CS230

Using GRUs to Predict Crop Yield

James Zheng
Department of Computer Science
Stanford University
jameszhe@stanford.edu

Abstract

The abstract should consist of 1 paragraph describing the motivation for your paper and a high-level explanation of the methodology you used/results obtained.

1 Introduction

Food insecurity is a problem that affects about 795 million people worldwide (FAO 2015). As the population continues to grow, and natural phenomena such as climate change and drought magnify the global food problem, it is essential that we develop solutions to adequately feed the increasing number of people that will inhabit this planet. In fact, this is one of the goals outlined in the 2030 Agenda for Sustainable Development of the United Nations (United Nations 2015).

One way to tackle this issue is to improve our ability to estimate crop yields. By more accurately monitoring crop harvests, information can be obtained that increases food production (Dodds and Bartram 2016).

A type of mechanism that hasn't been applied to this field is the gated recurrent unit (GRU) found in certain RNNs. By incorporating GRUs, which use two gates to control the flow of information and provide the ability to memorize information, it is possible that a model with higher accuracy can be developed to predict crop yield.

2 Related work

Existing methods of predicting crop yield tend to use survey data that includes information on factors such as the weather and soil properties. While these approaches are successful in theory, they can only be applied to regions that can provide the necessary data. In developing regions, this kind of data is scarce. One alternate solution is the use of remote sensing data, which is much easier to obtain. Past research has applied various types of neural network architectures, including CNNs and LSTMs, coupled with a dimensionality-reducing Gaussian Process layer to this data to predict crop yield (You et al. 2017). This was effective due to the high-dimensionality of the remote sensing data; having the Gaussian Process layer enabled the neural network to more accurately extract features from the data. Additional research expanded on this by using You's model (trained on the United States) for transfer learning on Argentina and Brazil, which also displayed promising results (Wang et al. 2018).

3 Dataset and Features

The dataset is a collection of 32x32x9 spectral band and temperature histograms over a county's harvest season, paired with its soybean yield over the season in metric tonnes per hectare. An example of one such histogram is shown in Figure 1. 857 U.S. counties, 135 Argentinian counties, and

CS230: Deep Learning, Fall 2021, Stanford University, CA. (LateX template borrowed from NIPS 2017.)

32 Brazilian counties are represented to produce 9049 U.S., 1615 Argentinian, and 384 Brazilian datapoints, respectively. This data is from 2005 to 2016 and collected through MODIS satellite images. A 60-20-20 train-validation-test split is used for each country.

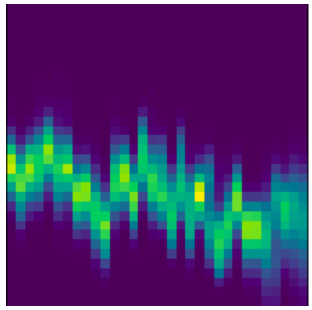

Figure 1: A band histogram from the dataset

This dataset was obtained from the SustainBench website at `https://sustainlab-group.github.io/sustainbench/docs/datasets/sdg2/crop_yield.html#download`.

4 Methods

Our neural network uses supervised learning to develop a deep learning model. Models are evaluated based on root mean squared error (RMSE) as well as R^2 metrics.

For baseline models, I included various ridge regression models as well as a linear regression model. Ridge regression models been shown to be effective for predicting crop yield (Bolton and Friedl 2013), so multiple such models with varying regularization strengths (alpha value α) are tested.

Once the baseline models are implemented, I will be training a recurrent neural network (RNN) model, specifically utilizing gated recurrent units (GRUs) to test on the crop yield data. This is a novel approach (to the best of my knowledge) in the field of crop yield prediction, and I will tune various hyperparameters to maximize the accuracy of the model. To build this model, I used Keras, a popular deep learning framework that is tightly integrated into TensorFlow.

To evaluate these models, I will be using the SustainBench package. This package provides dataloaders for all the datasets used in SustainBench, including the one that will be used in this project for crop yield prediction. Specifically, I will be looking at the data in the United States. Research from Wang et al. showed that transfer learning can be applied to other countries for crop yield prediction using CNNs, so theoretically, it may be possible that the results generated from this project can be used in developing regions as well.

