

4-2022

The Pandemic From Above: Estimating COVID-19 Cases Using Deep Learning and Satellite Imagery

John Hennin
William & Mary

Follow this and additional works at: <https://scholarworks.wm.edu/honorstheses>



Part of the [Data Science Commons](#), and the [Disease Modeling Commons](#)

Recommended Citation

Hennin, John, "The Pandemic From Above: Estimating COVID-19 Cases Using Deep Learning and Satellite Imagery" (2022). *Undergraduate Honors Theses*. William & Mary. Paper 1813.
<https://scholarworks.wm.edu/honorstheses/1813>

This Honors Thesis -- Open Access is brought to you for free and open access by the Theses, Dissertations, & Master Projects at W&M ScholarWorks. It has been accepted for inclusion in Undergraduate Honors Theses by an authorized administrator of W&M ScholarWorks. For more information, please contact scholarworks@wm.edu.



WILLIAM & MARY

CHARTERED 1693

WILLIAM & MARY

HONORS THESIS

The Pandemic From Above: Estimating COVID-19 Cases Using Deep Learning and Satellite Imagery

Author:

John HENNIN

Advisor:

Dan RUNFOLA

*A thesis submitted in fulfillment of the requirements for
Interdisciplinary Honors in the degree of Bachelors of Science in the
Data Science Program*

Accepted for Honors

Chair: Dr. Dan Runfola

Dr. Carrie Dolan

Dr. Anthony Stefanidis

Williamsburg, Virginia

April 27th, 2022

WILLIAM & MARY

Abstract

Data Science Program

Bachelors of Science

The Pandemic From Above: Estimating COVID-19 Cases Using Deep Learning and Satellite Imagery

by John HENNIN

Monitoring the spread of an outbreak of disease (such as COVID-19) is an important component of any coordinated pandemic response. Across the globe, our ability to conduct such monitoring - especially at early stages of the COVID-19 pandemic - was highly limited due to a lack of public reporting mechanisms. Today, the process of case data collection remains expensive and, in some regions, is subject to political considerations. Researchers have turned to some techniques leveraging Google Trends and Twitter data to overcome limitations in public data sources. Here, we provide another approach which leverages satellite information to provide estimates of case counts. Visible features in imagery - such as vehicles in parking lots, or temporary structures at hospitals - should convey some information about underlying disease spreads. We explore the use of a ResNet50V2-LSTM hybrid model that is trained on satellite information, seeking to predict case counts across counties in the USA using imagery alone. Using solely this imagery, this model produced an overall rate of error of 4.55%. We discuss advantages and drawbacks of this approach, and how some future directions may aid us in overcoming contemporary limits.

Contents

Abstract	i
1 Thesis	1
1 Introduction	1
2 Data and Methods	4
2.1 Covid-19 Case Data	4
2.2 Imagery Data	5
2.3 ResNet50V2 Transfer Learning Approach	5
2.4 ResNet50V2 to LSTM Approach	7
3 Results	9
3.1 Results: ResNet50V2 Transfer Learning Approach	9
3.2 Results: ResNet50V2 to LSTM Approach	9
4 Discussion	12
5 Conclusion	15
6 Acknowledgements	16
References	17

List of Figures

1.1	Images collected from the Sentinel-2 MSI Mission, 10m resolution.	2
1.2	Images of same location in Adams County, Iowa taken 90 days apart starting from 09-30-2020.	3
1.3	Quarterly case counts plotted against daily case counts (daily data has been scaled for visualization purposes).	4
1.4	ResNet50V2 input and first convolutional block architecture. Made using NN architecture visualization package VisualKeras (Gavrikov, 2020). The "Add" layer adds the outputs from the last MaxPooling2D and Conv2D layers.	6
1.5	ResNet50V2-LSTM Architecture. As before, the features from the ResNet50V2 have been reshaped for visualization purposes. . . .	7
1.6	Above are three random images from the dataset. Below these are two dimensional representations of their corresponding 2048 features.	9
1.7	Mappings of average case counts per county (upper left), average ResNet50V2-LSTM prediction per county (upper right), and the mean absolute error between them (lower left).	10
1.8	Mappings of average case counts per county (top) and average ResNet50V2-LSTM prediction per county (bottom) broken down by quantiles to better display variation.	11
1.9	Line plot of actual case trend line, ResNet50V2 trend line, and ResNet50V2-LSTM trend line.	12

