

Computational Challenges in Non-parametric Prediction of Bradycardia in Preterm

Infants

by

Sinjini Mitra

A Thesis Presented in Partial Fulfillment
of the Requirements for the Degree
Master of Science

Approved November 2020 by the
Graduate Supervisory Committee:

Antonia Papandreou-Suppappola, Co-Chair
Bahman Moraffah, Co-Chair
Pavan Turaga

ARIZONA STATE UNIVERSITY

December 2020

ABSTRACT

Infants born before 37 weeks of pregnancy are considered to be preterm. Typically, preterm infants have to be strictly monitored since they are highly susceptible to health problems like hypoxemia (low blood oxygen level), apnea, respiratory issues, cardiac problems, neurological problems as well as an increased chance of long-term health issues such as cerebral palsy, asthma and sudden infant death syndrome. One of the leading health complications in preterm infants is bradycardia - which is defined as the slower than expected heart rate, generally beating lower than 60 beats per minute. Bradycardia is often accompanied by low oxygen levels and can cause additional long term health problems in the premature infant.

The implementation of a non-parametric method to predict the onset of bradycardia is presented. This method assumes no prior knowledge of the data and uses kernel density estimation to predict the future onset of bradycardia events. The data is preprocessed, and then analyzed to detect the peaks in the ECG signals, following which different kernels are implemented to estimate the shared underlying distribution of the data. The performance of the algorithm is evaluated using various metrics and the computational challenges and methods to overcome them are also discussed.

It is observed that the performance of the algorithm with regards to the kernels used are consistent with the theoretical performance of the kernel as presented in a previous work. The theoretical approach has also been automated in this work and the various implementation challenges have been addressed.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support of many people. I would like to thank my advisers Dr. Antonia Papandreou-Suppappola and Dr. Bahman Moraffah for their guidance and support in helping me finishing this thesis. Thanks to Dr. Pavan Turaga for agreeing to be on my defense committee. A special thank you to Dr. Thomas Matthew Holeva for his support and advice throughout the duration of my graduate education. Thanks to my dear friend Brian Baddadah and his many lessons on LaTeX. Thanks to my co-workers who put up with me while I rambled endlessly about this concept. Thanks to my parents Bajradeb Mitra and Laily Mitra who not only support me in every endeavor but have taught me to pursue my ambition with determination and hard work. This thesis would not be possible without their love. Thank you to my brothers Rhitabrata and Dhritabrata Mitra who encourage me and keep me going. Finally, thank you to my friends Nicole Martin, Madeline Damasco, Shayla Puryear and Andrew Wharton for their encouragement and support.

TABLE OF CONTENTS

| | Page |
|---|------|
| LIST OF TABLES | v |
| LIST OF FIGURES | vi |
| 1 INTRODUCTION | 1 |
| 1.1 Motivation and Existing Methods | 1 |
| 1.2 Proposed Method | 2 |
| 1.3 Thesis Organization | 3 |
| CHAPTER | |
| 2 BACKGROUND IN NON-PARAMETRIC STATISTICS AND MACHINE LEARNING | 4 |
| 2.1 Machine Learning Methods | 5 |
| 2.1.1 Learning Scenarios | 7 |
| 2.1.2 Model Selection | 9 |
| 2.2 Bias Variance Decomposition | 12 |
| 2.3 Density Estimation | 15 |
| 2.3.1 Parametric Density Estimation | 18 |
| 2.3.2 Non-Parametric Density Estimation | 24 |
| 2.3.3 Kernel Density Estimator | 25 |
| 2.3.4 Mean Square Error of Kernel Estimators | 28 |
| 2.4 Cross Validation to Minimize Risk | 33 |
| 3 ALGORITHM AUTOMATION | 37 |
| 3.1 Data Collection | 37 |
| 3.2 Non-parametric Prediction | 38 |
| 3.3 Data Preprocessing | 40 |
| 3.4 Peak Detection | 41 |

| CHAPTER | Page |
|-------------------------------------|------|
| 3.5 Dataset Segmentation | 42 |
| 3.5.1 Training Set | 42 |
| 3.5.2 Validation Set | 42 |
| 3.5.3 Test Set | 43 |
| 3.5.4 Monte-Carlo Simulations | 43 |
| 3.5.5 Testing Mechanism..... | 44 |
| 4 EXPERIMENTS | 45 |
| 4.1 Evaluation Metrics..... | 47 |
| 4.2 Discussion..... | 53 |
| 5 FUTURE WORKS | 55 |
| REFERENCES | 56 |

LIST OF TABLES

| Table | Page |
|---|------|
| 2.1 The Basis Functions of the Different Kernels Implemented | 26 |
| 3.1 Number of Bradycardia Events and Duration of Ecg for Each Infant as Obtained From [1] | 41 |
| 4.1 Information Obtained from the Header File (.hea) for Each Infant from the Database..... | 46 |
| 4.2 A General Confusion Matrix | 52 |
| 4.3 Confusion Matrix for Gaussian Kernel Run on Infant 5 | 53 |
| 4.4 Average Test Error for Each Kernel as Calculated for Infant 5 Through Monte Carlo Simulation | 53 |
| 4.5 Calculated Epe for Epanechnikov, Cosine and Uniform Kernel Run on Infant 5 | 53 |
| 4.6 Calculated Metrics for Infant 5..... | 54 |

LIST OF FIGURES

| Figure | Page |
|--------|---|
| 2.1 | The Estimation and Approximation Error Are Represented in Blue and Green Respectively. These Errors Are the Distance of a Chosen Hypothesis h^* from h_1 and h_2 11 |
| 2.2 | Illustration of the Presence of Bias/Variance (or Both) in a Data Set... 13 |
| 2.3 | Figure Illustrating the Bias Variance Decomposition and the Optimal Complexity 14 |
| 2.4 | A Histogram Showing the Maximum Daily Temperatures at La Guardia Airport, New York. The Data Is Drawn from The <i>airquality</i> Dataset in R Which Contains the Daily Air Quality Measurements in New York, May-september, 1973..... 16 |
| 2.5 | All of the Implemented Kernels on a Common Coordinate System 27 |
| 2.6 | Squared Bias, Variance and Mean Squared Error (Solid Line) as Functions Of h As Seen In [2] 30 |
| 4.1 | Block Diagram Representation of the Preprocessing and Density Estimation Stage. The Bradycardia Onset Indices Can Be Found from The <i>.atr</i> File Available in the Database Used. 47 |
| 4.2 | A Segment of the Raw Ecg for Infant 5, Obtained Using The <i>rdsamp()</i> Command [1] 48 |
| 4.3 | Pan-Tompkin Output for an ECG Segment for Infant 5 50 |
| 4.4 | Estimate $\hat{p}_H(x)$ for Gaussian Kernel. 51 |
| 4.5 | Plots Showing the Kernel Density Estimate for Infant 5 Using the Gaussian Kernel with $h = 4.981171$ and $\mathcal{C}_k = 0.002179295$ and the Obtained Confidence Set \mathcal{B}_x . The Convex Hull for the Confidence Set Is Outlined in Black. 52 |

Chapter 1

INTRODUCTION

1.1 Motivation and Existing Methods

Infants born before 37 weeks of pregnancy are considered to be preterm. Typically, preterm infants have to be strictly monitored since they are highly susceptible to health problems like hypoxemia (low blood oxygen level), apnea, respiratory issues, cardiac problems, neurological problems as well as an increased chance of long-term health issues such as cerebral palsy, asthma and sudden infant death syndrome. One of the leading health complications in preterm infants is bradycardia - which is defined as the slower than expected heart rate, generally beating lower than 60 beats per minute. Bradycardia is often accompanied by low oxygen levels and can cause additional long term health problems in the premature infant. Certain therapeutic interventions, as presented in [3, 4], might be the most effective if intervention is done early in high-risk infants. However, in a study conducted with nineteen preterm infants (10 M/ 9 F) born between 25–33 weeks of gestation [5], it was found that even on the most sensitive setting of the oximeter, a significant number of bradycardias are not recorded. Given the severity of the condition in terms of long term effects and the lack of a comprehensive and robust mechanism to predict bradycardias, there has been significant research and development in this area in recent times.

Existing methods as mentioned in [6, 7], use a combination of signal detection using multivariate feature construction from multimodal measurements, followed by machine learning to generate predictive warnings to aid real-time therapeutic inter-

ventions. The approach outlined in [7] uses Gaussian Mixture Models to successfully train the model and subsequently generate the the warnings. The limited success of these methods might help health care professionals and alert clinicians in the short term and ultimately provide automatic therapeutic care to reduce the complexity of predicting preterm cardiorespiratory conditions. However, there still remains a need to present a more robust and reliable method for the prediction of bradycardia.

This is addressed partly in [8], which hypothesizes that the immature cardiovascular control system in preterm infants exhibits transient temporal instabilities in heart rate that can be detected as a precursor signal of bradycardia. Statistical features in heartbeat signals prior to bradycardia are extracted and evaluated to assess their utility for predicting bradycardia. The method uses point process analysis to generate real-time, stochastic measures from discrete observations of continuous biological mechanisms. This method achieves a false alarm rate of 0.79 ± 0.018 .

Although the process in [8] outperforms its predecessors, an alternate statistical method to predict the onset of bradycardia in preterm infants is presented in [9]. Unlike the parametric approach in [8], the method in [9] assumes no prior knowledge of the data and uses non-parametric methods to predict the future onset of bradycardia events with 95% accuracy. The data is modeled by first detecting the QRS complex in the ECG signals and then using kernel density estimator.

1.2 Proposed Method

The non-parametric estimation method for prediction of bradycardia mentioned in [9] is implemented in this project through the creation of an automated algorithm to

provide results consistent with the findings in the paper. The algorithm is created in R and MATLAB and uses non-parametric density estimation methods[10] to predict the onset of bradycardia within the proposed failure rate of 5%. Although the theoretical method is discussed in detail [9], the implementation of said method is presented in this document. The aim of the document is to analyze the performance of the algorithm as well as to evaluate its performance based on common metrics. This would bridge the gap between the mathematical model and real-time application of the model.

1.3 Thesis Organization

This Thesis is organized as follows. Chapter 2 provides a brief background on non-parametric statistics and machine learning. In Chapter 3, we discuss a non-parametric density estimation based method to detect the bradycardia, introduce a hypothesis testing and the process used for algorithm automation. In Chapter 4, we present the main automation process and experimental results. We then conclude by discussing the future works in Chapter 5.

Chapter 2

BACKGROUND IN NON-PARAMETRIC STATISTICS AND MACHINE LEARNING

This chapter seeks to provide an introductory background to the various statistical and machine learning concepts that have been used to automate the method presented in [9]. Some of the key concepts covered are empirical risk minimization, bias variance decomposition, leave one out cross validation and non-parametric density estimation.

Statistics is the field of mathematics which involves the summary and analysis of numerical data in large quantities. The field of statistics can be divided into two general areas: descriptive statistics and inferential statistics. Descriptive statistics is a branch of statistics in which data are only used for descriptive purposes and are not employed to make predictions. Thus, descriptive statistics consists of methods and procedures for presenting and summarizing data. The procedures most commonly employed in descriptive statistics are the use of tables and graphs, and the computation of measures of central tendency and variability [11].

Inferential statistics employs data in order to draw inferences (i.e., derive conclusions) or make predictions. Typically, in inferential statistics sample data are employed to draw inferences about one or more populations from which the samples have been derived. Typically (although there are exceptions) the ideal sample to employ in research is a random sample. In a random sample, each subject or object in the population has an equal likelihood of being selected as a member of that

sample. Machine learning draws heavily from statistics in not only classification and categorization of data but also for algorithm building and model selection[11].

2.1 Machine Learning Methods

Machine learning can be broadly defined as computational methods using experience to improve performance or to make accurate predictions. Here, experience refers to the past information available to the learner, which typically takes the form of electronic data collected and made available for analysis. This data could be in the form of digitized human-labeled training sets, or other types of information obtained via interaction with the environment. In all cases, its quality and size are crucial to the success of the predictions made by the learner. An example of a learning problem is how to use a finite sample of randomly selected documents, each labeled with a topic, to accurately predict the topic of unseen documents. Clearly, the larger is the sample, the easier is the task. But the difficulty of the task also depends on the quality of the labels assigned to the documents in the sample, since the labels may not be all correct, and on the number of possible topics [2, 12].

Machine learning consists of designing efficient and accurate prediction algorithms. Some critical measures of the quality of these algorithms are their time and space complexity [13, 14]. However, in machine learning, there is a need for an additional notion of sample complexity to evaluate the sample size required for the algorithm to learn a family of concepts. Since the success of a learning algorithm depends on the data used, machine learning is inherently related to data analysis and statistics. More generally, learning techniques are data-driven methods combining fundamental concepts in computer science with ideas from statistics, probability and optimization.

Classification is the problem of assigning a category to each item in the data set. For example, document classification consists of assigning a category such as politics,

business, sports, or whether to each document, while image classification consists of assigning to each image a category such as car, train, or plane. The number of categories in such tasks is often less than a few hundreds, but it can be much larger in some difficult tasks and even unbounded as in OCR, text classification, or speech recognition.

