Solar Flare Prediction With Recurrent Neural Networks

Jill M. Platts
AFRL/RISA, Rome, NY
John A. Marsh, Michael J. Reale, Christopher W. Urban
SUNY Polytechnic Institute, Utica, NY

ABSTRACT

As the star closest to Earth, the Sun offers a wealth of information on its own composition and behavior, as well as a basis for the composition and behavior of all stars. The Sun's violent magnetism gives rise to various solar activity, including solar flares. A form of space weather, strong solar flares can damage communications and expose astronauts to dangerous radiation. Monitored nearly 24 hours a day by various satellites, an amalgamation of the Sun's magnetic field properties, observed over a period of time, is theorized to be indicative of an upcoming flare. A popular choice when working with time series data, or sequences, Recurrent Neural Networks (RNNs) are excellent for solar flare forecasting models. RNNs are equipped with an internal memory and are able to understand sequences more effectively than other types of neural networks. Our work aims to prove the validity of using RNNs with multivariate time series data, related to the Sun's magnetic fields, to predict solar flares one day prior to occurrence. Predicting solar flares by class, one day prior to occurrence is also explored.

1. INTRODUCTION

Imposing order on chaos, magnetism plays a fundamental role in the universe. Spanning entire galaxy clusters, as well as between one galaxy and the next, magnetic fields have been detected across the cosmic web [26]. Locally, Earth's magnetic field creates a magnetosphere around the planet that acts as a shield. Combined with gases that make up the planet's atmosphere, Earth is able to endure the source of the closest and most powerful magnetic field, the Sun. [25]

1.1 THE SUN

As the closest star to Earth, understanding the Sun is key to understanding stars that are too far away to be observed in detail. Around 93 million miles away, the Sun's dynamic energy directly impacts the Earth's climate, weather and atmosphere. One of its three outer layers, the Sun's photosphere is an opaque, gas surface. Below the photosphere is the Sun's interior, including the convection zone, radiation zone and core. Above the photosphere is both the chromosphere and corona, respectively. [17]

The turbulent photosphere layer is made up of constantly moving gas currents. Granules, which resemble grains of rice, reveal rising gas currents, while neighboring dark areas reveal descending gas currents. Supergranules contain many granules and reveal both ascending gas currents and horizontally moving gas currents (from center to edges). Appearing temporarily, from a few hours to a few months, sunspots are darker, cooler blotches on the photosphere. Appearing over regions experiencing an increase in magnetic activity, sunspot numbers rise and fall over the course of each 11-year solar cycle. A larger number of sunspots, at any given time, generally indicates an increase in solar activity. [17]

1.2 ACTIVE REGIONS

A visible manifestation of the Sun's magnetic field, sunspots are a unique form of solar activity. Each sunspot indicates an area of the solar surface experiencing vertical magnetic flux (vertical flow of energy). However, sunspots are not distributed evenly across the photosphere; they are often concentrated in clusters, or what is commonly referred to as active regions. [21]

Scattered across the Sun's photosphere, the intense magnetic flux within active regions contributes to their clearly recognizable appearance. When a magnetograph is used to capture an image, a magnetogram, of the Sun's magnetic fields, active regions appear as white spots with magnetic field lines flowing towards the viewer (positive polarity) and black spots with magnetic field lines flowing away from the viewer (negative polarity) [27]. A magnetogram

can be seen in Fig. 1. Active regions can vary significantly in lifespan and size, but they can also grow to be as large as 100,000 km (over 62,000 miles) across and can live for several months [21].

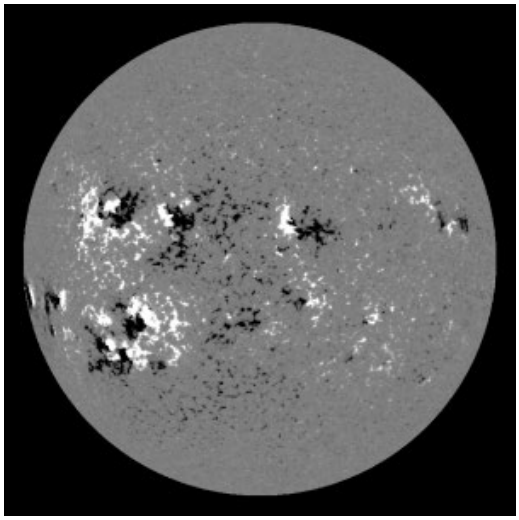

Fig. 1: A Solar Magnetogram [27].

Emerging quickly, the creation of both sunspots and active regions is thought to be caused by toroidal-shaped magnetic flux rising out of the convection zone and breaking through the photosphere in a fragmented form. It is further theorized that the intense magnetic flux originates from a global toroidal field generated by a dynamo somewhere under the convection zone [28]. [21]

1.3 SPACE WEATHER

Recently beginning its 25th solar cycle, with the 24th running from 2008 to 2019, the Sun is currently midlife (around 4.5 billion years away from its death) and experiencing less active solar cycles, resulting in less solar weather [4][7]. The main source of space weather we experience here on Earth, the Sun is constantly pushing charged particles into space, also known as solar wind [24].

One form of space weather, solar wind is generally easily pushed away by Earth's own, albeit weak, magnetic field. When solar wind pushes past Earth's magnetic field and into its atmosphere, auroras, sometimes referred to as polar lights or northern lights, occur [24]. More intense forms of space weather include solar flares and coronal mass ejections (CMEs). A CME can be described as a large, concentrated cloud of magnetized particles hurled into space, right through any surrounding solar wind [30]. In contrast, a solar flare is an enormous burst of radiation visible via optical light, as well as X-ray and ultraviolet (UV) [29]. A solar flare is the largest type of explosive event that directly impacts Earth. Additionally, although they are fundamentally different, solar flares and CMEs often accompany one another [6]. Despite our deceptively large distance from the Sun, space weather occurs quickly and can damage satellites, disrupt power grids and expose astronauts to intense radiation and unstable conditions [29].

1.4 SOLAR FLARES

Intense eruptions, solar flares release anywhere from 〖10〗^17 to 〖10〗^25 Joules of energy over time periods ranging from milliseconds to hours [16]. For scale, one large solar flare can release as much energy as the Earth's current population could consume in 100,000 years [17]. Occurring near or in active regions, solar flares are classified based on their brightness in the X-ray spectrum (wavelength range 0.1-0.8 nanometer), in order of increasing energy flux. Table 1 shows the accepted classification of A through X-class flares.