List of Tables

1.1	Model test results compared to statistics of actual data.	9
-----	---	---

List of Abbreviations

CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
LSTM	Long Short Term Memory Network
RGB	Red Green Blue (image bands)

Chapter 1

Thesis

1 Introduction

The speed and accuracy at which we currently collect Covid-19 case surveillance data directly impacts how effectively we can target and prevent the spread of the disease (The Centers for Disease Control and Prevention, 2021). In more developed countries, this has proven to be a costly and overly complex process, as multiple entities across various levels of jurisdiction are hastily collaborating to aggregate, anonymize, and release case data to both the public and policy-makers (Komenda, 2020; The Centers for Disease Control and Prevention, 2021). In less developed countries, this process is either nonexistent or overwhelmingly inaccurate (Milan and Treré, 2020). This has led to demand for supplementary alternatives that can deliver Covid-19 case estimates in a cost-effective and timely manner.

In order to fulfill this need, some researchers have been utilizing proxy data sources such as Google Trends and Twitter for their modeling (Sulyok, Ferenci, and Walker, 2021; Gharavi, Nazemi, and Dadgostari, 2020). These methods have demonstrated success when applied to select more developed countries as case studies (Skiera et al., 2020; Nindrea, 2020). However, there are considerable limitations to these approaches. For example, there are a number of nations that lack accurate Google Trend and Twitter data, driven by government restrictions, severe political bias, or very limited accessibility to Internet services (Ensafi and Crandall, 2015; Timoneda and Vera, 2021; Kemp, 2022).

To effectively estimate Covid-19 case counts on a comprehensive global level, researchers need to source data that has universal coverage, regardless of the

myriad of unique federal restrictions or developmental challenges that each nation may face. With this challenge in mind, this study proposes utilizing publicly available high-resolution satellite imagery and a computer vision-based deep learning model ensemble to estimate Covid-19 cases. With this technique, we are able to employ inexpensive, timely, and accessible data collection that covers the entire globe to satisfy the urgent nature of a pandemic.




FIGURE 1.1: Images collected from the Sentinel-2 MSI Mission, 10m resolution.

We chose to pursue computer vision-based modeling, specifically transfer learning convolutional neural networks (CNNs), as recent literature has demonstrated increasing confidence in the success of these models when trained on satellite imagery. These models have been applied to a wide variety of contexts, such as estimating road-quality, school tests scores, and daily economic consumption (Brewer et al., 2021; Runfola, 2021; Jean et al., 2016). It does so through a process unique to the CNN called convolution, where the model convolves a filter over an image and uses the products as potential abstractions of informative features, such as distinct lane markings in the road-quality example. The weights associated with these abstractions are adjusted as the model

trains. On top of the CNN, we investigate the use of a recurrent neural network (RNN), which is a deep learning model that incorporates a temporal dimension to better inform estimations. This RNN method is supported by recent studies that explored applications such as detecting floods, classifying land cover, and tracking urban development (Rahnemoonfar et al., 2018; Ienco et al., 2017; Papadomanolaki et al., 2019). Within this study, the model will estimate the Covid-19 case counts for a single period of time and a specific region, then use this estimation in its prediction for the same region during the next period of time.




FIGURE 1.2: Images of same location in Adams County, Iowa taken 90 days apart starting from 09-30-2020.

