lceil (N_C+1)(1-\alpha) \rceil}\}$ . Then, denoting  $\mathcal{C}(x) := \{y : s(x, y) \in \mathcal{Q}(\alpha)\}$ ,*

$$\mathcal{P}_{X', Y'}(Y' \in \mathcal{C}(X')) \geq 1 - \alpha \quad (\text{B.1})$$

*Proof.* Denote  $s_i = s(X_i, Y_i)$  for each  $i = 1, \dots, n$  and  $s' = s(X', Y')$ . We define  $t_i = \inf\{t \geq 0 : s_i \in \mathcal{A}_t\}$  and  $t' = \inf\{t \geq 0 : s' \in \mathcal{A}_t\}$ . We consider the case that  $\mathcal{P}_{X, Y}(t_i \neq t_j) = 1$ , that is, that the probability of ties is a probability measure 0 set. Without loss of generality, we then assume the scores are sorted according to the assumed pre-order, namely that  $s_1 \lesssim s_2 \lesssim \dots \lesssim s_{N_C}$ , or equivalently  $t_1 \leq t_2 \leq \dots \leq t_{N_C}$ . We then again have that  $\hat{t} = t_{\lceil (N_C+1)(1-\alpha) \rceil}$  if  $\alpha > 1/(N_C + 1)$  and  $\hat{t} = \infty$  otherwise. In the latter case, coverage is trivially satisfied. In the former case, we see

$$\mathcal{P}_{X, Y}(Y' \in \mathcal{C}(X')) = \mathcal{P}_{X, Y}(s(X, y) \lesssim s_{\lceil (N_C+1)(1-\alpha) \rceil}) = \mathcal{P}_{X, Y}(t' \leq t_{\lceil (N_C+1)(1-\alpha) \rceil}). \quad (\text{B.2})$$

By the assumed exchangeability of  $\mathcal{D}_C := \{(X_i, Y_i)\}_{i=1}^{N_C}$  and  $(X', Y')$ , we have that

$$\mathcal{P}_{X, Y}(t' \leq t_k) = \frac{k}{N_C + 1}, \quad (\text{B.3})$$

for any  $k$ . From here, we have the desired conclusion that

$$\mathcal{P}_{X, Y}(s(X, y) \lesssim s_{\lceil (N_C+1)(1-\alpha) \rceil}) = \left(\frac{1}{N_C + 1}\right) (\lceil (N_C + 1)(1 - \alpha) \rceil) \geq 1 - \alpha, \quad (\text{B.4})$$

completing the proof as desired.  $\square$

### B.3 Conformal Aggregation Methods

We now describe the methods from Gasparin and Ramdas [2024a] that were compared against experimentally, specifically the standard majority-vote  $\mathcal{C}^M$ , partially randomized thresholding  $\mathcal{C}^R$ , and fully randomized thresholding  $\mathcal{C}^U$  approaches. As discussed in 4.2.3, these methods all follow the structural form of 4.2 and largely differ in their choice of weights and thresholds. The standard majority-vote  $\mathcal{C}^M$  is the most natural choice, defined by

$$\mathcal{C}^M(x) := \left\{ y \mid \frac{1}{K} \sum_{k=1}^K \mathbb{1}[y \in \mathcal{C}_k(x)] > \frac{1}{2} \right\}. \quad (\text{B.5})$$

The randomized methods differ in that independent randomization is leveraged over the threshold, namely with:

$$\mathcal{C}^R(x) := \left\{ y \mid \frac{1}{K} \sum_{k=1}^K \mathbb{1}[y \in \mathcal{C}_k(x)] > \frac{1}{2} + \frac{U}{2} \right\} \quad (\text{B.6})$$

$$\mathcal{C}^U(x) := \left\{ y \mid \frac{1}{K} \sum_{k=1}^K \mathbb{1}[y \in \mathcal{C}_k(x)] > U \right\}, \quad (\text{B.7})$$

for  $U \sim \text{Unif}([0, 1])$ . Notably, all these methods retain the guarantees typical of conformal prediction.

## APPENDIX C

### Online Conformal Prediction: Enforcing Monotonicity via Online Optimization

#### C.1 Proofs

##### Proof Proposition 2

*Proof.* Since  $\tilde{w}_i > 0$  for all  $i$ , the map

$$w \mapsto D_{\text{KL}}(w \parallel \tilde{w}) = \sum_{i=1}^n w_i \log \frac{w_i}{\tilde{w}_i}$$

is strictly convex on  $\mathbb{R}_+^n$ . As  $\mathcal{X}$  is a nonempty closed convex set, the minimizer exists and is unique.

Consider the constrained problem

$$\min_{w \in \mathbb{R}^n} \sum_{i=1}^n w_i \log \frac{w_i}{\tilde{w}_i} \quad \text{s.t.} \quad w_i \geq \mu, \quad \sum_{i=1}^n w_i = 1.$$

Its Lagrangian is

$$\mathcal{L}(w, \lambda, \nu) = \sum_{i=1}^n w_i \log \frac{w_i}{\tilde{w}_i} + \lambda \left( \sum_{i=1}^n w_i - 1 \right) + \sum_{i=1}^n \nu_i (\mu - w_i),$$

where  $\nu_i \geq 0$ . Let  $w^*$  be the optimizer. By the KKT conditions,

$$\log \frac{w_i^*}{\tilde{w}_i} + 1 + \lambda - \nu_i = 0, \quad i = 1, \dots, n.$$

Hence

$$w_i^* = \tilde{w}_i e^{-1-\lambda+\nu_i}.$$

Now, if  $w_i^* > \mu$ , then complementary slackness gives  $\nu_i = 0$ , and therefore

$$w_i^* = c \tilde{w}_i, \quad \text{where } c := e^{-1-\lambda} > 0.$$

If instead  $w_i^* = \mu$ , then the lower-bound constraint is active. Thus, for every  $i$ ,

$$w_i^* = \max\{\mu, c \tilde{w}_i\}.$$

Finally, the equality constraint  $\sum_{i=1}^n w_i^* = 1$  implies that  $c > 0$  must satisfy

$$\sum_{i=1}^n \max\{\mu, c \tilde{w}_i\} = 1.$$

Since the left-hand side is continuous and strictly increasing in  $c$ , such a  $c$  exists and is unique. This proves the claim.  $\square$

**Proposition 3.** (Proposition 5 in Gibbs and Candès [2024]) *Let  $s$  be a random variable and assume that there exists a value  $q^*$  such that  $\mathbb{P}(s < q^*) = \alpha$ . Then, for any  $q$ ,*

$$\mathbb{E}[\rho_t(q, 1 - \alpha)] - \mathbb{E}[\rho_t(q^*, 1 - \alpha)] = \begin{cases} \mathbb{E}[(q - s)\mathbb{1}_{q < s \leq q^*}], & \text{if } q \leq q^* \\ \mathbb{E}[(s - q)\mathbb{1}_{q^* < s \leq q}], & \text{if } q^* \leq q \end{cases}$$