Table 1: Solar Flare Classes [21].

Class	W/m² between 1 & 8 Ångströms
A	**$<10^{-7}$**
B	**$\geq 10^{-7} < 10^{-6}$**
C	**$\geq 10^{-6} < 10^{-5}$**
M	**$\geq 10^{-5} < 10^{-4}$**
X	**$\geq 10^{-4}$**

Within active regions, solar flares often emerge near or directly above sunspots, particularly where sunspots with opposite polarities are in close contact [18]. While X-class flares can cause major blackouts on Earth, M-class flares are generally only capable of causing brief blackouts and C-class flares lack the strength to cause any significant adverse effects on Earth [21]. Currently, monitoring solar X-ray flux is the accepted method for detecting solar flare events. A solar flare event is identified by a rapid rise in X-ray flux, with a slower decline following [18]. Exhibiting the Sun's overall unpredictability, solar flare events can be both abrupt and short or gradual and long. The diversity of solar flare events continues to present a forecasting conundrum.

2. RELATED WORK

Due to being an active area of research, solar physics, and the general behavior and composition of the Sun, are not yet fully understood. Additionally, due to distance and technological limitations, our understanding of star mechanics is minimal. With much left to discover, the Sun, our closest star, is continuously studied. Regarding space weather, there have been many approaches to forecasting the Sun's behavior.

2.1 DATA SOURCES

Launched in 2010, the Helioseismic and Magnetic Imager (HMI) on the Solar Dynamics Observatory (SDO) has been a popular source for gathering data on the Sun's photosphere and magnetic field lines therein [1]. Nearly all recent work with space weather forecasting focuses on data from the HMI. Observing the Sun continuously, the HMI is one of three instruments on the SDO designed to observe and study the Sun. Recording four main types of data, including dopplergrams, continuum filtergrams, vector magnetograms and line-of-sight magnetograms, the HMI gathers over one terabyte of data a day, every day. [32]

Considering solar flares are the result of a sudden release of magnetic energy, data acquired from the HMI's magnetograms is thought to be the most current and best source for flare prediction models. Due to our understanding that solar flares occur near or in active regions [21], the team behind HMI created a data product focused on tracking active regions, called HARPs. Working to meet the demands of space weather research, the team behind HMI further improved their HARP product by re-adding certain data from HMI's full magnetograms and creating yet another data product called SHARPs. This newer, space weather-focused data product has served as the basis for nearly all most recent research on the application of neural networks to the solar flare forecasting problem.

2.2 RECENT APPROACHES

In *Predicting Solar Flares Using a Long Short-term Memory Network* [2], a popular RNN, called a Long Short-Term Memory (LSTM) network, is applied to a combination of 25 datapoints from HMI's SHARP data product and 15 calculated datapoints related to class-based flaring history. A binary classification approach, researchers sought to predict whether or not an AR would produce a C or M-class flare within the following 24hrs. LSTM ultimately performed well, outperforming past models on the multivariate, time series data, but researchers concluded that using fewer features may have achieved better performance.

In the most recent approach to solar flare prediction, *Predicting Solar Flares with Machine Learning: Investigating Solar Cycle Dependence* [1], researchers sought to predict whether or not an AR would produce any class of flare, with a focus on C and M-class flares. Using a LSTM, input to the model consisted of 24hr time series', using 20 SHARP parameters, and retaining the 12-minute cadence each SHARP parameter is regularly calculated at (a maximum of 120 times per AR, per day). The researchers concluded that this is a viable approach to solar flare prediction, however, significant performance variation is seen across different years, possibly due to the evolving solar cycle [17] or simply noisy data.

3. METHODOLOGY

Echoing the most recent research [1] on solar flare prediction, 20 key parameters of Helioseismic and Magnetic Imager's (HMI's) SHARP data product were selected to serve as multivariate feature input. Specifically, the hmi.sharp_720s SHARP series of definitive data was used [8]. Table 2 lists each of the 20 parameters, as well as their descriptions.

Calculated per day, every 12 minutes, SHARP summary parameters are derived from magnetograms and are a numerical representation of active region behavior on the Sun's photosphere. Moving with the rotation of the sun, active regions evolve over time due to magnetic energy. The behavior of ARs, at any given time, is thought to contain indicators for upcoming solar flare occurrences, therefore, it is regularly hypothesized [1][2], that AR characteristics can be used to predict solar flares.

Table 2: SHARP Parameters and Descriptions [1].

Parameter	Description
USFLUX	Total unsigned flux, in Mx
MEANGAM	Mean inclination angle, gamma, in deg
MEANGBT	Mean value of the total field gradient, in G $M^{m^{-1}}$
MEANGBZ	Mean value of the vertical field gradient, in G $M^{m^{-1}}$
MEANGBH	Mean value of the horizontal field gradient, in G $M^{m^{-1}}$
MEANJZD	Mean vertical current density, in mA m^{-2}
TOTUSJZ	Total unsigned vertical current, in A
MEANALP	Total twist parameter, alpha, in 1 $M^{m^{-1}}$
MEANJZH	Mean current helicity, in G^2 m^{-1}
TOTUSJH	Total unsigned current helicity, in G^2 m^{-1}
ABSNJZH	Absolute value of net current helicity, in G^2 m^{-1}
SAVNCPP	Sum of the absolute value of net currents per polarity, in A
MEANPOT	Mean photospheric excess magnetic energy density, in ergs cm^{-2}
TOTPOT	Total photospheric magnetic energy density, in ergs cm^{-2}
MEANSHR	Mean shear angle (measured using Btotal), in deg
SHRGT45	Percentage of pixels with a mean shear angle greater than 45°, in %
SIZE	Projected area of patch on image in microhemisphere
SIZE_ACR	Projected area of active pixels on image in microhemisphere
NACR	Number of active pixels in patch
NPIX	Number of pixels within patch

With the goal of supervised learning, once SHARP data was acquired there was a need to acquire additional information on solar flare occurrences, or data that could serve as the labels for the SHARP-based time series

features. Detecting and recording flares since 1975, NOAA's series of Geostationary Operational Environmental Satellites (GOES) are one of the best sources of information on past solar flare occurrences. Available on NOAA's website [10], the GOES yearly X-ray Sensor (XRS) reports provide a variety of useful information on the previous year's detected solar flares. From occurrence times to class label, the GOES reports, from 2011 to 2016, were used. Additionally, it should be noted that the GOES yearly reports include NOAA AR numbers, so data is easily linked to SHARP summary parameters, which also includes NOAA AR numbers.

