Spring 2023

Dynamic Discounted Satisficing Based Driver Decision Prediction in Sequential Taxi Requests

Sree Pooja Akula
Missouri University of Science and Technology

DYNAMIC DISCOUNTED SATISFICING BASED DRIVER DECISION PREDICTION

IN SEQUENTIAL TAXI REQUESTS

by

SREE POOJA AKULA

A THESIS

Presented to the Graduate Faculty of the

MISSOURI UNIVERSITY OF SCIENCE AND TECHNOLOGY

In Partial Fulfillment of the Requirements for the Degree

MASTER OF SCIENCE

in

COMPUTER SCIENCE

2023

Approved by:

Dr. Venkata Sriram Siddhardh Nadendla, Advisor
Dr. Tie Luo
Dr. Ardhendu Tripathy

ABSTRACT

Ridesharing platforms rely on connecting available taxi drivers to potential passengers to maximize their revenue. However, predicting the stopping decision made by every driver, i.e., the final task performed during a given day, is crucial to achieving this goal. Unfortunately, little research has been done on predicting drivers' stopping decisions, especially when they deviate from expected utility maximization behavior. This research proposes a ***D**ynamic **D**iscounted **S**atisficing* (DDS) heuristic to model and learn the task at which human agents will stop working for that day, assuming that the human agents are taking sequential decisions based on their preference order. We apply this approach to the problem of predicting the stopping decision of taxi drivers in a ridesharing platform. To estimate the model parameters and predict the stopping time, we propose an algorithm - Sampling Based Back Propagation Through Time(SBPTT) and evaluate it using real-time data from the Chicago taxi dataset. The proposed model consistently has better accuracy on simulated and real world data sets, when compared with discounted satisficing model.

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Dr. Venkata Sriram Siddhardh Nadendla, for his exceptional guidance, support, and encouragement throughout my research journey. Dr. Nadendla's extensive expertise, attention to detail, and dedication to excellence have been instrumental in shaping my research and personal growth. I am fortunate to have had such a mentor who believed in me and challenged me to strive for success. I would also like to extend my thanks to Dr. Tie Luo and Dr. Ardhendu Tripathy, who have served as my committee members. Their belief in my abilities and their challenging feedback have motivated me to strive for excellence and have undoubtedly improved the quality of my work. I would also like to extend my heartfelt appreciation to my family for their unconditional love and constant support. Their unwavering encouragement and belief in me have been the driving force behind my academic achievements.To my lab mates, I am incredibly grateful for their friendship, collaboration, and valuable support. Their insightful feedback, constructive criticism, and technical expertise have been integral to the success of my research project.

I would like to acknowledge the Department of Computer Science at Missouri S&T for providing me with the resources, opportunities, and education necessary to achieve my academic goals. The department has played a critical role in my academic journey, and I am grateful for the valuable experiences and knowledge gained during my time there. Finally, I would like to express my thanks to everyone who has contributed to my academic journey. Your support and encouragement, have been instrumental in my success, and I am deeply grateful for all that you have done for me.

TABLE OF CONTENTS

Page

LIST OF ILLUSTRATIONS

1. INTRODUCTION

Ridesharing platforms connect available taxi drivers to potential passengers whenever a service request is submitted to the platform. The use of ridesharing platforms has become increasingly popular over the past few years, with more and more people relying on them for transportation [1]. One of the biggest challenges faced by drivers on these platforms is the long hours required to meet their income goals [2]. As the hours of work accumulate, the level of cognitive fatigue experienced by drivers may increase, resulting in a decline in their decision-making abilities. This can affect their ability to navigate routes, make quick decisions in traffic, and assess passenger safety [3]. If the ridesharing platform can predict the stopping decision made by every driver, then it can improve driver's decision-making ability and reduce the likelihood of fatigue-related errors. Predicting drivers' final tasks can improve the performance of ridesharing platforms by assigning ride requests optimally to maximize the efficiency and average revenue of the platform.

The stopping decision of a driver may depend on task demands and cognitive fatigue levels of drivers and various external factors such as such as traffic conditions, passenger behavior and the length of their workday. However, To the best of our knowledge, there is little work on the analysis of driver's behavior (e.g. stopping time) on a ridesharing platform, especially when they exhibit behavioral deviations from expected utility maximization (EUM) behavior [4]. While traditional satisficing models [5, 6, 7] have been shown to model human decision-making under conditions of cognitive fatigue, they may not fully capture the dynamic nature of decision-making in the context of ridesharing platforms. Devaguptapu et al., in [8] predicts the stopping time of an agent using Discounted Satisficing heuristic where the agent's threshold discounts over time. However, the model was found to be inaccurate when dealing with human agents with greater discontent levels. For example, in a ridesharing platform, the threshold of a driver discounts over time and also changes across days.

To address these limitations, there is a need for a more dynamic approach to satisficing or discounted satisficing that can adapt to the changing demands of the ridesharing context. By accounting for the dynamic nature of decision-making, these models could help drivers make more informed decisions about when to take breaks, how much effort to expend on driving, and when to stop working altogether. In addition, dynamic discounted satisficing models could help ridesharing companies allocate resources more effectively. For example, they could optimize work schedules to reduce the likelihood of driver burnout and increase overall productivity. Overall, the development of dynamic discounted satisficing models has the potential to improve the safety and well-being of ridesharing drivers, as well as increase organizational productivity and improve outcomes. Therefore, this thesis focuses on the prediction of drivers' stopping decision to help improve the performance of ridesharing platforms.

The main contribution of this thesis is threefold. Firstly, it proposes a novel approach to model the sequential decisions made by drivers using *dynamic discounted satisficing* (DDS). By capturing the dynamics of the satisficing threshold, the DDS heuristic provides a more accurate representation of the decision-making process, which can lead to more effective way to predict a driver's stopping task. Secondly, the thesis develops a learning algorithm for stochastic neural networks to estimate the model parameters of dynamic discounted satisficing. This algorithm addresses the challenges involved in estimating model parameters to capture randomness in the decision-making process. Finally, the thesis validates the proposed model on the real-world dataset (Chicago taxi dataset)from a ride-sharing platform [9]. This validation demonstrates the practical applications of the proposed model and its effectiveness in predicting the stopping time of drivers.

2. LITERATURE REVIEW

Ridesharing applications [10], have revolutionized the transportation industry by providing a convenient and cost-effective alternative to traditional taxis. One of the key features of these applications is sequential taxi requests, which allow the driver to complete multiple rides every day. There are two types of ridesharing platforms depending on whether the driver has the autonomy to pick the request. For example, in traditional taxi services, drivers are expected to serve passengers who are assigned to them by the platform without any choice. On the other hand, modern ridesharing platforms such as Uber and Lyft match drivers with passengers while preserving the decision autonomy at both types of agents. By allowing drivers to pick up passengers according to requests, these platforms maximize the efficiency of each trip and reduce idle time between rides. In either setting [11, 12], a taxi driver services a sequence of ride requests until they decide to stop working for that day, as in the case of [13, 14, 15, 16]. The use of ridesharing platforms has become increasingly popular over the past few years, with more and more people relying on them for transportation [1]. This led to long hours for drivers on these platforms [2]. As the hours of work accumulate, the level of cognitive fatigue experienced by drivers may increase, resulting in a decline in their decision-making abilities. This can affect their ability to navigate routes, make quick decisions in traffic, and assess passenger safety [3]. If the ridesharing platform can predict the stopping decision made by every driver, then it can improve driver's decision-making ability and reduce the likelihood of fatigue-related errors.

Practical ridesharing platforms offer diverse recommendations to both passengers as well as drivers. For example, passengers are provided with wait-location recommendation to reduce the trip-cost [17] and/or plan ride suggestions [18] (e.g. compact car vs. luxury car vs. large SUV) to improve their experience. On the other hand, drivers are provided with ride choices, along with incentives if/when passengers' future activity is predicted in a location that has very few drivers. The success of such recommendations relies heavily on the accuracy of *network state information* (NSI) [19, 20] available at the platform. For example, Tal Altshuler et.al., in [21] predicted spatio-temporal utilization of ridesharing services from passenger activity models extracted from NSI. [22] extracted NSI from GPS data and identifed the passengers' demand hot area and proposed a taxi station optimization model by analyzing the time series distribution dynamic characteristics of passengers' temporal variation in certain land use types and taxi driver's searching behavior in connection with different activity spaces for different lengths of observation period. Charles C Macadam in [23] emphasizes the importance of including human characteristics in models of driver control behavior to accurately predict the performance of the driver-vehicle system. He identified physical limitations and unique attributes of human drivers and presented driver models commonly used for prediction. However, this paper does not address driver/passenger activity or the impact of driver's physical/cognitive state on system efficiency.

Ride-sharing platforms involve complex decision-making processes for drivers, who need to evaluate limited information about passengers and their destinations in a short amount of time. These decisions are subject to physical and cognitive limitations, such as time constraints and the ability to process and remember information. Simon [24] introduced the concept of bounded rationality to explain how humans deviate from the ideal of economic rationality, such as Expected Utility Maximization (EUM) [25], when making decisions under such limitations. One manifestation of bounded rationality is *Satisficing*. In the context of multi-armed bandit problems, where an agent must choose between multiple options with uncertain rewards over time, satisficing has been studied extensively. [26] showed that the satisficing model based on mean or instantaneous reward is equivalent to the exploration/exploitation trade-off problem in multi-armed bandit problems, and presented bounds on the agent's performance.

However, one assumption made by [27] is that the threshold for acceptability remains constant throughout the decision process. In reality, the threshold may change over time due to various factors, such as the driver's mood, fatigue, or experience. Understanding how the threshold changes and affects the decision-making process could help improve the performance and satisfaction of drivers and passengers on ride-sharing platforms. To address this issue, [8] proposed a *Discounted Satisficing* heuristic, where the agent's threshold discounts over time. This means that the decision-maker becomes less satisfied with a given level of utility as time goes on. The decision-maker still does not aim for the optimal outcome, but instead aims to achieve a level of utility that is at least as good as the discounted threshold. However, such a model is restrictive since people typically exhibit different discounting factors over days, making it infeasible to learn the model parameters on a daily basis in practice.