So, in particular, if  $s$  has a density  $p(\cdot)$  on  $[0, B]$  with  $p(x) \geq p > 0$  for all  $x \in [0, B]$ , then

$$\mathbb{E}[\rho_t(q, 1 - \alpha)] - \mathbb{E}[\rho_t(q^*, 1 - \alpha)] \geq \frac{p(q - q^*)^2}{2}$$

**Lemma C.1.1.** *Let  $s_1, \dots, s_T$  be an arbitrary sequence in  $[0, B]$ , then the EG update rule has the following regret bound.*

$$\sum_{t=1}^T g_t(w_t) - \sum_{t=1}^T g_t(w_t^*) \leq \frac{(1 + \log(1/\mu))}{\eta} \left( 1 + \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1 \right) + \eta G_\infty^2 T \quad (\text{C.1})$$

*Proof.* Since the projected exponentiated gradient update is exactly the mirror descent update with negative entropy, we may write

$$w_{t+1} = \arg \min_{w \in \Delta_{K+1}^\mu} \{\eta \langle \nabla g_t, w \rangle + D_{\text{KL}}(w \| w_t)\}.$$

We first prove the standard one-step mirror descent inequality.

**Step 1: One-step inequality.** Fix  $t$  and let  $w_t^* \in \Delta_{K+1}^\mu$

$$\begin{aligned} \eta \langle \nabla g_t, w_t - w_t^* \rangle &= \langle \nabla \phi(w_t) - \nabla \phi(w_{t+1}) - \eta \nabla g_t, w_t^* - w_{t+1} \rangle \\ &\quad + \langle \nabla \phi(w_{t+1}) - \nabla \phi(w_t), w_t^* - w_{t+1} \rangle \\ &\quad + \langle \eta \nabla g_t, w_t - w_{t+1} \rangle \end{aligned}$$

By optimality of  $w_{t+1}$  over the convex set  $\Delta_{K+1}^\mu$ ,

$$\langle \eta \nabla g_t + \nabla \phi(w_{t+1}) - \nabla \phi(w_t), w_t^* - w_{t+1} \rangle \geq 0.$$

Therefore,

$$\eta \langle \nabla g_t, w_t - w_t^* \rangle \leq \langle \eta \nabla g_t, w_t - w_{t+1} \rangle + \langle \nabla \phi(w_{t+1}) - \nabla \phi(w_t), w_t^* - w_{t+1} \rangle$$

Using the three-point identity for Bregman divergences,

$$\langle \nabla \phi(w_{t+1}) - \nabla \phi(w_t), w_t^* - w_{t+1} \rangle = D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) - D_{\text{KL}}(w_{t+1} \| w_t),$$

Hence

$$\eta \langle \nabla g_t, w_t - w_t^* \rangle \leq \eta \langle \nabla g_t, w_t - w_{t+1} \rangle + D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) - D_{\text{KL}}(w_{t+1} \| w_t).$$

Using Hölder's inequality,

$$\begin{aligned} \eta \langle \nabla g_t, w_t - w_t^* \rangle &\leq \eta \langle \nabla g_t, w_t - w_{t+1} \rangle + D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) - D_{\text{KL}}(w_{t+1} \| w_t) \\ &\leq \eta \|\nabla g_t\|_\infty \|w_t - w_{t+1}\|_1 + D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) - D_{\text{KL}}(w_{t+1} \| w_t) \\ &\leq \frac{\eta^2}{2} \|\nabla g_t\|_\infty^2 + \frac{1}{2} \|w_t - w_{t+1}\|_1^2 + D_{\text{KL}}(w_t^* \| w_t) \\ &\quad - D_{\text{KL}}(w_t^* \| w_{t+1}) - D_{\text{KL}}(w_{t+1} \| w_t) \\ &\leq \frac{\eta^2}{2} \|\nabla g_t\|_\infty^2 + D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) \end{aligned}$$

Since  $\|\nabla g_t\|_\infty \leq G_\infty$ ,

$$\eta \langle \nabla g_t, w_t - w_t^* \rangle \leq D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) + \frac{\eta^2 G_\infty^2}{2}.$$

**Step 2: Sum over time.** Summing over  $t = 1, \dots, T$  yields

$$\eta \sum_{t=1}^T \langle \nabla g_t, w_t - w_t^* \rangle \leq \sum_{t=1}^T \left( D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right) + \frac{\eta^2 G_\infty^2 T}{2}.$$

We rewrite the KL terms as

$$\begin{aligned} \sum_{t=1}^T \left( D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right) &= \sum_{t=1}^T \left( D_{\text{KL}}(w_t^* \| w_t) - D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) \right. \\ &\quad \left. + D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right) \\ &= D_{\text{KL}}(w_1^* \| w_1) - D_{\text{KL}}(w_{T+1}^* \| w_{T+1}) \\ &\quad + \sum_{t=1}^{T-1} \left( D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right). \end{aligned}$$

Since KL divergence is nonnegative,

$$-D_{\text{KL}}(w_{T+1}^* \| w_{T+1}) \leq 0,$$

so

$$\eta \sum_{t=1}^T \langle \nabla g_t, w_t - w_t^* \rangle \leq D_{\text{KL}}(w_1^* \| w_1) + \sum_{t=1}^{T-1} \left( D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right) + \frac{\eta^2 G_\infty^2 T}{2}.$$

**Step 3: KL is Lipschitz in the first argument on  $\Delta_{K+1}^\mu$ .** We have

$$D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) - D_{\text{KL}}(w_t^* \| w_{t+1}) \leq (1 + \log(1/\mu)) \|w_{t+1}^* - w_t^*\|_1.$$

Therefore,

$$\sum_{t=1}^{T-1} \left( D_{\text{KL}}(w_{t+1}^* \| w_{t+1}) - D_{\text{KL}}(w_t^* \| w_{t+1}) \right) \leq (1 + \log(1/\mu)) \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1.$$

**Step 4: Conclude the regret bound.** Combining the previous inequalities,

$$\eta \sum_{t=1}^T \langle \nabla g_t, w_t - w_t^* \rangle \leq D_{\text{KL}}(w_1^* \| w_1) + (1 + \log(1/\mu)) \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1 + \frac{\eta^2 G_\infty^2 T}{2}.$$

Dividing by  $\eta$  gives

$$\sum_{t=1}^T \langle \nabla g_t, w_t - w_t^* \rangle \leq \frac{D_{\text{KL}}(w_1^* \| w_1)}{\eta} + \frac{1 + \log(1/\mu)}{\eta} \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1 + \frac{\eta G_\infty^2 T}{2}.$$