The different stages of machine learning are described below:

- *Examples*: Items or instances of data used for learning or evaluation. This is also referred to as the data set.
- *Features*: The set of attributes, often represented as a vector, associated to an example.
- *Labels*: Values or categories assigned to examples. In classification problems, examples are assigned specific categories. In regression, items are assigned real-valued labels.
- *Hyperparameters*: Free parameters that are not determined by the learning algorithm, but rather specified as inputs to the learning algorithm.
- *Training sample*: Examples used to train a learning algorithm. The training sample varies for different learning scenarios as described in the next subsection.
- *Validation samples*: Examples used to tune the parameters of a learning algorithm when working with labeled data. The validation sample is used to select appropriate values for the learning algorithm's free parameters (hyperparameters).
- *Test Sample*: Examples used to evaluate the performance of a learning algorithm. The test sample is separate from the training and validation data and is

not made available in the learning stage. The learning algorithm predicts labels for the test sample based on features. These predictions are then compared with the labels of the test sample to measure the performance of the algorithm.

- *Loss function*: A function that measures the difference, or loss, between a predicted label and a true label. Denoting the set of all labels as y and the set of possible predictions as y' a loss function L is a mapping $L : y \times y' \rightarrow \mathbb{R}_+$. In most cases, $y' = y$ and the loss function is bounded, but these conditions do not always hold.
- *Hypothesis set*: A set of functions mapping features (feature vectors) to the set of labels y .

2.1.1 Learning Scenarios

The different learning scenarios for a learning algorithm differ in the types of training data available to the learner, the order and method by which training data is received and the test data used to evaluate the learning algorithm[12, 15].

- *Supervised learning* : The learner receives a set of labeled examples as training data and makes predictions for all unseen points. This is the most common scenario associated with classification, regression, and ranking problems.
- *Unsupervised learning* : The learner exclusively receives unlabeled training data, and makes predictions for all unseen points. Since in general no labeled example is available in that setting, it can be difficult to quantitatively evaluate the performance of a learner.
- *Semi-supervised learning* : The learner receives a training sample consisting of

both labeled and unlabeled data, and makes predictions for all unseen points. Semi-supervised learning is common in settings where unlabeled data is easily accessible but labels are expensive to obtain. Various types of problems arising in applications, including classification, regression, or ranking tasks, can be framed as instances of semi-supervised learning. The hope is that the distribution of unlabeled data accessible to the learner can help it achieve a better performance than in the supervised setting. The analysis of the conditions under which this can indeed be realized is the topic of much modern theoretical and applied machine learning research.

- *Transductive inference* : As in the semi-supervised scenario, the learner receives a labeled training sample along with a set of unlabeled test points. However, the objective of transductive inference is to predict labels only for these particular test points. Transductive inference appears to be an easier task and matches the scenario encountered in a variety of modern applications. However, as in the semi-supervised setting, the assumptions under which a better performance can be achieved in this setting are research questions that have not been fully resolved.
- *On-line learning* : In contrast with the previous scenarios, the online scenario involves multiple rounds where training and testing phases are intermixed. At each round, the learner receives an unlabeled training point, makes a prediction, receives the true label, and incurs a loss. The objective in the on-line setting is to minimize the cumulative loss over all rounds or to minimize the regret, that is the difference of the cumulative loss incurred and that of the best expert in hindsight. Unlike the previous settings just discussed, no distributional assumption is made in on-line learning. In fact, instances and their labels may

be chosen adversarially within this scenario.

- *Reinforcement learning* : The training and testing phases are also intermixed in reinforcement learning. To collect information, the learner actively interacts with the environment and in some cases affects the environment, and receives an immediate reward for each action. The object of the learner is to maximize his reward over a course of actions and iterations with the environment. However, no long-term reward feedback is provided by the environment, and the learner is faced with the exploration versus exploitation dilemma, since he must choose between exploring unknown actions to gain more information versus exploiting the information already collected.
- *Active learning* : The learner adaptively or interactively collects training examples, typically by querying an oracle to request labels for new points. The goal in active learning is to achieve a performance comparable to the standard supervised learning scenario (or passive learning scenario), but with fewer labeled examples. Active learning is often used in applications where labels are expensive to obtain, for example computational biology applications.

2.1.2 Model Selection

A key problem in learning algorithm is the selection of the hypothesis set \mathcal{H} . The choice of \mathcal{H} is subject to a trade-off that can be analyze using estimation and approximation errors[12, 13].

Let \mathcal{H} be a family of functions mapping \mathcal{X} to $\{-1, +1\}$. The excess error of a hypothesis h chosen from \mathcal{H} , that is the difference between its error $R(h)$ and the Bayes error R^* , can be decomposed as follows:

$$R(h) - R^* = \underbrace{\left(R(h) - \inf_{h \in \mathcal{H}} R(h) \right)}_{\text{estimation}} + \underbrace{\left(\inf_{h \in \mathcal{H}} R(h) - R^* \right)}_{\text{approximation}}$$

The first term is called the estimation error, the second term the approximation error. The estimation error depends on the hypothesis h selected. It measures the error of h with respect to the infimum of the errors achieved by hypotheses in \mathcal{H} . The approximation error measures how well the Bayes error can be approximated using \mathcal{H} . It is a property of the hypothesis set \mathcal{H} , a measure of its richness. For a more complex or richer hypothesis \mathcal{H} , the approximation error tends to be smaller at the price of a larger estimation error. Model selection consists of choosing \mathcal{H} with a favorable trade-off between the approximation and estimation errors.

Empirical Risk Minimization

Empirical Risk Minimization is a standard algorithm by which the estimation error can be bounded. ERM seeks to minimize the error on the training sample and the following discussion is adapted from [12]:

$$h_S^{\text{ERM}} = \operatorname{argmin}_{h \in \mathcal{H}} \widehat{R}_S(h)$$

where $\widehat{R}_S(h)$ is the calculated error on the training sample composed of data derived from the sample space S .

Let us consider that there is a joint probability distribution $P(x, y)$ over X and Y , and the training set consists of n instances $(x_1, y_1), \dots, (x_n, y_n)$ drawn i.i.d. from $P(x, y)$. It is also assumed there is a non-negative real-valued loss function $L(\hat{y}, y)$ which measures how different the prediction \hat{y} of a hypothesis is from the true outcome y . The hypothesis h is a function of x , therefore, $\hat{y} = h(x)$. The loss function can be




Figure 2.1: The Estimation and Approximation Error Are Represented in Blue and Green Respectively. These Errors Are the Distance of a Chosen Hypothesis h^* from h_1 and h_2

written as $L(\hat{y}, y) = L(h(x), y) = h(x) - y$ [12]. The risk associated with hypothesis $h(x)$ is then defined as the expectation of the loss function $L(h(x), y)$:

$$R(h) = \mathbf{E}[L(h(x), y)] = \int L(h(x), y)dP(x, y)$$

In general, the risk $R(h)$ cannot be computed because the distribution $P(x, y)$ is unknown to the learning algorithm. However, an approximation (called empirical risk) is calculated by averaging the loss function on the training set:

$$R_{\text{emp}}(h) = \frac{1}{n} \sum_{i=1}^n L(h(x_i), y_i),$$

where x_i and y_i are data points drawn from the training set of length n . Thus the learning algorithm defined by the ERM principle consists in solving the above optimization problem [12], [15]. The concept of the loss function $L(h(x), y)$ and risk is elaborated further in Section 2.3.4 and Section 2.4.

2.2 Bias Variance Decomposition

The loss function in Section 2.1.2 is the L_2 loss which is calculated using the calculated estimate and the true value of the estimate at a particular point. For any such estimation, there are a number of errors present. This section aims to analyze such errors and present a bound on the L_2 loss as observed in [15]. Given a data set $D = (x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ it is assumed these points are independent identically distributed (i.i.d) and drawn from the same unknown distribution $P(x, y)$. For non-parametric density estimation the parameters of $P(x, y)$ are unknown but it is still possible to estimate $P(x, y)$ and this estimate is denoted by f . If this f is used to design a hypothesis to classify our given data set according to some condition, there appears two sources of errors in the estimation of f as elaborated in [15]:

1. **Bias:** This kind of error is caused by the inability of f to estimate $P(x, y)$ correctly.

2. **Variance:** This kind of error is caused by the presence of random noise in the data which causes the estimate f to be inaccurate and vary from $P(x, y)$

Let us assume any arbitrary point z drawn from the distribution $P(x, y)$ such that it is not a sample data point in the dataset D . The estimate at z is represented by

$$\gamma = f(z) + \epsilon.$$

Note that γ is random. It is worth mentioning $f(z)$ is the value of the estimate of f at z and ϵ is added noise such that

$$\begin{aligned} \mathbb{E}[\gamma] &= f(z), \\ \text{Var}(\gamma) &= \text{Var}(\epsilon) \end{aligned} \tag{2.1}$$

A square loss is assumed, $|P(z) - f(z)|^2$, the difference between the value of the actual distribution P and the estimate f at z . The risk which is the expectation of the squared loss is computed. Risk, $R(z)$ is written as

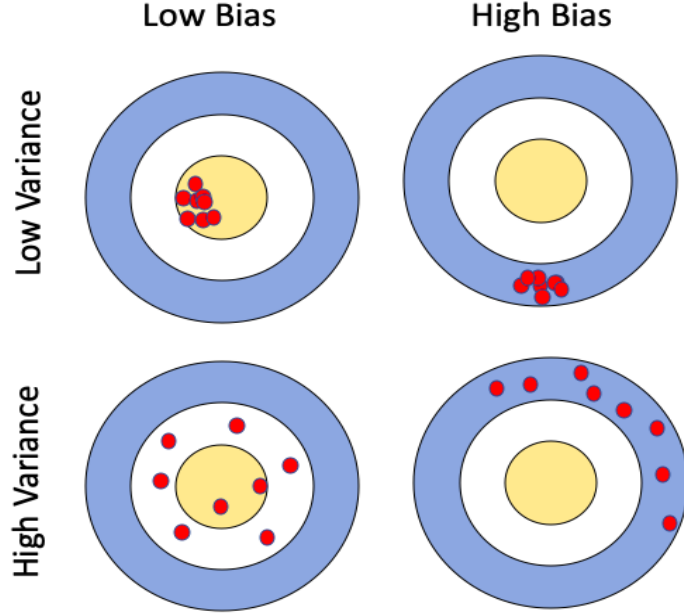


Figure 2.2: Illustration of the Presence of Bias/Variance (or Both) in a Data Set.

$$\begin{aligned}
R(z) &= \text{E} [(P(z) - \gamma)^2] \\
&= \text{E} [P(z)^2] + \text{E} [\gamma^2] - 2\text{E}[\gamma P(z)] \\
&= \text{Var}(P(z)) + \text{E}[P(z)]^2 + \text{Var}(\gamma) + \text{E}[\gamma]^2 - 2\text{E}[\gamma]\text{E}[P(z)] \quad (2.2) \\
&= (\text{E}[P(z)] - \text{E}[\gamma])^2 + \text{Var}(P(z)) + \text{Var}(\gamma) \\
&= (\text{E}[P(z)] - f(z))^2 + \text{Var}(P(z)) + \text{Var}(\epsilon)
\end{aligned}$$

$\text{E}[P(z)] - f(z)$ is defined to be the bias. If the bias is larger than zero, the estimator is said to be positively biased, if the bias is smaller than zero, the estimator is negatively biased, and if the bias is exactly zero, the estimator is unbiased. The higher the bias, the worse our estimate is a fit for the actual distribution $P(x, y)$. Bias is inherent to a particular model and a large amount of bias suggests underfitting [2, 12].




Figure 2.3: Figure Illustrating the Bias Variance Decomposition and the Optimal Complexity

$\text{Var}(P(z))$ is the variance of the estimator (kernel density, regression, etc.). The variance as the difference between the expected value of the squared estimator minus the squared expectation of the estimator. A large variance may indicate that our model is highly specialized to only one dataset, model is extremely complicated, which indicates overfitting.

$\text{Var}(\epsilon)$ is called irreducible error. It is an error that cannot be decomposed further and will always be present in the estimate. This is also called noise.

In real world scenarios, $P(x, y)$ is rarely known and especially in the case of non-parametric estimations, we lack knowledge about the parameters of $P(x, y)$. However,

in calculating the estimate at arbitrary points, as seen above, there exists three different sources of error. Therefore, it is necessary to minimize the risk of any estimation method so that the bias and variance are both minimized. There is however, a trade-off between the two [13, 15].

As seen in Figure 2.3, the most complex models overfit the data while the simple models underfit the data. Therefore the optimum complexity for a model is given by the values of bias and variance which minimize the total error.

2.3 Density Estimation

Density estimation is the process of reconstructing the probability density using a set of given data points or observations. In statistics, a random variable is a variable whose value depends on a random phenomenon [16]. For instance, the outcome of a coin toss, the event that it rains in Phoenix tomorrow or the sum of any two numbers between 0 and 1 are all random variables. A random variable can be either discrete (having specific values) or continuous (any value in a continuous range). Some outcomes of a random variable are more likely to occur (high probability density) and other outcomes are less likely to occur (low probability density). The overall shape of the probability density is referred to as a probability distribution, and the calculation of probabilities for specific outcomes of a random variable is performed by a probability density function (PDF).

From the PDF of certain data it can be judged whether a given observation is likely, or unlikely. One of the most popular and common ways to visualize the probability density function is the histogram (as seen in Figure 2.4). A histogram is a graph of the frequency distribution in which the vertical axis represents the count (frequency) and the horizontal axis represents the possible range of the data values. A histogram is created by dividing up the range of the data into a small number of intervals or




Figure 2.4: A Histogram Showing the Maximum Daily Temperatures at La Guardia Airport, New York. The Data Is Drawn from The *airquality* Dataset in R Which Contains the Daily Air Quality Measurements in New York, May-september, 1973

bins. The number of observations falling in each interval is counted. This gives a frequency distribution [17].