3. DYNAMIC DISCOUNTED SATISFICING

Dynamic Discounted Satisficing (DDS) is a novel decision heuristic which models dynamic thresholds in satisficing behavior when agents deviate from optimal decisions due to cognitive and/or informational limitations. Specifically, DDS captures the cognitive atrophy dynamics within the agent over the course of a day due to changing fatigue levels across days. Traditional models such as satisficing and discounted satisficing assume that an agent always begins with the same threshold and/or discounting factor across days. However, this is not observed in practice because people start their working day with a different target and discounting factor each day. For example, a driver may start the day with a fresh mind today, but can start working tomorrow with either higher fatigue levels due to lack of sleep, or a higher target when he/she realizes additional domestic expenses in the near future.

Consider a ridesharing platform where a driver is presented with an indefinite sequence of ride requests, until he/she decides to stop working for the day. In the d^{th} day, assume that the driver serves T_d rides, for which he/she obtains a utility sequence $\{u_{d,1}, \cdots, u_{d,T_d}\}$ for tasks $k = 1, \cdots, T_d$ respectively. In other words, the driver's accumulated utility on day d after completing t tasks can be computed as

$$U_{d,t} = \sum_{k=1}^{t} u_{d,k}. \tag{3.1}$$

Henceforth, for simplicity, we ignore the subscript T_d in U_{d,T_d} and denote the total accumulated utility as U_d. In practice, the driver typically exhibits two types of dynamics within their decision behavior:

- the constant deterioration of threshold (as modeled in Eqn. (3.3)) within a given day due to increasing weariness over time, and

- the evolution of the initial target (as modeled in Eqn. (3.4)) and fatigue rate (as in Eqn. (3.5)) across days.

Let $\mathbb{P}_{\mathcal{S}}(x)$ denote a projection operator that projects the input argument x onto the set $\mathcal{S}$, i.e.,

$$\mathbb{P}_{\mathcal{S}}(x) = \begin{cases} x_L, & \text{if } x \leq x_L \triangleq \inf \mathcal{S}, \\ x, & \text{if } x \in \mathcal{S}, \\ x_U, & \text{if } x \geq x_L \triangleq \sup \mathcal{S}, \end{cases} \tag{3.2}$$

In other words, the projection operator returns the closest value to x that is also in $\mathcal{S}$, or the boundary value of $\mathcal{S}$ if x is outside of the range defined by $\mathcal{S}$. Then, DDS can be formally defined in the following manner:

Definition 1. *A driver exhibits <u>dynamic discounted satisficing</u> heuristic, if there exists four real numbers $a_1, a_2, b_1, b_2 \in \mathbb{R}$, one positive real number $\lambda \in \mathbb{R}_+$, one bounded real number $\beta \in (0, 1]$, and two arrays of random numbers $\epsilon_d \sim \mathcal{N}(0, 1)$ and $\eta_d \sim \mathcal{N}(0, 1)$ for $d \in \mathbb{N}$, such that his/her final ride count t^* is given by*

$$t^* = \mathit{minimize} \quad \left\{ t \in \mathcal{T}_d \mid U_{d,t} = \sum_{k=1}^{t} u_{d,k} \geq \beta_d^{t-1} \cdot \lambda_d \right\} \tag{3.3}$$

where the dynamics of initial target λ_d and the discounting factor β_d are given by

$$\lambda_d = \mathbb{P}_{[0,\infty]} \Big(a_1 \lambda_{d-1} + a_2 \cdot U_{d-1} + \epsilon_d \Big) \tag{3.4}$$

and

$$\beta_d = \mathbb{P}_{[0,1]} \Big(b_1 \beta_{d-1} + b_2 \cdot e^{-T_{d-1}} + \eta_d \Big) \tag{3.5}$$

respectively.

Illustrative Example: In order to understand the differences between satisficing, discounted satisficing, and dynamic discounted satisficing heuristics, we consider an illustrative example of a ride-sharing platform. Specifically, we examine the case of a driver who has the same set of taxi rides across three days, where the utility vector follows an exponential distribution.

$$utility\ vector = [6.04, 4.34, 2.27, 0.41, 9.33, 12.96, 1.27, 4.62, 4.04, 17.41].$$

The exponential distribution is a suitable choice for modeling the driver's utilities because it is a continuous probability distribution that is commonly used to model the events in a Poisson process. In the context of our example, the exponential distribution captures the randomness and unpredictability of the driver's utility from each ride.

