University of Texas at Arlington

MavMatrix

2020 Spring Honors Capstone Projects

Honors College

5-1-2020

DISABLED PARKING PASS DETECTOR FOR LICENSE PLATE RECOGNITION PARKING SYSTEM USING A CONVOLUTION NEURAL NETWORK AND TRANSFER LEARNING

Thomas Perappadan

Follow this and additional works at: https://mavmatrix.uta.edu/honors_spring2020

Recommended Citation

Perappadan, Thomas, "DISABLED PARKING PASS DETECTOR FOR LICENSE PLATE RECOGNITION PARKING SYSTEM USING A CONVOLUTION NEURAL NETWORK AND TRANSFER LEARNING" (2020). *2020 Spring Honors Capstone Projects*. 44.
https://mavmatrix.uta.edu/honors_spring2020/44

This Honors Thesis is brought to you for free and open access by the Honors College at MavMatrix. It has been accepted for inclusion in 2020 Spring Honors Capstone Projects by an authorized administrator of MavMatrix. For more information, please contact leah.mccurdy@uta.edu, erica.rousseau@uta.edu, vanessa.garrett@uta.edu.

Copyright © by Thomas Antony Perappadan 2020

All Rights Reserved

DISABLED PARKING PASS DETECTOR FOR LICENSE PLATE RECOGNITION PARKING SYSTEM USING A CONVOLUTION NEURAL NETWORK AND TRANSFER LEARNING

by

THOMAS ANTONY PERAPPADAN

Presented to the Faculty of the Honors College of

The University of Texas at Arlington in Partial Fulfillment

of the Requirements

for the Degree of

HONORS BACHELOR OF SCIENCE IN COMPUTER SCIENCE

THE UNIVERSITY OF TEXAS AT ARLINGTON

May 2020

ACKNOWLEDGMENTS

I would like to acknowledge members of the UTA faculty who have contributed to my learning experience and understanding of key fundamental Computer Science concepts that inspired me to work on this project and take an interest in the field of Computer Vision.

I would like to extend a special thank you to Dr. Shawn Gieser for being my advisor for this project and providing support and guidance when needed. Would also like to acknowledge Dr. Vassilis Athitsos, whose computer vision and machine learning classes helped augment my knowledge in this field and who was a great resource for me to learn from.

May 06, 2020

ABSTRACT

DISABLED PARKING PASS DETECTOR FOR LICENSE PLATE

RECOGNITION PARKING SYSTEM USING A

CONVOLUTION NEURAL NETWORK

AND TRANSFER LEARNING

Thomas Antony Perappadan, B.S Computer Science

The University of Texas at Arlington, 2020

Faculty Mentor: Shawn Gieser

The License Plate Recognition Advanced Parking System is a fully automated parking enforcement system that is able to identify license plate information of cars as they pull into a parking lot and track these cars to their individual parking spots after they have entered. This allows an administrator to view exactly where each car is parked and enforce parking rules down to a per spot level accuracy. Computer vision techniques are used to read license plates as cars enter the parking lot and additional cameras will track each car to its individual spot. The disabled parking pass component of this project is the feature to detect disabled parking passes on cars as they enter and validate the cars that can park in disabled parking spots using this information. This will be implemented using a

convolutional neural network and computer vision techniques to detect and identify these passes in a video input.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

CHAPTER 1

INTRODUCTION

1.1 The Problem

Parking enforcement is a logistical problem that can be dealt with in multiple ways. Traditionally, cars displayed tags that an enforcement officer would check to verify that the vehicle had permission to be in a lot at a specific time. More recently, companies have created License Plate recognition systems that can automatically capture a car's license plate information and verify it with an internal database to validate that car's parking privileges in a specified lot. Most of these systems involve a camera mounted to a vehicle that drives through a parking lot to validate each car's license plate. Some systems make use of cameras at entrances and exits to enforce parking rules in a parking lot/garage. One of the shortcomings of existing systems is that they do not allow for per-spot-level tracking of the cars that are parked in a lot. It is harder to enforce constraints like special reserved permits and handicap parking automatically without knowing exactly where a car is parked.