Finally, since  $\nabla g_t \in \partial g_t(w_t)$  and  $g_t$  is convex,

$$g_t(w_t) - g_t(w_t^*) \leq \langle \nabla g_t, w_t - w_t^* \rangle.$$

Then

$$\sum_{t=1}^T (g_t(w_t) - g_t(w_t^*)) \leq \frac{D_{\text{KL}}(w_1^* \| w_1)}{\eta} + \frac{1 + \log(1/\mu)}{\eta} \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1 + \frac{\eta G_\infty^2 T}{2}.$$

To bound the initial divergence, note that for any  $u, w \in \Delta_{K+1}^\mu$ ,

$$D_{\text{KL}}(u \| w) = \sum_{i=1}^N u_i \log \frac{u_i}{w_i} \leq \sum_{i=1}^N u_i \log \frac{1}{\mu} = \log(1/\mu).$$

Thus,

$$D_{\text{KL}}(w_1^* \| w_1) \leq 1 + \log(1/\mu),$$

which gives the simplified bound.  $\square$

**Lemma C.1.2.** (Theorem 10.1. in Hazan [2016]) Let  $s_1, \dots, s_T$  be an arbitrary sequence in  $[0, B]$ , then the projected gradient on  $\mathcal{Q}$  update rule has the following regret bound.

$$\sum_{t=1}^T f_t(q_t) - \sum_{t=1}^T f_t(q_t^*) \leq \frac{3D^2}{\eta} \left( 1 + \sum_{t=1}^{T-1} \|q_{t+1}^* - q_t^*\|_1 \right) + \eta G_2^2 T \quad (\text{C.2})$$

Where  $D$  is the diameter of  $\mathcal{Q}$  and  $\|\nabla f\|_2 \leq G_2$ .

### Proof Theorem 6.4.1

*Proof.* We can apply Proposition 3 and Lemma C.1.1 so

$$\begin{aligned}
\sum_{t=1}^T \frac{p \|q_t - q_t^*\|_2^2}{2} &= \sum_{t=1}^T \sum_{i=1}^K \frac{p (q_{t,i} - q_{t,i}^*)^2}{2} \\
&\leq \sum_{t=1}^T \sum_{i=1}^K \mathbb{E} [\rho_t(q_{t,i}, 1 - \alpha_i)] - \mathbb{E} [\rho_t(q_{t,i}^*, 1 - \alpha_i)] \\
&= \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=1}^K (\rho_t(q_{t,i}, 1 - \alpha_i) - \rho_t(q_{t,i}^*, 1 - \alpha_i)) \right] \\
&= \mathbb{E} \left[ \sum_{t=1}^T g_t(w_t) - g_t(w_t^*) \right] \\
&\leq \frac{(1 + \log(1/\mu))}{\eta} \left( 1 + \sum_{t=1}^{T-1} \|w_{t+1}^* - w_t^*\|_1 \right) + \eta(BK)^2 T
\end{aligned}$$

□

### Proof Theorem 6.4.2

*Proof.* We can apply Proposition 3 and Lemma C.1.2 so

$$\begin{aligned}
\sum_{t=1}^T \frac{p \|q_t - q_t^*\|_2^2}{2} &= \sum_{t=1}^T \sum_{i=1}^K \frac{p (q_{t,i} - q_{t,i}^*)^2}{2} \\
&\leq \sum_{t=1}^T \sum_{i=1}^K \mathbb{E} [\rho_t(q_{t,i}, 1 - \alpha_i)] - \mathbb{E} [\rho_t(q_{t,i}^*, 1 - \alpha_i)] \\
&= \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=1}^K (\rho_t(q_{t,i}, 1 - \alpha_i) - \rho_t(q_{t,i}^*, 1 - \alpha_i)) \right] \\
&= \mathbb{E} \left[ \sum_{t=1}^T f_t(q_t) - f_t(q_t^*) \right] \\
&\leq \frac{3B^2K}{\eta} \left( 1 + \sum_{t=1}^{T-1} \|q_{t+1}^* - q_t^*\|_1 \right) + \eta KT
\end{aligned}$$

□

## BIBLIOGRAPHY

- Ali Abbasi, Parvin Sadeghi, and Zahra Taghizadeh Rahmat Abadi. Characterization of microplastics in digestive tract of commercial fish species from the oman sea. Marine Pollution Bulletin, 197:115769, 2023.
- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. Advances in neural information processing systems, 24, 2011.
- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Information fusion, 76:243–297, 2021.
- Naoki Abe and Philip M Long. Associative reinforcement learning using linear probabilistic concepts. In ICML, pages 3–11. Citeseer, 1999.
- Mary B Alatis and Gerhard P Hancke. A review on challenges of autonomous mobile robot and sensor fusion methods. IEEE Access, 8:39830–39846, 2020.
- A. Alfaro-Núñez, D. Astorga, L. Cáceres-Farías, L. Bastidas, C. Soto Villegas, K. Macay, and J. H. Christensen. Microplastic pollution in seawater and marine organisms across the tropical eastern pacific and galápagos. Scientific Reports, 12:3502, 2022.
- Steve Allen, Deonie Allen, Vernon R Phoenix, Gaël Le Roux, Pilar Durántez Jiménez, Anaëlle Simonneau, Stéphane Binet, and Didier Galop. Atmospheric transport and deposition of microplastics in a remote mountain catchment. Nature geoscience, 12(5):339–344, 2019.
- Anastasios Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I Jordan. Uncertainty sets for image classifiers using conformal prediction. arXiv preprint arXiv:2009.14193, 2020.
- Anastasios N Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511, 2021.
- Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. Foundations and Trends in Machine Learning, 16(4):494–591, 2023.
- Anastasios N Angelopoulos, Rina Foygel Barber, and Stephen Bates. Online conformal prediction with decaying step sizes. arXiv preprint arXiv:2402.01139, 2024.