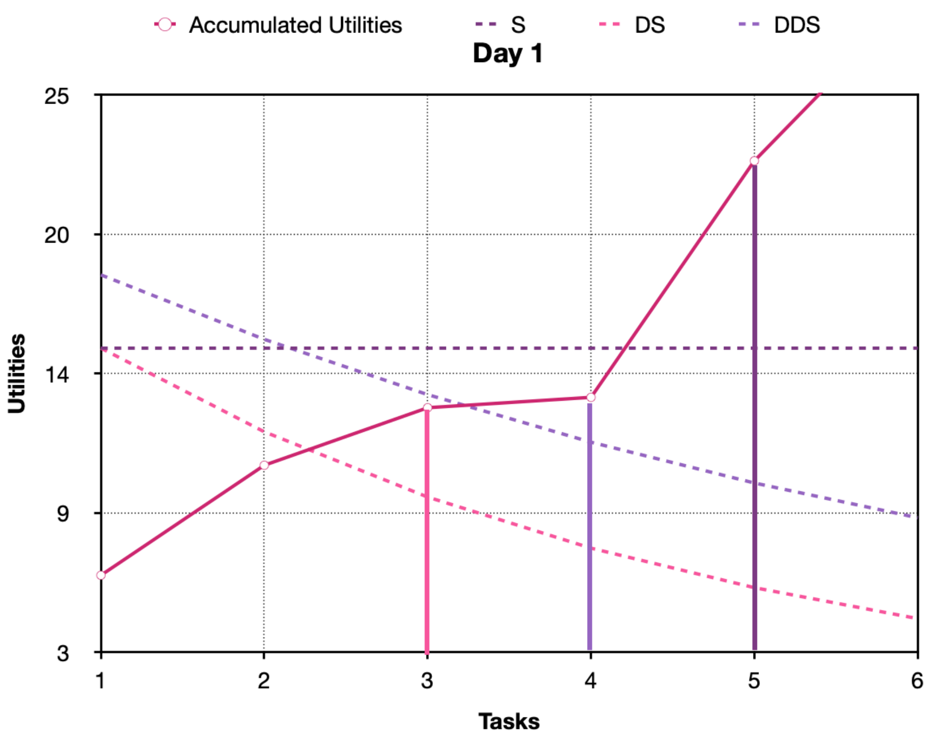

Figure 3.1. Illustrative Example: Day 1

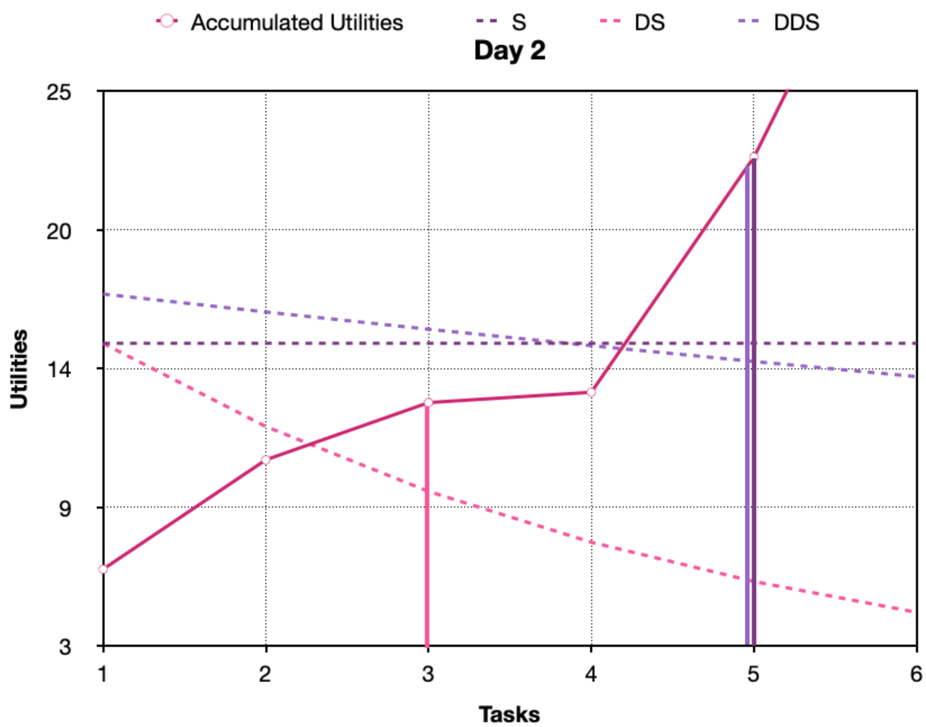

Figure 3.2. Illustrative Example: Day 2

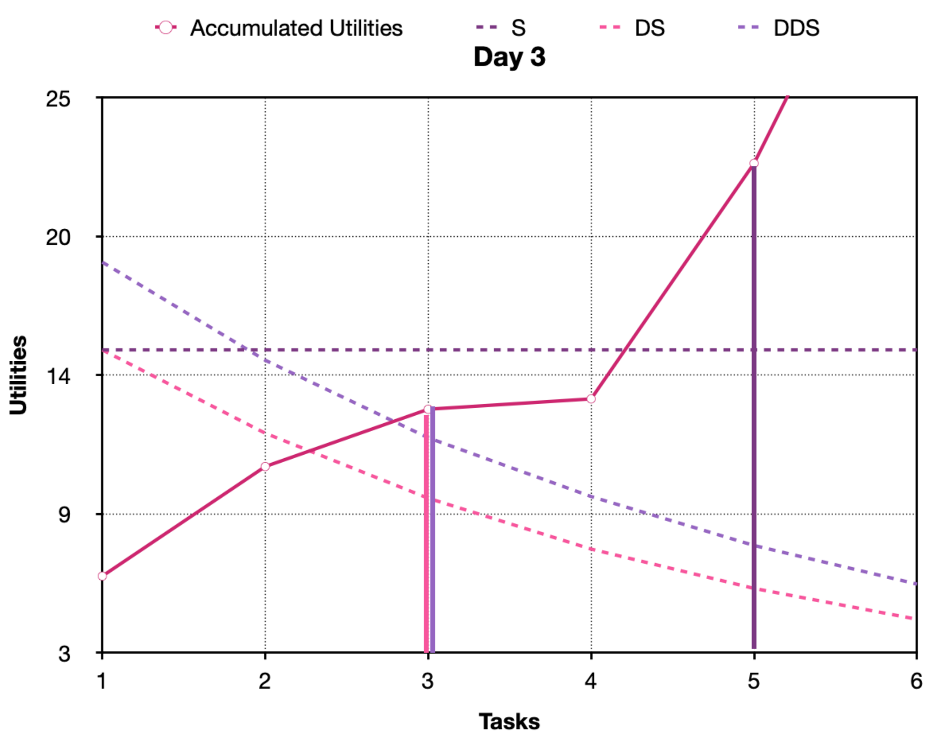

Figure 3.3. Illustrative Example: Day 3

To develop the various satisficing models, we initialize their parameters as follows: $\lambda_0 = 15$ and $\beta_0 = 0.78$, along with $a_1 = 0.68$, $a_2 = 0.37$, $b_1 = 1.2$, and $b_2 = 0.23$. In the satisficing (S) model as shown in Figure 3.1, the threshold λ_s is fixed at 15 for all days, yielding a constant stopping task of $t_s = 5$ for all days. In the discounted satisficing (DS) model as shown in Figure 3.2, both the threshold λ_{ds} and discounting factor β_{ds} are held constant at 15 and 0.78, respectively, resulting in a consistent stopping task of $t_{ds} = 3$ for all days.

In the dynamic discounted satisficing (DDS) model as shown in Figure 3.3, the threshold and discounting factor vary across days, as determined by the functions presented in Eqn. 3.4 and Eqn. 3.5. On day 1, the threshold and discounting factor for DDS are $\lambda_{dds1} = 17.89$ and $\beta_{dds1} = 0.8579$, respectively, resulting in a stopping task of $t_{dds1} = 4$. On day 2, the threshold and discounting factor for DDS are $\lambda_{dds2} = 16.957$ and $\beta_{dds2} = 0.958$, respectively, resulting in a stopping task of $t_{dds2} = 5$. On day 3, the threshold and discounting factor for DDS are $\lambda_{dds3} = 18.475$ and $\beta_{dds2} = 0.798$, respectively, resulting in a stopping task of $t_{dds3} = 3$.

Notably, the satisficing and discounted satisficing models yield the same stopping time every day respectively, when the utility vector remains constant. However in the case of dynamic discounted satisficing the model produces different stopping times for each day due to the changing threshold and discounting factor.The dynamic nature of the DDS model is particularly relevant for understanding the decision-making behavior of taxi drivers, who often face a variety of factors that can affect their stopping behavior on any given day. For example, traffic conditions, weather, and personal factors such as fatigue or mood can influence how long a driver is willing to continue working.

As a result, even if a taxi driver performs similar tasks each day, their stopping task may vary depending on various external factors. For instance, let us assume that the taxi driver has an upcoming rent payment that they needs to be made. In this situation, the driver may be more motivated to work for longer hours to earn as much money as possible to cover their expenses. This increased motivation might cause the DDS model to adjust the threshold value upwards, reflecting the driver's higher need for income. Additionally, the driver may be more willing to accept rides that are further away or require more effort, as these can result in higher payouts that would bring them closer to their financial goal. In this case, the higher threshold value would likely result in a higher stopping task Figure 3.2 for the driver, as they become more willing to continue working for longer periods of time to achieve their financial target. On the other hand, let us assume that the taxi driver did not enough sleep the previous night due to personal reasons. As a result, they begin their workday feeling fatigued and less motivated than usual. In this situation, the DDS model might adjust the discounting factor downwards to reflect the driver's immediate need for rest and recuperation. This would result in the driver placing less weight on future rewards and prioritizing their current well-being over earning more money. The lower discounting factor would likely result in a lower stopping task Figure 3.3 for the driver. In all of these scenarios, the DDS model adjusts the threshold and discounting factor based on external factors that affect the driver's motivation and perceived utility. By accounting for these factors, the DDS model provides a more nuanced and accurate representation of how satisficing behavior plays out in practice.

4. MODELING DYNAMIC DISCOUNTED SATISFICING USING STOCHASTIC NEURAL NETWORKS

4.1. MODEL ARCHITECTURE

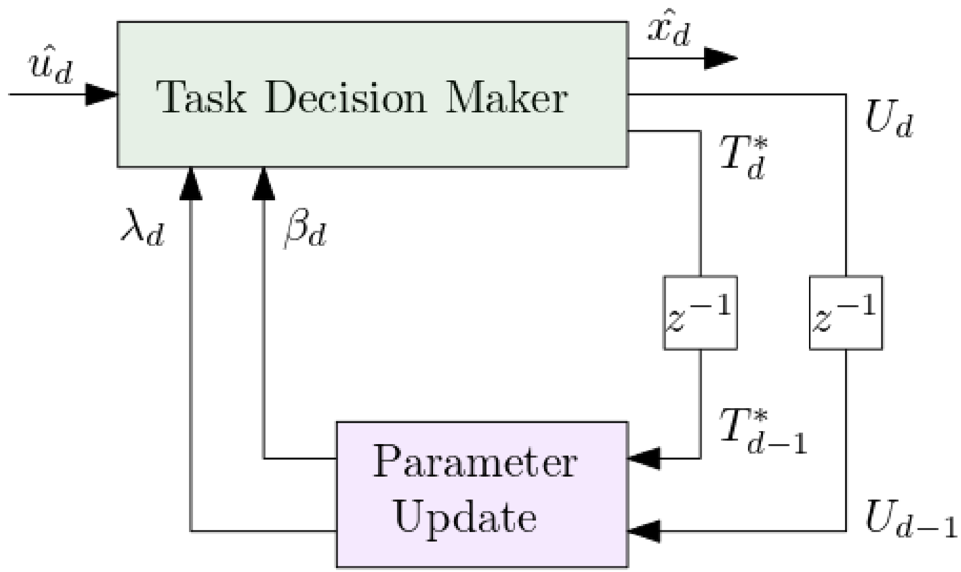

Figure 4.1. Model Architecture

In this thesis, we model Dynamic Discounted Satisficing using the combination of classical statistical modeling techniques and data-driven systems as discussed in [28]. Classical statistical modeling techniques provide a solid foundation for decision-making, while data-driven systems capture the dynamic and complex nature of real-world decision-making, resulting in a more accurate and effective model.

The Model Architecture in Figure 4.1 represents dynamic discounted satisficing heuristic as the sequential decision-making strategy employed by a driver. This architecture incorporates an adaptive decision-making approach through the *Task Decision Maker* and *Parameter Update Network*.

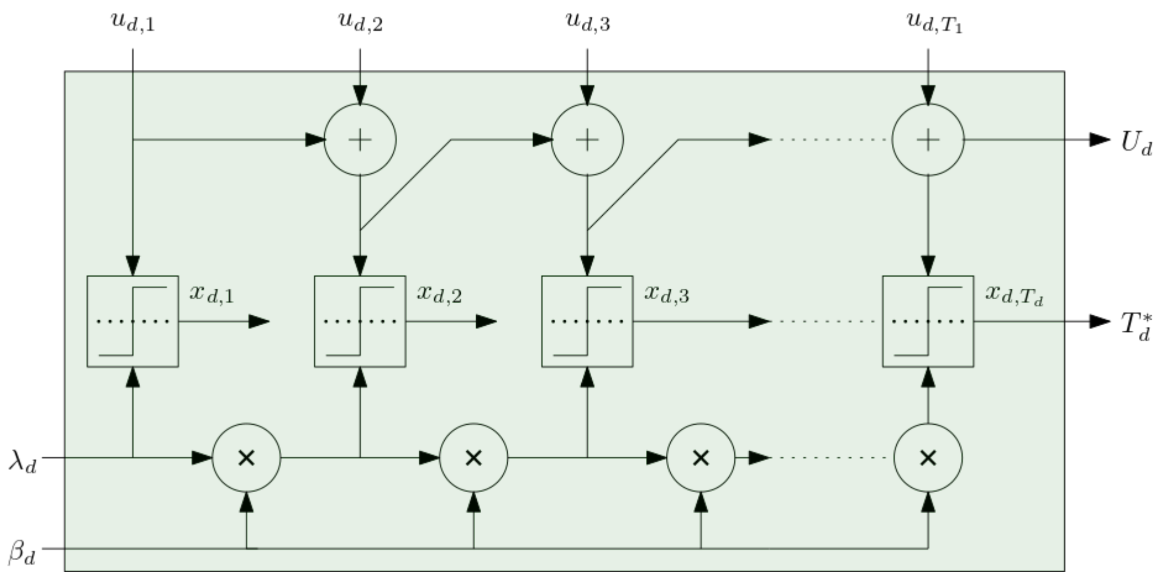

Figure 4.2. Task Decision Maker - Neural Network Architecture

Let $x_{d,t}$ denote the probability of driver choosing to accept the t^{th} ride request on day d. Note that, according to the DDS heuristic stated in Definition 1, the driver continues to accept the ride requests as long as the difference between the discounted threshold $\beta_d^{t-1} \cdot \lambda_d$ and the accumulated utility $U_{d,t}$ up to task t on day d is non-negative.

The decision-making process for the Task Decision Maker (TDM) is shown in Figure 4.2. The TDM considers the driver's utility vector and the dynamically updated parameters λ_d and β_d to make decisions on each task. The stopping task decision is computed by considering the accumulated utility as shown in Equation (3.3). Let $x_{d,t}$ denote the probability of the driver choosing to accept the t^{th} ride request on day d. According to the DDS heuristic stated in Definition 1, the driver continues to accept ride requests as long as the difference between the discounted threshold $\beta_d^{t-1} \cdot \lambda_d$ and the accumulated utility $U_{d,t}$ up to task t on day d is non-negative.