1.2 Automated Parking System Solution

This system allows an administrator to view exactly where each car is parked and enforce parking rules down to a per-spot-level accuracy. This will be done using computer vision to read license plates as cars enter the parking lot and additional cameras that will track each car to its individual spot. The additional Honors credit component of this project that will be implemented is the feature to detect handicap passes on cars as they enter and validate the cars that can park in handicapped spots using this information. This will require an additional component to the image recognition algorithm that will be able to identify handicap passes that are displayed on cars. It will involve implementing and training a convolutional neural network that will be able to detect and identify these disabled passes and automatically validate their parking.

CHAPTER 2

LITERATURE REVIEW

2.1 License Plate Recognition Software and Existing Companies

The current solutions that exist in the market are mainly aimed at parking enforcement and is marketed towards the administrators rather than the users. For example, UTA is currently using a service called NuPark which implements basic functionality related to enforcing parking permits including cameras mounted onto a car or cameras mounted to a fixed location [3]. The customer wanted a similar product that integrates blockchain technology and we decided to add the functionality of being more user oriented and being able to monitor entire parking lots at a per spot level. The research done on existing products yielded a host of other companies that offer very similar if not the exact same features to NuPark [1, 3, 5, 6]. The plan for the Senior Design project is to try and integrate some open source projects that already exist and use computer vision to implement license plate recognition. The work that this team is interested in includes the Open ALPR project on GitHub [4] and the DIY ALPR project that implements license plate recognition on the Raspberry Pi [2].

2.2 Disabled Pass Detection

There are various methods already described in literature that employ computer vision techniques to try and detect specific objects. The most basic form of object detection is template matching and there have been various attempts to improve on this method and create faster algorithms, such as a very fast template matching method [7]. However, recent advances in machine learning technology and the availability of large datasets have allowed new and more efficient methods to develop, such as convolutional neural networks (CNNs).

2.2.1 Convolutional Neural Networks (CNNs)

CNNs combine template matching with machine learning to greatly improve detection accuracies. Template matching requires us to know exactly what template is being searched for to obtain meaningful results. However, CNNs can employ template matching at the initial layers of the network and then learn the templates that need to be extracted from an image by using a large amount of training data. CNNs have been used to detect and classify objects for a wide variety of applications such as vehicle detection [8], facial recognition, food type recognition [9], etc. The main drawback to this method is the large amount of training data required to train this network and enable it to learn the right features to look for in an image. There are methods to augment a smaller dataset such as applying different transforms to the same image [10]. Since this application requires near-real time detection of the template, faster Region-based CNNs (R-CNNs) are of particular interest as they allow for fast recognition in an image [11].

2.2.2 Transfer Learning

Transfer Learning is the method that will be used to achieve the detection of disabled parking passes. Transfer Learning makes use of other CNN models that have been trained for recognizing a completely different set of objects. Even though they recognize different objects, the base training layers can be reused [12]. The idea here is that the basic templates that the model learn to extract still applies to the new set of objects and the learning can be transferred to the new model. This allows us to use significantly fewer training examples compared to the number required to build a model from scratch. There is also the added advantage of the significant reduction in training time as a lot of the layers are already trained and frozen.

CHAPTER 3

METHODOLOGY

3.1 Data Collection

One of the main challenges for this project was data collection due to complications caused by the COVID-19 pandemic. The original plan was to collect training data by driving around the parking lots and taking pictures of cars with a handicap sign hanging from the rear-view mirror. This did not seem feasible after stay at home orders were issued, so custom template was created instead (Appendix A). This template was used to collect a data set comprised of around 100 images that have the template in various orientations and backgrounds.