Lastly, to increase the validity of the overall data, while verifying the authenticity and accuracy of the GOES yearly flare reports, a third data source was used. A department of NOAA dedicated to monitoring space weather, the Space Weather Prediction Center (SWPC), generates daily solar event reports. Labeled as XRA events, or an X-ray event observed by an orbiting GOES, solar flares are classified soon after occurrence. However, lacking the formal NOAA AR numbers, it is not possible to link SWPC's daily reports directly with SHARP data. Therefore, daily reports were used only to verify GOES yearly reports when generating the overall dataset for this research.

SWPC's daily flare reports were compared with the GOES yearly flare reports for a thoroughly accurate set of solar flare occurrences. Serving as labels for the SHARP-derived multivariate, time series dataset, flare events acquired from both the GOES yearly reports and SWPC's daily reports provided a solid dataset of solar flare activity from 2011 to 2016.

Once all necessary data was acquired from SHARP [8], GOES [10] and SWPC [9], feature data (SHARP data) and label data (GOES and SWPC combined data) generated groups by day and active region. If (for each group) there was at least one flare occurrence the following day, the group was selected and saved. At the same time, the corresponding flare class, or label, was selected and saved. All data was downsampled. Downsampling [19] was necessary for this specific machine learning problem due to imbalanced classes (solar flare classes do not occur at an equal rate over the course of a year or several years). Additionally, in the train, validation, test split a random state was used to ensure a reproducible split across all solar flare forecasting models. Stratify was also used to force an even split of solar flare class occurrences across all sets, further balancing the overall dataset.

4. MODEL ARCHITECTURE

Multiple RNN-based machine learning models were created and tuned to the solar flare prediction problem. Particularly useful for sequence datasets, RNN-based model architectures work as regression classification hybrids (predicting values off multivariate sequences and classifying). Attempting to forecast a solar flare occurrence 24hrs in advance, the two specific approaches used for this research are described in Section 4.2.

4.1 RECURRENT NEURAL NETWORKS

The basic form of a RNN uses its internal state to achieve 'memory', however, unlike more advanced forms of a RNN, the basic form does not have gates. It should be noted that the basic form of a RNN is implementable via Keras [39] as a SimpleRNN. A Long Short-Term Memory (LSTM) is a more complex version of a basic RNN. To extend 'memory' length, a LSTM uses gates. A LSTM neuron has a split state, which can be thought of as the long-term state and short-term state, respectively. The long-term state moves through forgot and addition gates, dropping and adding memories. The long-term state is then sent out to the next neuron, however, it is also copied and filtered by the output gate. This produces the short-term state. Prior to this, the previous short-term state and the current input vector move through four fully connected layers, adjusting the composition of the long-term state as needed. Simplified, an LSTM neuron is designed to capture long-term patterns, making it particularly useful for sequential, or time series data. [20]

A compact version of a LSTM, a Gated Recurrent Unit (GRU) is a fairly new type of RNN. Mirroring a LSTM, a GRU was specifically designed to be a less complex, more efficient version of a LSTM for less intensive tasks. However, it mirrors an LSTM's usage of gates. Recombining the split states seen in a LSTM, a GRU only has one state. This is used to add and remove from the long-term 'memory'. When a new memory must be stored, the location where it will be stored is erased first. Growing in popularity, GRUs often perform equally as well as LSTMs.

With the goal of comparing and evaluating the performance of RNNs on solar flare forecasting, the three main types of RNNs are all utilized. Effective at recognizing patterns in sequential data, current RNNs still have limited 'memory' capabilities, often experiencing learning difficulties with sequences that contain more than 100 timesteps. Additionally, all RNNs tend to suffer from unstable gradients, often resulting in the need for regularization. [20]

4.2 PREDICTION METHODS

Based on previous research [1][2], most solar flare forecasting models simply seek to predict the next-day occurrence of any flare class. In *Predicting Solar Flares with Machine Learning: Investigating Solar Cycle Dependence* [1], an LSTM-based model is applied to an unbalanced dataset of SHARP features and GOES labels. This approach is used to predict all flares greater than or equal to A-class, as well as to predict flares greater than or equal to C-class and greater than or equal to M-class. The target label is whatever maximum flare class an active region is predicted to produce the next day, after the end of a 24-hour time sequence. In slight contrast, *Predicting Solar Flares Using a Long Short-term Memory Network* [2] uses a LSTM-based model to predict flares greater than or equal to M-class, greater than or equal to 5.0 strength or higher M-class and greater than or equal to C-class. Again, the target label is whatever maximum flare class an AR is predicted to produce the next day, after the end of a 24-hour time sequence.

Two approaches, or two questions regarding flare occurrences, were developed and attempted for this research. First, it is important to note, due to the lack of easy detection of A-class flares, they were ignored. Additionally, X-class flares occur rarely, resulting in minimal data, so they were also ignored. Lastly, sequences from the remaining flare classes, including B, C and M-class were downsampled in an effort to achieve a balanced dataset for better classification performance. Negative, or no flare, sequences were also downsampled.

The first approach combined C and M-class sequences with a balanced amount of no flare sequences. This dataset was then fed to one-layer RNNs, including a basic RNN, LSTM and GRU. The dataset was then fed to multi-layer RNNs, including a basic RNN, LSTM and GRU. All six models attempted to forecast whether or not a C or M-class flare would occur the next day, after a 24-hour period ended.

The second approach combined B, C and M-class sequences. An experimental solution to the solar flare prediction problem, all no flare data, or negative samples, were removed from the dataset. Fed to one-layer RNNs and multi-layer RNNs, including basic RNNs, LSTMs and GRUs, the goal of this approach was to evaluate a RNN's ability to learn the subtleties between magnetic field precursors for different class strengths. The six models attempted to forecast whether a B, C, or M-class flare would occur in the next-day, 24-hour period after a 24-hour period ended.