- C. F. Araujo, M. M. Nolasco, A. M. P. Ribeiro, and P. J. A. Ribeiro-Claro. Identification of microplastics using raman spectroscopy: Latest developments and future prospects. Water Research, 142:426–440, 2018.
- Felipe Areces, Christopher Mohri, Tatsunori Hashimoto, and John Duchi. Online conformal prediction via online optimization. In Forty-second International Conference on Machine Learning, 2025. URL <https://openreview.net/forum?id=KwGc2JUIDK>.
- Kamiar Asgari and Michael J Neely. Projection-free non-smooth convex programming. arXiv preprint arXiv:2208.05127, 2022.
- Miriam Ayer, H Daniel Brunk, George M Ewing, William T Reid, and Edward Silverman. An empirical distribution function for sampling with incomplete information. The annals of mathematical statistics, pages 641–647, 1955.
- Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. The Annals of Statistics, 51(2):816–845, 2023.
- R. J. Barnes, M. S. Dhanoa, and S. J. Lister. Standard normal variate transformation and de-trending of near-infrared diffuse reflectance spectra. Applied Spectroscopy, 43(5): 772–777, 1989. doi: 10.1366/0003702894202201.
- Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivald conformal prediction. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, Advances in Neural Information Processing Systems, 2022. URL <https://openreview.net/forum?id=QNjyrDBx6tz>.
- Melanie Bergmann, Sophia Mützel, Sebastian Pimpke, Mine B. Tekman, Jan Trachsel, and Gunnar Gerdt. White and wonderful? microplastics prevail in snow from the alps to the arctic. Science Advances, 5(8):eaax1157, 2019. doi: 10.1126/sciadv.aax1157.
- Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In International Conference on Machine Learning, pages 2337–2363. PMLR, 2023.
- Gilberto Binda, Gabriela Kalčíková, Ian John Allan, Rachel Hurley, Elisabeth Rødland, Davide Spanu, and Luca Nizzetto. Microplastic aging processes: Environmental relevance and analytical implications. TrAC Trends in Analytical Chemistry, 172:117566, 2024.
- Erik Blasch, Tien Pham, Chee-Yee Chong, Wolfgang Koch, Henry Leung, Dave Braines, and Tarek Abdelzاهر. Machine learning/artificial intelligence for sensor data fusion—opportunities and challenges. IEEE Aerospace and Electronic Systems Magazine, 36(7): 80–93, 2021.
- Jan S. Böke, Jürgen Popp, and Christoph Krafft. Optical photothermal infrared spectroscopy with simultaneously acquired raman spectroscopy for two-dimensional microplastic identification. Scientific Reports, 12(1):18785, 2022. doi: 10.1038/s41598-022-23318-2.

- J. Brahney, M. Hallerud, E. Heim, M. Hahnenberger, and S. Sukumaran. Plastic rain in protected areas of the united states. Science, 368:1257–1260, 2020.
- Ramon F Brena, Antonio A Aguilera, Luis A Trejo, Erik Molino-Minero-Re, and Oscar Mayora. Choosing the best sensor fusion method: A machine-learning approach. Sensors, 20(8):2350, 2020.
- Bruker. OPUS Package: SEARCH & IDENT, n.d. Accessed 2026-05-10.
- Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In Ricard Gavaldà, Gábor Lugosi, Thomas Zeugmann, and Sandra Zilles, editors, Algorithmic Learning Theory, pages 23–37, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-04414-4.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. Foundations and Trends® in Machine Learning, 5(1):1–122, 2012.
- Sébastien Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In Sanjoy Dasgupta and David McAllester, editors, Proceedings of the 30th International Conference on Machine Learning, volume 28 of Proceedings of Machine Learning Research, pages 258–265, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <https://proceedings.mlr.press/v28/bubeck13.html>.
- L. Cabernard, L. Roscher, C. Lorenz, G. Gerdtz, and S. Primpke. Comparison of raman and fourier transform infrared spectroscopy for the quantification of microplastics in the aquatic environment. Environmental Science & Technology, 52:13279–13288, 2018.
- Jesús Carrete, Hadrián Montes-Campos, Ralf Wanzenböck, Esther Heid, and Georg KH Madsen. Deep ensembles vs committees for uncertainty estimation in neural-network force fields: Comparison and application to active learning. The Journal of Chemical Physics, 158(20), 2023.
- Kani Chen, Inchi Hu, and Zhiliang Ying. Strong consistency of maximum quasi-likelihood estimators in generalized linear models with fixed and adaptive designs. The Annals of Statistics, 27(4):1155–1163, 1999.
- Abhilash Chenreddy and Erick Delage. End-to-end conditional robust optimization. arXiv preprint arXiv:2403.04670, 2024.
- Victor Chernozhukov, Kaspar Wüthrich, and Zhu Yinchu. Exact and robust conformal inference methods for predictive machine learning with dependent data. In Conference On learning theory, pages 732–749. PMLR, 2018.
- C Anela Choy, Bruce H Robison, Tyler O Gagne, Benjamin Erwin, Evan Firl, Rolf U Halden, J Andrew Hamilton, Kakani Katija, Susan E Lisin, Charles Rolsky, et al. The vertical distribution and biological transport of marine microplastics across the epipelagic and mesopelagic water column. Scientific reports, 9(1):7843, 2019.

- Liz Clark, Rebecca Allen, Zara LR Botterell, Beatriz Callejo, Brendan J Godley, Clare Henry, David Santillo, and Sarah E Nelms. Using citizen science to understand floating plastic debris distribution and abundance: A case study from the north cornish coast (united kingdom). Marine pollution bulletin, 194:115314, 2023.
- Madeline E Clough, Eduardo Ochoa Rivera, Rebecca L Parham, Andrew P Ault, Paul M Zimmerman, Anne J McNeil, and Ambuj Tewari. Enhancing confidence in microplastic spectral identification via conformal prediction. Environmental Science & Technology, 58(49):21740–21749, 2024.
- Win Cowger, Andrew Gray, Silke H Christiansen, Hannah DeFron, Ashok D Deshpande, Ludovic Hemabessiere, Eunah Lee, Leonid Mill, Keenan Munno, Barbara E Ossmann, et al. Critical review of processing and classification techniques for images and spectra in microplastic research. Applied Spectroscopy, 74:989–1010, 2020.
- Win Cowger, Zacharias Steinmetz, Andrew Gray, Kimberly Munno, Jennifer Lynch, Hannah Hapich, Sebastian Primpke, Hannah De Fron, Chelsea Rochman, and Orla Herodotou. Microplastic spectral classification needs an open source community: Open specy to the rescue! Analytical Chemistry, 93(21):7543–7548, 2021. doi: 10.1021/acs.analchem.1c00123.
- Samuel J Cusworth, William J Davies, Martin R McAinsh, Andrew S Gregory, Jonathan Storkey, and Carly J Stevens. Agricultural fertilisers contribute substantially to microplastic concentrations in uk soils. Communications Earth & Environment, 5(1):7, 2024.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In COLT, volume 2, page 3, 2008.
- Hannah De Fron, Rachel Rubinovitz, and Chelsea M. Rochman.  $\mu$ ATR-FTIR spectral libraries of plastic particles (FLOPP and FLOPP-e) for the analysis of microplastics. Analytical Chemistry, 93(48):15878–15885, 2021a. doi: 10.1021/acs.analchem.1c02549.
- Hannah De Fron, Rebecca Rubinovitz, and Chelsea M. Rochman.  $\mu$ atr-ftir spectral libraries of plastic particles (flopp and flopp-e) for the analysis of microplastics. Analytical Chemistry, 93:15878–15885, 2021b.
- Rémy Degenne, Pierre Menard, Xuedong Shang, and Michal Valko. Gamification of pure exploration for linear bandits. In Hal Daumé III and Aarti Singh, editors, Proceedings of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pages 2432–2442. PMLR, 13–18 Jul 2020.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- John Denker and Yann LeCun. Transforming neural-net output levels to probability distributions. Advances in neural information processing systems, 3, 1990.