Using Transfer Learning is especially beneficial in this scenario as this test set is very limited. It is also taken into consideration that using a custom template could cause issues with detecting the actual passes in the future, but the model can easily be retrained in the future using more data to get better real-world results. Since this is a classifier that only needs to distinguish whether the sign is present or not (binary classifier), training images that do not contain the template are also required. A dataset of random articles was collected using a camera and some pictures from open source online repositories.

3.2 Training the Custom Model

After researching the various available networks to use for Transfer Learning, the ResNet50 network was chosen. Transfer Learning works by using the weights from this pre-built model that is trained using huge datasets to recognize hundreds of classes. Thus, this model already has the ability to extract useful information from the input images. The top layer of this network was reconfigured by removing its output layer and implementing a custom fully-connected layer and output layer that can be trained to detect the specific template that is needed for this application. See table 3.1 for the model summary from this training stage.

Table 3.1: ResNet50 CNN Network Training Summary

Layer (type)	Output Shape	Param #
model_1 (Model)	(None, 204800)	23587712
dense_1 (Dense)	(None, 512)	104858112
dropout_1 (Dropout)	(None, 512)	0
dense_2 (Dense)	(None, 512)	262656
dropout_2 (Dropout)	(None, 512)	0
dense_3 (Dense)	(None, 1)	513

Total params: 128,708,993; Trainable params: 105,121,281; Non-trainable params: 23,587,712

3.3 Measuring Performance

Expressing the performance of any classification model with a single percentage for accuracy does not portray the full picture as there could be a number of other factors at play that affect this number in non-intuitive ways. The following measures can be used to identify how well the model is performing and to decide whether changing hyperparameters has a positive effect or not on this model.

Table 3.2: Initial Performance Measurements

	Predicted Positive	Predicted Negative
Actual Positive	TP = 96	FN = 246
Actual Negative	FP = 75	TN = 261

Legend:
True Positives (TP) - The number of template images identified correctly
False Positives (FP) – The number of images incorrectly identified as a handicap pass
True Negatives (TN) – The number of images correctly identified as not a parking pass
False Negatives (FN) – The number of handicap passes not identified
Accuracy = (TP+FP)/(TP+FP+TN+FN)

CHAPTER 4

DISCUSSION

4.1 Choosing a Transfer Learning Model

There are multiple models to choose from when trying to pick a Transfer Learning model to base the CNN model on. From my research into existing models, the ones that stood out to be most state of the art are VGG [13], InceptionV3 [14] and ResNet5 [15]. ResNet5 was chosen as it seems to have faster training speeds along with better accuracies for image classification. This works by increasing the depth, which results in fewer parameters than widening the network [15].

4.1.1 Initial Results

The initial model that was trained kept the previous convolutional layers constant from the ResNet5 model and trained the fully connected layers that ResNet5 was outputting to. This did not yield very good results (Table 3.1). The training accuracy was high, but the testing accuracy seemed dismal and this usually suggests over-fitting where the model learns the training data itself and fails to extract meaningful results, which give bad testing accuracies.