Both machine learning solutions attempted were unique from the most recent research [1][2][3] regarding solar flare forecasting. With a wide variety of approaches possible, there's currently no true solution for predicting solar flares in advance, therefore, any research on this problem contributes to advancing towards a high-performance solution.

5. ANALYSIS AND RESULTS

Five methods are used to evaluate the RNN models designed for this research, including confusion matrix, precision, recall, F1 score and accuracy. Table 3 shows all formulas used to evaluate the RNN models, along with their corresponding formulas. A popular evaluation method for classification models, a confusion matrix provides an easy-to-understand visualization of model performance. Providing the number of true positive (TP), false positive (FP), false negative (FN) and true negative (TN) values, a confusion matrix represents each actual class as a row and each predicted class as a column. Another useful metric is precision, which gives the accuracy of positive predictions. Recall, or the true positive rate (TPR), is also useful, giving the ratio of positive instances correctly detected. Another excellent metric, the F1 score is the harmonic mean (more weight given to lower values) of precision and recall. If precision and recall are both high the F1 score will be high. Lastly, accuracy is generally always calculated for every machine learning model. [20][22]

Table 3: Metric Formulas [20][22].

Metric	**Formula**
Precision	True Positive / (True Positive + False Positive)
Recall	True Positive / (True Positive + False Negative)
F1	2 x ((Precision x Recall) / (Precision + Recall))
Accuracy	(True Positive + True Negative) / Total

The first set of three models (one SimpleRNN, one LSTM and one GRU) sought to predict whether or not a flare would occur in the next 24hrs after a 24-hour time sequence. A 'F' was used to represent a flare occurring and a 'N' was used to represent no flare occurring. Only C and M-class flares were included, along with negative sequences (a sequence where no flare occurred the following day). Using a balanced dataset, which spanned from January 1st, 2011 to December 31, 2016, both the LSTM and GRU models performed equally well, with the SimpleRNN lagging behind. Figures 2, 3 and 4 show the confusion matrices for the binary, single layer SimpleRNN, LSTM and GRU models, respectively.

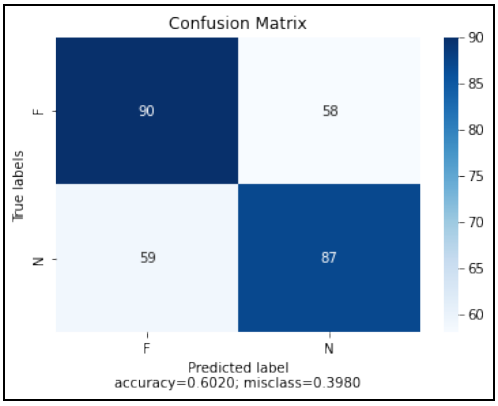

Fig. 2: Binary Classification with SimpleRNN.

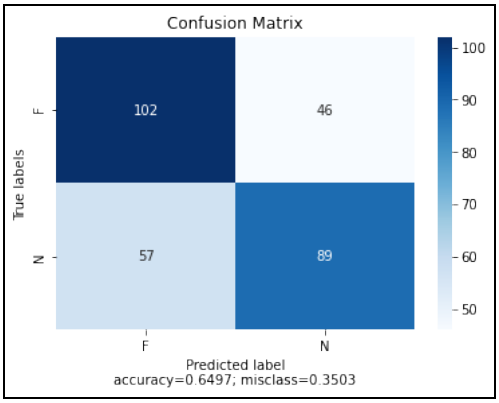

Fig. 3: Binary Classification with LSTM.

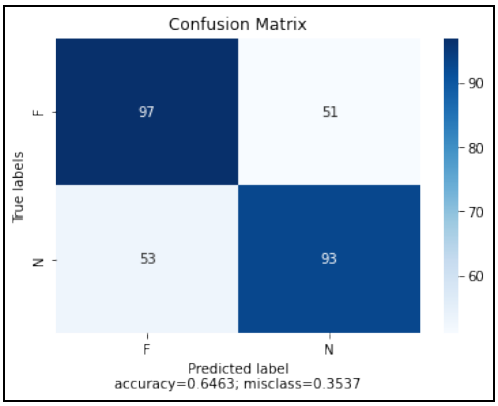

Fig. 4: Binary Classification with GRU.

An extension of the machine learning solution above, the same dataset was fed into a binary classifier with two RNN layers. Performed with all three RNN types (SimpleRNN, LSTM and GRU), the goal was to evaluate performance with multiple RNN layers (added complexity). This second set of models sought to again predict whether or not a flare would occur in the next 24hrs after a 24-hour time sequence. Again, a 'F' was used to represent a flare occurring and a 'N' was used to represent no flare occurring. Only C and M-class flares were included, along with negative sequences (a sequence where no flare occurred the following day). Using a balanced dataset that spanned from January 1st, 2011 to December 31, 2016, increased model complexity appeared to improve the performance of the GRU, with the LSTM and SimpleRNN lagging behind. Figures 5, 6 and 7 show the confusion matrices for the binary, multi-layer SimpleRNN, LSTM and GRU models, respectively.

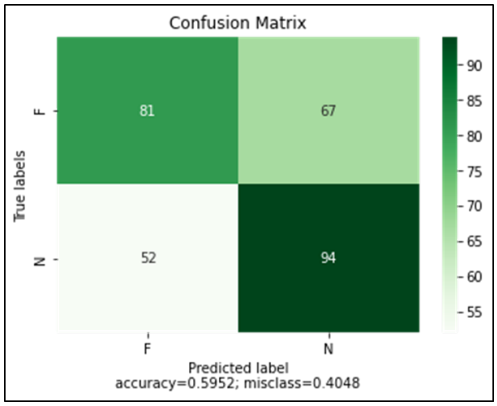

Fig. 5: Binary, Multi-Layer Classification with SimpleRNN.

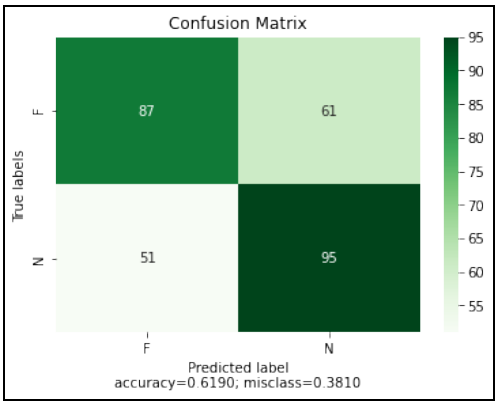

Fig. 6: Binary, Multi-Layer Classification with LSTM.