The shape of the probability density function across the domain for a random variable is referred to as the probability distribution and common probability distributions have names, such as uniform, normal, exponential, and so on. Given a random variable the aim is the estimation of a density of its probabilities. Therefore, it follows that one would want to know what the probability density looks like. However, in most cases the distribution of the random variable is not known. This is because there is no available knowledge of all possible outcomes for a random variable, only a small set of observations. However, the distribution can be estimated and this is called density estimation. In density estimation, the small sample of observed

variables are used to estimate the overall probability distribution [10].

Generalizing further to the case of any distribution of a set of random variables, let $D = X_1, X_2, \dots, X_n$ be independent identically distributed (i.i.d.) real valued random variables that share a common distribution. The density of this distribution, denoted by $p(x)$, is a function on \mathbb{R} from $[0, +\infty)$, but is unknown. An estimator of $p(x)$ is a function $x \mapsto p_n(x) = p_n(x, X_1, \dots, X_n)$ measurable with respect to the observations $\mathbf{X} = (X_1, \dots, X_n)$.

As described earlier, inferential statistics is aimed at making inferences about the larger population from which the samples are drawn. The main goals of inferential statistics are: parameter estimation, data prediction and model comparison. There are two main approaches in inferential statistics:

- *Frequentist*: The frequentist school only uses conditional distributions of data given specific hypotheses. The presumption is that some hypothesis (parameter specifying the conditional distribution of the data) is true and that the observed data is sampled from that distribution. In particular, the frequentist approach does not depend on a subjective prior that may vary from one investigator to another. In frequentist approaches, only repeatable random events have probabilities. These probabilities are equal to the long-term frequency of occurrence of the events in question. No probability is attached to hypotheses or to any fixed but unknown values [18].
- *Bayesian*: In contrast, the Bayesian school models uncertainty by a probability distribution over hypotheses. The ability to make inferences depends on the degree of confidence in the chosen prior, and the robustness of the findings to alternate prior distributions may be relevant and important. Bayesian approaches associate probabilities to any event or hypotheses. Probabilities are

also attached to non-repeatable events[16].

Frequentist measures like p -values and confidence intervals continue to dominate research, especially in the life sciences. However, in the current era of powerful computers and big data, Bayesian methods have undergone an enormous renaissance in fields like machine learning and genetics. In this method a Bayesian approach is used to estimate the unknown density function.

2.3.1 Parametric Density Estimation

If it is known a priori that $p(x)$ belongs to a parametric family $\{g(x, \theta) : \theta \in \Theta\}$, where $g(\cdot, \cdot)$ is a given function, and Θ is a subset of \mathbb{R}^k with a fixed dimension k independent of n , the estimator for $p(x)$ is equivalent to the estimation of the finite-dimensional parameter θ . This is a parametric problem of estimation [12, 16].

If the shape of the unknown distribution follows well-known sets/classes of distributions it can be estimated using certain parameters (like the mean, median, standard deviation etc.). For instance, the normal distribution has two parameters: the mean and the standard deviation. Given these two parameters, the probability distribution function is now known. These parameters can be estimated from data by calculating the sample mean and sample standard deviation. This process is called parametric density estimation [19]. There are two ways in which the parameter can be estimated:

Maximum Likelihood Estimation : In this case, the parameter θ is unknown but fixed. Given the data, θ is chosen such that it maximizes the probability of obtaining the samples that have already been observed. the following discussion has been adapted from [19]. The density $p(x)$ is completely specified by parameter $\theta = [\theta_1, \theta_2, \dots, \theta_k]$. If $p(x)$ is Gaussian with $N(\mu, \sigma^2)$ then $\theta = [\mu, \sigma^2]$. Since $p(x)$ depends on θ , it can be denoted by $p(x | \theta)$, where $p(x | \theta)$ is not a conditional density

but only demonstrates dependence. If $p(x)$ is $N(\mu, \sigma^2)$ then $D = X_1, X_2, \dots, X_n$ are i.i.d. samples from $N(\mu, \sigma^2)$ and,

$$p(D | \theta) = \prod_{k=1}^{k=n} p(x_k | \theta) = F(\theta) \quad (2.3)$$

Equation (2.3) is called the likelihood of θ with respect to the observations D .

The value of θ that maximizes the likelihood function $p(D | \theta)$ is given by

$$\hat{\theta} = \arg \max_{\theta} (p(D | \theta)) \quad (2.4)$$

Instead of maximizing $p(D | \theta)$ it is often easier to maximize $\ln(p(D | \theta))$. Since log is a monotonic function Equation (2.4) is rewritten as,

$$\arg \max_{\theta} (p(D | \theta)) = \arg \max_{\theta} (\ln p(D | \theta))$$

Therefore,

$$\hat{\theta} = \arg \max_{\theta} (\ln p(D | \theta)) = \arg \max_{\theta} \left(\ln \prod_{k=1}^{k=n} p(x_k | \theta) \right) = \arg \max_{\theta} \left(\sum_{k=1}^n \ln p(x_k | \theta) \right) \quad (2.5)$$

Let us consider the Gaussian parameter mentioned before. It is assumed that $p(x | \mu)$ is $N(\mu, \sigma^2)$ where σ^2 is known but μ is unknown and needs to be estimated. Therefore, $\theta = \mu$ for this problem and using Equation (2.5) it can be concluded that:

$$\begin{aligned} \hat{\mu} &= \arg \max_{\mu} \left(\sum_{k=1}^n \ln p(x_k | \mu) \right) \\ &= \arg \max_{\mu} \left(\sum_{k=1}^n \ln \left(\frac{1}{\sqrt{2\pi}\sigma} \exp \left(-\frac{(x_k - \mu)^2}{2\sigma^2} \right) \right) \right) \\ &= \arg \max_{\mu} \sum_{k=1}^n \left(-\ln \sqrt{2\pi}\sigma - \frac{(x_k - \mu)^2}{2\sigma^2} \right) \end{aligned}$$

For easier notation, the previous equation is rewritten as $M(\mu)$, and subsequently,

$$\arg \max_{\mu} (M(\mu)) = \arg \max_{\mu} \sum_{k=1}^n \left(-\ln \sqrt{2\pi\sigma} - \frac{(x_k - \mu)^2}{2\sigma^2} \right) \quad (2.6)$$

Taking the derivative of Equation (2.6),

$$\frac{d}{d\mu} (M(\mu)) = \sum_{k=1}^n \frac{1}{\sigma^2} (x_k - \mu) = 0$$

Simplifying further,

$$\sum_{k=1}^n x_k - n\mu = 0$$

and finally,

$$\hat{\mu} = \frac{1}{n} \sum_{k=1}^n x_k \quad (2.7)$$

As seen in Equation (2.7), the maximum likelihood estimator of the mean is just the average value of the observed samples, D . Equation (2.7) makes intuitive sense since in general, one would assume that the mean of a set of data is the numerical average.

Bayesian Estimation : In this method, the observed data is fixed and different values of θ are assumed. Therefore, unlike the maximum likelihood approach, θ is now the random variable. The following discussion has been adapted from [20]. The aim is the estimation of θ given $D = X_1, X_2, \dots, X_n$. It is assumed that θ is continuous. The posterior probability distribution of θ is given by $p(\theta | I)$ and it should be normalized such that

$$\int_{-\infty}^{\infty} p(\theta | I) d\theta = 1 \quad (2.8)$$

$p(\mathbf{D} | \theta, I)$ is the sampling distribution for D given the model implied by I and θ . Baye's Theorem tells us that posterior distribution for θ should be

$$p(\theta | \mathbf{D}, I) = \frac{p(\mathbf{D} | \theta, I)p(\theta | I)}{p(\mathbf{D} | I)} \quad (2.9)$$

It is true that Equation (2.9) is also normalized such that,

$$\int_{-\infty}^{\infty} p(\theta | \mathbf{D}, I) d\theta = 1$$

The denominator $p(\mathbf{D} | I)$ in Equation (2.9) can be evaluated as

$$p(\mathbf{D} | I) = \int_{-\infty}^{\infty} p(\mathbf{D}, \theta | I) d\theta = \int_{-\infty}^{\infty} p(\mathbf{D} | \theta, I)p(\theta | I) d\theta \quad (2.10)$$

Since Equation (2.10) only depends on D and not the parameter θ , one can say that Equation (2.9) is actually

$$p(\theta | \mathbf{D}, I) \propto p(\mathbf{D} | \theta, I)p(\theta | I)$$

where \propto implies that there is some numerical constant M that equates the terms on the left of \propto to the terms on the right. The value of M does not depend on θ but may depend on D .

$$p(\theta | \mathbf{D}, I) = M(\mathbf{D})p(\mathbf{D} | \theta, I)p(\theta | I)$$

M can be calculated using

$$\int_{-\infty}^{\infty} p(\theta | \mathbf{y}, I) d\theta = 1$$

The Gaussian case is considered where only μ is unknown ($\theta = \mu$). Therefore, there is a need to establish the posterior probability of μ denoted by $p(\mu)$. Assuming that

$$p(x | \mu) \sim N(\mu, \sigma^2), \quad (2.11)$$

and,

$$p(\mu) \sim N(\mu_0, \sigma_0^2), \quad (2.12)$$

where σ, μ_0, σ_0 are known to us.

Here, $p(\mu)$ is akin to $p(\theta | I)$ (the posterior probability of θ). The aim is to find $p(\theta | \mathbf{D}, I)$ using Equation (2.9). By assumption of independence,

$$p(\mathbf{D} | \theta, I) = \prod_{k=1}^{k=n} p(x_k | \theta) \quad (2.13)$$

To solve for Equation (2.9), Equation (2.11), Equation (2.12) and Equation (2.13) are used.

$$p(\mu | \mathbf{D}, I) = \alpha \prod_{k=1}^n p(x_k | \mu) p(\mu) \quad (2.14)$$

where α is the scale parameter as discussed previously and is independent of μ . As x_k is normally distributed, $p(x_k | \mu)$ and $p(\mu)$ are updated with the specific equations:

$$p(x_k | \mu) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{x_k - \mu}{\sigma} \right)^2 \right] \quad (2.15)$$

$$p(\mu) = \frac{1}{(2\pi\sigma_0^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{\mu - \mu_0}{\sigma_0} \right)^2 \right] \quad (2.16)$$

Substituting Equation (2.15) and Equation (2.16) in Equation (2.14),

$$p(\mu | \mathbf{D}) = \alpha \prod_{k=1}^n \frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{x_k - \mu}{\sigma} \right)^2 \right] \frac{1}{(2\pi\sigma_0^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{\mu - \mu_0}{\sigma_0} \right)^2 \right]$$

$$p(\mu | \mathbf{D}) = \alpha \prod_{k=1}^n \frac{1}{(2\pi\sigma^2)^{1/2}} \frac{1}{(2\pi\sigma_0^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{\mu - \mu_0}{\sigma_0} \right)^2 - \frac{1}{2} \left(\frac{x_k - \mu}{\sigma} \right)^2 \right]$$

Simplifying further,

$$p(\mu | \mathbf{D}) = \alpha' \exp \sum_{k=1}^n \left(-\frac{1}{2} \left(\frac{\mu - \mu_0}{\sigma_0} \right)^2 - \frac{1}{2} \left(\frac{x_k - \mu}{\sigma} \right)^2 \right)$$

where

$$\alpha' = \frac{1}{(2\pi\sigma^2)^{1/2}} \frac{1}{(2\pi\sigma_0^2)^{1/2}}$$

On further simplification,

$$p(\mu | \mathbf{D}) = \alpha'' \exp \left[-\frac{1}{2} \left(\frac{n}{\sigma^2} + \frac{1}{\sigma_0^2} \right) \mu^2 - 2 \left(\frac{1}{\sigma^2} \sum_{k=1}^n x_k + \frac{\mu_0}{\sigma_0^2} \right) \mu \right] \quad (2.17)$$

Comparing Equation (2.17) to the Gaussian distribution in the standard form:

$$p(\mu | \mathbf{D}) = \frac{1}{(2\pi\sigma_n^2)^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{\mu - \mu_n}{\sigma_n} \right)^2 \right] \quad (2.18)$$

where,

$$\mu_n = \left(\frac{n\sigma_0^2}{n\sigma_0^2 + \sigma^2} \right) \bar{x}_n + \frac{\sigma^2}{n\sigma_0^2 + \sigma^2} \mu_0 \quad (2.19)$$

and,

$$\sigma_n^2 = \frac{\sigma_0^2 \sigma^2}{n\sigma_0^2 + \sigma^2} \quad (2.20)$$

It is important to note that for Gaussian random variables the variance and the mean are the required parameters to estimate the underlying distribution. Therefore, it is important to emphasize both μ_n and σ_n^2 .

In order to find $p(x | \mathbf{D})$ the following equation is used

$$p(x | \mathbf{D}) = \int p(x, \theta | \mathbf{D}) d\theta = \int p(x | \theta) p(\theta | \mathbf{D}) d\theta \quad (2.21)$$

$p(\mu | \mathbf{D})$ is given by Equation (2.18) and $p(x | \mu)$ is given by Equation (2.11).