Figure 4.1: Training Accuracy vs Epochs

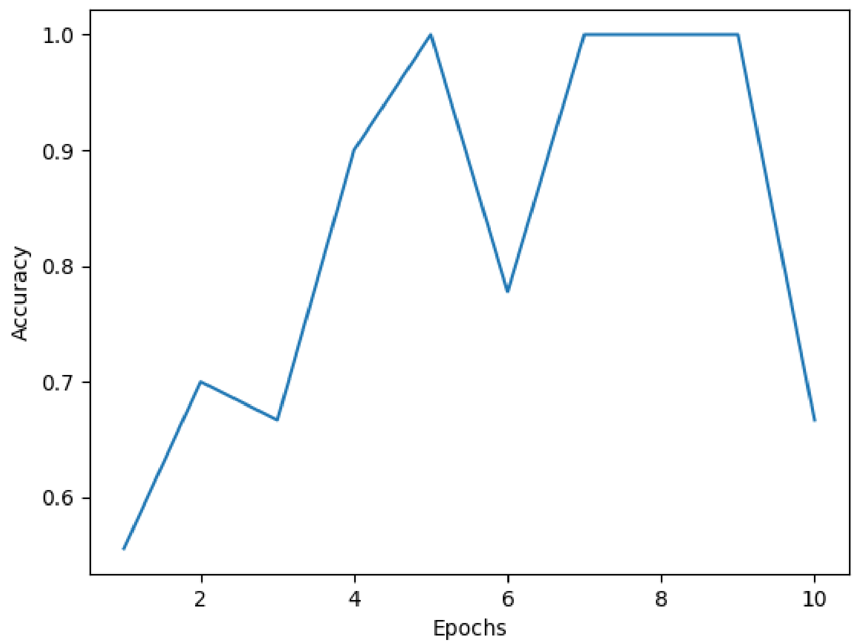

Figure 4.2 Validation Accuracy vs Epochs

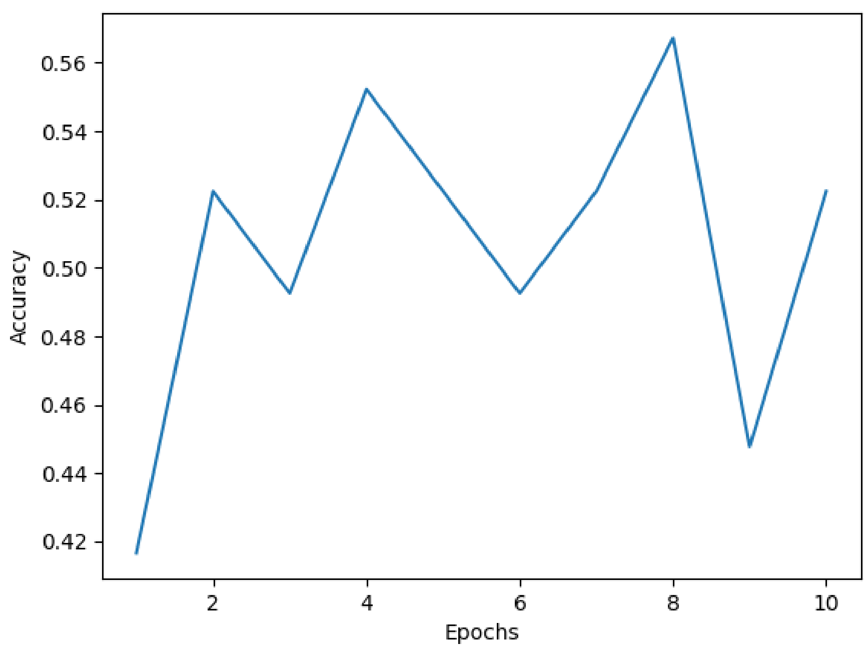

4.1.2 Improving the Model's Performance

One of the reasons for getting this result was the low number of training images that were available to use to train this model. However, before trying to add more training images to the dataset, the next step taken to improve the model's performance was to unfreeze some of the pre-trained layers that came with the ResNet50 model to fine-tune the performance of this Transfer Learning model. Doing this resulted in the following accuracies and results.