- Thomas G Dietterich. Ensemble methods in machine learning. In International workshop on multiple classifier systems, pages 1–15. Springer, 2000.
- Ming Dong, Qing Zhang, Xia Xing, Wei Chen, Zhen She, and Zhi Luo. Raman spectra and surface changes of microplastics weathered under natural environments. Science of the Total Environment, 739:139990, 2020. doi: 10.1016/j.scitotenv.2020.139990.
- Ming Dong, Zhen She, Xiaoping Xiong, Gang Ouyang, and Zhi Luo. Automated analysis of microplastics based on vibrational spectroscopy: Are we measuring the same metrics? Analytical and Bioanalytical Chemistry, 414(11):3359–3372, 2022. doi: 10.1007/s00216-022-03951-6.
- Kevin Dowd. Too good to be true? the (in) credibility of the uk inflation fan charts. Journal of Macroeconomics, 29(1):91–102, 2007.
- Karin Duswald, Verena Pichler, Verena Kopatz, Thomas Limberger, Verena Karl, Daniel Hennerbichler, Roman Zimmerleiter, Wolfgang Wadsak, Matthias Hettich, Elisabeth S. Gruber, et al. Detection of unlabeled polystyrene micro- and nanoplastics in mammalian tissue by optical photothermal infrared spectroscopy. Analytical Chemistry, 97(31):16714–16722, 2025. doi: 10.1021/acs.analchem.4c05400.
- Bradley Efron. Bootstrap methods: Another look at the jackknife. The Annals of Statistics, 7(1):1–26, 1979.
- I. Evangelou, S. Bucci, and A. Stohl. Atmospheric microplastic emissions from land and ocean. Nature, 2026. doi: 10.1038/s41586-025-09998-6.
- Louis Faury, Marc Abeille, Clément Calauzènes, and Olivier Fercoq. Improved optimistic algorithms for logistic bandits. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 2020.
- Louis Faury, Marc Abeille, Kwang-Sung Jun, and Clément Calauzènes. Jointly efficient and optimal algorithms for logistic bandits. In International Conference on Artificial Intelligence and Statistics, pages 546–580. PMLR, 2022.
- Shai Feldman, Stephen Bates, and Yaniv Romano. Calibrated multiple-output quantile regression with representation learning. Journal of Machine Learning Research, 24(24): 1–48, 2023.
- Tanner Fiez, Lalit Jain, Kevin G Jamieson, and Lillian Ratliff. Sequential experimental design for transductive linear bandits. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
- Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. Advances in neural information processing systems, 23, 2010.
- J. P. G. L. Frias and R. Nash. Microplastics: Finding a consensus on the definition. Marine Pollution Bulletin, 138:145–147, 2019.

- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 998–1027. PMLR, 23–26 Jun 2016.
- Matteo Gasparin and Aaditya Ramdas. Conformal online model aggregation. *arXiv preprint arXiv:2403.15527*, 2024a.
- Matteo Gasparin and Aaditya Ramdas. Merging uncertainty sets via majority vote. *arXiv preprint arXiv:2401.09379*, 2024b.
- Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.
- Isaac Gibbs and Emmanuel J Candès. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.
- R. C. Hale, M. E. Seeley, M. J. La Guardia, L. Mai, and E. Y. Zeng. A global perspective on microplastics. *Journal of Geophysical Research: Oceans*, 125:e2018JC014719, 2020.
- Marc Hallin, Davy Paindaveine, and Marianna Šíman. Multivariate quantiles and multiple-output regression quantiles: From  $\ell_1$  optimization to halfspace depth. *Annals of Statistics*, 38:635–669, 2010.
- Marton Havasi, Rodolphe Jenatton, Stanislav Fort, Jeremiah Zhe Liu, Jasper Snoek, Balaji Lakshminarayanan, Andrew M Dai, and Dustin Tran. Training independent subnetworks for robust prediction. *arXiv preprint arXiv:2010.06610*, 2020.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- Atsuhiko Isobe, Nicolette T. Buenaventura, Sarah Chastain, Suchana Chavanich, Andrés Cózar, Marie DeLorenzo, Paul Hagmann, Hirofumi Hinata, Nikolai Kozlovskii, Amy L. Lusher, et al. An interlaboratory comparison exercise for the determination of microplastics in standard sample bottles. *Marine Pollution Bulletin*, 146:831–837, 2019. doi: 10.1016/j.marpolbul.2019.07.033.
- Natalia P Ivleva. Chemical analysis of microplastics and nanoplastics: challenges, advanced methods, and perspectives. *Chemical reviews*, 121(19):11886–11936, 2021.
- Jenna R. Jambeck, Roland Geyer, Chris Wilcox, Theodore R. Siegler, Miriam Perryman, Anthony Andrady, Ramani Narayan, and Kara Lavender Law. Plastic waste inputs from land into the ocean. *Science*, 347(6223):768–771, 2015. doi: 10.1126/science.1260352.
- Yassir Jedra and Alexandre Proutiere. Optimal best-arm identification in linear bandits. *Advances in neural information processing systems*, 33, 2020.
- John Wiley & Sons, Inc. KnowItAll Analytical Edition Software. <https://sciencesolutions.wiley.com/knowitall-analytical-edition-software/>, n.d. Accessed 2026-05-10.

- Kwang-Sung Jun, Lalit Jain, Blake Mason, and Houssam Nassif. Improved confidence bounds for the linear logistic model and applications to bandits. In International Conference on Machine Learning, pages 5148–5157. PMLR, 2021.
- Melissa R Jung, F David Horgen, Sara V Orski, Viviana Rodriguez, Kathryn L Beers, George H Balazs, T Todd Jones, Thierry M Work, Kayla C Brignac, Sarah-Jeanne Royer, et al. Validation of atr ft-ir to identify polymers of plastic marine debris, including those ingested by marine organisms. Marine pollution bulletin, 127:704–716, 2018.
- Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In ICML, volume 12, pages 655–662, 2012.
- La Daana K Kanhai, Katarina Gardfeldt, Thomas Krumpfen, Richard C Thompson, and Ian O’Connor. Microplastics in sea ice and seawater beneath ice floes from the arctic ocean. Scientific reports, 10(1):5004, 2020.
- Hideaki Kano, Junya Honda, Kentaro Sakamaki, Kentaro Matsuura, Atsuyoshi Nakamura, and Masashi Sugiyama. Good arm identification via bandit feedback. Machine Learning, 108:721–745, 2019.
- Abbas Kazerouni and Lawrence M. Wein. Best arm identification in generalized linear bandits. Operations Research Letters, 49(3):365–371, 2021. ISSN 0167-6377. doi: <https://doi.org/10.1016/j.orl.2021.03.011>. URL <https://www.sciencedirect.com/science/article/pii/S0167637721000523>.
- Linglong Kong and Ivan Mizera. Quantile tomography: Using quantiles with multivariate data. Statistica Sinica, pages 1589–1610, 2012.
- S. Kotar, R. McNeish, C. Murphy-Hagan, V. Renick, C.-F. T. Lee, C. Steele, A. Lusher, C. Moore, E. Minor, J. Schroeder, et al. Quantitative assessment of visual microscopy as a tool for microplastic research: Recommendations for improving methods and reporting. Chemosphere, 308:136449, 2022. doi: 10.1016/j.chemosphere.2022.136449.
- R. Kozloski, W. Cowger, and M. M. Arienzo. Moving toward automated  $\mu$ FTIR spectra matching for microplastic identification: Addressing false identifications and improving accuracy. Microplastics and Nanoplastics, 4(1):27, 2024. doi: 10.1186/s43591-024-00106-5.
- Arun K Kuchibhotla and Richard A Berk. Nested conformal prediction sets for classification with applications to probation data. The Annals of Applied Statistics, 17(1):761–785, 2023.
- A. K ppler, D. Fischer, S. Oberbeckmann, G. Schernewski, M. Labrenz, K.-J. Eichhorn, and B. Voit. Analysis of environmental microplastics by vibrational microspectroscopy: Ftir, raman or both? Analytical and Bioanalytical Chemistry, 408:8377–8391, 2016.
- Tor Lattimore and Csaba Szepesv ri. Bandit algorithms. Cambridge University Press, 2020.

