

RESEARCH INTO THE USE OF MACHINE LEARNING AND NEURAL NETWORKS ON PERSONALISED NUMBER PLATES

Project Proposal

Abstract

Machine learning has a heavy use in many different areas of the modern world. This document discusses an investigation into the use of machine learning with regards to personalized registration plates for vehicles, by analysing the choices made by previous personalised plate holders.

Lee Rodger
1809707

Table of Contents

Table of Figures	2
<i>1 – Introduction</i>	<i>3</i>
<i>2 – Background Research.....</i>	<i>4</i>
2.1 – Supervised Learning	4
2.2 – Unsupervised Learning	5
2.3 – Neural Networks	5
2.3.1 – Real World Applications of Neural Networks	6
2.4 – Applying AI Ideas.....	6
<i>3 – Aims & Objectives</i>	<i>7</i>
3.1 – Project Aim.....	7
3.2 – Objectives	7
<i>4 – Research Methods</i>	<i>8</i>
4.1 – Strategy and Approach	8
4.2 – Risk Analysis.....	9
4.3 – Data Evaluation	9
<i>5 – Software Development Methods.....</i>	<i>10</i>
<i>6 – Resources and Planning</i>	<i>11</i>
<i>7 – Summary</i>	<i>12</i>
<i>References.....</i>	<i>13</i>

Table of Figures

Figure 1 - An example of a decision tree taken from Rokach and Maimon [3]	4
Figure 2 - An example of data after a k-means clustering algorithm has been applied. Taken from Dey [2].	5
Figure 3 - Software Development Lifecycle taken from Synotive - https://www.synotive.com/blog/software-development-client-questionnaire	10
Figure 4 - Gantt chart showing each stage of development and research throughout the project	11

1 – Introduction

In modern day computing, artificial intelligence plays a key role in many people's day to day lives. We can find many examples of AI from autonomous vehicles, to electric toothbrushes as demonstrated in an article by Vincent James [1]. We are progressing into a world where we use the power of computing to solve many of our day-to-day problems as well as revolutionise the industry with new technologies.

The aim of his project is to research and investigate machine learning through neural networks and apply these ideas to a system whereby a user will receive customised registration plate suggestions based upon his/her name. The suggestions brought back will adapt as the system is used more regularly and provide results based on choices made by previous users. The data will be used within a neural network so that suggestions will be based on what has been learnt by others' preferences.

A successful outcome for this project would enable us to see how viable machine learning can be in any number of given scenarios. An organisation could adapt the use of neural networks in order to fit the needs for any service which they use/provide.

2 – Background Research

Dey describes machine learning as a means to teach machines how to handle data more effectively [2]. In his paper, he describes a number of different machine learning techniques, each of which can be attributed to a learning style. The two main categories of learning are labelled 'supervised' and 'unsupervised', however there are a vat variety of learning categories which all hold valuable techniques for learning with different types of data.

2.1 – Supervised Learning

This approach to machine learning always requires human intervention at some level. The algorithms must be given datasets for 'train' and 'test' data. The train dataset contains an output variable which needs to be predicted [2]. A common implementation of a supervised learning algorithm is a decision tree. These trees work by grouping attributes and then sorting each one based on its value.

Rokach and Maimon describe decision trees as "a classifier expressed as a recursive pattern of the instance space" [3]. Data flows through the decision tree to reach an eventual end-node, whereby a 'decision' has been made based on the data.

An example of such tree can be seen below in *Figure 1*. This tree looks at the data given and checks it against three scenarios. The first step checks the age of the subject within the data, if greater than 30, the traversal of the tree ends there. If not, the gender is then checked. If female, the final check is carried out, 'Last R'. This is a simple model of a supervised form of machine learning which constantly needs to be fed data and is built upon decisions made by humans beforehand. The tree essentially sorts the data it is given and produces a result based upon pre-determined attribute values.




Figure 1 - An example of a decision tree taken from Rokach and Maimon [3]

2.2 – Unsupervised Learning

Unsupervised learning algorithms learn new features about data as it is introduced. Comparisons are made between new and old data and it continues to learn as more and more data is introduced.