5 Experiments/Results/Discussion

For my baseline ridge regression models, I used regularization strengths of $\alpha = 0.25, 0.5, 1, 2, 4$. Additionally, I tested a model that utilized linear regression. The results of those models, which includes their RMSE and R^2 values, are shown in Table 1 below:

Table 1

Model	**RMSE**	R^2
Ridge ($\alpha = 0.25$)	9.373	0.149
Ridge ($\alpha = 0.5$)	9.373	0.149
Ridge ($\alpha = 1.0$)	9.373	0.149
Ridge ($\alpha = 2.0$)	9.374	0.149
Ridge ($\alpha = 4.0$)	9.374	0.149
Linear Regression	9.369	0.150

We see that the results of each of the baseline models are similar, with the RMSE values all within the range 9.369 to 9.374.

For the RNN model, several hyperparameters were tuned in order to minimize loss, including the number of hidden layers, the number of nodes per layer, mini-batch size, and other decisions such as whether to include batch normalization. All models were trained for 60 epochs.

The model that was finally settled on includes 3 GRU layers followed by a batch normalization layer and a dense layer. The number of units per GRU layer decreased from 256 to 128 to 64. See Figure 2 below for an illustration of the RNN:

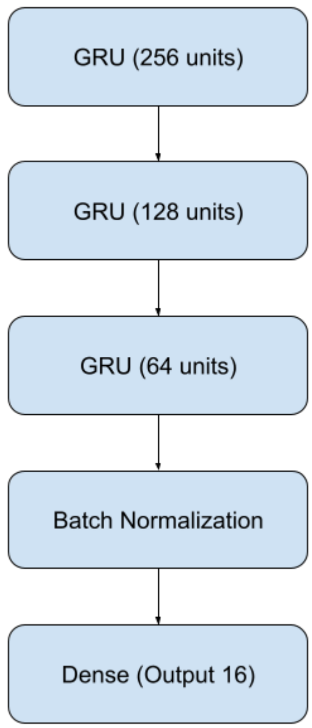

Figure 2: RNN architecture

The results of the SustainBench evaluator show that the RMSE value of the RNN model using GRU layers is 10.850, which is higher than the baseline models.

See the graphs below for a visual of the the loss and accuracy throughout the training of the model:

6 Conclusion/Future Work

Based on my research, using RNNs that incorporate GRUs are not necessarily better at predicting crop yield from histogram data than baseline models such as ridge regression are. The RMSE of the RNN that I developed turned out to be greater than those of the baseline models, indicating that it performed worse a predicting crop yield.

In the future, possible directions to continue this research include testing out other hyperparameters in conjunction with GRUs. These can include additional variations in the number of layers or number of nodes, or a different number of epochs, batch size, optimization algorithm, learning rate, etc.