Therefore, Equation (2.21) can be rewritten as:

$$p(x | \mathbf{D}) = \int p(x | \mu) p(\mu | \mathbf{D}) d\mu$$

$$p(x | \mathbf{D}) = \int \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x - \mu}{\sigma}\right)^2\right] \frac{1}{\sqrt{2\pi}\sigma_n} \exp\left[-\frac{1}{2}\left(\frac{\mu - \mu_n}{\sigma_n}\right)^2\right] d\mu$$

Substituting the Equation (2.19) and Equation (2.20),

$$p(x | \mathbf{D}) = \frac{1}{2\pi\sigma\sigma_n} \exp\left[-\frac{1}{2}\frac{(x - \mu)^2}{\sigma^2 + \sigma_n^2}\right] \int \exp\left[-\frac{1}{2}\frac{\sigma^2 + \sigma_n^2}{\sigma^2\sigma_n^2}\left(\mu - \frac{\sigma_n^2\bar{x}_n + \sigma^2\mu_n}{\sigma^2 + \sigma_n^2}\right)^2\right] d\mu$$

Therefore, $p(x | \mathbf{D})$ is distributed normally as,

$$p(x | D) \sim N(\mu_n, \sigma^2 + \sigma_n^2)$$

Thus, it is observed that for Bayesian parametric estimation, the underlying distribution is estimated using the posterior probability and conditional densities obtained from the given data.

2.3.2 Non-Parametric Density Estimation

In some cases however, a priori knowledge about $p(x)$ is not known. In this case it is usually assumed that $p(x)$ belongs to some class \mathcal{P} of densities [2, 12]. For example, \mathcal{P} can be the set of all the continuous probability densities on \mathbb{R} . In this

case, parametric density estimation is not feasible and alternative methods must be used to estimate the density $p(x)$. The distributions still have parameters, but they cannot be controlled or used in the process of estimation since they are unknown. One of the most common non-parametric density approaches is the kernel density estimator (KDE) [2] which is elaborated in the next section.

2.3.3 Kernel Density Estimator

A kernel is a mathematical weighting function that returns a probability for a given value of a random variable. The kernel effectively smooths or interpolates the probabilities across the range of outcomes for a random variable such that the sum of probabilities equals one [11].

Smoothing parameter (h) : The smoothing parameter or bandwidth controls how wide the probability mass is spread around a particular point as well as controlling the smoothness or roughness of a density estimate. In other words, the bandwidth controls the number of samples or window of samples used to estimate the probability for a new point. It is often denoted by h . It is important to note that this h is different from the hypothesis h mentioned in Section 2.1.2 and henceforth the use of h in this document denotes the bandwidth of a kernel. A large bandwidth may result in a rough density with few details, whereas a small window may have too much detail but not be general enough to cover any new unseen samples. Therefore, it is important to select the correct bandwidth h . In this project, the best value of h is learned through leave one out cross validation[21] elaborated in Section 2.4.

The contribution of samples in a window is governed by one of the many kernel basis functions denoted by $K(\cdot)$. A list of the kernel basis functions used in this project can be found in Table 2.1.

| Kernel | Equation |
|------------------------------|--|
| Gaussian | $K(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2}$ |
| Epanechnikov (parabolic) | $K(u) = \frac{3}{4}(1 - u^2)$, support: $ u \leq 1$ |
| Uniform (rectangular window) | $K(u) = \frac{1}{2}$, support: $ u \leq 1$ |
| Cosine | $K(u) = \frac{\pi}{4} \cos\left(\frac{\pi}{2}u\right)$, support: $ u \leq 1$ |

Table 2.1: The Basis Functions of the Different Kernels Implemented

The derivation for the kernel density estimator mentioned in the following discussion has been adapted from [2]. Generalizing to the case of a set of random variables $D = X_1, X_2, X_3, \dots, X_n$ which are independent identically distributed (i.i.d) and have a probability density $p(x)$, for any small $h > 0$, $p(x)$ can be calculated as

$$p(x) \approx \frac{F(x+h) - F(x-h)}{2h}, \quad (2.22)$$

such that $x \in \mathcal{R}$ and $F(x)$ is defined as

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n I(X_i \leq x)$$

where $I()$ is the indicator function. By the strong law of large numbers,

$$F_n(x) \rightarrow F(x), \quad \forall x \in \mathcal{R}$$

which allows us to replace $F(x)$ with $F_n(x)$ in Equation (2.22),

$$\hat{p}_n^R(x) = \frac{F_n(x+h) - F_n(x-h)}{2h} \quad (2.23)$$

Equation (2.23) is the estimator of $p(x)$ called the Rosenblatt Estimator and it can be rewritten as,




Figure 2.5: All of the Implemented Kernels on a Common Coordinate System

$$\hat{p}_n^R(x) = \frac{1}{2nh} \sum_{i=1}^n I(x-h < X_i \leq x+h) = \frac{1}{nh} \sum_{i=1}^n K_0\left(\frac{X_i - x}{h}\right),$$

where n is the total number of samples in D , h is the bandwidth of the kernel, K_0 is the selected kernel basis function and x is the point at which the kernel density estimate is to be calculated. Generalizing further the equation can be simplified to

$$\hat{p}_n(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{X_i - x}{h}\right) \quad (2.24)$$

Equation (2.24) is the kernel density estimator and $\int K(u)du = 1$, where K is called the kernel and h is the bandwidth associated with a particular estimator [18].

Multivariate Density Estimation

This application is concerned with multivariate density estimation. Consider that X_1, \dots, X_n constitute an i.i.d. q -vector ($X_i \in \mathbb{R}^q$, for some $q > 1$) having a common

PDF $f(x) = f(x_1, x_2, \dots, x_q)$. Let X_{is} denote the s th component of $X_i (s = 1, \dots, q)$. Using the product kernel function (separability of kernels) constructed from the product of univariate kernel functions, the PDF $f(x)$ is estimated by

$$\hat{f}(x) = \frac{1}{nh_1 \dots h_q} \sum_{i=1}^n K\left(\frac{X_i - x}{h}\right)$$

where

$$K\left(\frac{X_i - x}{h}\right) = k\left(\frac{X_{i1} - x_1}{h_1}\right) \times \dots \times k\left(\frac{X_{iq} - x_q}{h_q}\right),$$

and where $k(\cdot)$ is a univariate kernel function.

Similar to the univariate case, the optimal smoothing parameter h_{opt} should balance the squared bias and variance term, i.e., $h_s^4 = O((nh_1 \dots h_q)^{-1})$ for all s . Thus, $h_s = c_s n^{-1/(q+4)}$ for some positive constant $c_s (s = 1, \dots, q)$. It is assumed that the data is independently distributed over all i . The multivariate kernel density estimator is capable of capturing general dependence among the different components of X_i [2, 16].

2.3.4 Mean Square Error of Kernel Estimators

A basic measure of accuracy of the estimator is the Mean Square Error (MSE) also called the mean squared risk at any arbitrary point $x_0 \in \mathbb{R}$. According to [2], the MSE is defined as

$$\text{MSE} = \text{MSE}(x_0) \triangleq \mathbf{E}_p [(\hat{p}_n(x_0) - p(x_0))^2] \quad (2.25)$$

where \mathbf{E}_p denotes the expectation with respect to the random distribution of X_1, X_2, \dots, X_n denoted by p

$$\mathbf{E}_p [(\hat{p}_n(x_0) - p(x_0))^2] \triangleq \int \dots \int (\hat{p}_n(x_0, x_1, \dots, x_n) - p(x_0))^2 \prod_{i=1}^n [p(x_i) dx_i]$$

We have

$$\text{MSE} = b^2(x_0) + \sigma^2(x_0)$$

where

$$b(x_0) = \mathbf{E}_p[\hat{p}_n(x_0)] - p(x_0)$$

and

$$\sigma^2(x_0) = \mathbf{E}_p[(\hat{p}_n(x_0) - \mathbf{E}_p[\hat{p}_n(x_0)])^2].$$

The quantities $b(x_0)$ and $\sigma^2(x_0)$ are called the bias and variance of the estimator at a particular point x_0 . As seen in the previous section, the error of an estimator can be decomposed into a sum of its bias and variance at a particular point.

To better understand how the MSE depends on the bias and variance, each of these terms are analyzed separately as seen in [2]

It is known that the MSE at a random point x is

$$\begin{aligned} \text{MSE}(\hat{p}_n(x)) &\equiv \mathbf{E}\{[\hat{p}_n(x) - p(x_0)]^2\} \\ &= \text{var}(\hat{p}_n(x)) + [\mathbf{E}(\hat{p}_n(x)) - p(x_0)]^2 \\ &\equiv \text{var}(\hat{p}_n(x_0)) + [\text{bias}(\hat{p}_n(x))]^2 \end{aligned}$$

To analyze the bias the Taylor Series expansion formula[2] is used. For an univariate function $g(x)$ that is m times differentiable,

$$\begin{aligned} g(x) = &g(x_0) + g^{(1)}(x_0)(x - x_0) + \frac{1}{2!}g^{(2)}(x_0)(x - x_0)^2 + \\ &\dots + \frac{1}{(m-1)!}g^{(m-1)}(x_0)(x - x_0)^{m-1} + \frac{1}{m!}g^{(m)}(\xi)(x - x_0)^m, \end{aligned}$$

where

$$g^{(s)}(x_0) = \left. \frac{\partial^s g(x)}{\partial x^s} \right|_{x=x_0},$$




Figure 2.6: Squared Bias, Variance and Mean Squared Error (Solid Line) as Functions Of h As Seen In [2]

and ξ lies between x and x_0 .

The bias term bias ($\hat{p}_n(x)$) can be simplified as

$$\begin{aligned}
 &= \mathbb{E} \left\{ \frac{1}{nh} \sum_{i=1}^n k \left(\frac{X_i - x}{h} \right) \right\} - p(x) \\
 &= h^{-1} \mathbb{E} \left[k \left(\frac{X_1 - x}{h} \right) \right] - p(x) \\
 &\quad \text{(by identical distribution)} \\
 &= h^{-1} \int p(x_1) k \left(\frac{x_1 - x}{h} \right) dx_1 - p(x) \\
 &= h^{-1} \int p(x + hv) k(v) h dv - p(x)
 \end{aligned}$$

$$\begin{aligned}
& \text{(change of variable : } x_1 - x = hv) \\
& = \int \left\{ p(x) + p^{(1)}(x)hv + \frac{1}{2}p^{(2)}(x)h^2v^2 + O(h^3) \right\} k(v)dv - p(x) \\
& = \left\{ p(x) + 0 + \frac{h^2}{2}p^{(2)}(x) \int v^2k(v)dv + O(h^3) \right\} - p(x) \\
& = \frac{h^2}{2}p^{(2)}(x) \int v^2k(v)dv + O(h^3)
\end{aligned}$$

where the $O(h^3)$ term comes from

$$(1/3!)h^3 \left| \int f^{(3)}(\tilde{x})v^3k(v) \right| dv \leq Ch^3 \int |v^3k(v)| dv = O(h^3),$$

where C is a positive constant, and \tilde{x} lies between x and $x + hv$.

It has been assumed that $p(x)$ is three times differentiable, however, one can also assume that it is two times differentiable and subsequently weaken the initial assumption. In that case, bias is written as

$$\begin{aligned}
\text{bias}(\hat{p}_n(x)) &= E(\hat{p}_n(x)) - p(x) \\
&= \frac{h^2}{2}p^{(2)}(x) \int v^2k(v)dv + o(h^2)
\end{aligned}$$

Next analyzing the variance term, where $\text{var}(\hat{p}_n(x)) =$

$$\begin{aligned}
& \text{var} \left[\frac{1}{nh} \sum_{i=1}^n k \left(\frac{X_i - x}{h} \right) \right] \\
&= \frac{1}{n^2h^2} \left\{ \sum_{i=1}^n \text{var} \left[k \left(\frac{X_i - x}{h} \right) \right] + 0 \right\} \\
& \quad \text{(by independence)} \\
&= \frac{1}{nh^2} \text{var} \left(k \left(\frac{X_1 - x}{h} \right) \right)
\end{aligned}$$

$$\begin{aligned}
& \text{(by identical distribution)} \\
&= \frac{1}{nh^2} \left\{ \mathbb{E} \left[k^2 \left(\frac{X_1 - x}{h} \right) \right] - \left[\mathbb{E} \left(k \left(\frac{X_1 - x}{h} \right) \right) \right]^2 \right\} \\
&= \frac{1}{nh^2} \left\{ \int p(x_1) k^2 \left(\frac{x_1 - x}{h} \right) dx_1 - \left[\int p(x_1) k \left(\frac{x_1 - x}{h} \right) dx_1 \right]^2 \right\} \\
&= \frac{1}{nh^2} \left\{ h \int p(x + hv) k^2(v) dv - \left[h \int p(x + hv) k(v) dv \right]^2 \right\} \\
&= \frac{1}{nh^2} \left\{ h \int [p(x) + p^{(1)}(\xi)hv] k^2(v) dv - O(h^2) \right\} \\
&= \frac{1}{nh} \left\{ p(x) \int k^2(v) dv + O \left(h \int |v| k^2(v) dv \right) - O(h) \right\} \\
&= \frac{1}{nh} \{ \kappa f(x) + O(h) \}
\end{aligned}$$

where,

$$\kappa = \int k^2(v) dv$$

Thus it is seen that by the correct selection of h the MSE for $p(x)$ can be minimized. By choosing $h = cn^{1/\alpha}$ the conditions for the consistent estimation of $p(x)$ are satisfied. However, the question still remains regarding what values of c and α should be used. For a given sample size n , if h is too small the resulting estimator will have a small bias but a large variance. On the other hand if h is too large, the estimator will have a large bias but a small variance. So to minimize $\text{MSE}(\hat{p}_n(x))$ the bias and variance terms have to be balanced [2, 13].