- Robin Lenz, Kristina Enders, Colin A. Stedmon, David M. A. Mackenzie, and Torkel Gissel Nielsen. A critical assessment of visual identification of marine microplastic using raman spectroscopy for analysis improvement. Marine Pollution Bulletin, 100(1):82–91, 2015. doi: 10.1016/j.marpolbul.2015.09.026.
- Yifeng Li, Fang-Xiang Wu, and Alioune Ngom. A review on machine learning principles for multi-view biological data integration. Briefings in bioinformatics, 19(2):325–340, 2018.
- Michael W. Liemohn. Data Analysis for the Geosciences: Essentials of Uncertainty, Comparison, and Visualization. Wiley, 2024.
- J. Lim, J. Seo, and D. Shin. Reducing spectral confusion in microplastic analysis: A U-Net deep learning approach. Analytical Chemistry, 97(34):18432–18443, 2025. doi: 10.1021/acs.analchem.5c00584.
- Lars Lindemann, Matthew Cleaveland, Gihyun Shim, and George J Pappas. Safe planning in dynamic environments using conformal prediction. IEEE Robotics and Automation Letters, 8(8):5116–5123, 2023.
- Yufeng Liu and Yichao Wu. Simultaneous multiple non-crossing quantile regression estimation using kernel constraints. Journal of nonparametric statistics, 23(2):415–437, 2011.
- Andrea Locatelli, Maurilio Gutzeit, and Alexandra Carpentier. An optimal algorithm for the thresholding bandit problem. In International Conference on Machine Learning, pages 1690–1698. PMLR, 2016.
- Adrián López-Rosales, Borja Ferreiro, José Andrade, María Fernández-Amado, Miguel González-Pleiter, Purificación López-Mahía, Roberto Rosal, and Soledad Muniategui-Lorenzo. A reliable method to determine airborne microplastics using quantum cascade laser infrared spectrometry. Science of the Total Environment, 913:169678, 2024.
- Charles Lu, Andréanne Lemay, Ken Chang, Katharina Höbel, and Jayashree Kalpathy-Cramer. Fair conformal predictors for applications in medical imaging. In Proceedings of the AAAI conference on artificial intelligence, volume 36, pages 12008–12016, 2022.
- Charles Lu, Yaodong Yu, Sai Praneeth Karimireddy, Michael Jordan, and Ramesh Raskar. Federated conformal predictors for distributed uncertainty quantification. In Federated Learning and Analytics in Practice: Algorithms, Systems, Applications, and Opportunities, 2023. URL <https://openreview.net/forum?id=qEPYJNHNC0>.
- Rui Luo and Zhixin Zhou. Weighted aggregation of conformity scores for classification. arXiv preprint arXiv:2407.10230, 2024.
- Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. Advances in neural information processing systems, 31, 2018.
- Rémi Marta Soare, Alessandro Lazaric. Best-arm identification in linear bandits. Advances in neural information processing systems, 27, 2014.

- Blake Mason, Lalit Jain, Subhojyoti Mukherjee, Romain Camilleri, Kevin Jamieson, and Robert Nowak. Nearly optimal algorithms for level set estimation. In International Conference on Artificial Intelligence and Statistics, pages 7625–7658. PMLR, 2022.
- Kleopatra Miserli, Christos Lykos, Angelos G Kalampounias, and Ioannis Konstantinou. Screening of microplastics in aquaculture systems (fish, mussel, and water samples) by ftir, scanning electron microscopy–energy dispersive spectroscopy and micro-raman spectroscopies. Applied Sciences, 13(17):9705, 2023.
- Annette M. Molinaro, Richard Simon, and Ruth M. Pfeiffer. Prediction error estimation: A comparison of resampling methods. Bioinformatics, 21(15):3301–3307, 2005. doi: 10.1093/bioinformatics/bti499.
- Vanessa Morgado, Carla Palma, and Ricardo JN Bettencourt da Silva. Microplastics identification by infrared spectroscopy–evaluation of identification criteria and uncertainty by the bootstrap method. Talanta, 224:121814, 2021.
- Kimberly Munno, Hannah De Frond, Brianna O’Donnell, and Chelsea M. Rochman. Increasing the accessibility for characterizing microplastics: Introducing new application-based and spectral libraries of plastic particles (SLoPP and SLoPP-E). Analytical Chemistry, 92(3):2443–2451, 2020. doi: 10.1021/acs.analchem.9b03626.
- Veronica Nava, Maria Luce Frezzotti, and Barbara Leoni. Raman spectroscopy for the analysis of microplastics in aquatic systems. Applied Spectroscopy, 75(11):1341–1357, 2021.
- Shanshan Niu, Rui Liu, Qing Zhao, S. Gagan, Andrea Dodero, Qi Ying, Xiaohui Ma, Zhen Cheng, Swarup China, Manjula Canagaratna, et al. Quantifying the chemical composition and real-time mass loading of nanoplastic particles in the atmosphere using aerosol mass spectrometry. Environmental Science & Technology, 58(7):3363–3374, 2024. doi: 10.1021/acs.est.3c10286.
- Eduardo Ochoa Rivera and Ambuj Tewari. Near optimal pure exploration in logistic bandits. arXiv preprint arXiv:2410.20640, 2024.
- Eduardo Ochoa Rivera, Yash Patel, and Ambuj Tewari. Conformal prediction for ensembles: Improving efficiency via score-based aggregation. In The Thirty-Ninth Annual Conference on Neural Information Processing Systems, 2025.
- Eduardo Ochoa-Rivera, Yash Patel, and Ambuj Tewari. Conformal prediction for ensembles: Improving efficiency via score-based aggregation. In The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2026. URL <https://openreview.net/forum?id=sNTqqdPVSv>.
- Nicole E. Olson, Y. Xiao, Z. Lei, and Andrew P. Ault. Simultaneous optical photothermal infrared (O-PTIR) and raman spectroscopy of submicrometer atmospheric particles. Analytical Chemistry, 92(14):9932–9939, 2020. doi: 10.1021/acs.analchem.0c01495.