An example of this algorithm can be seen within an implementation known as ‘K-Means Clustering’. This technique automatically creates groups or ‘clusters’, and items which share similar characteristics are grouped together. Dey tells us that the name ‘k-means’ is given because the number of clusters which will be determined is unknown at the beginning [2]. Below is an example of a k-means clustering result, whereby we can see that the data has been classified based on similar characteristics [Figure 2].




Figure 2 - An example of data after a k-means clustering algorithm has been applied. Taken from Dey [2].

2.3 – Neural Networks

The primary focus of this research is around the implementation of neural networks. Gurney describes these networks as “an interconnected assembly of simple processing elements, units and nodes, whose functionality is loosely based on the animal neuron” [4]. Neural networks are particularly useful in cases whereby it may be a struggle to classify data. We can see neural networks being used in fields such as facial and speech recognition, whereby over continued use, the application using this network will become better at making decisions and smarter at understanding different dialects. These algorithms allow machines to learn independently, and therefore become better over time without there necessarily being a need to implement changes to the system.

Zwass tells us that neural networks are typically trained under supervised learning [5]. A network requires a level of training before it is ready to begin working alone on data it is given. Firstly, it must be told what it is looking for in the data so it can begin to categorise what it sees. For example, a speech recognition neural network would need to be trained against a number of different accents so that it can begin to learn words being spoken in different ways. A paper published by Microsoft details the extensive use of deep neural networks with regards to speech recognition [6]. The network analyses voice data for use with speech recognition within applications and devices.

Neural networks have the ability to be as complicated as necessary when it comes to sorting out the data they are given. A more complex neural network will require a larger amount of compute time, so it is important to find the balance between accuracy and speed. It would be unwise to create a very complex network for an application which does not require a high level of accuracy.

Unfortunately for developers, according to Rojas, neural networks require a large amount of computation time [7]. This can be problematic for developers who want to create quick-to-use applications that have the power of artificial intelligence which neural networks provide. However, the benefits outweigh the computational expense in most cases.

2.3.1 – Real World Applications of Neural Networks

In today's world, we are always exposed to this type of learning without even knowing it. Many homes are now adopting 'smart' devices such as Google Home and Amazon Alexa. These devices collect voice data continually to adapt to its specific users as well as learn different styles in which users may speak.

The Alexa neural network plays a big focus on personalised data. In a paper by Kim, Kim, Kumar and Sarikaya of Amazon [8], it is explained that Alexa utilises a "Personalized attention method". This is particularly interesting with regards to the research being carried out for this project, as a large focus on that on neural networks are personalisation.

2.4 – Applying AI Ideas

By looking at these modern-day examples, we can see just how complex the tasks typically are for neural networks, and we can see how computationally expensive they can be. This may often make people question the application of a neural network within simple software applications, but this study aims to prove that there is a need for these complex systems in small tasks. On a grander scale, this application may be used hundreds of thousands of times a year and having a system which can speed up the rate of producing these plates will aid heavily in the long-run.

3 – Aims & Objectives

3.1 – Project Aim

This project overall aims develop a system for personalised registration plates, whereby a user can enter a name and have a personalised number plate for a vehicle is generated based around the name entered. The system should incorporate machine learning and utilise a neural network in order to generate these personalised registration plates. By definition of a neural network, the application should be able to generate more desirable results for users throughout an extensive use of the system.

3.2 – Objectives

The aim has been outlined in order to determine the success of the project. The area of machine learning and artificial intelligence is extremely broad, therefore, to avoid losing focus on the ultimate goal, a number of objectives have been set to work towards the final goal.

The three main objectives of this project are as follows:

- Develop a Personalised Number Plate Generation Application
- Implement a Neural Network System onto the Personalisation Application
- Conduct a User Study to determine if the implementation of a neural network has a positive impact

To determine whether or not a neural network implementation of a personalisation system is viable, we need an application by which we can begin to test these ideas. Firstly, an application which generates personalised suggestions should be created. This gives us a means of comparison when looking into the benefits of the AI implementation.