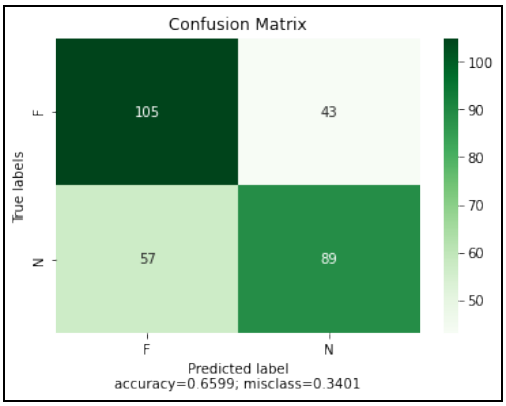

Fig. 7: Binary, Multi-Layer Classification with GRU.

Veering into the second proposed machine learning solution to solar flare forecasting, a third group of models were created. Three versions (SimpleRNN, LSTM, GRU) sought to predict whether or not a flare would occur in the next 24hrs after a 24-hour time sequence. Only B, C and M-class flares were used, with all negative sequences excluded. Using a balanced dataset that spanned from January 1st, 2011 to December 31, 2016, all performed similarly, with the LSTM slightly outperforming the other two RNNs. Figures 8, 9 and 10 show the confusion matrices for the multi-class, single layer SimpleRNN, LSTM and GRU models, respectively.

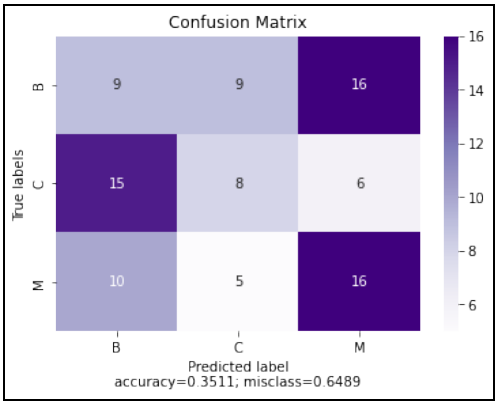

Fig. 8: Multi-Class Classification with SimpleRNN.

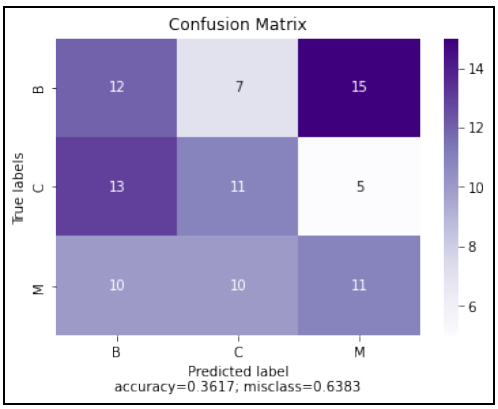

Fig. 9: Multi-Class Classification with LSTM.

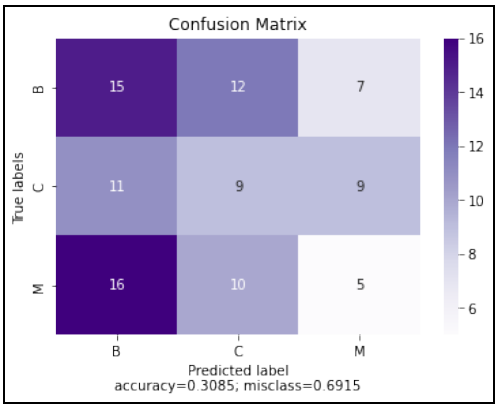

Fig. 10: Multi-Class Classification with GRU.

An extension of the machine learning solution above, the same dataset was fed into a multi-class classifier with two RNN layers. Performed with all three RNN types (SimpleRNN, LSTM and GRU), the goal was to evaluate forecasting performance with additional model complexity. Similar to all the other models, this group of models also worked to predict whether or not a flare would occur in the next 24hrs after a 24-hour time sequence. Only B, C and M-class flares were used, with all negative sequences excluded. Using a balanced dataset that spanned from January 1st, 2011 to December 31, 2016, the GRU and SimpleRNN performed significantly better with the added complexity, while the LSTM performed significantly worse. Figures 11, 12 and 13 show the confusion matrices for the multi-class, multi-layer SimpleRNN, LSTM and GRU models, respectively.

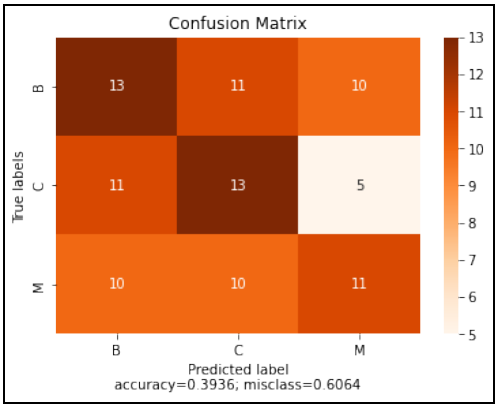

Fig. 11: Multi-Class, Multi-Layer Classification with SimpleRNN.

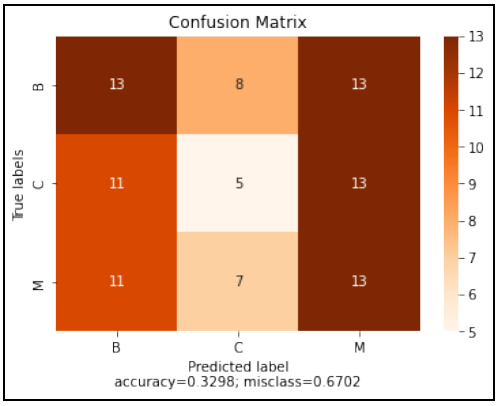

Fig. 12: Multi-Class, Multi-Layer Classification with LSTM.