- Davy Paindaveine and Miroslav Šiman. On directional multiple-output quantile regression. Journal of Multivariate Analysis, 102(2):193–212, 2011.
- Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In European conference on machine learning, pages 345–356. Springer, 2002.
- Rebecca L. Parham, Abbygail M. Ayala, L. Meagher, Madeline E. Clough, Eduardo Ochoa Rivera, J. H. Shi, Ambuj Tewari, Anne J. McNeil, and Andrew P. Ault. Identifying microplastics in laboratory and atmospheric aerosol mixtures via optical photothermal infrared and raman microspectroscopy. Analytical Chemistry, 97(33):18136–18143, 2025. doi: 10.1021/acs.analchem.5c02968.
- Yash Patel, Sahana Rayan, and Ambuj Tewari. Conformal contextual robust optimization. arXiv preprint arXiv:2310.10003, 2023.
- E. Pięta, N. Piergies, K. Chrabąszcz, A. Banas, K. Banaś, M. K. F. Lo, A. Panek, W. M. Kwiatek, M. B. H. Breese, and K. Pogoda. Photothermal infrared imaging of nanoplastics in human cells with nanoscale resolution. ACS Applied Materials & Interfaces, 2025. doi: 10.1021/acsami.5c16048.
- Joana C Prata, João P da Costa, Isabel Lopes, Anthony L Andrady, Armando C Duarte, and Teresa Rocha-Santos. A one health perspective of the impacts of microplastics on animal, human and environmental health. Science of the Total Environment, 777:146094, 2021.
- Joana C. Prata, Jorge Padrão, Muhammad Tariq Khan, and Tony R. Walker. Do’s and don’ts of microplastic research: a comprehensive guide. Water Emerging Contaminants & Nanoplastics, 3(2), 2024. ISSN 2831-2597. doi: 10.20517/wecn.2023.61. URL <https://www.oaepublish.com/articles/wecn.2023.61>.
- Craig B. Prater, Mustafa Kansiz, and Ji-Xin Cheng. A tutorial on optical photothermal infrared (O-PTIR) microscopy. APL Photonics, 9(9), 2024. doi: 10.1063/5.0219983.
- Sebastian Primpke, Claudia Lorenz, Richard Rascher-Friesenhausen, and Gunnar Gerdts. An automated approach for microplastics analysis using focal plane array (fpa) ftir microscopy and image analysis. Analytical Methods, 9(9):1499–1511, 2017.
- Sebastian Primpke, Marisa Wirth, Claudia Lorenz, and Gunnar Gerdts. Reference database design for the automated analysis of microplastic samples based on fourier transform infrared (ftir) spectroscopy. Analytical and bioanalytical chemistry, 410(21):5131–5141, 2018.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In International conference on machine learning, pages 8748–8763. PMLR, 2021.

- Rahul Rahaman et al. Uncertainty quantification and deep ensembles. Advances in neural information processing systems, 34:20063–20075, 2021.
- Eva Regnier. Public evacuation decisions and hurricane track uncertainty. Management Science, 54(1):16–28, 2008.
- Allen Z Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, et al. Robots that ask for help: Uncertainty alignment for large language model planners. arXiv preprint arXiv:2307.01928, 2023.
- Gerrit Renner, Torsten C Schmidt, and Jürgen Schram. A new chemometric approach for automatic identification of microplastics from environmental compartments based on ft-ir spectroscopy. Analytical chemistry, 89(22):12045–12053, 2017.
- Gerrit Renner, Astrid Nellessen, Andreas Schwiers, Martin Wenzel, Torsten C. Schmidt, and Jürgen Schram. Data preprocessing & evaluation used in the microplastics identification process: A critical review & practical guide. TrAC Trends in Analytical Chemistry, 111: 229–238, 2019. doi: 10.1016/j.trac.2018.12.004.
- Herbert Robbins. Some aspects of the sequential design of experiments. Bulletin of the American Mathematical Society, pages 527–535, 1952.
- Peter J Rousseeuw and Anja Struyf. Computing location depth and regression depth in higher dimensions. Statistics and Computing, 8:193–203, 1998.
- Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized banditxs. Mathematics of Operations Research, 35(2):395–411, 2010.
- Yoan Russac, Louis Fauray, Olivier Cappé, and Aurélien Garivier. Self-concordant analysis of generalized linear bandits with forgetting. In International Conference on Artificial Intelligence and Statistics, pages 658–666. PMLR, 2021.
- Robert E Schapire et al. A brief introduction to boosting. In Ijcai, volume 99, pages 1401–1406. Citeseer, 1999.
- Stefan Schnabel and Wolfhard Janke. A simple algorithm for uniform sampling on the surface of a hypersphere. arXiv preprint arXiv:2204.14004, 2022.
- John H. Seinfeld and Spyros N. Pandis. Atmospheric Chemistry and Physics: From Air Pollution to Climate Change. John Wiley & Sons, 2 edition, 2006.
- Robert Serfling. Quantile functions for multivariate analysis: approaches and applications. Statistica Neerlandica, 56(2):214–232, 2002.
- Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. Journal of Machine Learning Research, 9(3), 2008.

- J. A. Smolen, G. E. Moore, N. D. Perez, and K. L. Wooley. Adaptable microplastic classification using similarity learning on  $\mu$ FTIR spectra collected from  $\mu$ FTIR focal plane array imaging. Proceedings of the National Academy of Sciences, 122(42):e2509745122, 2025. doi: 10.1073/pnas.2509745122.
- Marta Soare. Sequential resource allocation in linear stochastic bandits. In Sequential Resource Allocation in Linear Stochastic Bandits, 2015.
- Young Kyoung Song, Sang Hee Hong, Soeun Eo, and Won Joon Shim. A comparison of spectroscopic analysis methods for microplastics: Manual, semi-automated, and automated fourier transform infrared and raman techniques. Marine Pollution Bulletin, 173:113101, 2021. ISSN 0025-326X. doi: <https://doi.org/10.1016/j.marpolbul.2021.113101>. URL <https://www.sciencedirect.com/science/article/pii/S0025326X21011358>.
- M. Stone. Cross-validators choice and assessment of statistical predictions. Journal of the Royal Statistical Society: Series B, 36(2):111–147, 1974. doi: 10.1111/j.2517-6161.1974.tb00994.x.
- Mahesh Subedar, Ranganath Krishnan, Paulo Lopez Meyer, Omesh Tickoo, and Jonathan Huang. Uncertainty-aware audiovisual activity recognition using deep bayesian variational inference. In Proceedings of the IEEE/CVF international conference on computer vision, pages 6301–6310, 2019.
- Shiliang Sun. A survey of multi-view machine learning. Neural computing and applications, 23:2031–2038, 2013.
- Sophia Huiwen Sun and Rose Yu. Conformal prediction for time-series forecasting with change points. In The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025. URL <https://openreview.net/forum?id=HgLaVgCpCl>.
- R. Talbot and H. Chang. Microplastics in freshwater: A global review of factors affecting spatial and temporal variations. Environmental Pollution, 292:118393, 2022.
- Thermo Fisher Scientific Inc. FTIR Analysis Software, n.d. Accessed 2026-05-10.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 25(3-4):285–294, 1933.
- Junjiao Tian, Wesley Cheung, Nathaniel Glaser, Yen-Cheng Liu, and Zsolt Kira. Uno: Uncertainty-aware noisy-or multimodal fusion for unanticipated input degradation. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 5716–5723. IEEE, 2020.
- A. Tirkey and L. S. B. Upadhyay. Microplastics: An overview on separation, identification and characterization of microplastics. Marine Pollution Bulletin, 170:112604, 2021.
- Vladimir G Trunov and Vladimir V V'yugin. Online aggregation of conformal predictive systems. In Conformal and Probabilistic Prediction with Applications, pages 430–449. PMLR, 2023.