It is worth noting that the typical scope of the RNN studies in the context of satellite imagery is largely limited to recognizing changes in physical geographic structures, whether that be bodies of water, forest cover, or buildings. By applying the RNN method to the novel context of disease surveillance, we hope to demonstrate the substantial potential of this modeling technique, as well as highlight a relationship between the geographical structure within a high-resolution satellite image and the underlying characteristics of the represented population.

In this piece we validate the aforementioned approach using the United States as a case study. We compare the merits of a basic CNN, a transfer learning CNN model, and then the final transfer learning CNN-RNN stacked ensemble for this

context. Finally, we discuss possible applications of our final model both within the current Covid-19 pandemic, as well as any future pandemics that may arise.

2 Data and Methods

2.1 Covid-19 Case Data

For the Covid-19 case data, we used a US county-level dataset provided by The New York Times (The New York Times, 2021) and supplemented any missing counties with the Johns Hopkins University Covid-19 Global data repository (Dong, 2020). Both data sources collaborated with departments of health at either the state or county level, which led to certain counties being excluded between them and made it necessary to cross-reference the data. At the time the data was collected for modeling, it contained 264,075 daily observations of 3006 counties from 01/27/2020 to 10/01/2021 and included measures of confirmed cases and deaths attributed to Covid-19.




FIGURE 1.3: Quarterly case counts plotted against daily case counts (daily data has been scaled for visualization purposes).

To help the model generalize, it is necessary to normalize the cases by bucketing them into larger time periods. This process smooths out dramatic spikes in county-specific case counts. We aggregated the data by summing all confirmed cases for each county within 90 day periods. As seen in figure 2.1, the daily trend line in Covid-19 cases was apparently preserved by the quarterly aggregation.

2.2 Imagery Data

We are using satellite imagery data collected by the Sentinel 2 project (Monitoring, 2020, 2021), under which two satellites, SENTINEL-2A and SENTINEL-2B, image the globe every 5 days. For each observation in the case data, three (RGB) bands are downloaded with the Google Earth Engine API (Gorelick et al., 2017). The region covered is $5km^2$, with the most populated city within each US county at its center. The most populated city of the county serves as a proxy for the remaining area of the region. Additionally, by bringing the city to the forefront, we are avoiding the case in which our imagery is mostly comprised of farmland or nature and now the model is more likely to capture key features, such as hospitals, temporary structures that may be associated with hospitals, and general activity metrics such as large volumes of cars on roads. Each band has a spatial resolution of 10m, which means each pixel in an image corresponds to a 10 meter by 10 meter area on the ground. The granularity advantage of Sentinel 2 imagery is especially apparent when compared to similar contemporary missions, such as the Landsat 8 (launched 2013) and Landsat 9 (launched 2021) satellites each with a maximum resolution capacity of 30m and 15m, respectively (Science, 2013; Science, 2021). This unique high resolution allows for sophisticated sub-pixel analysis, which is especially important in the analysis of population-based behaviors such as disease spread.

For example, recent literature has demonstrated significant changes in the pixel values of parking lots depending on the amount of cars within the lot (Liu, 2021; Vincent and Besson, 2019). Within our context, if a hospital parking lot is captured in an image, the change in pixel values may indicate how busy the hospital is, which in turn could be indicative of the state of the disease spread. This potential relationship between physical changes in a region and the number of Covid-19 cases within the same region is what we aim to establish by using the 10m satellite imagery offered by Sentinel 2.

2.3 ResNet50V2 Transfer Learning Approach

Aggregating our data from daily to quarterly reduced the number of observations from 264,075 to 23,368. In order to effectively train and validate the model, we allocated 75% of the observations to a training set and the remaining 25% to




FIGURE 1.4: ResNet50V2 input and first convolutional block architecture. Made using NN architecture visualization package VisualKeras (Gavrikov, 2020). The "Add" layer adds the outputs from the last MaxPooling2D and Conv2D layers.

a validation set. Therefore, our model would have to train on 17,526 observations, which is a relatively small dataset for CNNs and could cause inadequate training. Research has shown that a strategy called transfer learning can help mitigate this issue (Brewer et al., 2021; Barman et al., 2019). In computer vision, transfer learning is the use of a CNN that has been pre-trained on another dataset. The advantage of this is the ability to parameterize the weights on a much larger dataset than your own in order to utilize the pre-trained model's elementary ability to identify significant features in an image such as shapes or edges. The final layer of the CNN is then adjusted for the user's specific modeling needs, whether it be classification for a number of classes or a regression estimation.

