Developing a Classification Algorithm for Pedestrians and Street Signs

Aviv Zohman, Darron Kotoyan, Eric Wasserman

Abstract

Our grand challenge is to lower traffic accidents by using autonomous vehicles. To achieve this in the scope of our GCI project we created a machine learning algorithm that classifies street signs and pedestrians. This grand challenge contributes to the problem of car accidents. We wrote a machine learning algorithm that is trained to detect object and people on the road with the help of fastAI. Our grand challenge highlights machine learning, an essential mechanism to the furtherance of science and society, and its implementation and efficacy.

Introduction

- Our grand challenge is to create a machine learning algorithm that classifies street signs and pedestrians with 95% accuracy
- Car accidents cause 35,000 annual deaths, 3 million annual injuries, and create $230.6 billion in annual costs. Autonomous vehicles are a foreseeable solution to these problems.
- Since we are software engineers by trade, there is a large knowledge gap necessary to conceptualize the major parts of an autonomous vehicle. We therefore decided to focus on the software aspect for our grand challenge by implementing a classification model.

Methods

- We built and characterized a machine learning algorithm that is trained using images we downloaded from Google.
- We downloaded images of pedestrians and of street signs, as shown in Figure 1.
- The model requires enough computer power that we couldn't train the model using a CPU. We purchased an inexpensive GPU and trained our model with the GPU's aid.
- Our model was trained with four epochs, each epoch training in ~1 second, demonstrated in Figure 2.

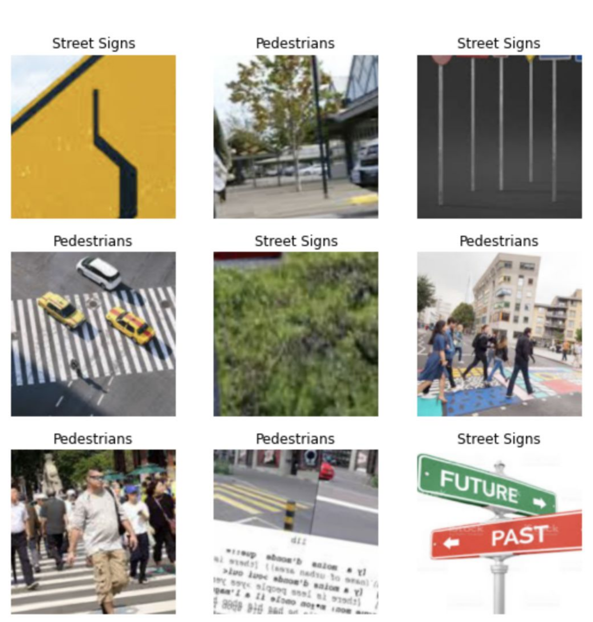

Figure 1. Examples of the pictures that were downloaded and used to train the model

epoch	train_loss	valid_loss	error_rate	time
0	0.358020	0.401790	0.125000	00:08
1	0.354501	0.528958	0.125000	00:07
2	0.313141	0.354825	0.101562	00:07
3	0.273332	0.379199	0.109375	00:07

Figure 2. Data table shows four trials with predictive and actual error time generated

Results

- Our confusion matrix, Figure 3.a, is our confusion matrix. From our confusion matrix we can calculate the following results:
  - Accuracy: ~89%
  - Precision: ~92%
  - Sensitivity: ~88%

Our data tells us that our machine learning algorithm is more precise when detecting a street sign (91%), than pedestrians (88%). Overall, it has an accuracy score of .89 (89%) and very good recall scores which indicates that the algorithm is correctly classifying a large total of relevant results.

	precision	recall	f1-score	support
Pedestrians	0.88	0.93	0.90	70
Street Signs	0.91	0.84	0.88	58
accuracy			0.89	128
macro avg	0.89	0.89	0.89	128
weighted avg	0.89	0.89	0.89	128

Figure 3. Data table showing the precision, recall, f1-score,and support for pedestrians and street signs

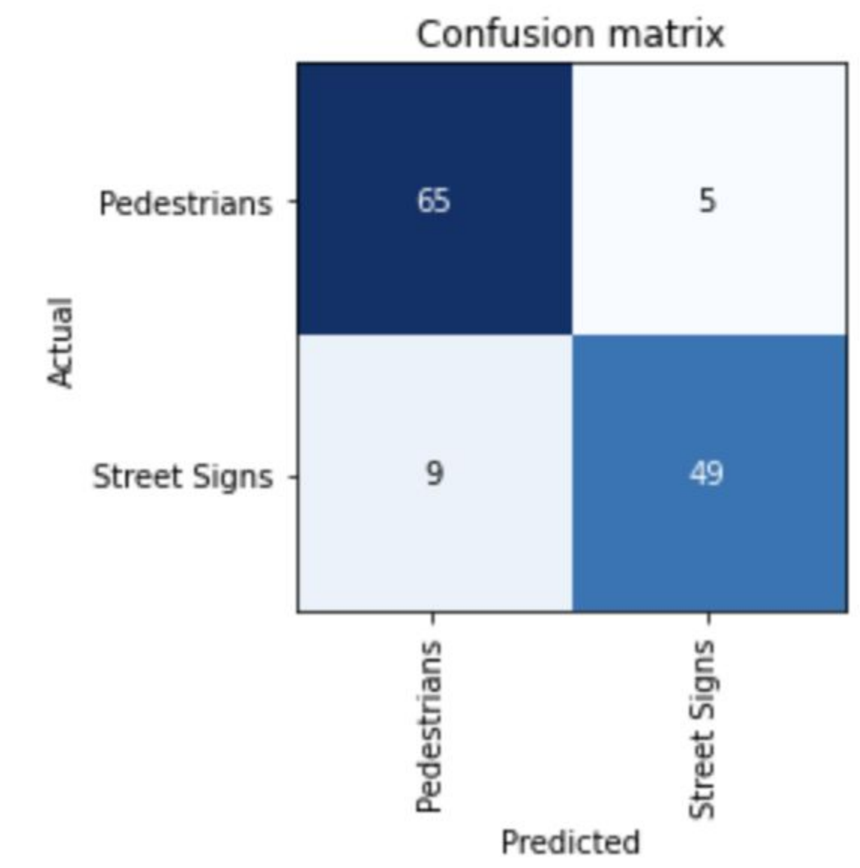

Figure 4. Confusion matrix between Pedestrians and Street signs generated from data

Conclusion

- The model correctly predicts street signs with 89% accuracy
- The model correctly predicts pedestrians with 89% accuracy

An accurate and wide breadth machine learning algorithm is essential to autonomous vehicles. Our accurate and low breadth model illustrates the implications machine learning has on autonomous vehicle safety and efficacy. The next steps are to expand the model classifications and investigate this algorithm's application to autonomous vehicles.

Acknowledgements

- Fast.AI
- Brenna Gormally
- Kevin Brice
- Jose Raul Gonzalez

References

- Howard, Jeremy, et al. 2018. *fastai*. GitHub. https://github.com/fastai/fastai

Author Contributions

- Aviv Zohman developed and trained the classification model and contributed to the poster. Eric Wasserman founded the GPU necessary to train the model and contributed to the poster. Darron Kotoyan contributed to the poster.