After the first objective, we can begin to implement a neural network to see the results the application generates over time given n number of entries. We can study the data and see what sorts of suggestions the application begins to give, whilst also comparing with results from the basic version of the application.

Once we have a working version of the application both with and without the neural network implementation, we can begin to look into what users feel about the results which are given from both applications. Over extended use, we can begin to determine if the results give a more 'desirable' outcome, indicating whether or not the system works given an extended period of use. If the application is able to deliver more desirable results with the neural network implementation than without at a reduced number of attempts, we can determine just how effective this implementation would be in a real-world scenario.

4 – Research Methods

When conducting research for this project, it is important that we aim to achieve the most reliable results possible, to determine the validity of the outcome. We hope to see that the use of neural networks does indeed aid in helping to suggest personalised registration plates. If this does indeed work, we may be able to implement these ideas into many other areas, speeding up the process by which people receive desirable results, and the improving the users' perceived quality of these results overall.

4.1 – Strategy and Approach

When conducting research for this project in the form of a user study, it must be ensured that the users are not led to be biased before they begin to use the systems. The applications, both with and without the neural network, will need to look and feel identical so as to not sway the user to favouring one system of the other unfairly.

Each participant will be in front of two identical software application. The only difference being that one will contain the neural network and the other will not. The participant will not be told which version of the software contains a neural network, therefore not influencing a biased result.

The software (both versions) will allow the user to enter their first name (or any name of their choosing), and a personalised vehicle registration plate will be generated based around the name entered. Each participant will enter the same name 10 times on each system.

Each user should rate their personalised results on a scale from 'Not at all desirable' to 'Very Desirable' for each plate that is generated. This should be split over a Likert scale [9], of 0-5 points. This scale has been chosen as it allows to gain a user's opinion of the results whilst mapping them to a numerical value. Since we are only interested in the opinion of the result rather than an analysis of the experience, a quantitative measure is ideal. This data can then be mapped and trends can be analysed to establish if the system was successful in producing more desirable results when using a neural network.

It should also be noted how many users will have used the system prior to voting, giving an indication as to more desirable results appearing over time, as the system has been able to learn what sorts of results are more or less desirable.

At the end of the user studies, it should be clear through analysis of the results if the software with the neural network was able to produce more desirable results than the other.

4.2 – Risk Analysis

There are no immediate risks for the users involved in the study due to the simplicity of the experiment and lack of hazardous equipment. However, it is not to be ignored that there are a number of risks with reliability of data when it comes to this type of research. For example, in the real world, this system would be used to create and sell number plates. This means that each plate would have a price which would be a very important factor in desirability for the user. To combat this issue, we could develop a further study with price attributes being accounted for and determine if the results have suffered any change.

For a real-world implementation, a user should be able to enter a price range for which they are willing to spend. This would of course alter the results and would require a deeper neural network to count for the adding factors.

The data retrieved from the study may also not be entirely accurate as not everyone will have the same perspective when it comes to what they find to be desirable. For example, a plate generated may not be attractive for one user but may be highly desirable for another. These may lead to an inconsistency in the results. This could be avoided by conducting multiple repeat tests for the same participants and measuring results based on each individual separately. This however would mean that the network would only cater for one individual and may be compromised when used by other users.

4.3 – Data Evaluation

Once the study has been completed, all of the results will be collected and plotted on line-graphs where we should instantly obtain some visible results. We would expect the graphs representing the application with no neural network to produce randomly dotted levels of satisfaction with no clear trend in the data.

We would then expect to see a definitive trend in the data on the graph representing the application with the implementation of a neural network. With an already trained network, we would expect to see higher levels of satisfaction earlier on, but also should be able to notice that users who used the system after others had used it, should have higher results of desirability overall, since the network will be continually trained throughout the study.

Over time, we should expect to see less anomalous results appearing in the data as the network should become continually better at delivering desirable results. This would further support the claim that neural networks can be beneficial within a system such as this one.