For the purposes of this study, we started with a ResNet50V2 model (He et al., 2016) pretrained on the ImageNet dataset which contains 1.28 million training images of 1000 different classes (Deng et al., 2009). We resized each image to 256 pixels by 256 pixels for three RGB color bands and rescaled the pixel values by inverse the maximum pixel value of 255. Since the weights were trained using a dataset for classification, we had to explicitly set the model parameters for regression. At this stage, the model was ready to estimate the number of Covid-19 cases for a region given its satellite image.

2.4 ResNet50V2 to LSTM Approach

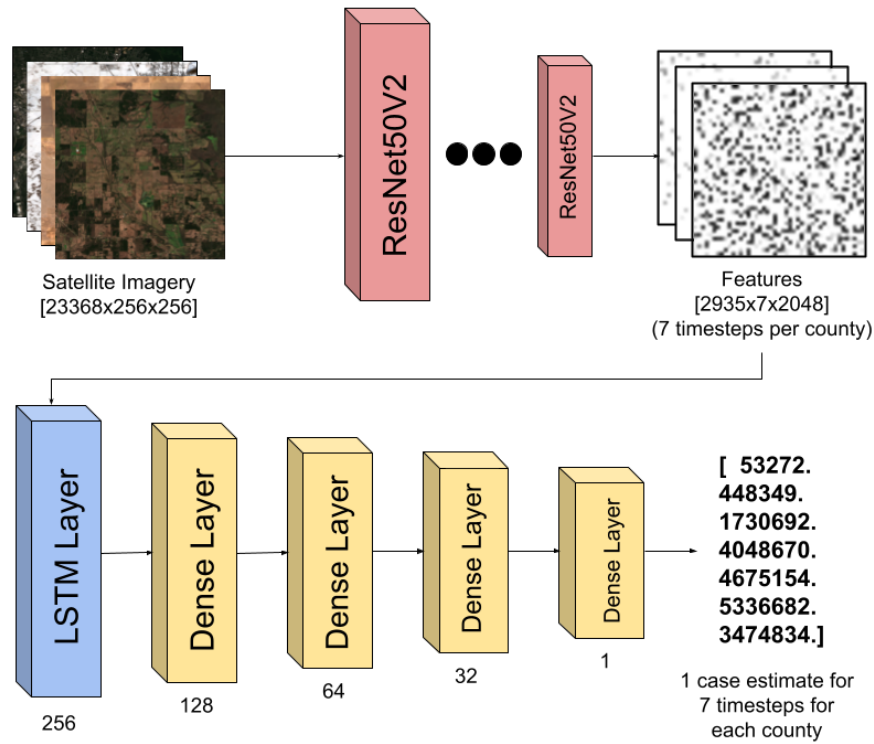


FIGURE 1.5: ResNet50V2-LSTM Architecture. As before, the features from the ResNet50V2 have been reshaped for visualization purposes.

Disease case counts at an instance in time is often not an independent figure. Typically, there is some correlation between the number of cases at one point in time and the number of cases at another point in time for the same area, whether it be due to a relatively fixed population size, the simultaneous spread of the disease to new people and the recovery of the previously diseased, etc. Therefore, it is helpful to consider previous confirmed case figures in the process of estimating future cases. Sequential relationships like the one described are often modeled with a special type of neural network called a Recurrent Neural Network (RNN). RNNs can serve a number of functions, such as estimating the emotional sentiment of a Tweet, generating sheet music based on previous works, and detecting outliers in data (Yuan and Zhou, 2015; Dua et al., 2020; Williams et al., 2002).