The optimal choice of h should satisfy

$$d\text{MSE}(\hat{p}_n(x))/dh = 0$$

It can be shown that the value of h which minimizes MSE at a point x is given by

$$h_{opt} = c(x)n^{-1/5}$$

where

$$c(x) = \left\{ \kappa p(x) / [\kappa_2 p^{(2)}(x)]^2 \right\}^{1/5}$$

It is important to note that $\text{MSE}(\hat{p}_n(x))$ is a calculation at a particular point x , and that the value of h minimizing the MSE is different depending on the choice of x . For instance, the value of h that minimizes the error at a point at the tail end of a distribution is different from the value of h that would minimize MSE if x was the average [2].

The main concern is the minimization of error and subsequently the risk while computing the estimators using kernel density, globally - that is, for all x in the support of $p(x)$. In this case, the optimal h is obtained minimizing the mean integrated square error (MISE) which is discussed in the next section.

2.4 Cross Validation to Minimize Risk

This section explains the method by which the optimal bandwidth h_{opt} is determined for the kernel density estimator as seen in [2, 21]. In the previous section, the MSE was defined in Equation (2.25) as

$$\text{MSE} = \text{MSE}(x_0) \triangleq \mathbf{E}_p [(\hat{p}_n(x_0) - p(x_0))^2]$$

which is calculated at a fixed arbitrary point x_0 . However, in analyzing the risk of the estimator, the global risk must be analyzed. The global risk is defined as the Mean Integrated Squared error which can be calculated as follows

$$\text{MISE} \triangleq \mathbf{E}_p \int (\hat{p}_n(x) - p(x))^2 dx \quad (2.26)$$

The MISE can be rewritten as $\text{MISE}(h)$ to denote that it is a function of the bandwidth h and the ideal value of h can be defined as

$$h_{\text{id}} = \arg \min_{h>0} \text{MISE}(h) \quad (2.27)$$

Since $\text{MISE}(h)$ depends on the unknown value of p (which is not known), a different approach to find the h_{id} must be employed. In this case, an unbiased estimation of the risk through leave one out cross validation can be used. Instead of minimizing $\text{MISE}(h)$ the approximately unbiased estimator of $\text{MISE}(h)$ is minimized.

It is important to note that

$$\text{MISE}(h) = \mathbf{E}_p \int (\hat{p}_n(x) - p(x))^2 dx = \mathbf{E}_p \left[\int (\hat{p}_n(x))^2 dx - 2 \int \hat{p}_n(x)p(x) dx \right] + \int (p(x))^2 dx \quad (2.28)$$

The last term does not depend on h and therefore can be excluded from the minimization. The first term is the expectation with respect to the distribution of $X_1, X_2, X_3, \dots, X_n$ denoted by p . Therefore, the minimizer for h_{id} seeks to minimize the term

$$J(h) \triangleq \mathbf{E}_p \left[\int (\hat{p}_n(x))^2 dx - 2 \int \hat{p}_n(x)p(x) dx \right] \quad (2.29)$$

where $J(h)$ is the risk function. The second term,

$$\int \hat{p}_n(x)p(x) dx \quad (2.30)$$

can be written as

$$\mathbf{E}_p[\hat{p}(x)],$$

where \mathbf{E}_x denotes the expectation with respect to \mathbf{X} and not with respect to the random variables \mathbf{X}_i . Therefore, this term is the expectation of $\hat{p}_n(x)$ with respect to

$p(x)$. Rewriting the integral in Equation (2.30) as

$$\hat{p}_{n,-i}(x) = \frac{1}{(n-1)h} \sum_{j \neq i} K\left(\frac{X_j - x}{h}\right) \quad (2.31)$$

which is the leave one out kernel estimator for $f(X_i)$. Finally, the first term in Equation (2.28), $\int (\hat{p}_n(x))^2 dx$ can be estimated as

$$\int (\hat{p}_n(x))^2 dx = \frac{1}{n^2 h^2} \sum_{i=1}^n \sum_{j=1}^n \int k\left(\frac{X_i - x}{h}\right) k\left(\frac{X_j - x}{h}\right) dx = \frac{1}{n^2 h} \sum_{i=1}^n \sum_{j=1}^n \bar{k}\left(\frac{X_i - X_j}{h}\right) \quad (2.32)$$

where $\bar{k}(v) = \int k(u)k(v-u)du$ is the twofold convolution kernel derived from $k(\cdot)$. If

$$k(v) = \exp(-v^2/2) / \sqrt{2\pi},$$

a standard normal kernel, then

$$\bar{k}(v) = \exp(-v^2/4) / \sqrt{4\pi}$$

, a normal kernel (i.e. normal PDF) with mean zero and variance 2, which follows since two independent $N(0, 1)$ random variables sum to a $N(0, 2)$ random variable[2].

For this project, the convolution kernel is computed analytically in the programming language R. Combining Equation (2.31) and Equation (2.32), the leave one out cross validation estimator which minimizes h is obtained as

$$CV_f(h) = \frac{1}{n^2 h} \sum_{i=1}^n \sum_{j=1}^n \bar{k}\left(\frac{X_i - X_j}{h}\right) - \frac{2}{n(n-1)h} \sum_{i=1}^n \sum_{j \neq i, j=1}^n k\left(\frac{X_i - X_j}{h}\right). \quad (2.33)$$

This equation can be easily generalized for multivariate estimation

$$CV_f(h_1, \dots, h_q) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \bar{K}_h(X_i, X_j) - \frac{2}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i, j=1}^n K_h(X_i, X_j),$$

where

$$K_h(X_i, X_j) = \prod_{s=1}^q h_s^{-1} k\left(\frac{X_{is} - X_{js}}{h_s}\right)$$

$$\bar{K}_h(X_i, X_j) = \prod_{s=1}^q h_s^{-1} \bar{k}\left(\frac{X_{is} - X_{js}}{h_s}\right),$$

and $\bar{k}(v)$ is the two fold convolution kernel based upon $k(\cdot)$, where $k(\cdot)$ is the univariate kernel function.

Note that h_{id} can be found by using different numerical search algorithms. For this particular problem, a range of values for h is defined such that $h = h_1, h_2, h_3, \dots, h_n$. For every value of h , $CV_f(h)$ is computed and then the value of h for which $CV_f(h)$ is minimum [2, 21] is obtained.

Thus, $CV_f(h)$ yields an unbiased estimator of $MISE(h)$, which is independent of h . This means that the functions $MISE(h)$ and $\mathbf{E}_{\mathbf{p}}[CV_f(h)]$ have the same minimizers. In turn, the minimizers of $\mathbf{E}_{\mathbf{p}}[CV_f(h)]$ can be approximated by those of the function $CV(\cdot)$ which can be computed from the observations $X_1, X_2, X_3, \dots, X_n$. This means that

$$h_{CV} = \arg \min CV(h) = h_{\text{id}}.$$

As discussed in this chapter, given a data set of random variables, the main aim is to estimate the density of these points through non-parametric methods by using the kernel density estimator. Four kernel basis functions are chosen: Gaussian, Epanechnikov, uniform (boxcar/rectangular window) and cosine. The best value of h for each of these methods is learned through leave one out cross validation. The chosen value of h minimizes the loss function and subsequently the global risk of the chosen estimation method. The mathematics of the exact process employed is discussed in the next chapter.

Chapter 3

ALGORITHM AUTOMATION

In this chapter, the problem statement and the specific application of the concepts and methods discussed in the previous chapter are discussed. While Chapter 2 provided a general overview of the proposed method and concepts employed, this chapter analyzes the specific applications of the concepts previously discussed to the chosen data set. Referring to the mathematical model presented in [9] this chapter seeks to relate the theoretical concepts to the practical implementation.

Following the method presented in [9], the ECG data is obtained from the Preterm Infant Cardio-Respiratory Signals (PICS) database [1, 22]. This database contains simultaneous ECG and respiration recordings of ten preterm infants collected from the Neonatal Intensive Care Unit (NICU) of the University of Massachusetts Memorial Healthcare. Statistical features based on linear estimates of heart rate are used to predict episodes of bradycardia.

3.1 Data Collection

As presented in [9], ten pre-term infants were studied, with post-conceptional age of $29 \frac{3}{7}$ to $34 \frac{2}{7}$ weeks (mean: $31 \frac{1}{7}$ weeks) and study weights of 843 to 2100 grams (mean: 1468 grams). The infants were spontaneously breathing room air and did not have any congenital or perinatal infections or health complications. A single channel of a 3-lead electrocardiogram (ECG) signal was recorded at 500 Hz for 20-70 hours per infant. In absence of an ECG channel, a compound ECG signal was recorded (250Hz) [1].

Each ECG signal is pre-processed and segmented such that each segment contains both normal and unhealthy beats. R-peak information is extracted using the Pan-Tompkin algorithm[23] and estimate the non specific probability density function. After setting a desired level of false alarm to be tolerated by the system, a threshold region is found using the estimated density and a generated threshold plane. R-peaks from further along in time are used to test against this threshold region to determine the onset of near-term bradycardia [9].

3.2 Non-parametric Prediction

Once the R-peaks have been obtained, a non-parametric method is employed to determine the kernel-based probability density function which is used to predict the onset of bradycardia in the future [2, 9].

Assuming the number of R peaks in an ECG segment is N , the R-tuple is defined as $x = (t_n, R_n)$, $n=1, \dots, N$ and $x_n \in \mathcal{X}$, where t_n is the sample at which the peak occurs and R_n is the amplitude of the peak. Therefore, the set of R tuples $X = x_1, x_2, \dots, x_N$ is i.i.d. and drawn from an unknown distribution $p(x)$. As discussed in Chapter 2, using the kernel density estimator, one can estimate $p(x)$ as

$$\hat{p}_H(x) = \frac{1}{N} \sum_{n=1}^N K_H(x - x_n)$$

for some positive definite bandwidth matrix H . without loss of generality, it is assumed that $H = h^2 I_2$, where I_2 is the (2x2) identity matrix. The best value of h is learned using leave one out cross validation to ensure unbiased estimation [9, 21].

In order to construct the hypothesis set, a probability of false alarm P_{FA} is declared such that the hypothesis testing produces $(1 - P_{FA})$ confidence. For this application, $P_{FA} = 5\%$ and therefore the hypothesis testing produces 95% confidence with which the future onset of bradycardia is predicted. The null hypothesis is defined as \mathcal{H}_0 as

the hypothesis that the density of the next R-tuple x_{N+1} is the same as that of the previous R-tuples in the set \mathcal{X}_N ; that is, $\mathcal{H}_0 : x_{N+1} = x$, for all possible values of $x \in \mathcal{X}$ [9]. The aim is to construct a confidence set $\mathcal{A}_{\mathcal{X}}$, that consists of all values X_N , such that the probability of the next R-tuple, x_{N+1} , belonging to this set satisfies the following condition

$$\Pr(x_{N+1} \in \mathcal{A}_{\mathcal{X}}) \geq (1 - P_{\text{FA}})$$

The kernel density estimator $\hat{p}_H^a(x)$ based on the augmented data set $X_N \cup \{x\}$ for a fixed value of $x \in \mathcal{X}$. is then used The rank of $\hat{p}_H^a(x_1), \dots, \hat{p}_H^a(x_{N+1})$ is uniformly distributed under the null hypothesis. Thus, for each value of x , the p-value η_x is given by,

$$\eta_x = \frac{1}{N+1} \sum_{n=1}^N \mathbb{I}(\hat{p}_H^a(x_n) \leq \hat{p}_H^a(x)),$$

where \mathbb{I} is the indicator function. Therefore, the confidence set $(1 - P_{\text{FA}})$ is defined as

$$\mathcal{A}_{\mathcal{X}} = \{x : \eta_x \geq P_{\text{FA}}\}$$

$\mathcal{A}_{\mathcal{X}}$ is distribution free [24] and only determined using the finite set X_N [25]. This method is impractical since X_N are merely limited observations drawn from the original unknown distribution. Therefore, a larger set which contains $\mathcal{A}_{\mathcal{X}}$ and can be easily constructed named $\mathcal{B}_{\mathcal{X}}$ is defined. This set is easier to compute and preserves accuracy [9].

Additionally, $y_n = \hat{p}_H(x_n)$ for $n = 1, \dots, N$ is also defined and it is assumed that y_i is sorted in ascending order, that is $y_1 \leq \dots \leq y_N$. The prediction set $\mathcal{B}_{\mathcal{X}}$ can be constructed as

$$\mathcal{B}_{\mathcal{X}} = \{x : \hat{p}_H(x) \geq \mathcal{C}_k\}, \quad (3.1)$$

The threshold plane \mathcal{C}_k can be compute as

$$\mathcal{C}_k = y_k - (K_H(0)/N|H|^{1/2})$$

where $k = \lceil (N + 1)P_{FA} \rceil$.