- Renukanandan Tumu, Matthew Cleaveland, Rahul Mangharam, George J Pappas, and Lars Lindemann. Multi-modal conformal prediction regions by optimizing convex shape templates. arXiv preprint arXiv:2312.07434, 2023.
- U.S. Bureau of Labor Statistics. Consumer price index for all urban consumers: All items in u.s. city average [CPIAUCSL]. FRED, Federal Reserve Bank of St. Louis, n.d. URL <https://fred.stlouisfed.org/series/CPIAUCSL>. Retrieved March 2026.
- Michal Valko, Rémi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In International Conference on Machine Learning, pages 46–54. PMLR, 2014.
- Vladimir Vovk. Conditional validity of inductive conformal predictors. In Asian conference on machine learning, pages 475–490. PMLR, 2012.
- Vladimir Vovk, Alexander Gammerman, and Craig Saunders. Machine-learning applications of algorithmic randomness. In Proceedings of the Sixteenth International Conference on Machine Learning, pages 444–453. Morgan Kaufmann, 1999.
- Vladimir Vovk, Alex Gammerman, and Glenn Shafer. Algorithmic Learning in a Random World. Springer, 2005.
- VV V’yugin and VG Trunov. Online aggregation of conformal forecasting systems. Journal of Communications Technology and Electronics, 68(Suppl 2):S239–S253, 2023.
- C. Wang, J. Zhao, and B. Xing. Environmental source, fate, and toxicity of microplastics. Journal of Hazardous Materials, 407:124357, 2021a.
- Po-An Wang, Ruo-Chun Tzeng, and Alexandre Proutiere. Fast pure exploration via frank-wolfe. Advances in Neural Information Processing Systems, 34:5810–5821, 2021b.
- F. Watteau, M.-F. Dignac, A. Bouchard, A. Revallier, and S. Houot. Microplastic detection in soil amended with municipal solid waste composts as revealed by transmission electronic microscopy and pyrolysis/GC/MS. Frontiers in Sustainable Food Systems, 2, 2018. doi: 10.3389/fsufs.2018.00081.
- Jana Weisser, Teresa Pohl, Michael Heinzinger, Natalia P Ivleva, Thomas Hofmann, and Karl Glas. The identification of microplastics based on vibrational spectroscopy data—a critical review of data analysis routines. TrAC Trends in Analytical Chemistry, 148: 116535, 2022.
- Jun-Li Xu, Kevin V. Thomas, Zisheng Luo, and Aoife A. Gowen. FTIR and raman imaging for microplastics analysis: State of the art, challenges and prospects. TrAC Trends in Analytical Chemistry, 119:115629, 2019.
- Xiaoqiang Yan, Shizhe Hu, Yiqiao Mao, Yangdong Ye, and Hui Yu. Deep multi-view learning methods: A review. Neurocomputing, 448:106–129, 2021.

- Yachong Yang and Arun Kumar Kuchibhotla. Selection and aggregation of conformal prediction sets. Journal of the American Statistical Association, pages 1–13, 2024.
- Ye Yuan, Guangxu Xun, Kebin Jia, and Aidong Zhang. A multi-view deep learning method for epileptic seizure detection using short-time fourier transform. In Proceedings of the 8th ACM international conference on bioinformatics, computational biology, and health informatics, pages 213–222, 2017.
- Ye Yuan, Guangxu Xun, Kebin Jia, and Aidong Zhang. A multi-view deep learning framework for eeg seizure detection. IEEE journal of biomedical and health informatics, 23(1): 83–94, 2018.
- Margaux Zaffran, Olivier Féron, Yannig Goode, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In International conference on machine learning, pages 25834–25866. PMLR, 2022.
- Cha Zhang and Yunqian Ma. Ensemble machine learning, volume 144. Springer, 2012.
- Delong Zhang, Chen Li, Cheng Zhang, Mikhail N. Slipchenko, Gregory Eakins, and Ji-Xin Cheng. Depth-resolved mid-infrared photothermal imaging of living cells and organisms with submicrometer spatial resolution. Science Advances, 2(9):e1600521, 2016. doi: 10.1126/sciadv.1600521.
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3836–3847, 2023a.
- Xu Zhang, Chunhong Yu, Peng Wang, and Chunping Yang. Microplastics and human health: unraveling the toxicological pathways and implications for public health. Frontiers in public health, 13:1567200, 2025a.
- Y. Zhang, P. Wu, R. Xu, X. Wang, L. Lei, A. T. Schartup, Y. Peng, Q. Pang, X. Wang, L. Mai, et al. Plastic waste discharge to the global ocean constrained by seawater observations. Nature Communications, 14(1):1372, 2023b. doi: 10.1038/s41467-023-37108-5.
- Y. Zhang, J. H. Slade, A. P. Ault, and A. W. H. Chan. An atmospheric chemistry perspective on airborne micro- and nanoplastic particles. Environmental Science & Technology, 59(16):7810–7819, 2025b. doi: 10.1021/acs.est.5c03264.
- J. Zhao, H. Lui, D. I. McLean, and H. Zeng. Automated autofluorescence background subtraction algorithm for biomedical raman spectroscopy. Applied Spectroscopy, 61(11): 1225–1232, 2007. doi: 10.1366/000370207782597003.
- Jing Zhao, Xijiong Xie, Xin Xu, and Shiliang Sun. Multi-view learning overview: Recent progress and new challenges. Information Fusion, 38:43–54, 2017.
- Hui Zou and Ming Yuan. Composite quantile regression and the oracle model selection theory. The Annals of Statistics, 36:1108–1126, 2008.