For our intent, we chose to use an Equal Unit Size Many-to-Many LSTM RNN (from now on referred to as LSTM), which is trained on multiple inputs and returns the same number of outputs, where each output recursively informs the estimation for the following input. Since our imagery data consists of between 5-7 images of the same location across periods of 90 days, we can input a set of these images and an LSTM would return a set of equal size containing estimated case counts, one for each image. For example, if the first image corresponds to a count of 200 confirmed cases but the model estimated 250, when the model is trained on the second image in the set it now considers the second image's pixel data, the previous image's correct case count of 200, and the magnitude of error.

Unfortunately, this isn't as simple as taking the image data and sending it through an LSTM. First, passing every pixel into a LSTM would result in tremendous levels of noise, as individual pixels are generally not representative of features of interest. One pixel does not generally represent one car, or one structure; rather, it is groups of pixels we are interested in. Second, passing a raw image forward through an LSTM would require flattening it into a 1D array, which would degrade the spatial relationship between the pixels that help characterize an image. Thus, researchers have found it useful to combine a CNN and LSTM for sequential images. Notable examples include another Twitter sentiment analysis model that first predicts the text in a screenshot of a Tweet and automatic image caption generators (Basiri et al., 2021; Adnan et al., 2021).

We chose to combine the aforementioned ResNet50V2 model with a LSTM sequential model. By removing the last layer of the ResNet50V2 model, we can extract the image features from the penultimate step in the forward pass. For each image there are 2048 feature values that contain an abstract representation of the original image. These features may represent intensities (i.e., more bright objects in an image), or shapes (i.e., more sharp edges); the specific nature of the features represented is contingent on the model training. By using these features within the LSTM, we attempt to capture the key features that are most important for predicting COVID-19, while still retaining the capabilities provided by a LSTM structure.




FIGURE 1.6: Above are three random images from the dataset. Below these are two dimensional representations of their corresponding 2048 features.

3 Results

Results				
	\bar{x}	σ	MAE	MAE as % of \bar{x}
Actual	587088.03	2453954.29		
ResNet50V2	385382	876439.25	132092.86	25.53
ResNet50V2 & LSTM	560379.44	346318.7	26708.63	4.55

TABLE 1.1: Model test results compared to statistics of actual data.

3.1 Results: ResNet50V2 Transfer Learning Approach

The standalone ResNet50V2 model achieved an overall mean absolute error of 132,092.86 ($\sigma = 876,439.25$). As a percentage of the mean case count, the model scored 25.53%.

3.2 Results: ResNet50V2 to LSTM Approach

The ResNet50V2 to LSTM model resulted in an overall mean absolute error (MAE) of 26,708.63 ($\sigma = 346,318.7$). The MAE as a percentage of the mean case

count (587,088.03, $\sigma = 2,453,954.29$) is only 4.55%. Notably, the standard deviation of the ResNet50V2 predictions was slightly more comparable to the actual case data standard deviation, but the ResNet50V2 to LSTM model achieved a significantly better MAE and MAE as a percentage of the mean case count. Therefore, the addition of the temporal dimension seems to greatly improve the accuracy of our model's predictions.