The set $\mathcal{B}_{\mathcal{X}} \supset \mathcal{A}x$ satisfies

$$\Pr(x_{N+1} \in \mathcal{A}_{\mathcal{X}}) \geq (1 - P_{FA}) \implies \Pr(x_{N+1} \in \mathcal{B}_{\mathcal{X}}) \geq (1 - P_{FA}).$$

In other words, the prediction set $\mathcal{B}_{\mathcal{X}}$ is the projection of the estimated density which is above the threshold \mathcal{C}_k [9]. Once this prediction set has been obtained, any R-tuple $x_m = (t_m, R_m)$ is predicted with 95% confidence to be the onset of bradycardia if

$$x_m \notin \mathcal{B}_{\mathcal{X}}, m > N. \tag{3.2}$$

The proposed model employs unsupervised learning since the training data is unlabeled and predictions are made on all unseen points. As with all machine learning approaches, the data is segmented into training, validation and testing sets. However, before this can be done, the ECG signals must be analyzed and preprocessed.

3.3 Data Preprocessing

Each ECG signal is preprocessed by first removing baseline wander, which is a low frequency of around 0.5-0.6Hz. It is removed by using a high pass filter with cut off frequency between 0.5 and 0.6Hz. There are two other values present in the ECG header files obtained from the database:

- *Gain*: This is a floating-point number that specifies the difference in sample values that would be observed if a step of one physical unit occurred in the original analog signal. For ECGs, the gain is usually roughly equal to the R-wave amplitude in a lead that is roughly parallel to the mean cardiac electrical axis. If the gain is zero or missing, this indicates that the signal amplitude is not calibrated. For this dataset, every ECG recording has its own value of gain [1].

| Parameters | Infant 1 | Infant 2 | Infant 3 | Infant 4 | Infant 5 | Infant 6 | Infant 7 | Infant 8 | Infant 9 | Infant 10 |
|--------------------|----------|----------|----------|----------|----------|----------|----------|----------|----------|-----------|
| Bradycardia events | 77 | 72 | 80 | 66 | 72 | 56 | 34 | 28 | 97 | 40 |
| Duration (hours) | 45.6 | 43.8 | 43.7 | 46.8 | 48.8 | 48.6 | 20.3 | 24.6 | 70.3 | 45.1 |

Table 3.1: Number of Bradycardia Events and Duration of Ecg for Each Infant as Obtained From [1]

- *Base*: The base is a floating point number that specifies the counter value corresponding to sample 0. For this database, the base counter is 16 which means that sample 0 is shifted to the 16th position in the files [1].

The ECG signals are adjusted for the gain and base values for each file. These values can be obtained from the header file for each infant contained in the database.

Once the baseline wander has been removed, each signal is segmented into events. The number of events is equal to the number of bradycardia events for that infant which is obtained from the annotation file. For instance, the annotation file for infant 5 contains the onset of 72 bradycardia events. Therefore the cleaned ECG signal for this infant is segmented into 72 events where each event contains some samples before and after the bradycardia onset. For this project, 5000 samples before and 2500 samples after the onset of bradycardia were chosen randomly to be included in each event. Therefore, every event has 7501 samples.

3.4 Peak Detection

The R-peaks are detected by running the Pan-Tompkin algorithm on each event. Each detected peak x_n is a tuple (t_n, R_n) for every event for a particular infant. These peaks are used as the data to run the rest of the model for bradycardia prediction. The R peaks are stored in .csv files and are used to generate the training, validation and testing set for the algorithm.

3.5 Dataset Segmentation

Once the peaks are detected, the entire data set of peaks is randomly shuffled and then segmented into training, validation and testing set according to some pre-determined demarcations. To this end, a percentage of total points to be allocated to each set is determined. The sets are generated using MATLAB following which the training set is used for kernel density estimation. The output of the kernel estimator is used to calculate C_k and to generate the confidence set \mathcal{B}_x according to Equation (3.1) and Equation (3.2).

3.5.1 Training Set

The training set is used to train the model - that is, this is the set used for kernel density estimation of the original probability density function (pdf) [?]. The training set contains the location and amplitude of R peaks as detected by the Pan Tompkin algorithm. Before the set can be used for kernel density estimation, the location of R peaks is scaled by the sampling frequency and normalized. The training set is also used to generate the points on the x and y axis (evaluation points) where the kernel density estimation is to be calculated. The results from the estimation are used to compute the threshold C_k and to generate the confidence set \mathcal{B}_x .

3.5.2 Validation Set

Before the estimator can be run on the training set, the leave one out cross validation is used on the validation set to learn the best value of h . The validation set is used to prevent reuse of data points for both training and cross validation and subsequently prevent bias. The leave one out cross validation is run for the chosen kernel [9].

3.5.3 Test Set

The testing condition based on the prediction set is defined according to Equation (3.1). A point m from the test set is said to be the onset of bradycardia if

$$x_m \notin \mathcal{B}_x, m > N. \quad (3.3)$$

The condition $m > N$ is maintained since it is assumed that the points in the test set are separate from the training data (which contains N points).

3.5.4 Monte-Carlo Simulations

To obtain the average error Monte-Carlo simulations for different demarcations of the entire event data set. For instance, the data set is split into: a) 60%, 20%, 20%, b) 70%, 20%, 10% and c) 70%, 10%, 20% to form three separate training, validation and testing sets respectively.

The annotation file of every infant contains the onset of bradycardia while the Pan-Tompkin output contains the location of the peaks. It is reasonably assumed that the peak locations that occur after the bradycardia onset are the bradycardia peaks [1]. The question then arises how far into the future from the onset of bradycardia should the algorithm search to locate the bradycardia peak. It is assumed that it needs to search k samples after the bradycardia onset. Since each event was generated by taking 2500 samples after the onset m , $k < 2500$ otherwise the end of each event set would be reached

In theory, $m < k < m + 2500$. Therefore, to calculate the appropriate value of k , a random number between u is generated such that $u \in (0, 1]$. The value of k is then calculated as

$$m < k < m + \text{ceil}(u * 1500).$$

For a test point, a peak x_n is said to be a bradycardia peak if it satisfies the

condition for k outlined above. The estimated prediction error (EPE) is calculated as

$$\text{EPE} = \frac{\text{Number of False Alarms}}{\text{Number of R-tuples tested}}.$$

Since a lower value of EPE indicates that lesser false alarms are generated, the test error is a measure of performance of this method [9]. For a fixed value of P_{FA} the total test error for each infant is calculated in Chapter 4.

3.5.5 Testing Mechanism

Once \mathcal{B}_x has been generated using \mathcal{C}_k , the convex hull of \mathcal{B}_x is calculated to determine its boundary. Every point from the test set is then used to assess whether it belongs to \mathcal{B}_x by determining its position with respect to the boundary. If a point lies outside the boundary, it is said to lie outside \mathcal{B}_x and therefore be the onset of bradycardia [9].

The implementation of the test statistic as seen in Equation (3.2) depends on the creation of the Equation (3.1). Starting from the raw ECG signal, the signal is pre-processed, segmented into events and the R-peaks are detected using the Pan-Tompkin algorithm. The peaks are then stored as tuples and used to create the training set on which the kernel density estimator is run. The validation set is used to estimate h and finally the test error is calculated by evaluating the performance of the algorithm on the test set. This is elaborated in more detail in the Chapter 4.

Chapter 4

EXPERIMENTS

In this section, the simulation results of the process discussed in Chapter 3 are presented. First each ECG signal is pre-processed by accounting for the gain of the recording machine and the baseline wander. Once this is done, each ECG signal is segmented into events. Each event contains N peaks which are detected using the Pan-Tompkin algorithm. The simulation results are categorized infant to infant - i.e., all the steps of the algorithm are executed for the events generated by each infant and calculate an average test error by running Monte Carlo simulations for each infant. Therefore, for a particular infant, if there are M total events, the Pan-Tompkin is run M times to detect the peaks in each event. The output of the Pan-Tompkin is stored in a large matrix and then shuffled randomly to generate the training, validation and test sets.

The output of the Pan-Tompkin is a tuple (t_n, r_n) where t_n is the location of the peak and r_n (where $n = 1, 2, \dots, N$) is the amplitude at that location [23]. The set of tuples that constitute the training set is the input to the kernel density estimator where each tuple represents an instance of a random variable [9]. It is assumed that the tuples are independent and identically drawn from the same underlying distribution. In the first step of the automation the aim is to estimate this distribution using the kernel density estimator. The initial kernel density estimate is calculated using the built in *density()* function in R. However, in calculation of the threshold \mathcal{C}_k and creation of the confidence set \mathcal{B}_x a custom function designed for this purpose called *kerneldensity()* is used.

Once the baseline wander has been removed, the peaks and their locations are

| Infant | F_s (Hz) | Base | Gain |
|--------|------------|------|-----------|
| 1 | 250 | 16 | 800.6597 |
| 2 | 500 | 16 | 1220.7707 |
| 3 | 250 | 16 | 1140.7954 |
| 4 | 500 | 16 | 834.3036 |
| 5 | 500 | 16 | 800.6597 |
| 6 | 500 | 16 | 800.6141 |
| 7 | 500 | 16 | 1283.8528 |
| 8 | 500 | 16 | 1420.7631 |
| 9 | 500 | 16 | 800.6597 |
| 10 | 500 | 16 | 800.4159 |

Table 4.1: Information Obtained from the Header File (.hea) for Each Infant from the Database.

detected using the Pan-Tompkin algorithm. Following which they are randomly shuffled and used to create the training, validation and testing sets according to some predetermined ratio. Then the training set is used as input to the kernel density estimator.

The estimator output is used to calculate the threshold \mathcal{C}_k and to create the confidence set \mathcal{B}_x according to the following equation

$$\mathcal{B}_x = \{x : \hat{p}_H(x) \geq \mathcal{C}_k\}.$$

This means that all of the values of $\hat{p}_H(x) < \mathcal{C}_k$ are rejected. This gives us the desired confidence set. The convex hull of this shape is calculated to establish its boundary. This is important because for a point x_m to be considered the onset of bradycardia

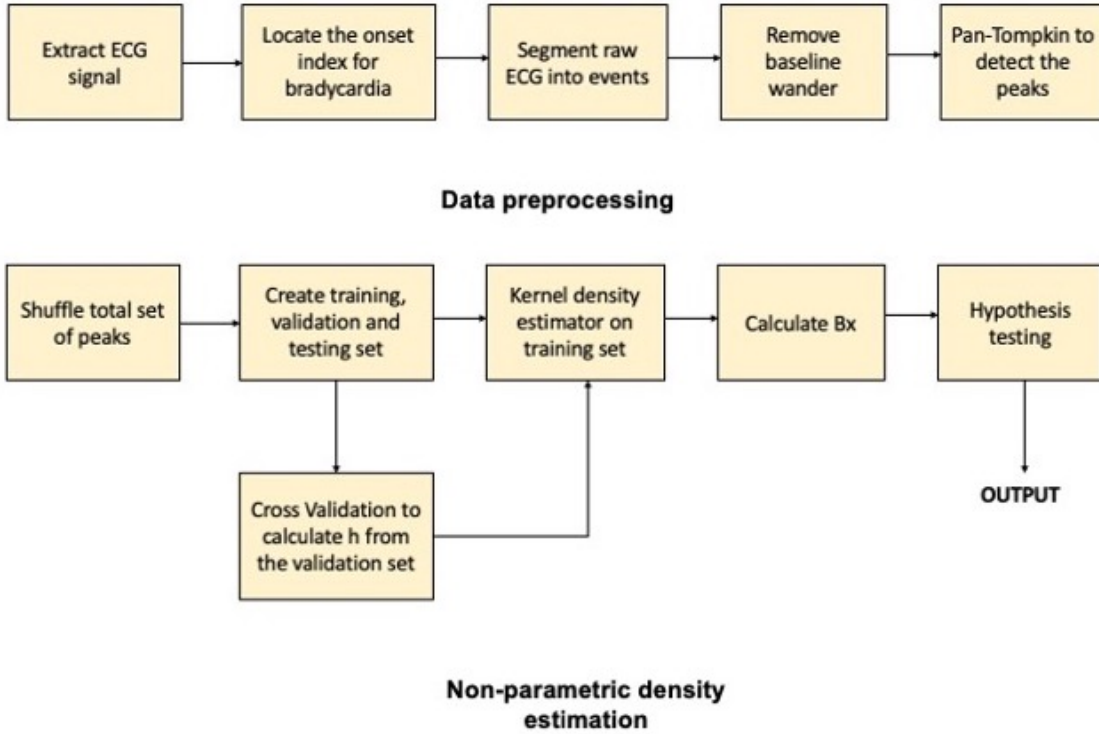


Figure 4.1: Block Diagram Representation of the Preprocessing and Density Estimation Stage. The Bradycardia Onset Indices Can Be Found from The .atr File Available in the Database Used.

it has to satisfy the following equation

$$x_m \notin \mathcal{B}_x, m > N.$$

Therefore, if the points in the generated test set lie within the convex hull of \mathcal{B}_x , they are not considered to be the onset of bradycardia whereas all points lying outside \mathcal{B}_x are considered to be the opposite.