Figure 4.3: Improved Training Accuracy vs Epochs

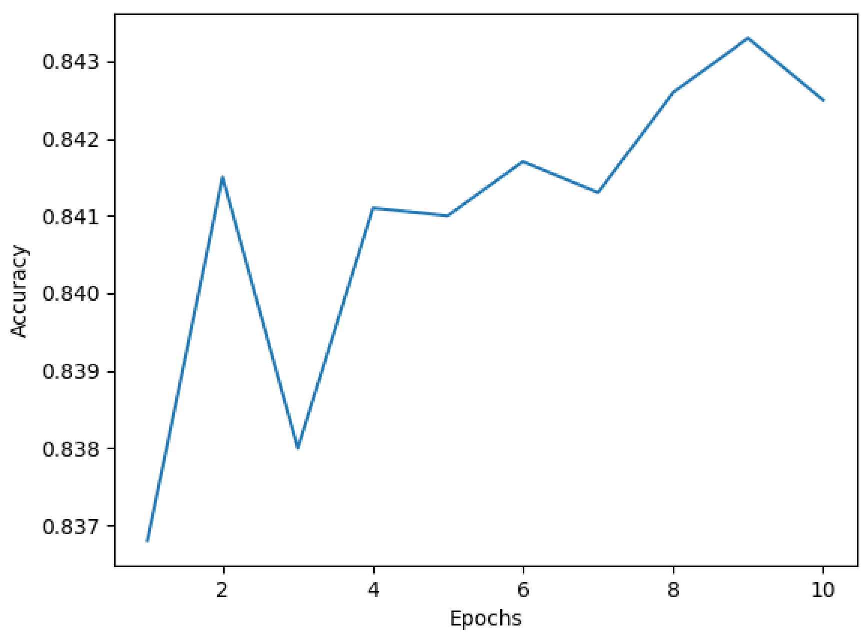

Figure 4.4: Improved Validation Accuracy vs Epoch

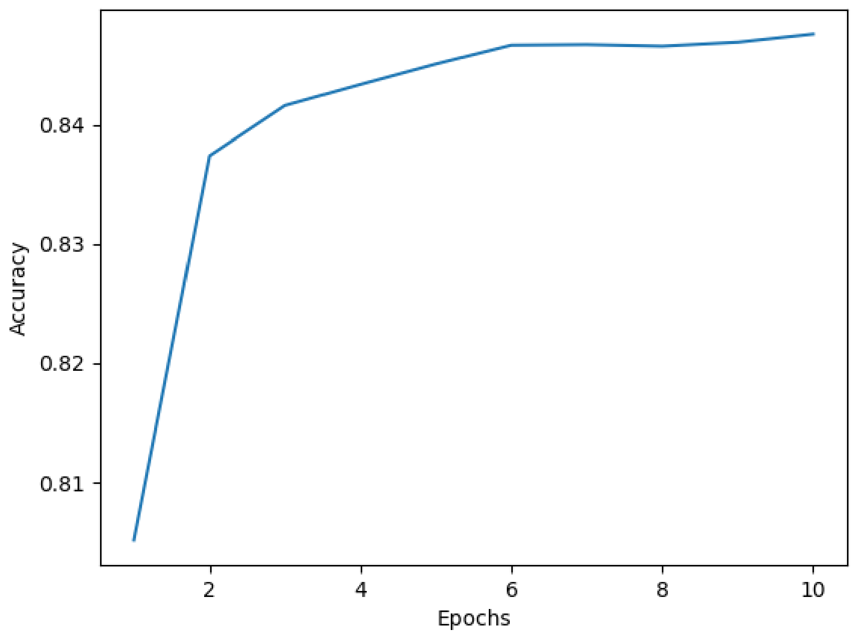

CHAPTER 5

CONCLUSION

The Transfer Learning technique to develop a handicap parking pass detector has proven effective to integrate into the larger system. The accuracy in this experiment was limited by the size of the training dataset and the availability of a variety of real-world images to train on. This accuracy can be further improved upon by augmenting the image data set and then further tuning the hyperparameters to detect the parking pass signs.

Transfer Learning techniques greatly help speed up the training phase for image recognition and classification problems and are effective in taking information that was previously learned by other models to extract key base-level information about an image and then use that learning to classify our own custom set of images instead of the dataset and classes that the model was originally trained on. If the new dataset is different from the original dataset, only re-training the top layer might not produce the best results as seen in the initial implementation. A better accuracy was achieved by allowing some of the lower levels to be retrained to fine tune the weights and extract the ideal features that are needed for the classification. In this case, adding an extra hidden layer on top of the previously trained ResNet5 model and then allowing some of the hidden layers to be retrained to fit the current dataset increased the accuracy and produced actionable results that could be used in this application.