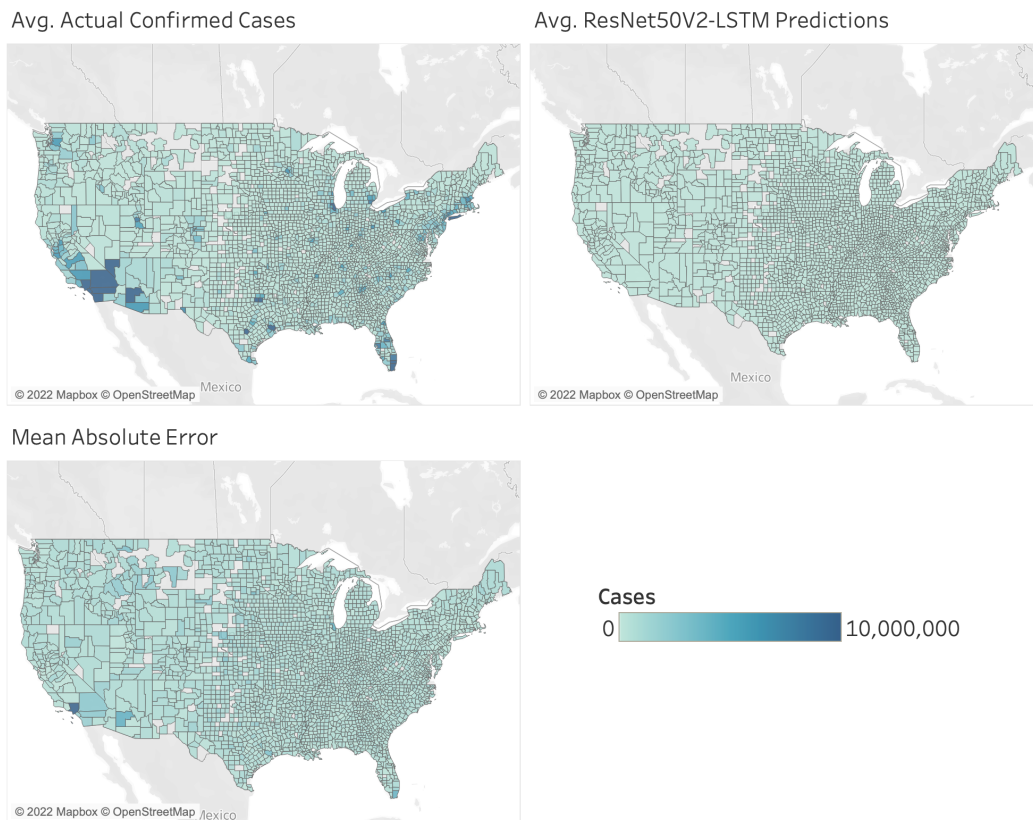


FIGURE 1.7: Mappings of average case counts per county (upper left), average ResNet50V2-LSTM prediction per county (upper right), and the mean absolute error between them (lower left).

From 1.7 and 1.8, the average actual confirmed cases can vary significantly from county to county. The Southwest region, specifically Los Angeles county and the counties adjacent to it, had the greatest average 90-day aggregated active case counts, at about 10 million. Counties in the state of Florida, along the Northeast coast, and in the Midwest region all had similar counts to this 10 million mark. As you can see from the upper right map, the predictions made by