4.1 Evaluation Metrics

To evaluate the performance of the kernel density estimator, the following metrics are used [14]:

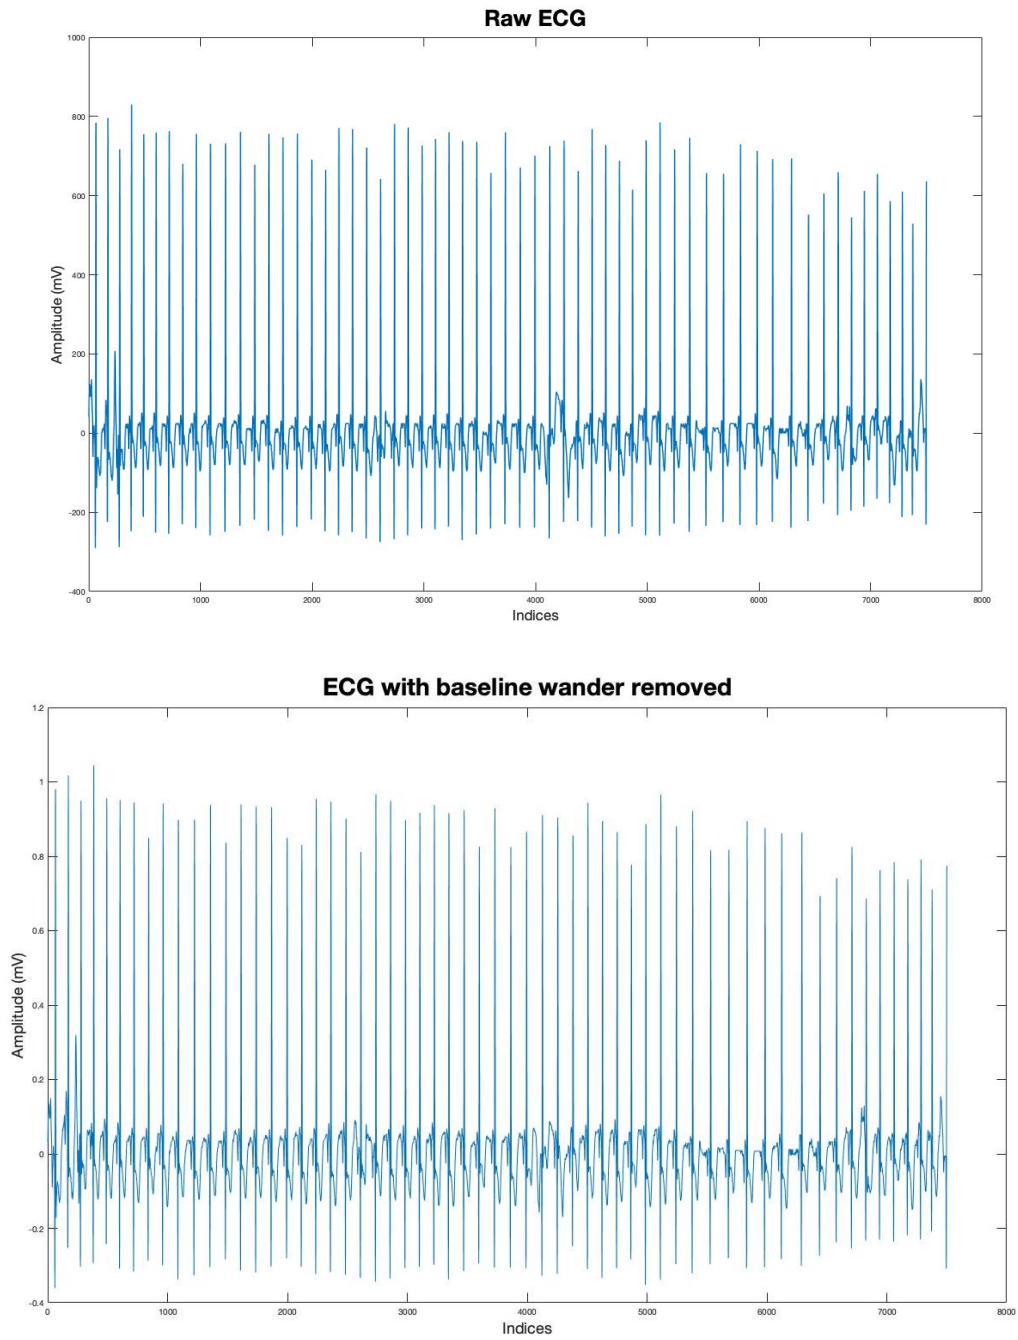


Figure 4.2: A Segment of the Raw Ecg for Infant 5, Obtained Using The `rdsamp()` Command [1]

- *Test error* : The test error is defined as the ratio of the number of false classifications and the number of tuples tested.
- *Sensitivity(Recall)* : Sensitivity is defined as the ability of an algorithm to predict a positive outcome when the actual outcome is positive. In this case, it is the ability of the algorithm to predict the onset of bradycardia correctly.
- *Specificity (Precision)* : The ability of an algorithm to not predict a positive outcome when the outcome is not positive. In this case, specificity is the ability of the algorithm to classify the healthy heartbeat correctly.
- *False Discovery Rate* : The ratio of all false positive classifications to the total number of all positive classifications. As the name indicates this metric points to all discoveries that are classified as positive but are actually negative.
- *False Omission Rate* : As the name suggests, this score is the ratio of all false negative classifications to the sum of all negative classifications. This score displays how many positive instances are falsely omitted as being negative.
- *Accuracy* : The ratio of correctly predicted observations to the total observations. This metric is an evaluation of how accurately the algorithm predicts the desired outcome.
- *F1 score* : The weighted average of Precision and Recall. This score takes both false positives and false negatives into account.

These metrics can be decided from the confusion matrix as seen in Table 4.2 and Table 4.3. The confusion matrix is the tabular representation of each combination of prediction and actual value. Confusion matrices require a binary outcome. In this case, the positive case is that the point is the onset of bradycardia i.e., it lies outside

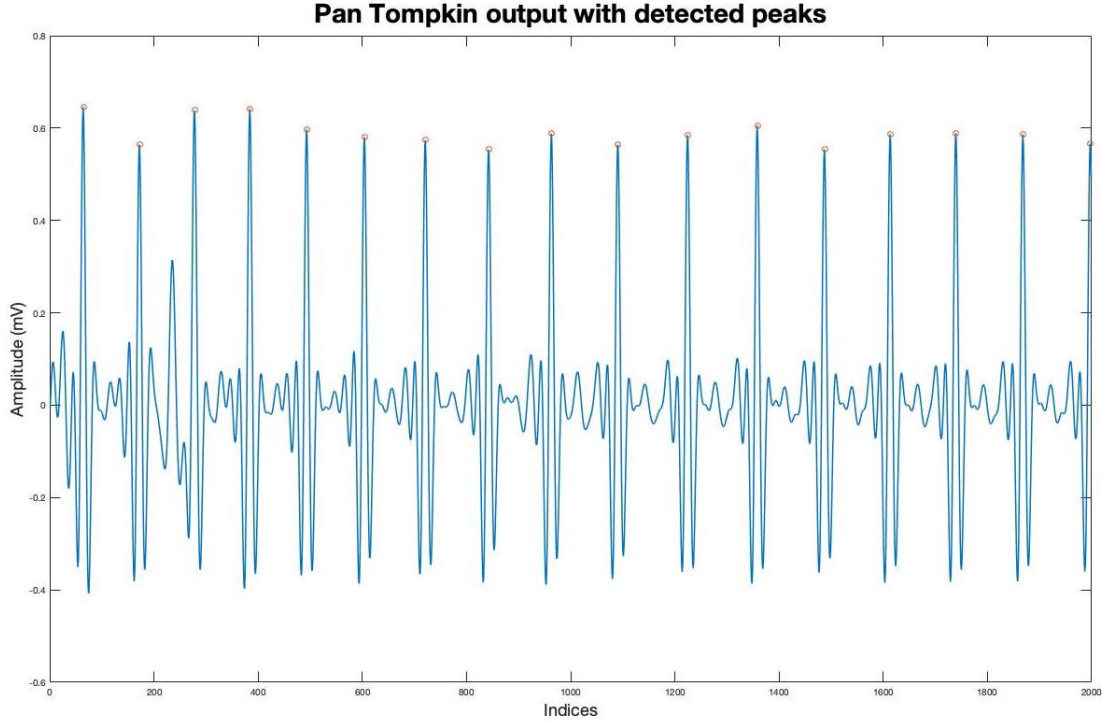


Figure 4.3: Pan-Tompkin Output for an ECG Segment for Infant 5

the convex hull (denoted by 0) and it is not the onset (denoted by 1). The average test error for infant 5 is calculated in Table 4.4.

In Table 4.4, **Error1** = average test error calculated when data is split into the training/validation/testing set in the ratio of .7/.1/.2. For **Error2** the split is .7/.2/.1 and for **Error3** the split is .6/.2/.2. This error is the result of Monte Carlo simulations run on the data. 20 such simulations are run to calculate the average test error.

On the other hand, Table 4.5 is the result of running the Gaussian kernel on the data only once. The total number of tuples in the test set for infant 5= 468. Therefore, the estimated predictive error (EPE) is calculated as:

$$\text{EPE} = \frac{\text{Number of False Alarms}}{\text{Number of R-tuples tested}}$$

where, Number of False classifications = FP + FN . From Table 4.2, FN = 9 and FP




Figure 4.4: Estimate $\hat{p}_H(x)$ for Gaussian Kernel.

= 26, therefore,

$$\text{EPE} = \frac{(26 + 9)}{468} = 7.478\%.$$

As observed this value is very close to the average test error for Gaussian kernels presented in Table 4.4.

The value of EPE for other kernels is shown in Table 4.5. These findings are consistent with the average test error for each kernel presented in Table 4.4. To evaluate the performance of this algorithm, additional metrics are calculated from

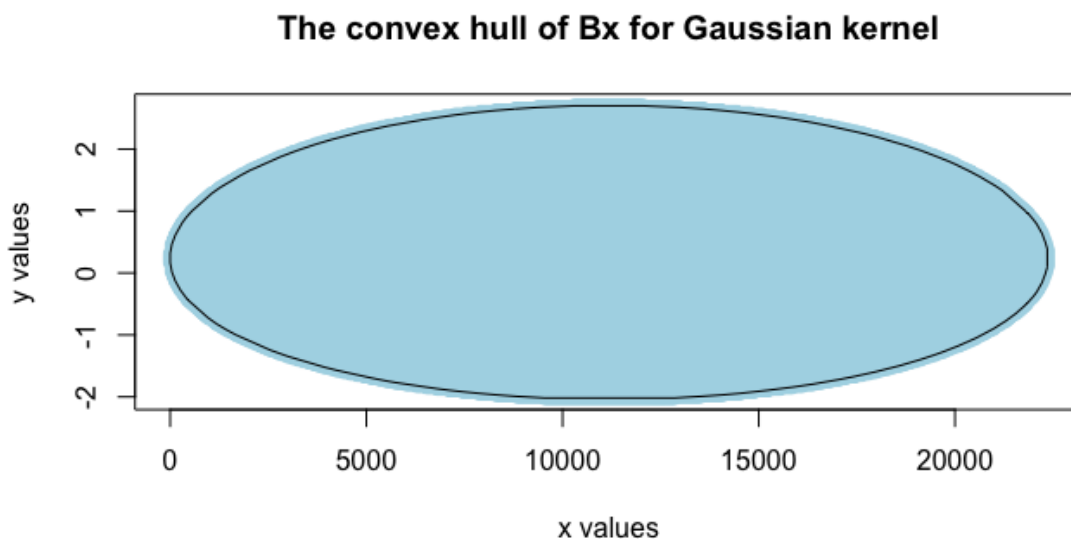


Figure 4.5: Plots Showing the Kernel Density Estimate for Infant 5 Using the Gaussian Kernel with $h = 4.981171$ and $\mathcal{C}_k = 0.002179295$ and the Obtained Confidence Set \mathcal{B}_x . The Convex Hull for the Confidence Set Is Outlined in Black.

| Classifications | ActuallyPositive | ActuallyNegative |
|--------------------------|-------------------------|-------------------------|
| PredictedPositive | True Positives (TP) | False Positives (FP) |
| PredictedNegative | False Negatives (FN) | True Negatives (TN) |

Table 4.2: A General Confusion Matrix

Table 4.2.

Finally, all metrics mentioned in Section 4.1 are evaluated for the Gaussian Kernel and presented in Table 4.6.

| Classifications | Actually Positive | Actually Negative |
|---------------------------|-------------------|-------------------|
| Predicted Positive | 18 | 26 |
| Predicted Negative | 9 | 415 |

Table 4.3: Confusion Matrix for Gaussian Kernel Run on Infant 5

| Kernel | h | C_k | Error1 | Error2 | Error3 |
|--------------|----------|------------------------|------------|------------|------------|
| Gaussian | 4.981171 | 0.002179295 | 7.90000005 | 7.6266667 | 7.95026705 |
| Epanechnikov | 10.52929 | 1.02×10^{-05} | 6.9500001 | 6.5284985 | 6.8285046 |
| Cosine | 10.23305 | 1.26×10^{-07} | 7.4 | 7.2599999 | 7.4782377 |
| Uniform | 2.675163 | 0.004355385 | 8.4401727 | 8.19666675 | 8.45300395 |

Table 4.4: Average Test Error for Each Kernel as Calculated for Infant 5 Through Monte Carlo Simulation

| Kernel | Number of points tested | Test error |
|--------------|-------------------------|------------|
| Epanechnikov | 468 | 6.41025% |
| Cosine | 468 | 7.2649% |
| Uniform | 468 | 8.5470% |

Table 4.5: Calculated Epe for Epanechnikov, Cosine and Uniform Kernel Run on Infant 5

4.2 Discussion

From the results of the Monte Carlo simulations in Table 4.4, it is observed that the Epanechnikov kernel performs the best with respect to the EPE. This is because

| Metric | Formula | Value |
|----------------------|-----------------------------|--------------|
| Sensitivity | $\frac{TP}{TP+FN}$ | 0.75 |
| Precision | $\frac{TP}{TP+FP}$ | 0.4091 |
| False Discovery Rate | $\frac{FP}{TP+FP}$ | 0.5909 |
| False Omission Rate | $\frac{FN}{FN+TN}$ | 0.0212 |
| Accuracy | $\frac{TP+TN}{TP+TN+FP+FN}$ | 0.9252 |
| F ₁ score | $\frac{TP}{TP+0.5(FP+FN)}$ | 0.5070 |

Table 4.6: Calculated Metrics for Infant 5

the Epanechnikov kernel is optimal with respect to MSE. Due to the MSE minimization employed in this estimation process, it ensures the best performance of the Epanechnikov kernel. On the contrary, the uniform kernel performs the worst out of the four kernels studied. This makes sense since the uniform kernel assigns equal weight to all points in its support.