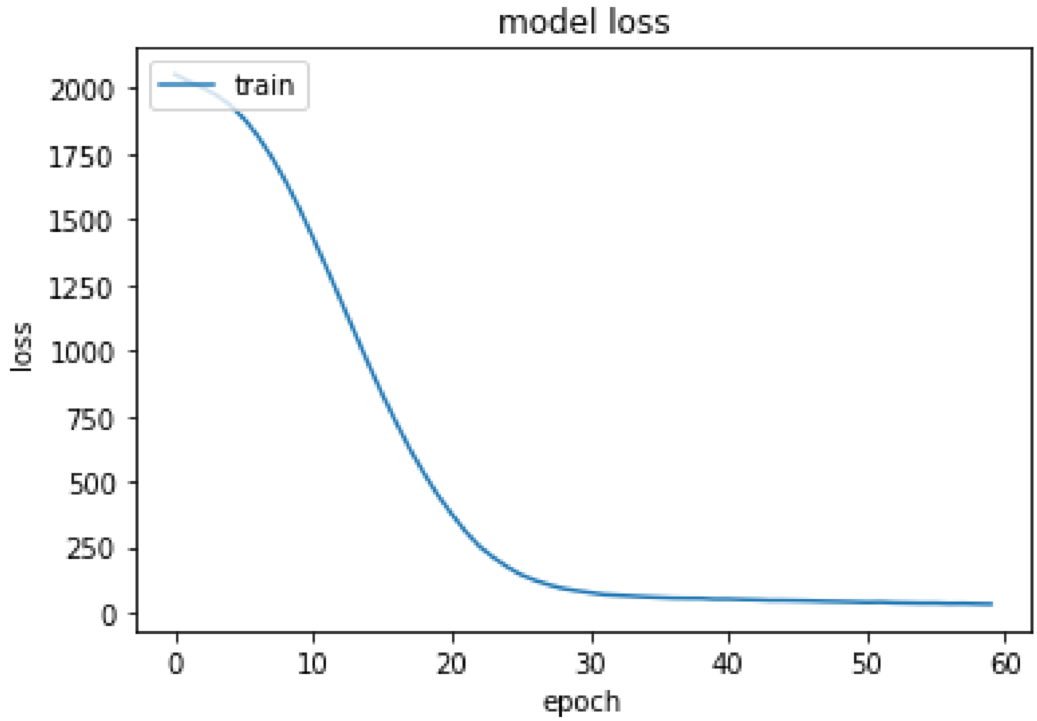

Figure 3: RNN Loss

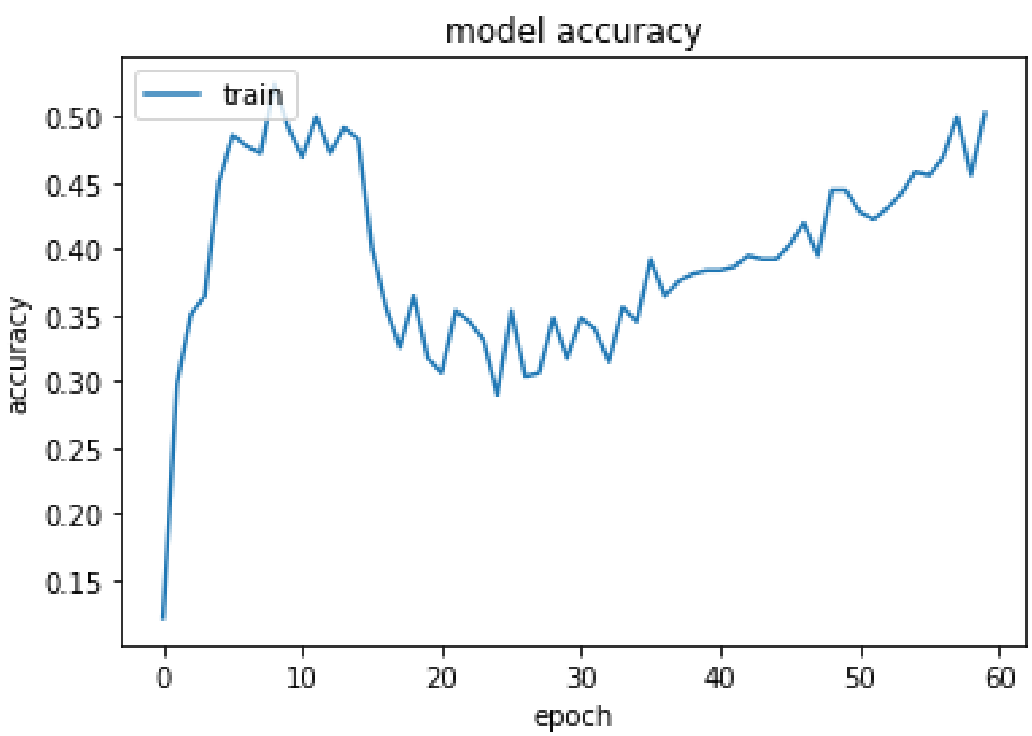

Figure 4: RNN Accuracy

Additionally, it may be interesting to incorporate long short-term memory (LSTM) layers into the model along with GRUs. LSTM layers are a similar type of RNN layer that involve three gates instead of the two seen in GRUs.

References

[1] A. X. Wang, C. Tran, N. Desai, D. Lobell, and S. Ermon. Deep transfer learning for crop yield prediction with remote sensing data. In Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, COMPASS '18, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450358163. doi: 10.1145/3209811.3212707. URL https://doi.org/10.1145/3209811.3212707.

[2] Bolton, D. K., and Friedl, M. A. 2013. Forecasting crop yield using remotely sensed vegetation indices and crop phenology metrics. Agricultural and Forest Meteorology 173:74–84.

[3] Dodds, F., and Bartram, J. 2016. The Water, Food, Energy and Climate Nexus: Challenges and an Agenda for Action. Routledge

[4] FAO. 2015. The state of food insecurity in the world. meeting the 2015 international hunger targets: Taking stock of uneven progress.

[5] J. You, X. Li, M. Low, D. Lobell, and S. Ermon. Deep Gaussian Process for Crop Yield Prediction Based on Remote Sensing Data. In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17, page 4559-4565. AAAI Press, 2017. URL https://aaai.org/ocs/index.php/AAAI/592AAAI17/paper/view/14435.

[6] United Nations, G. A. 2015. Transforming our world: the 2030 agenda for sustainable development. New York: United Nations.