The Parameter Update Network updates λ_d and β_d as shown in Equation (3.5) and (3.4). The initial threshold for a driver on day d, denoted by λ_d, is a positive real number, which depends on the previous day λ_{d-1}, the total accumulated utility of the driver up to the previous day U_{d-1}, and ϵ_d with a standard normal distribution $\epsilon_d \sim \mathcal{N}(0, 1)$, as shown

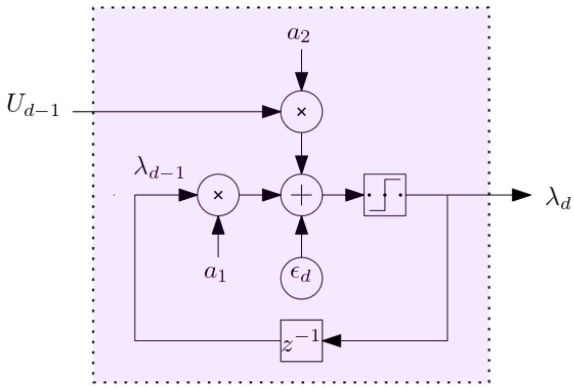

(a) Modeling Dynamics in Satisficing Threshold

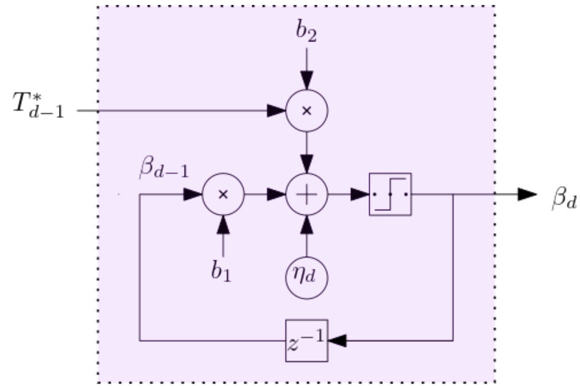

(b) Modeling Dynamics in Discounting Factor

Figure 4.3. Neural Network Models to Characterize Dynamics in DDS Parameters

in Figure 4.3a. The neural network model used is a noisy perceptron, which introduces randomness into the model. The dynamics of the parameter update network are updated , where a_1 and a_2 are the model parameters. The model uses the rectified linear unit (ReLU) activation function, which maps input values to the set [0, ∞). The discounting behavior of the driver over time is determined by the discounting factor β_d, where $\beta_d \in (0, 1]$. The value of β_d on day d, as shown in Figure 4.3b, depends on the previous day's discounting factor, β_{d-1}, the stopping task of the driver on the previous day, T_{d-1}, and η_d with a standard normal distribution, $\eta_d \sim \mathcal{N}(0, 1)$. The neural network model used is a noisy perceptron, where η_d introduces randomness into the model. The dynamics of the parameter update network are updated , where b_1 and b_2 are the model parameters. The model uses the sigmoid activation function, which maps input values to the set [0,1).

In this stochastic neural networks, η_d and ϵ_d represent the randomness in λ_d and β_d iterations. This randomness can be interpreted as a way of capturing the inherent variability in human decision-making, as human behavior is often influenced by various random factors such as emotions and other environmental factors. By incorporating this randomness into the parameter update network, the model can adapt to the dynamic nature of human decision-making and capture the variability in decision-making behavior over time.

4.2. PROPOSED TRAINING ALGORITHM AND PERFORMANCE METRICS

In order to effectively train our neural networks on sequential data, we propose a novel approach called Sampling-based Backpropagation Through Time (SBPTT). This method is inspired by the conventional Backpropagation Through Time (BPTT) algorithm [29].

Algorithm 1 Sampling based Back Propagation Through Time(SBPTT)

```
Inputs:
    Û_{D×T}, x*_{D×T}, ε ~ N(0,1), η ~ N(0,1)
Initialize:
    Model parameters w = [a_1, a_2, b_1, b_2]
    Learning rate α = 0.01
    Initial gradient ∂L/∂w = 0
for d = D to 1 do
    for r = 1 to R do
        Perform forward propagation to compute predicted output: x̂_{d,r} = f(Û_{d×T}, x*_{D×T}, ε_r, η_r)
        Compute loss between predicted and target output: L_{d,r} = L(x̂_{d,r}, x*_{D×T})
        Compute the gradients of the loss with respect to the network parameters: ∂L_{d,r}/∂w
    end for
    Compute the mean gradients of the loss across R: ∂L/∂w = (1/R) Σ_{r=1}^{R} ∂L_{d,r}/∂w
    Update the network parameters using gradient descent: w ← w − α ∂L/∂w
end for
```

The SBPTT algorithm (Algorithm 1) works by inputting a sequence $\hat{\boldsymbol{U}}_{D\times T}$ into the network one time step at a time. The network then generates a prediction $\hat{x}_{d,r}$ for each time step. In Algorithm 1, we perform forward propagation to compute the predicted output $\hat{x}_{d,r}$ for each random sample r by passing the input sequence $\hat{\boldsymbol{U}}_{d\times T}$ through the network, incorporating random variables ϵ_r and η_r sampled from a normal distribution. Subsequently, we calculate the loss $L_{d,r}$ between the predicted output $\hat{x}_{d,r}$ and the true output $x^*_{D\times T}$. The gradients of the loss with respect to the network parameters are then computed for each random sample, and the average gradients $\frac{\partial L}{\partial w}$ are obtained by averaging the gradients over all the random samples. Finally, we update the model parameters w using gradient descent with a learning rate α, and repeat this process for all time steps d

until the model converges to a satisfactory level of accuracy on the training data. Notably, by incorporating randomness through ϵ_r and η_r, we account for the inherent variability and randomness in human behavior, potentially enhancing the overall performance and generalization capabilities of the model.

5. VALIDATION ON SIMULATION DATA

To validate the performance of the model, we generated a simulated data for a driver for 50 epochs over 500 days with λ_d, β_d. Task-utilities $\hat{\boldsymbol{U}}_{D\times T}$ are randomly generated from an exponential distribution with $scale = 10$ and $y^*_{D\times T}$ is the expected binary output, where the value is set to 1 if the driver performs a task and the value is set to 0 otherwise.

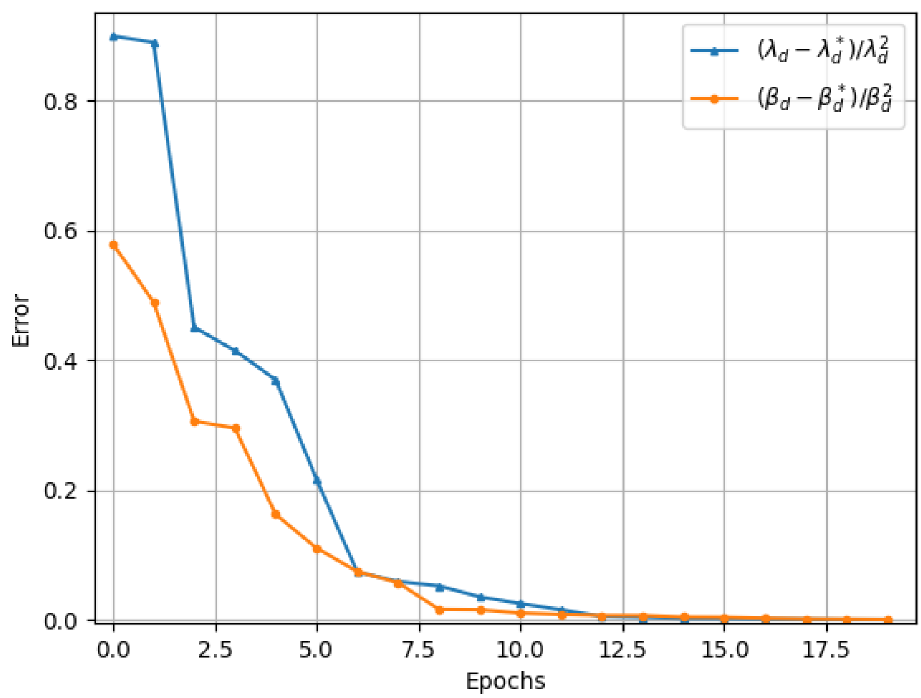

Figure 5.1. Error performance in estimating λ and β across epochs when R = 1.

Lambda and Beta errors graph as shown in Figure 5.1, Figure 5.2, Figure 5.3 illustrates that the estimation error of λ converges to zero consistently as β converges. This can be attributed to the fact that the dynamic thresholds of the driver with higher values of β generally deteriorate at a much slower rate, thereby revealing about the model parameters in their choices.

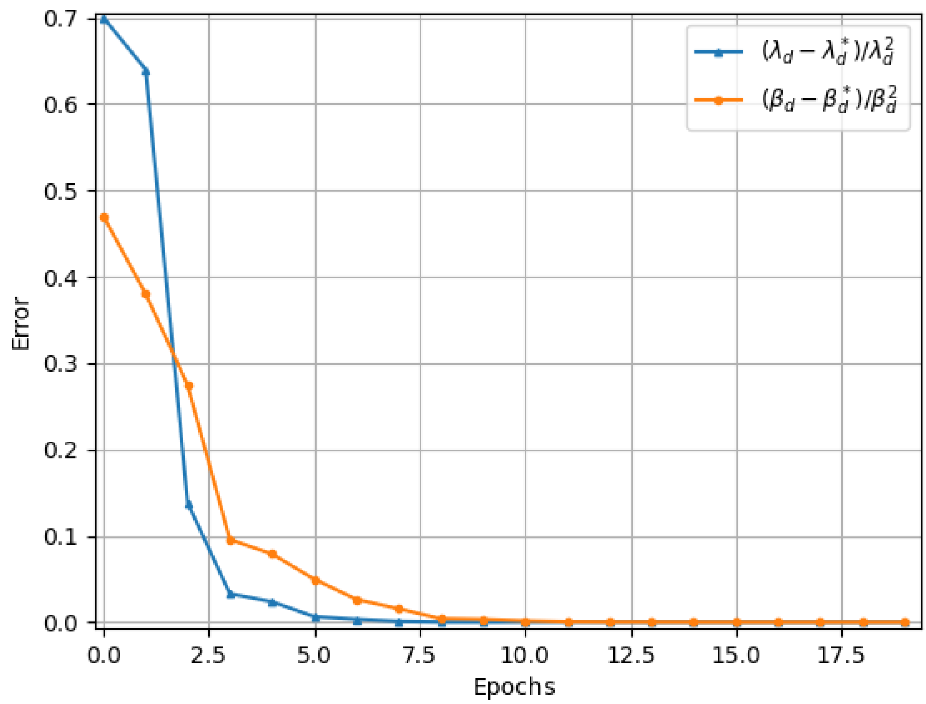

Figure 5.2. Error performance in estimating λ and β across epochs when R = 8.