APPENDIX A

CUSTOM TEMPLATE TO REPRESENT HANDICAP PARKING PASSES

REFERENCES

[1] "LPR Enforcement," LPR Enforcement AIMS Parking Management Software. [Online]. Available: https://aimsparking.com/features-benefits/lpr-enforcement. [Accessed: 05-Feb-2020].

[2] J Franklin-Hodge "DIY ALPR," GitHub. [Online]. Available: https://github.com/jfh01/diy-alpr. [Accessed: 05-Feb-2020].

[3] "NuPark," LPR Technology. [Online]. Available https://www.passportinc.com/lpr/. [Accessed: 05-Feb-2020].

[4] "Open ALPR," GitHub. [Online]. Available: https://github.com/openalpr/openalpr. [Accessed: 05-Feb-2020].

[5] "Perceptics," License Plate Readers. [Online]. Available: https://www.perceptics.com/products/license-plate-readers/. [Accessed: 05-Feb-2020].

[6] "Operations Security," License Plate Recognition automates Parking Management. [Online]. Available https://ops-com.com/parking-security-platform/license-plate-recognition/. [Accessed: 05-Feb-2020].

[7] H. Schweitzer, J. W. Bell, and F. Wu, "Very Fast Template Matching," Computer Vision — ECCV 2002 Lecture Notes in Computer Science, pp. 358–372, 2002.

[8] D. He, C. Lang, S. Feng, X. Du, and C. Zhang, "Vehicle detection and classification based on convolutional neural network," Proceedings of the 7th International Conference on Internet Multimedia Computing and Service - ICIMCS '15, 2015.

[9] Y. Kawano and K. Yanai, “Food image recognition with deep convolutional features,” Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct Publication - UbiComp '14 Adjunct, 2014.

[10] F. Chollet, “The Keras Blog,” The Keras Blog ATOM. [Online]. Available: https://blog.keras.io/building-powerful-image-classification-models-using-very-little-data.html. [Accessed: 08-May-2020].

[11] S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 6, pp. 1137–1149, 2017.

[12] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, “A Survey on Deep Transfer Learning,” arXiv:1808.01974 [cs, stat], Aug. 2018, Accessed: May 08, 2020. [Online]. Available: http://arxiv.org/abs/1808.01974.

[13] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” arXiv:1409.1556 [cs], Apr. 2015, Accessed: May 09, 2020. [Online]. Available: http://arxiv.org/abs/1409.1556.

[14] J. Fu, “Application of Modified Inception-ResNet and CondenseNet in Lung Nodule Classification,” in Proceedings of the 3rd International Conference on Computer Engineering, Information Science & Application Technology (ICCIA 2019), Chongqing City, China, 2019, doi: 10.2991/iccia-19.2019.28.

[15] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, Jun. 2016, pp. 770–778, doi: 10.1109/CVPR.2016.90.

BIOGRAPHICAL INFORMATION

Thomas Perappadan is a senior majoring in Computer Science at the University of Texas at Arlington (UTA). He has developed an interest in computer vision research during his college career and has worked on developing computer vision software at the University of Texas at Arlington Research Institute that is able to capture a person's portrait using a camera and transform it into a sketch drawing that can be drawn on a piece of paper using a mechanical Drawbot in under four minutes. This was being developed for an outreach program at the Perot museum in Dallas by the research institute. Thomas is also actively involved in the Drone Club at UTA that he helped start and is responsible for building the drone club's fleet of three drones from scratch with ambitions to eventually equip these drones with software that can coordinate multi-drone flights simultaneously. Thomas plans on joining Microsoft as a Software Engineer after graduating from UTA with an Honors Bachelor of Science degree in Computer Science.