If another study were to be conducted after the initial one, it should be expected that the results for desirability should begin in the same region as the results from the last few tests in the first study. These results should continue to improve until the network has reached its 'peak' and may not be able to get better without being introduced to more complex data and without being made into a 'deeper' neural network.

5 – Software Development Methods

The software will be designed and implemented using an agile methodology [10]. The software product will go through a number of iterations, adding further levels of functionality over time, but always having a level of usability at each iteration. This development model allows for a working product to be produced in a less amount of time compared to traditional approaches.

Throughout development however, the software development lifecycle will be taken into consideration when implementing new features.




Figure 3 - Software Development Lifecycle taken from Synotive - <https://www.synotive.com/blog/software-development-client-questionnaire>

By following this, it will be far easier to develop new features and implement them over time through a number of iterations. Each stage of the lifecycle shows how a software application is created which the cycle starting again at each new feature being implemented.

Even though the software will be developed by a single individual, a scrum methodology will be adhered to where each feature will be developed using short sprints to make sure that time is being managed effectively. This is to ensure that by the time of the user study, the software is fully ready and the neural network is working correctly.

6 – Resources and Planning

The software implementation will require an IDE in order to develop the application for the user study. The user study itself will require one or more computers for the participant to be able to use the software.

To ensure that the project is completed within the allocated period of time, it is essential that the project is time-managed well to allow for any unforeseen issues arising. This may involve time to fix bugs which will inevitably occur during the development process.




Figure 4 - Gantt chart showing each stage of development and research throughout the project

A good portion of time has been allocated for research and testing so that a strong familiarity of the ideas and technology can be established. This will ensure that the best possible version of the neural network can be developed and implemented successfully and therefore allow for a more successful project outcome overall.

This chart has been designed with each stage estimated at slightly longer period than originally thought. This is to allow time for debugging and any other problems which may arise. This will hopefully lead to a more successful project in the long-run.

7 – Summary

The project, by the end, should be able to determine whether or not the implementation of a neural network has a beneficial impact on a system such as the one in question. It will also give good insight into the field of artificial intelligence and may open up possibilities of alternative technologies for achieving a similar goal in future work.

Neural networks can be a difficult concept to get right when it comes to implementing a new software design, however, this study should show how viable they are within modern day applications and therefore will be able to give us indication as to just how dependent we will become on these technologies in the future.

References

- [1] J. Vincent, "The State of AI in 2019," The Verge, 28 January 2019. [Online]. Available: <https://www.theverge.com/2019/1/28/18197520/ai-artificial-intelligence-machine-learning-computational-science>. [Accessed 25 May 2019].
- [2] A. Dey, "Machine Learning Algorithms: A Review," *International Journal of Computer Science and Information Technologies*, vol. 7, no. 3, pp. 1174-1179, 2016.
- [3] L. Rokach and O. Maimon, *Data Mining and Knowledge Discovery Handbook* (Second Edition), New York: Springer Science and Business Media, LLC, 2010.
- [4] K. Gurney, *An Introduction to Neural Networks*, London: UCL Press, 1997.
- [5] V. Zwass, "Neural Network," Britannica, [Online]. Available: <https://www.britannica.com/technology/neural-network>. [Accessed 31 May 2019].
- [6] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath and B. Kingsbury, "Deep Neural Networks for Acoustic Modeling in Speech Recognition," *IEEE Signal Processing Magazine*, pp. 82-97, November 2012.
- [7] R. Rojas, "Neural Networks," Springer-Verlag, Berlin, 1996.
- [8] Y.-B. Kim, D. Kim, A. Kumar and R. Sarikaya, "Efficient Large-Scale Neural Domain Classification with Personalized Attention," Amazon.
- [9] A. Joshi, S. Kale, S. Chandel and D. K. Pal, "Likert Scale: Explored and Explained," *British Journal of Applied Science & Technology*, vol. 7, no. 4, pp. 396-403, 2015.
- [10] Agile Alliance, "Agile 101," Agile Alliance, 2019. [Online]. Available: <https://www.agilealliance.org/agile101/>. [Accessed 31 May 2019].