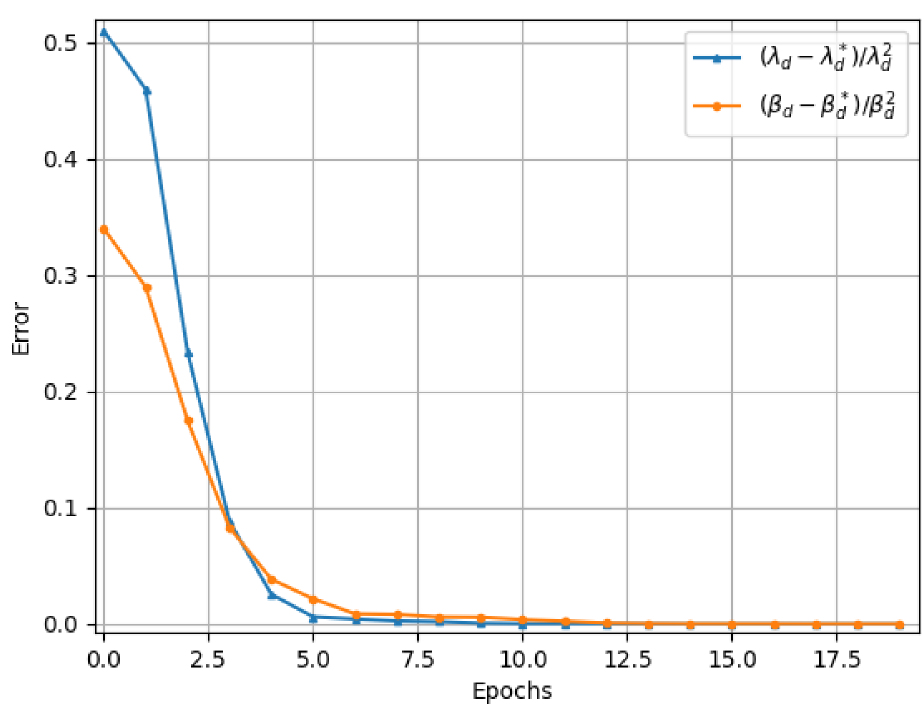

Figure 5.3. Error performance in estimating λ and β across epochs when R = 32.

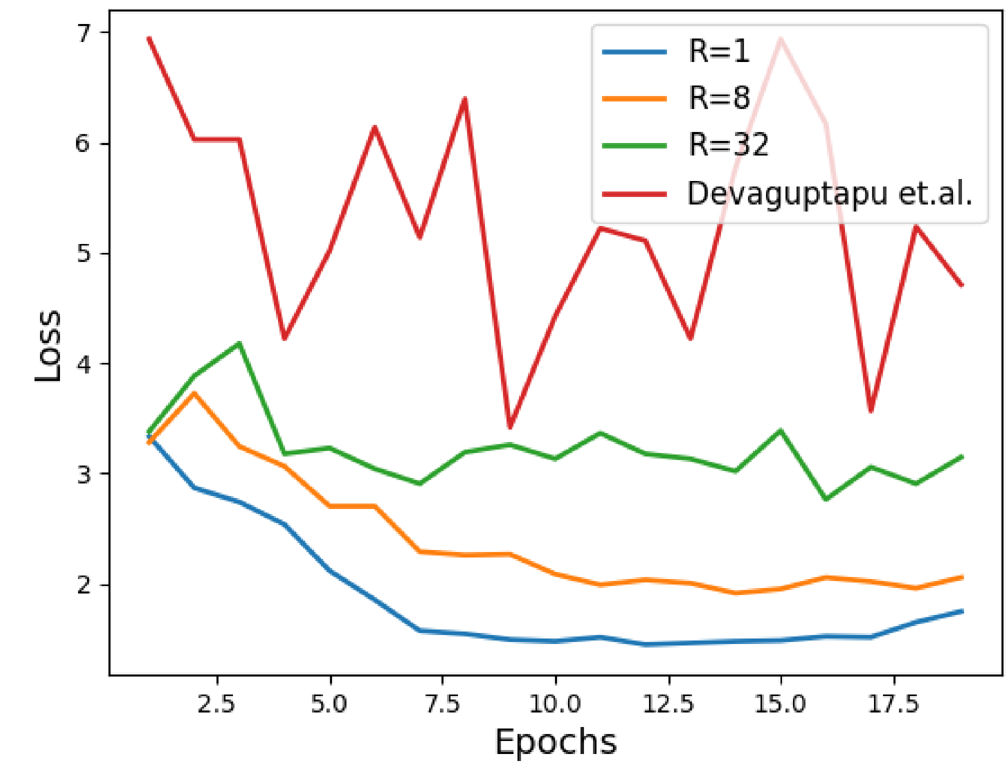

(a) Training Loss across epochs.

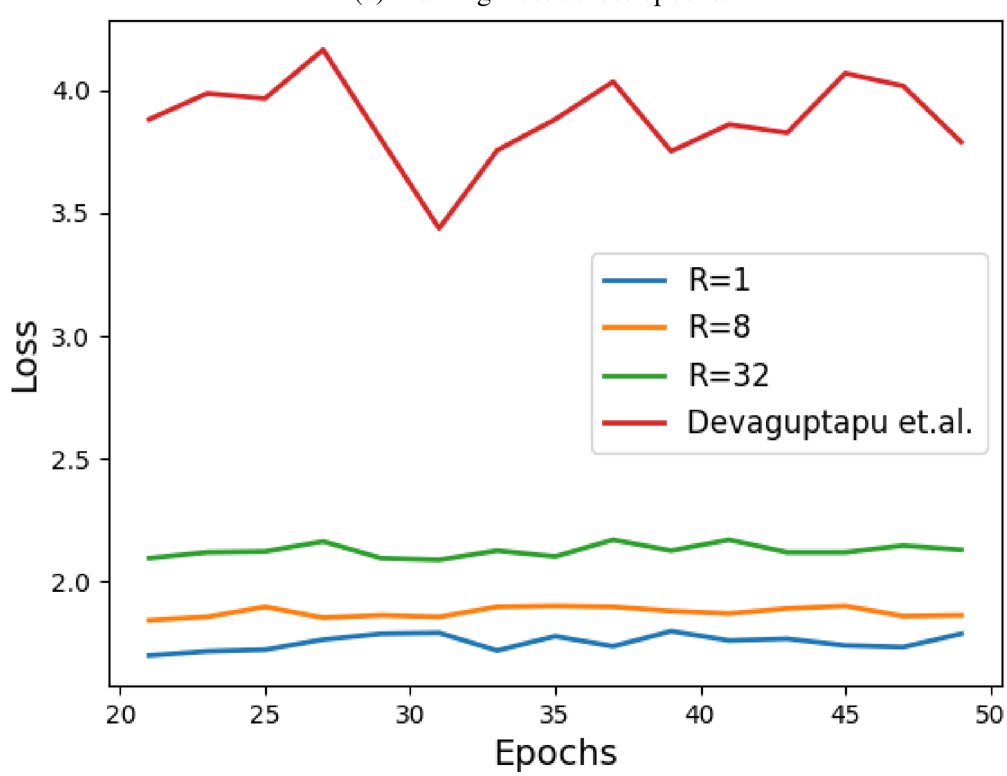

(b) Testing Loss across epochs.

Figure 5.4. Loss across Epochs on Simulation data.

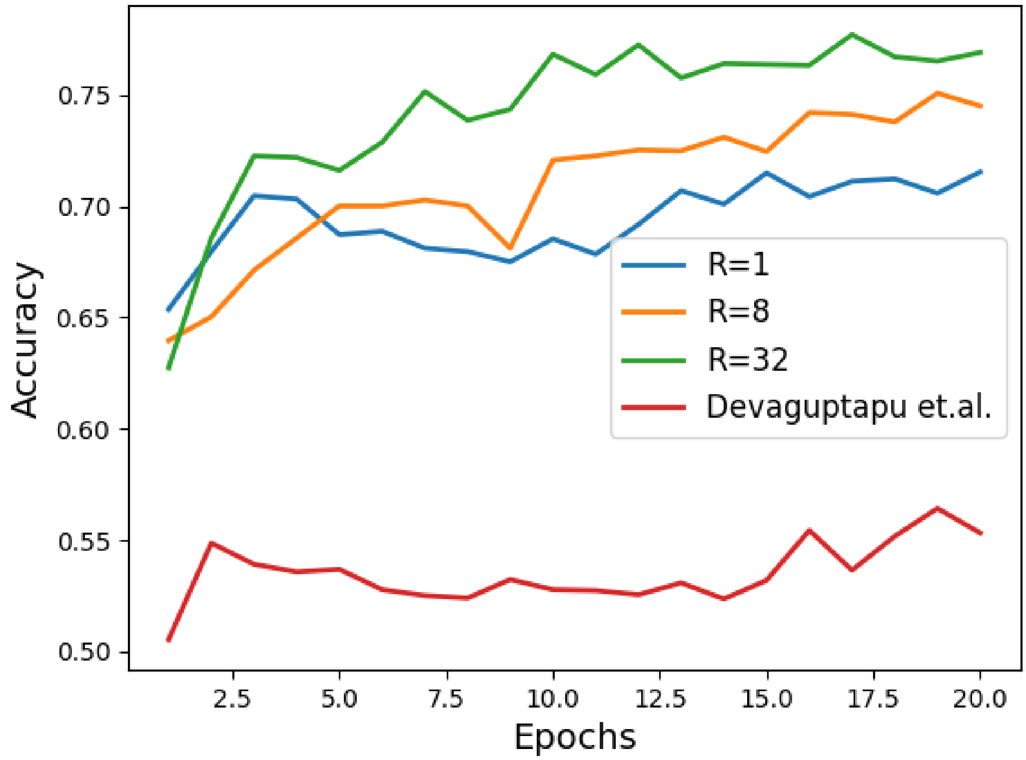

(a) Training Accuracy across epochs.