These findings are consistent with Table 4.5 where the Epanechnikov kernel performs the best and the uniform kernel performs the worst in terms of the test error. The accuracy of the method is 92.52% as seen in Table 4.6. The code for this implementation can be found in [26].

FUTURE WORKS

The previous sections have highlighted the motivations for the work presented in this document, the related theories implemented and the results obtained from this process. As seen in Chapter 4, the algorithm is evaluated using various performance metrics and it is important to note that the average testing error calculated is slightly higher than the proposed $P_{FA} = 5\%$ in [9].

The method of non-parametric density estimation can also be extended to other areas of health care such as early detection autism, MRI imaging, early detection of cancer, study of cardio-respiratory diseases in adults etc. In addition to advances in nonparametric modeling, Bayesian nonparametric modeling has found many applications in the field [27, 28].

The results obtained can also be used to study the effects of early detection of bradycardia in preterm infants. Outside of the health care industry - multiple object tracking using Bayesian modeling [27, 29], non-parametric bayesian estimation of noise density for non-linear filters [30], mining of spatio-temporal behaviour on social media using non-parametric modeling [31]; are just a few applications where the discussed method can be extended. In short, the method presented can be used for any applications that employ non-parametric density estimation and the calculation of a Type I error.

REFERENCES

- [1] A. H. Gee, R. Barbieri, D. Paydarfar, and P. Indic, “Predicting bradycardia in preterm infants using point process analysis of heart rate,” *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 9, pp. 2300–2308, 2017.
- [2] A. B. Tsybakov, *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- [3] A. Thommandram, J. M. Eklund, C. McGregor, J. E. Pugh, and A. G. James, “A rule-based temporal analysis method for online health analytics and its application for real-time detection of neonatal spells,” in *2014 IEEE International Congress on Big Data*. IEEE, 2014, pp. 470–477.
- [4] Ö. N. Onak, Y. S. Dogrusoz, and G. W. Weber, “Evaluation of multivariate adaptive non-parametric reduced-order model for solving the inverse electrocardiography problem: a simulation study,” *Medical & Biological Engineering & Computing*, vol. 57, no. 5, pp. 967–993, 2019.
- [5] B. Iglesias, M. J. Rodríguez, E. Aleo, E. Criado, J. Martínez-Orgado, and L. Arzuza, “3-lead electrocardiogram is more reliable than pulse oximetry to detect bradycardia during stabilisation at birth of very preterm infants,” *Archives of Disease in Childhood-Fetal and Neonatal Edition*, vol. 103, no. 3, pp. F233–F237, 2018.
- [6] J. R. Williamson, D. W. Bliss, and D. Paydarfar, “Forecasting respiratory collapse: theory and practice for averting life-threatening infant apneas,” *Respiratory physiology & neurobiology*, vol. 189, no. 2, pp. 223–231, 2013.
- [7] M. Lucchini, N. Burtchen, W. P. Fifer, and M. G. Signorini, “Multi-parametric cardiorespiratory analysis in late-preterm, early-term, and full-term infants at birth,” *Medical & biological engineering & computing*, vol. 57, no. 1, pp. 99–106, 2019.
- [8] A. H. Gee, R. Barbieri, D. Paydarfar, and P. Indic, “Predicting bradycardia in preterm infants using point process analysis of heart rate,” *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 9, pp. 2300–2308, 2016.
- [9] S. Das, B. Moraffah, A. Banerjee, S. K. S. Gupta, and A. Papandreou-Suppappola, “Bradycardia prediction in preterm infants using nonparametric kernel density estimation,” in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, 2019, pp. 1309–1313.
- [10] B. Lantz, *Machine learning with R*. Packt publishing ltd, 2013.
- [11] E. T. Jaynes, *Probability theory: the logic of science*.
- [12] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*, 2018.

- [13] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [14] M. Kuhn, K. Johnson *et al.*, *Applied predictive modeling*. Springer, vol. 26.
- [15] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [16] K. P. Murphy, *Machine learning: a probabilistic perspective*, 2012.
- [17] D. Freedman and P. Diaconis, “On the histogram as a density estimator: L 2 theory,” *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 57, no. 4, pp. 453–476, 1981.
- [18] B. W. Silverman, *Density estimation for statistics and data analysis*. CRC press, 1986, vol. 26.
- [19] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*. CRC press, 2013.
- [20] C. D. Manning, H. Schütze, and P. Raghavan, *Introduction to information retrieval*. Cambridge university press, 2008.
- [21] A. Celisse, “Optimal cross-validation in density estimation with the l^2 -loss,” *The Annals of Statistics*, vol. 42, pp. 1879–1910, 2014.
- [22] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff *et al.*, “PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals,” *Circulation*, vol. 101, pp. e215–e220, 2000. [Online]. Available: <http://physionet.mit.edu/physiobank/database/picsdb>
- [23] J. Pan and W. J. Tompkins, “A real-time QRS detection algorithm,” *IEEE Transactions on Biomedical Engineering*, vol. 32, pp. 230–236, 1985.
- [24] J. Lei, J. Robins, and L. Wasserman, “Distribution-free prediction sets,” *Journal of the American Statistical Association*, vol. 108, pp. 278–287, 2013.
- [25] G. Shafer and V. Vovk, “A tutorial on conformal prediction,” *Journal of Machine Learning Research*, vol. 9, pp. 371–421, 2008.
- [26] M. Sinjini, “Prediction of bradycardia in preterm infants,” <https://github.com/Sinjini15/Master-sThesis2020>, 2020.
- [27] B. Moraffah and A. Papndreou-Suppopola, “Bayesian nonparametric modeling for predicting dynamic dependencies in multiple object tracking,” *arXiv preprint arXiv:2004.10798*, 2020.
- [28] J. K. Ghosh and R. Ramamoorthi, *Bayesian nonparametrics*. Springer Science & Business Media, 2003.

- [29] B. Moraffah and A. Papandreou-Suppappola, “Random infinite tree and dependent Poisson diffusion process for nonparametric Bayesian modeling in multiple object tracking,” in *International Conference on Acoustics, Speech, and Signal Processing*, 2019, pp. 5217–5221.
- [30] E. Özkan, S. Saha, F. Gustafsson, and V. Šmídl, “Non-parametric bayesian measurement noise density estimation in non-linear filtering,” in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011, pp. 5924–5927.
- [31] X. Du, Y. Pei, W. Duivesteijn, and M. Pechenizkiy, “Exceptional spatio-temporal behavior mining through bayesian non-parametric modeling,” *Data Mining and Knowledge Discovery*, vol. 34, 01 2020.
- [32] R. J. Martis, U. R. Acharya, and L. C. Min, “Ecg beat classification using pca, lda, ica and discrete wavelet transform,” *Biomedical Signal Processing and Control*, vol. 8, no. 5, pp. 437–448, 2013.
- [33] A. Khazaei and A. Ebrahimzadeh, “Classification of electrocardiogram signals with support vector machines and genetic algorithms using power spectral features,” *Biomedical Signal Processing and Control*, vol. 5, no. 4, pp. 252–263, 2010.
- [34] T. A. O’Brien, K. Kashinath, N. R. Cavanaugh, W. D. Collins, and J. P. O’Brien, “A fast and objective multidimensional kernel density estimation method: fastkde,” *Computational Statistics & Data Analysis*, vol. 101, pp. 148–160, 2016.
- [35] S. Asaeedi, F. Didehvar, and A. Mohades, “ α -concave hull, a generalization of convex hull,” *Theoretical Computer Science*, vol. 702, pp. 48–59, 2017.
- [36] J.-S. Park and S.-J. Oh, “A new concave hull algorithm and concaveness measure for n-dimensional datasets.”
- [37] P. Krauthausen and U. D. Hanebeck, “Regularized non-parametric multivariate density and conditional density estimation,” in *2010 IEEE Conference on Multisensor Fusion and Integration*, 2010, pp. 180–186.
- [38] N. Montazeri Ghahjaverestan, S. Masoudi, M. B. Shamsollahi, A. Beuchée, P. Pladys, D. Ge, and A. I. Hernández, “Coupled hidden markov model-based method for apnea bradycardia detection,” *IEEE Journal of Biomedical and Health Informatics*, vol. 20, no. 2, pp. 527–538, 2016.
- [39] J. Cruz, A. I. Hernandez, S. Wong, G. Carrault, and A. Beuchee, “Algorithm fusion for the early detection of apnea-bradycardia in preterm infants,” in *2006 Computers in Cardiology*, 2006, pp. 473–476.
- [40] F. Portet, F. Gao, J. Hunter, and S. Sripada, “Evaluation of on-line bradycardia boundary detectors from neonatal clinical data,” in *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2007, pp. 3288–3291.

- [41] D. K. Ravish, K. J. Shanthi, N. R. Shenoy, and S. Nisargh, “Heart function monitoring, prediction and prevention of heart attacks: Using artificial neural networks,” in *2014 International Conference on Contemporary Computing and Informatics (IC3I)*, 2014, pp. 1–6.
- [42] C. Choi, Y. Kim, and K. Shin, “A pd control-based qrs detection algorithm for wearable ecg applications,” in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2012, pp. 5638–5641.
- [43] E. R. Adams and A. Choi, “Using neural networks to predict cardiac arrhythmias,” in *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2012, pp. 402–407.
- [44] B. Moraffah, “Bayesian nonparametric modeling and inference for multiple object tracking,” Ph.D. dissertation, Arizona State University, 2019.
- [45] C. M. Bishop, *Pattern recognition and machine learning*. springer, 2006.
- [46] C. J. Upton, A. D. Milner, and G. M. Stokes, “Episodic bradycardia in preterm infants,” *Archives of Disease in Childhood*, vol. 67, pp. 831–834, 1992.
- [47] A. H. Gee, J. Chang, J. Ghosh, and D. Paydarfar, “Bayesian online changepoint detection of physiological transitions,” in *IEEE Engineering in Medicine and Biology Society*, 2018, pp. 45–48.
- [48] H. Truong, “Predicting adverse outcomes in preterm infants using early bedside monitor data,” 2018, thesis for Bachelor of Science in Physics, College of William and Mary, Williamsburg, VA.
- [49] A. H. Gee, R. Barbieri, D. Paydarfar, and P. Indic, “Predicting bradycardia in preterm infants using point process analysis of heart rate,” *IEEE Transactions on Biomedical Engineering*, vol. 64, pp. 2300–2308, 2017.
- [50] S. M. Mahmud, H. Wang, and Y. Kim, “Accelerated prediction of bradycardia in preterm infants using time-frequency analysis,” in *Int. Conference on Computing, Networking and Communications*, 2019, pp. 468–472.
- [51] A. Hosseini and M. Sarrafzadeh, “Unsupervised prediction of negative health events ahead of time,” *arXiv preprint arXiv:1901.11168*, 2019.
- [52] S. Blackburn, “Problems of preterm infants after discharge,” *Journal of Obstetric, Gynecologic, & Neonatal Nursing*, vol. 24, no. 1, pp. 43–49, 1995.
- [53] J. M. Perlman and J. J. Volpe, “Episodes of apnea and bradycardia in the preterm newborn: impact on cerebral circulation,” *Pediatrics*, vol. 76, no. 3, pp. 333–338, 1985.
- [54] G. Pichler, B. Urlesberger, and W. Müller, “Impact of bradycardia on cerebral oxygenation and cerebral blood volume during apnoea in preterm infants,” *Physiological measurement*, vol. 24, no. 3, p. 671, 2003.

- [55] C. F. Poets, R. S. Roberts, B. Schmidt, R. K. Whyte, E. V. Asztalos, D. Bader, A. Bairam, D. Moddemann, A. Peliowski, Y. Rabi *et al.*, “Association between intermittent hypoxemia or bradycardia and late death or disability in extremely preterm infants,” *Jama*, vol. 314, no. 6, pp. 595–603, 2015.
- [56] A. Mittal and N. Paragios, “Motion-based background subtraction using adaptive kernel density estimation,” in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, vol. 2. Ieee, 2004, pp. II–II.
- [57] B. Moraffah, “Inference for multiple object tracking: A Bayesian nonparametric approach,” *arXiv preprint arXiv:1909.06984 cs.LG*, 2019.