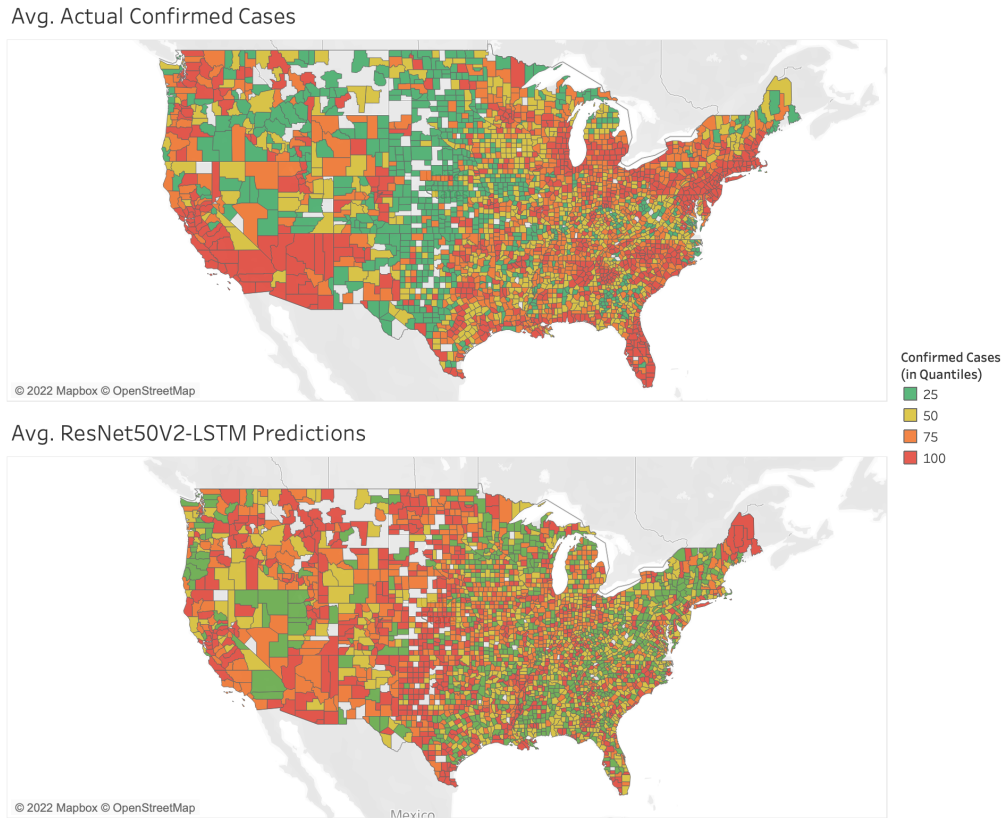


FIGURE 1.8: Mappings of average case counts per county (top) and average ResNet50V2-LSTM prediction per county (bottom) broken down by quantiles to better display variation.

our model did not effectively capture these extreme cases. Now looking at the Mean Absolute Error map, we see that Los Angeles county, shaded the darkest in the bottom left of the map, is the most notable example of this discrepancy, as the MAE for the county is about the same as its average actual confirmed cases figure.

The remaining majority of the counties have 90-day aggregated active case counts of between 0 to 1,000,000. The final model's predictions of these counties are much more accurate than the predictions of the outliers, as we can see from a much smaller MAE in the lower left map. However, the model seems to generalize in its predictions for these counties as they are on a significantly different scale than the outliers.

In 1.9 we see the value of the LSTM addition to the ResNet50V2 base model.




FIGURE 1.9: Line plot of actual case trend line, ResNet50V2 trend line, and ResNet50V2-LSTM trend line.

The ResNet50V2 alone predicts about the same total case counts across our 7 quarterly timestamps. This does not mean that the ResNet50V2 model is predicting the same value for each observation. In fact, the standard deviation for the ResNet50V2 model is more than two times the standard deviation for the final hybrid model (see 1.1). Instead, the plot demonstrates that the ResNet50V2 predictions are not influenced by time. The ResNet50V2-LSTM model is able to incorporate the temporal dimension of the disease spread data. Consequently, the trend of the predictions for this model is much more representative of the actual case trend line.

4 Discussion

As you can see from 1.1, the final model's standard deviation is very small compared to the standard deviation of the case data. This means the model tends to make estimations more close to the mean. In potential future implementations, to avoid the bias due to an imbalanced scale in data, multiple models that cover different stages in pandemic may help provide more variance in estimates. This would enable it to account for unique situations like spikes in cases or significant drops in cases following an effective vaccination. This tactic would also allow for more granular training data, such as imagery every 5 days as opposed to our model's 90 days rate, since the data no longer spans the entirety of a pandemic.

Additionally, it should be expanded to more countries than the United States to further improve the variety of case counts and imagery. The biggest bottleneck in this study was the time it took to download the imagery, but for the purposes of future implementations that is a small issue. If the model is built in anticipation of future pandemics, the temporal expense of downloading imagery can be afforded so long as the entire process of data collecting and then model training finishes before the new pandemic begins.

In practice, the early stage model would be trained on early stage Covid-19 data and would manifest as a similar form to the final ResNet50V2-LSTM model of this study. In accordance, the later stage models would be trained on data from their respective time periods. However, for these models, as the anticipated pandemic unfolds and more information besides the immediately available satellite imagery becomes available, there is room for improvement in accuracy.