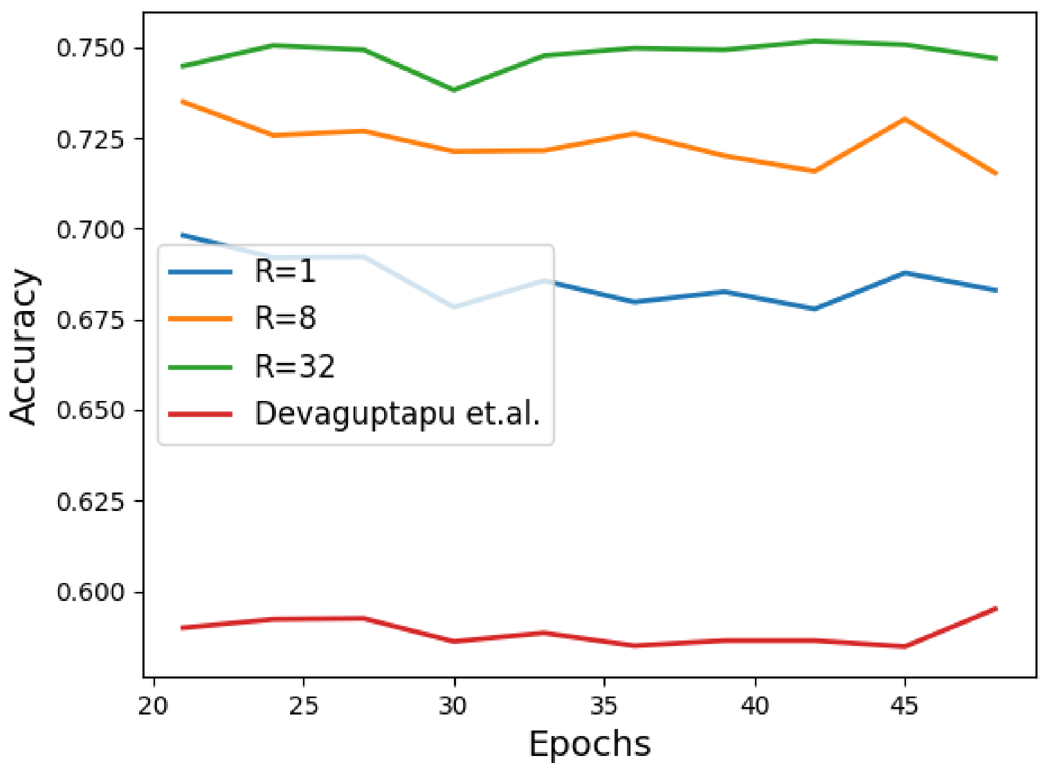

(b) Testing Accuracy across epochs.

Figure 5.5. Accuracy across Epochs on Simulation data.

The training and testing loss as shown in Figure 5.4 illustrates that the model's performance improves as R value increases. Training and Testing loss decreased by almost 70% when compared between $R = 32$ and Discounted Satisficing model. Similarly, training and testing accuracy as shown in Figure 5.5 also illustrates that the model's performance improves as R value increases.

6. VALIDATION ON REAL DATA

6.1. DATASETS AND PREPROCESSING

The dataset used in this study is obtained from the City of Chicago's Open Data Portal [9], and consists of taxi trips taken in the city during the year 2022. It contains detailed information about each taxi trip, including attributes such as trip start and end timestamps, trip durations, trip distances, fare amounts, payment types, and more. The dataset covers a time period of 2013 to the present day, and is updated monthly, providing researchers with a comprehensive and up-to-date resource for studying taxi usage in Chicago. To prepare the data for training our stochastic neural network model, we split the dataset into a 40% training set and a 60% test set for model evaluation.

The preprocessing steps involve converting the attribute 'Trip Start Timestamp' column to a datetime object and creating a new column with only the date information. The data is then grouped by date, and the values from the attribute 'Trip Total' column, which represents the total amount paid for each trip, are extracted and stored as a list. A new dataframe is constructed with the 'Trip Total' column and the grouped values, which is transposed to obtain the 'Trip Total' values as columns. To further prepare the data for model training, we create two files - one with the grouped 'Trip Total' values as input $\hat{\boldsymbol{U}}_{D\times T}$ and another with the same values as the expected output. The input file is padded with the average of all 'Trip Total' values to fill rest, and the expected output file is transformed into a binary format where 'Trip Total' values are set to 1's and rest are set to 0's , which represents $y^*_{D\times T}$.

For training model, we utilize forward propagation to compute the predicted output for each random sample by passing the input sequence through the network, incorporating random variables sampled from a normal distribution. Subsequently, we calculate the binary cross-entropy loss between the predicted output and the true output. The gradients of the loss with respect to the network parameters are then computed for each random sample, and the average gradients are obtained by averaging the gradients over all the random samples. Finally, we update the model parameters using the gradient descent algorithm with a learning rate of $\alpha = 0.01$, and repeat this process for all time steps until the model converges to a satisfactory level of accuracy on the training data. Notably, a significant aspect of our approach is the incorporation of randomness through the use of random variables in the model. This allows us to effectively account for the inherent variability and randomness in human behavior, potentially enhancing the overall performance and generalization capabilities of our model. The stochastic nature of the model enables it to effectively capture uncertainties and variations in the data, making it well-suited for predicting the stopping time of taxi drivers in the City of Chicago based on the total amount paid (utility) for each trip.

6.2. RESULTS AND DISCUSSION

This section presents the results obtained from applying Algorithm 1 to the Chicago Taxi dataset [9]. The model was trained on data from 10 different drivers, and the average loss and accuracy of the model were computed for both the training and testing datasets. To evaluate the performance of the proposed model across different R values, we conducted a comparison analysis.

In addition, the Discounted Satisficing model proposed by Devaguptapu et al. in [8] was also trained on the same datasets of the 10 drivers. We compared the loss and accuracy of this model with those of SBPTT to assess their relative performance.

Figure 6.1a shows the average training loss graph, which illustrates that the loss consistently decreases as R increases, improving the model's performance. This trend aligns with the idea that increased random samples allow for better optimization of the model parameters, resulting in reduced loss and improved model performance. Moreover, the test loss graph Figure 6.1b shows that the model performs well until reaching a certain

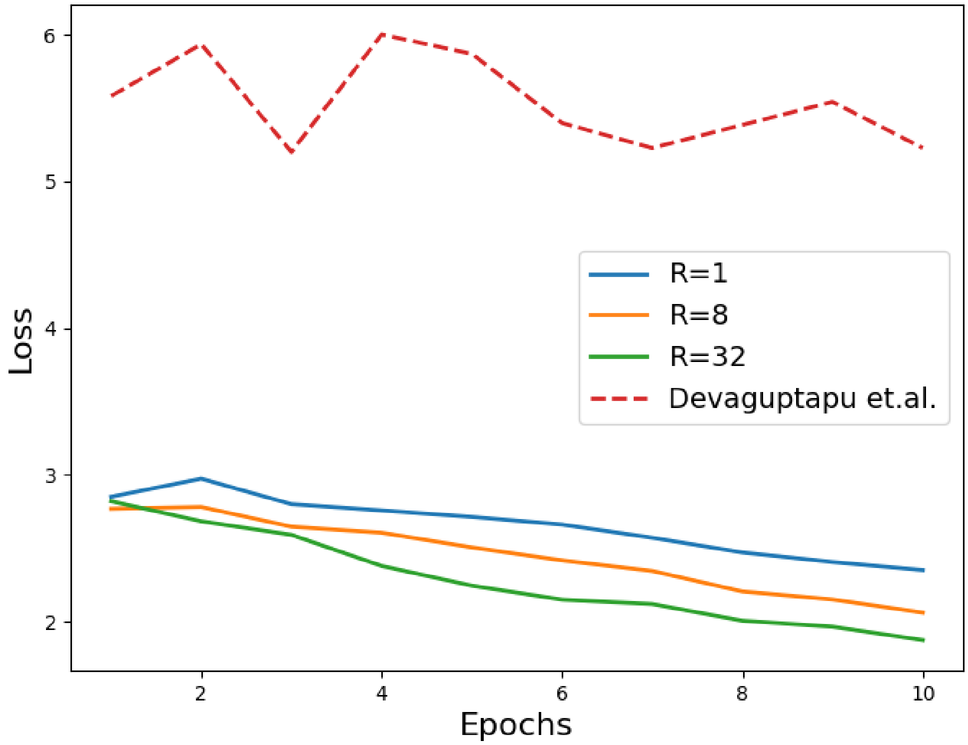

(a) Training Loss across epochs.