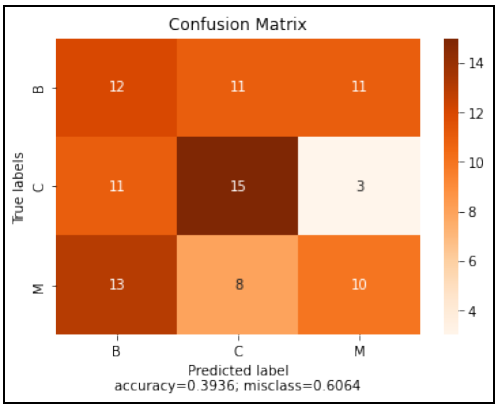

Fig. 13: Multi-Class, Multi-Layer Classification with GRU.

Table 4: Binary Metrics.

BINARY METRICS				
	Precision	**Recall**	**F1**	**Accuracy**
LSTM-2015 [1]	.6700	.6250	N/A	.8140
LSTM-15_18 [1]	.7020	.5300	N/A	.8000
SimpleRNN_1_bin	.6020	.6020	.6020	.6020
LSTM_1_bin	.6504	.6494	.6490	.6497
GRU_1_bin	.6462	.6462	.6462	.6463
Simple_RNN_2_bin_layer	.5964	.5956	.5944	.5952
LSTM_2_bin_layer	.6197	.6193	.6188	.6190
GRU_2_bin_layer	**.6612**	**.6595**	**.6589**	**.6599**

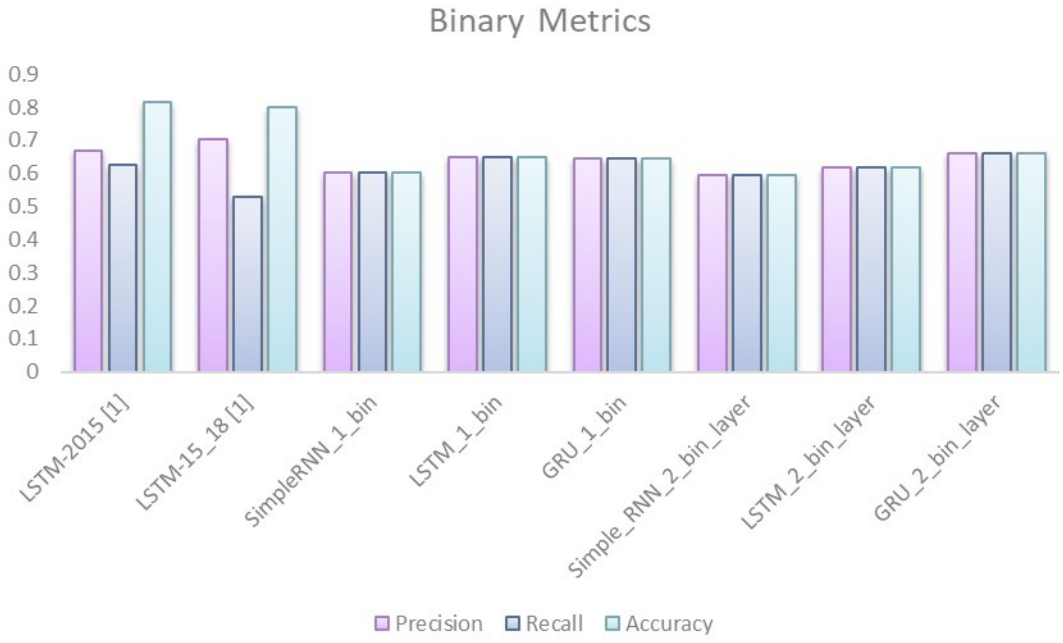

Fig. 14: Binary Metrics.

Table 5: Multi-Class Metrics.

MULTI-CLASS METRICS				
	Precision	**Recall**	**F1**	**Accuracy**
SimpleRNN_3_mc	.3498	.3522	.3474	.3511
LSTM_3_mc	.3635	.3624	.3629	.3617
GRU_3_mc	.2952	.3043	.2957	.3085
SimpleRNN_4_mc_layer	.3959	.3952	.3937	.3936
LSTM_4_mc_layer	.3183	.3247	.3174	.3298
GRU_4_mc_layer	**.3971**	**.3976**	**.3942**	**.3936**

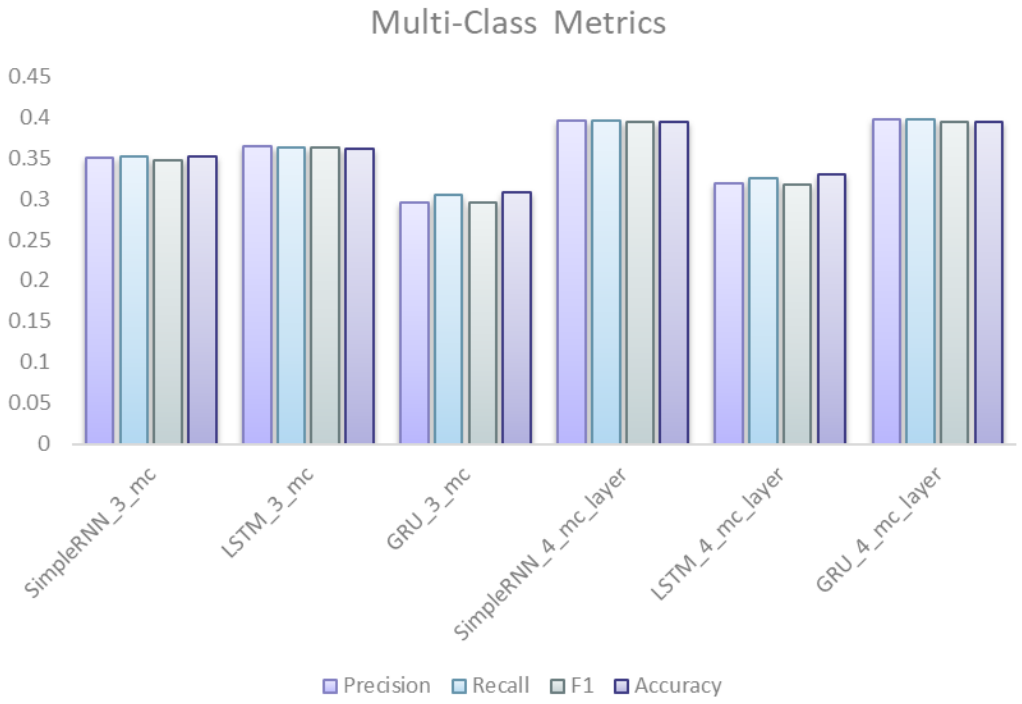

Fig. 15: Multi-Class Metrics.