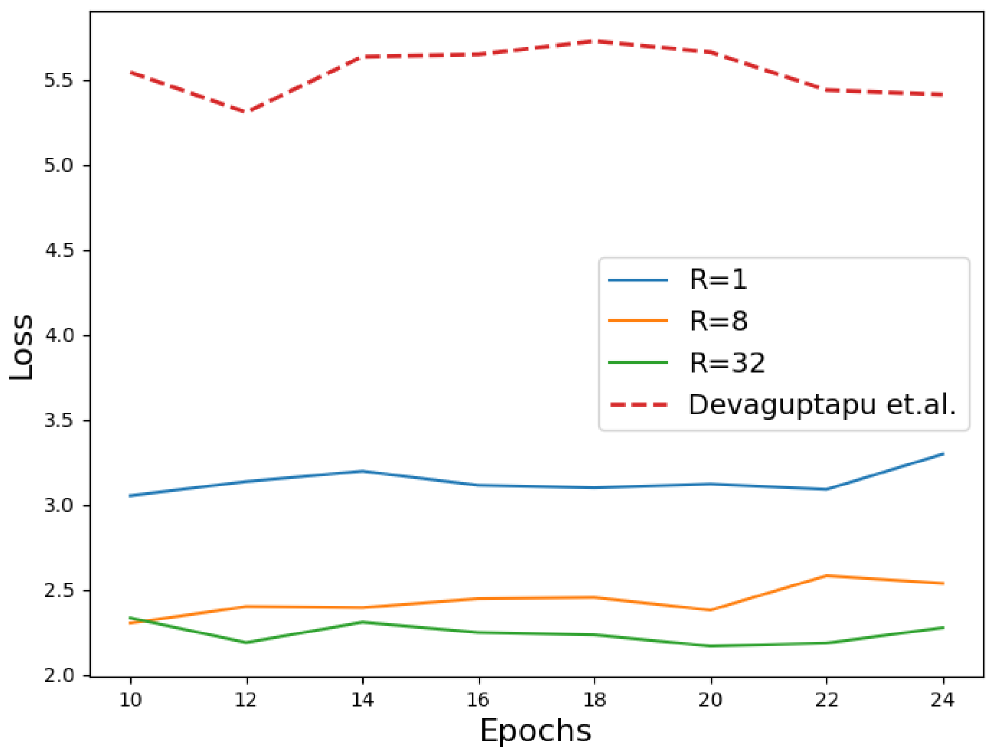

(b) Testing Loss across epochs.

Figure 6.1. Loss across Epochs.

number of epochs, after which it needs to be retrained. Furthermore, comparing the training and testing loss of the Discounted Satisficing model presented in [8], in graphs Figure 6.1a and Figure 6.1b, we can see that the loss for our proposed DDS model has decreased by almost 65%.

Figure 6.2 shows the average accuracy graph, which illustrates that the model's accuracy improves as R increases. The accuracy is calculated as the percentage of correct predictions with respect to the target output $y^*_{d\times T}$. This suggests that the model becomes more accurate in predicting the target output with a larger number of random samples during training. As the model is exposed to more diverse data points, it can adjust its parameters accordingly, leading to enhanced accuracy in its predictions. Furthermore, comparing the training and testing accuracy of the Discounted Satisficing model presented in [8], in Figure 6.2a and Figure 6.2b, we can see that the accuracy for our proposed DDS model has increased by almost 35%.

In addition, we analyzed the average λ (Figure 6.3a) and β (Figure 6.3b) of 10 drivers across epochs for different R values. These graphs show that as R increases, the average λ and β values become more stable across epochs, indicating that the model's parameters become better optimized with more diverse training data. This stability in the parameters suggests that the model becomes more robust and less sensitive to the specific training data it is exposed to. In summary, the results demonstrate that increasing the value of R has a positive impact on both the loss and accuracy of the model. The trend of decreasing loss and increasing accuracy with higher values of R indicates that the model's performance improves as it learns from more diverse training data. These findings suggest that increasing the number of samples, as represented by the parameter R, can lead to better model performance and improved accuracy in predicting the target output. Furthermore, we show that Dynamic Discounted Satisficing model has higher accuracy in predicting the stopping decision of a driver when compared with Discounted Satisficing model.

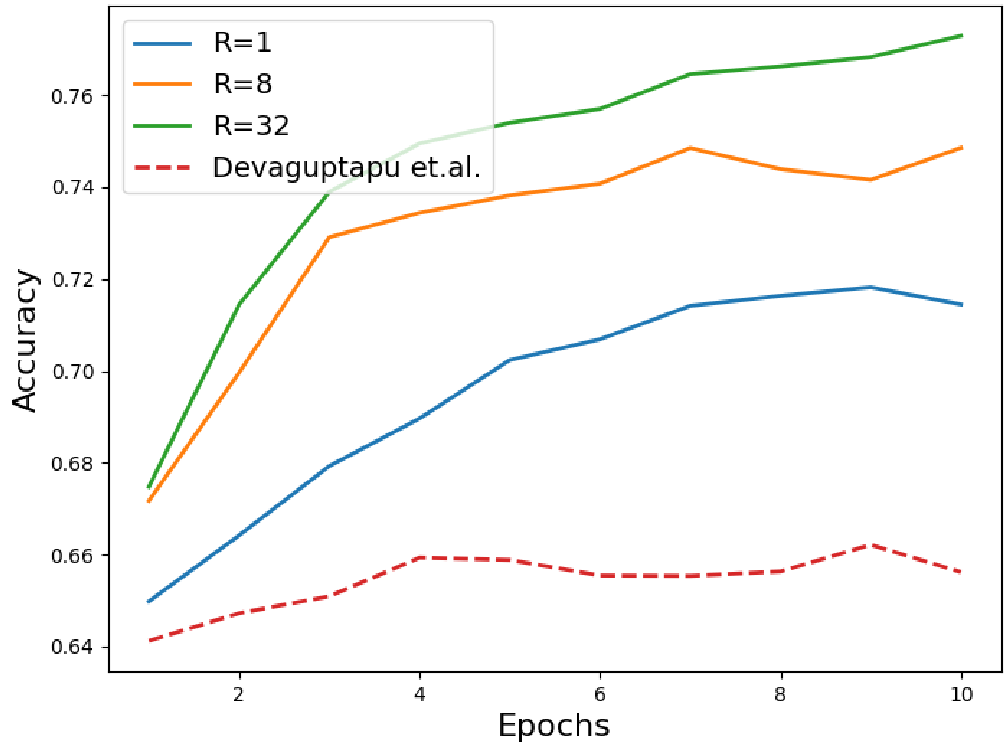

(a) Training Accuracy across epochs.

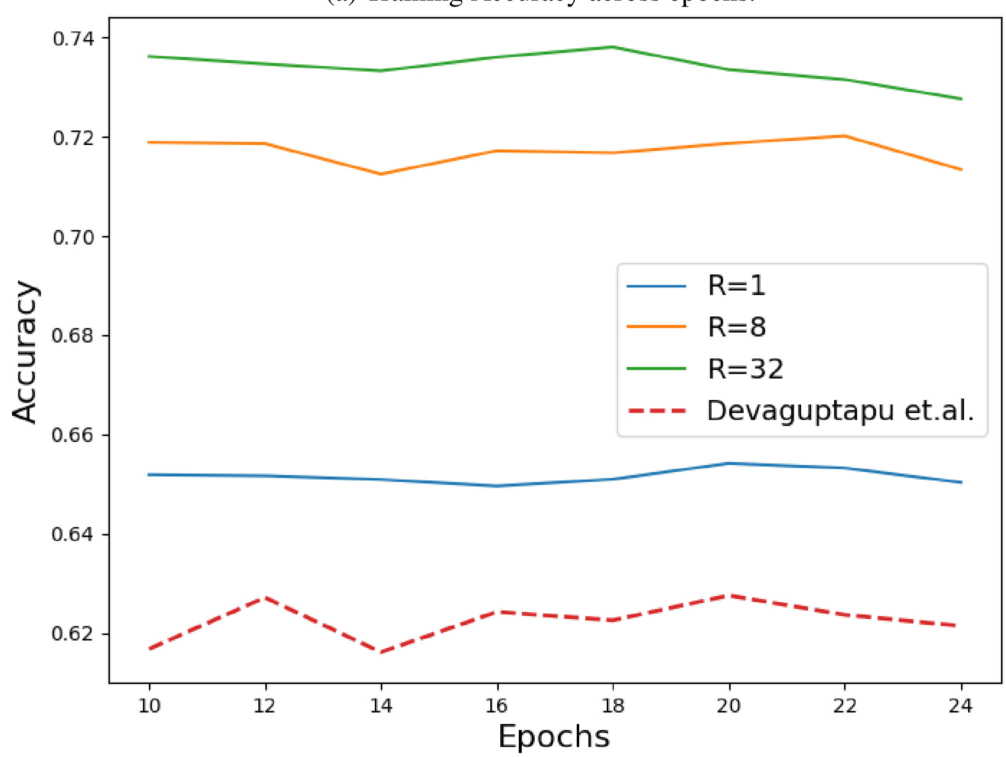

(b) Testing Accuracy across epochs.

Figure 6.2. Accuracy across Epochs.

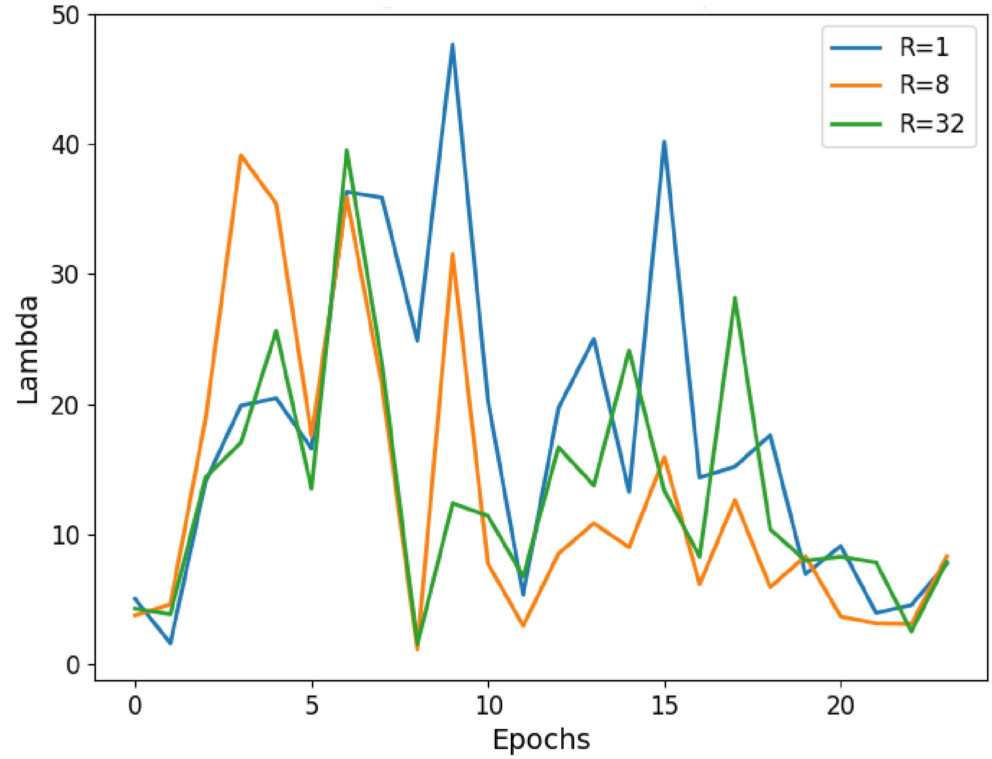

(a) Average λ for 10 drivers across Epochs

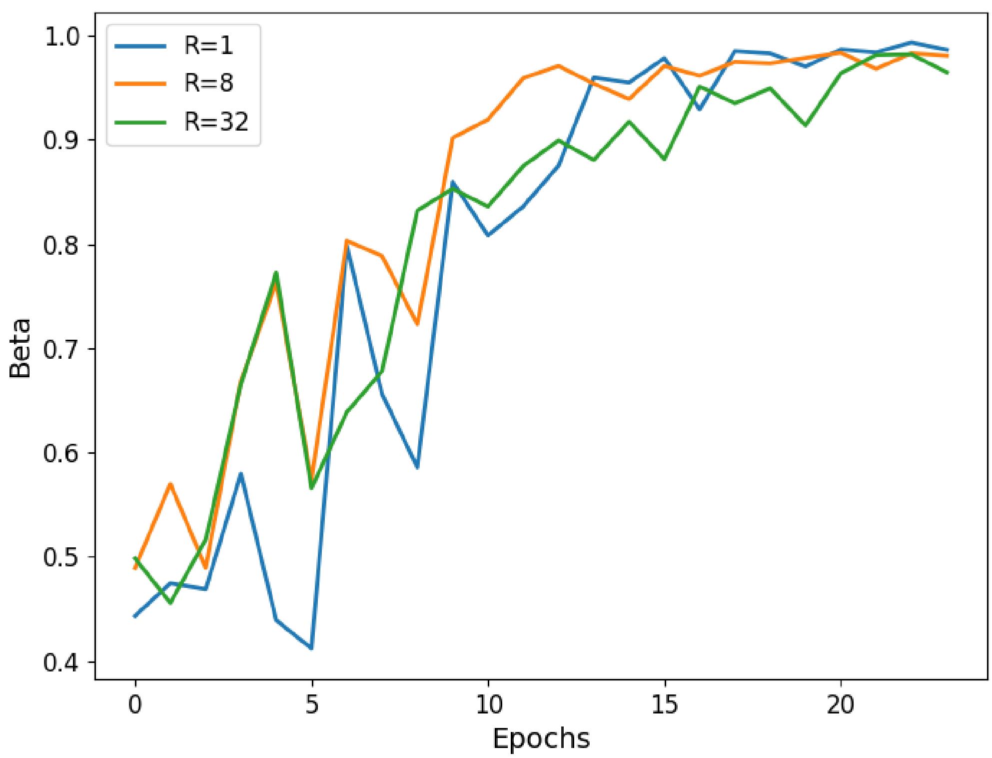

(b) Average β for 10 drivers across Epochs

Figure 6.3. Average λ and β across epochs

7. CONCLUSION AND FUTURE WORK

In this study, we developed a custom neural network model to predict the stopping time of taxi drivers in the City of Chicago based on the total amount paid for each trip they make. Our model achieved an accuracy of 35% and outperformed a baseline model [8] that predicted a constant value as the stopping time. In the future, we can improve the model's performance by incorporating additional features such as weather data, traffic congestion data, and the taxi driver's demographics. We can also explore other machine learning algorithms to compare their performance with our custom neural network model. Additionally, we can evaluate the model's performance on a more datasets to further validate its effectiveness in predicting the stopping time of taxi drivers in the City of Chicago.

REFERENCES

[1] Robert Hahn and Robert Metcalfe. The ridesharing revolution: Economic survey and synthesis. *More equal by design: economic design responses to inequality*, 4, 2017.

[2] Hao Yi Ong, Daniel Freund, and Davide Crapis. Driver positioning and incentive budgeting with an escrow mechanism for ride-sharing platforms. *INFORMS Journal on Applied Analytics*, 51(5):373–390, 2021.

[3] Massimiliano Gastaldi, Riccardo Rossi, and Gregorio Gecchele. Effects of driver task-related fatigue on driving performance. *Procedia-Social and Behavioral Sciences*, 111: 955–964, 2014.

[4] Richard C Green and Sanjay Srivastava. Expected utility maximization and demand behavior. *Journal of Economic Theory*, 38(2):313–323, 1986.

[5] E Frank Harrison and Monique A Pelletier. Managerial attitudes towards strategic decisions: Maximizing versus satisficing outcomes. *Management Decision*, 1997.

[6] Paul Reverdy, Vaibhav Srivastava, and Naomi Ehrich Leonard. Satisficing in multi-armed bandit problems. *IEEE Transactions on Automatic Control*, 62(8):3788–3803, 2016.

[7] Olena Kaminska, Allan L McCutcheon, and Jaak Billiet. Satisficing among reluctant respondents in a cross-national context. *Public Opinion Quarterly*, 74(5):956–984, 2010.

[8] Mounica Devaguptapu. *On predicting stopping time of human sequential decision-making using discounted satisficing heuristic*. Missouri University of Science and Technology, 2020.

[9] City of Chicago. Taxi trips: City of chicago: Data portal, Mar 2023. URL `https://data.cityofchicago.org/Transportation/Taxi-Trips/wrvz-psew`.

[10] Niels Agatz, Alan Erera, Martin Savelsbergh, and Xing Wang. Sustainable passenger transportation: Dynamic ride-sharing. 2010.

[11] Álvaro Aguilera-García, Juan Gomez, Guillermo Velázquez, and Jose Manuel Vassallo. Ridesourcing vs. traditional taxi services: Understanding users' choices and preferences in spain. *Transportation Research Part A: Policy and Practice*, 155: 161–178, 2022.

[12] Maria Vega-Gonzalo, Álvaro Aguilera-García, Juan Gomez, and José Manuel Vassallo. Traditional taxi, e-hailing or ride-hailing? a gsem approach to exploring service adoption patterns. *Transportation*, pages 1–40, 2023.

[13] Erik Ferguson. The rise and fall of the american carpool: 1970–1990. *Transportation*, 24(4):349–376, 1997.

[14] Kalon L Kelley. Casual carpooling—enhanced. *Journal of Public Transportation*, 10 (4):119–130, 2007.

[15] Catherine Morency. The ambivalence of ridesharing. *Transportation*, 34:239–253, 2007.

[16] Nelson D Chan and Susan A Shaheen. Ridesharing in north america: Past, present, and future. *Transport reviews*, 32(1):93–112, 2012.

[17] Chengcheng Dai. Ridesharing recommendation: Whether and where should i wait? In *Web-Age Information Management: 17th International Conference, WAIM 2016, Nanchang, China, June 3-5, 2016, Proceedings, Part I 17*, pages 151–163. Springer, 2016.

[18] Michael K Svangren, Mikael B Skov, and Jesper Kjeldskov. Passenger trip planning using ride-sharing services. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2018.

[19] Pengfeng Shu, Ying Sun, Yifan Zhao, and Gangyan Xu. Spatial-temporal taxi demand prediction using lstm-cnn. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, pages 1226–1230. IEEE, 2020.

[20] Jinmao Zhang, Huanchang Chen, and Yiming Fang. Taxiint: Predicting the taxi flow at urban traffic hotspots using graph convolutional networks and the trajectory data. *Journal of Electrical and Computer Engineering*, 2021:1–9, 2021.

[21] Tal Altshuler, Yaniv Altshuler, Rachel Katoshevski, and Yoram Shiftan. Modeling and prediction of ride-sharing utilization dynamics. *Journal of Advanced Transportation*, 2019:1–18, 2019.

[22] Xiaowei Hu, Shi An, and Jian Wang. Taxi driver's operation behavior and passengers' demand analysis based on gps data. *Journal of advanced transportation*, 2018, 2018.

[23] Charles C Macadam. Understanding and modeling the human driver. *Vehicle system dynamics*, 40(1-3):101–134, 2003.

[24] James G March Herbert A Simon. *Organizations*. 1958.

[25] Barry Schwartz, Andrew Ward, John Monterosso, Sonja Lyubomirsky, Katherine White, and Darrin R Lehman. Maximizing versus satisficing: happiness is a matter of choice. *Journal of personality and social psychology*, 83(5):1178, 2002.

[26] Aditya Mahajan and Demosthenis Teneketzis. Multi-armed bandit problems. *Foundations and applications of sensor management*, pages 121–151, 2008.

[27] P. Reverdy, V. Srivastava, and N. E. Leonard. Satisficing in multi-armed bandit problems. *IEEE Transactions on Automatic Control*, 62(8):3788–3803, Aug 2017. ISSN 0018-9286. doi: 10.1109/TAC.2016.2644380.

[28] Nir Shlezinger, Jay Whang, Yonina C Eldar, and Alexandros G Dimakis. Model-based deep learning. *Proceedings of the IEEE*, 2023.

[29] Paul J Werbos. Backpropagation Through Time: What It Does and How To Do It. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.

VITA

Sree Pooja Akula worked as a Graduate Research and Teaching assistant under Dr. Venkata Sriram Siddhardh Nadendla in the Computer Science department at Missouri University of Science and Technology (Missouri S&T). She received her Master's Degree in Computer Science from Missouri S&T in July 2023. She received her Bachelor's degree in Computer Science and Engineering from GRIET(India) in 2020.