Metric calculations, for all models, indicated learning to some degree. Table 4 lists the metrics for the binary classification models, in addition to the metrics from [1] (LSTM-2015 and LSTM-15_18). Fig. 14 visualizes metrics for the binary classification models in the form of a bar chart. It should be noted that [1] used an unbalanced dataset to apply a LSTM to 2015's SHARP and GOES data, called LSTM-2015, and to 2015 – 2018's SHARP and GOES data, called LSTM-15_18 [1]. Despite [1] using an unbalanced dataset, models in this research (using a balanced dataset) are comparable, and sometimes exceed the metric scores seen in [1]. The GRU_2_bin_layer (binary, multi-layered GRU) model performed particularly well.

The multi-class classification models did not perform as well as the binary classification models. Table 5 lists the metrics for the multi-class models, while Fig. 15 visualizes the metrics with a bar chart. Due to this being a completely new approach to solar flare prediction, there is no previous research to compare metrics to. Nonetheless, it can still be observed that all multi-class models showed some ability to learn. The GRU_4_mc_layer (multi-class, multi-layered GRU) model performed particularly well.

The newest form of a RNN, a Gated Recurrent Unit (GRU) was designed to be more efficient and more effective than a LSTM. Based on the metric scores for all twelve models, GRU multi-layered models consistently performed the best. Achieving higher metric scores, compared to SimpleRNN and LSTM, it is easy to recognize the value in applying GRU-based RNNs to the solar flare prediction problem.

6. CONCLUSION AND FUTURE WORK

Exploring the viability of RNN application to sequence data, which attempts to grasp the behavior of photospheric active regions to forecast solar flares, has yielded some optimistic results. Using a complex, multivariate dataset, the RNN models were all able to learn to some degree. Additionally, binary metric scores were either comparable to or exceeded previous research [1]. GRU-based RNNs were shown to be particularly promising when applied to solar flare forecasting.

Based on previous research [1][2], solar flare forecasting is a two-prong problem. One part of the problem focuses on compiling the right dataset [3] and the other part of the problem focuses on building the right machine learning model. Both parts of the problem offer opportunities for additional research.

One potential avenue for future research focuses on the limited availability of photospheric magnetic field data combined with the limited occurrence of significant (strong) solar flares, such as X-class flares, over the course of an average year. The generative functionality of Generative Adversarial Networks (GANs) [40] may produce more effective solar flare forecasting models, particularly when combined with the 'memory' of RNNs. Proven effective at a variety of generative tasks, GANs have gained in popularity in recent years due to their impressive performance.

Refocusing on the application of RNNs to solar flare forecasting, time dependence stands out as an area that should be further analyzed. Reducing the lofty goal of prediction models for next-day forecasting to a smaller time window, such as 6-hour forecasting or 8-hour forecasting, etc. may yield better results.

Another area that should be explored further is general composition of the dataset. Whether suffering from the curse of dimensionality (too many features) or simply a lack of data, the various parameters of the SHARP data product should be more thoroughly analyzed. Selecting or calculating additional features or more beneficial features, as researchers in *Predicting Solar Flares Using a Long Short-term Memory Network* [2] did for flare history features, beyond the commonly used SHARP data product parameters may also be a worthwhile endeavor.

The ever-changing center of our solar system, research concerning the Sun is ongoing. Space weather is one specific area that continues to gain attention. Monitored nearly 24 hours a day, the Sun's composition and behavior still holds many mysteries. Working to mitigate the harmful effects of space weather, various forecasting methods are actively being tested, particularly to predict solar flare occurrences. With the imposing goal of one-day, or multi-day, advance forecasting, solar flares present a difficult problem that may someday be solved with a machine learning or deep learning solution.

7. REFERENCES

[1] X. Wang et al., "Predicting Solar Flares with Machine Learning: Investigating Solar Cycle Dependence," Astrophys. J., vol. 895, no. 1, p. 3, May 2020, doi: 10.3847/1538-4357/ab89ac.

[2] H. Liu, C. Liu, J. T. L. Wang, and H. Wang, "Predicting Solar Flares Using a Long Short-term Memory Network," Astrophys. J., vol. 877, no. 2, p. 121, Jun. 2019, doi: 10.3847/1538-4357/ab1b3c.

[3] R. A. Angryk et al., "Multivariate Time Series Dataset for Space Weather Data Analytics," Sci. Data, vol. 7, no. 1, pp. 1–13, 2020, doi: 10.1038/s41597-020-0548-x.

[4] M. Bartels, "The sun has begun a new solar weather cycle. It should be pretty quiet, scientists say.," Space.com, 15-Sep-2020. [Online]. Available: https://www.space.com/sun-weather-solar-cycle-25-began-quiet-predictions.html. [Accessed: 01-Dec-2020].

[5] M. Bartels, "Stunning new sun images show our star's popcorn-like magnetic field structure," Space.com, 03-Sep-2020. [Online]. Available: https://www.space.com/amazing-sun-images-show-magnetic-field-gregor-telescope.html. [Accessed: 01-Dec-2020].

[6] C. Gohd, “The sun fires off its biggest solar flare in more than 3 years,” Space.com, 30-Nov-2020. [Online]. Available: https://www.space.com/sun-fires-off-solar-flare-m4-worst-three-years. [Accessed: 01-Dec-2020].

[7] P. Sutter, “Will our solar system survive the death of our sun?,” Space.com, 28-Oct-2020. [Online]. Available: https://www.space.com/solar-system-fate-when-sun-dies. [Accessed: 01-Dec-2020].

[8] Spaceweather Helioseismic and Magnetic Imager (HMI) Active Region Patch (SHARP), Joint Science Operations Center (JSOC), hmi.sharp_720s from 2011 to 2016. [Online]. Available: jsoc.stanford.edu/doc/data/hmi/sharp/sharp.htm

[9] Solar and Geophysical Daily Event Reports, National Oceanic and Atmospheric Administration (NOAA) Space Weather Prediction Center (SWPC), 2011 to 2016. [Online]. Available: ftp://ftp.swpc.noaa.gov/pub/warehouse/

[10] Geostationary Operational Environmental Satellite (GOES) X-ray Sensor (XRS) Yearly Reports, National Oceanic and Atmospheric Administration (NOAA) National Geophysical Data Center (NGDC), 2011 to 2016. [Online]. Available: https://www.ngdc.noaa.gov/stp/space-weather/solar-data/solar-features/solar-flares/x-rays/goes/xrs/

[11] NOAA National Environmental Satellite, Data, and Information Service (NESDIS). [Online]. Available: https://www.nesdis.noaa.gov/. [Accessed: 01-Dec-2020].

[12] “GOES Primary X-ray flux Interactive Timeline,” Community Coordinated Modeling Center (CCMC) and NASA. [Online]. Available: https://ccmc.gsfc.nasa.gov/support/ILWS/MATERIALS/iswa_wiki/index.php/GOES_Primary_X-ray_flux_Interactive_Timeline.html. [Accessed: 01-Dec-2020].

[13] G. Shen, Q. Tan, H. Zhang, P. Zeng, and J. Xu, “Deep learning with gated recurrent unit networks for financial sequence predictions,” Procedia Comput. Sci., vol. 131, pp. 895–903, 2018, doi: 10.1016/j.procs.2018.04.298.

[14] R. Fu, Z. Zhang, and L. Li, “Using LSTM and GRU neural network methods for traffic flow prediction,” Proc. - 2016 31st Youth Acad. Annu. Conf. Chinese Assoc. Autom. YAC 2016, pp. 324–328, 2017, doi: 10.1109/YAC.2016.7804912.

[15] M. Mudassir, S. Bennbaia, D. Unal, and M. Hammoudeh, “Time-series forecasting of Bitcoin prices using high-dimensional features: a machine learning approach,” Neural Comput. Appl., vol. 6, 2020, doi: 10.1007/s00521-020-05129-6.

[16] B. W. Carroll and D. A. Ostlie, An introduction to modern astrophysics. San Francisco: Pearson Addison Wesley, 2014.

[17] Moché Dinah L., Astronomy: a self-teaching guide, 7th ed. Hoboken, NJ: John Wiley and Sons, Inc.

[18] D. J. Mullan, Physics of the sun: a first course. Boca Raton: CRC, 2010.

[19] P. C. Bruce, A. G. Bruce, and P. Gedeck, Practical statistics for data scientists: 50+ essential concepts using R and Python. Beijing: O'Reilly, 2020.

[20] A. Géron, Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: concepts, tools, and techniques to build intelligent systems. Sebastopol, CA: O'Reilly Media, Inc., 2019.

[21] J. H. Thomas and N. O. Weiss, Sunspots and starspots. Cambridge: Cambridge University Press, 2012.

[22] I. Gnanadass, "Prediction of Gestational Diabetes by Machine Learning Algorithms," in IEEE Potentials, vol. 39, no. 6, pp. 32-37, Nov.-Dec. 2020, doi: 10.1109/MPOT.2020.3015190.

[23] J. Jeong, “The Most Intuitive and Easiest Guide for Recurrent Neural Network,” Medium, 06-Feb-2019. [Online]. Available: https://towardsdatascience.com/the-most-intuitive-and-easiest-guide-for-recurrent-neural-network-873c29da73c7. [Accessed: 01-Dec-2020].

[24] “What Is Space Weather?,” NASA, 17-Nov-2020. [Online]. Available: https://spaceplace.nasa.gov/spaceweather/en/. [Accessed: 01-Dec-2020].

[25] Y. Cendes, “Untangling the magnetic universe,” Astronomy.com, 30-Aug-2019. [Online]. Available: https://astronomy.com/magazine/2019/08/untangling-the-magnetic-universe. [Accessed: 01-Dec-2020].

[26] N. Wolchover, “The Hidden Magnetic Universe Begins to Come Into View,” Quanta Magazine. [Online]. Available: https://www.quantamagazine.org/the-hidden-magnetic-universe-begins-to-come-into-view-20200702/. [Accessed: 01-Dec-2020].

[27] “Magnetograms - NSO - National Solar Observatory,” 09-Nov-2018. [Online]. Available: https://nso.edu/data/nisp-data/magnetograms/. [Accessed: 01-Dec-2020].

[28] G. Ruediger, “Solar dynamo.” [Online]. Available: http://www.scholarpedia.org/article/Solar_dynamo. [Accessed: 01-Dec-2020].

[29] R. Garner, “Solar Storm and Space Weather - Frequently Asked Questions,” 19-Mar-2015. [Online]. Available: https://www.nasa.gov/mission_pages/sunearth/spaceweather/index.html. [Accessed: 01-Dec-2020].

[30] H. Zell, “The Difference Between Flares and CMEs,” 10-Feb-2015. [Online]. Available: https://www.nasa.gov/content/goddard/the-difference-between-flares-and-cmes. [Accessed: 01-Dec-2020].

[31] “Sun,” 10-Nov-2020. [Online]. Available: https://solarsystem.nasa.gov/solar-system/sun/overview/. [Accessed: 01-Dec-2020].

[32] Helioseismic and Magnetic Imager for SDO. [Online]. Available: http://hmi.stanford.edu/. [Accessed: 01-Dec-2020].

[33] drms. [Online]. Available: https://pypi.org/project/drms/. [Accessed: 01-Dec-2020].

[34] Pandas. [Online]. Available: https://pandas.pydata.org/. [Accessed: 01-Dec-2020].

[35] Scikit. [Online]. Available: https://scikit-learn.org/stable/. [Accessed: 01-Dec-2020].

[36] NumPy. [Online]. Available: https://numpy.org/. [Accessed: 01-Dec-2020].

[37] Project Jupyter. [Online]. Available: https://jupyter.org/. [Accessed: 01-Dec-2020].

[38] TensorFlow. [Online]. Available: https://www.tensorflow.org/. [Accessed: 01-Dec-2020].

[39] Keras. [Online]. Available: https://keras.io/. [Accessed: 01-Dec-2020].

[40] “A Beginner's Guide to Generative Adversarial Networks (GANs),” Pathmind. [Online]. Available: https://wiki.pathmind.com/generative-adversarial-network-gan. [Accessed: 01-Dec-2020].

[41] IBM Cloud Education, “What are Recurrent Neural Networks?,” IBM. [Online]. Available: https://www.ibm.com/cloud/learn/recurrent-neural-networks. [Accessed: 01-Dec-2020].