As a pandemic progresses, important public health variables gradually become widely available, like vaccination rates when a vaccine is created, certain regional mandates such as wearing masks, or the R-naught (R_0) value which measures how many people are expected to be infected by one contagious person. With this in mind, the predictions of the aforementioned later stage ResNet50V2-LSTM models could be sent through an additional basic regression model that could factor in these extra variables. Not only is cost effective and timely satellite imagery data still utilized with this method, but epidemiological variables could easily be incorporated to suit the data availability of the target region.

Often the difficulty in justifying computer vision modeling techniques to public health experts is the very novelty of the technique. It is difficult to identify the key features in an image that help inform the model's predictions. Consequently, there is an inherent clash between the lack of transparency in computer vision modeling and public health experts that are trying to identify key data that may be utilized in containing its spread. The amalgamation of our ResNet50V2-LSTM model and a regression model with key public health variables could function as a justifiable compromise and bridge the gap between computer vision and typical public health pandemic response techniques. Additional exploration into model explainability techniques which can help to identify key drivers in imagery could also be of high value.

A benefit of a satellite-based approach is that it can be employed before some of these public health figures are publicly available or in practice. If it is so early in a pandemic that there are no measures in place—such as social distancing, a vaccine, mask mandates—then a model that is trained on those variables will be ineffective. The model of this study, trained on satellite imagery data during the Covid-19 pandemic, circumvents that issue and could be effective in the early days of future pandemics, where public information used for epidemiological modeling is not available to practitioners.

It is possible that the model is mostly estimating the size and presence of population centers, rather than inherently Covid-19 cases, or this is at least considered in the estimation process. In further exploration, diversifying the training data by expanding the geographic scope to more rural areas could help ease the model's potential dependency on these population centers. Additionally, an expansion to both more distinctly rural and urban areas would help the model account for high and low extremes such as Los Angeles and Loving County, respectively. As noted in 3.2, the model's current form has difficulty adequately capturing the extremes. This is because a significant majority of the confirmed case training data is close to the mean. Diversifying the training data with a geographic expansion and, consequently, increasing the standard deviation, will help the model capture those extremes.

With a modeling context as omnipresent as the Covid-19 pandemic, it seems necessary to address the potential socioeconomic biases that influence the estimation process. We theorized that the model may utilize physical features such as hospitals, temporary structures that may be associated with hospitals, and general activity metrics such as large volumes of cars on roads. However, in less wealthy but possibly highly populated counties, these features may be less common as there could be a deficiency in funding, even if the case counts are relatively the same. Further, the case counts may even be higher in these areas, as they could lack the necessary resources to help prevent the spread of the disease. In this case, the model would be biased to predict higher case counts for wealthier counties.

5 Conclusion

In this piece, we aimed to utilize deep learning for estimating Covid-19 cases in contexts when alternative surveillance modeling techniques are limited by inaccurate or unavailable data. We chose satellite imagery as a data source unaffected by censorship, polarized political climates, nor regional technological capabilities. We then compared the merits of a CNN model against a CNN-LSTM hybrid model, using the United States as a case study. The CNN model estimated the imagery alone, but the CNN-LSTM model was able to consider the imagery temporally. As a result, the CNN-LSTM model had a much lower loss, achieving 4.55% MAE as a percent of the mean. The model has successfully demonstrated that Covid-19 cases can be estimated using satellite imagery, but we believe further development and testing is necessary before the model is generalizable for other pandemics.

6 Acknowledgements

I acknowledge William Mary Research Computing for providing computational resources and/or technical support that have contributed to the results reported within this paper.

I would also like to thank the faculty and students of the William & Mary geolab (geolab.wm.edu) for their feedback and support. Specifically, I would like to thank Jason Kristine Lin, Elias Halperin Wolman, Laura Kathleen Mills, Joseph David O'Brien, Maeve Naughton-Rockwell, Sean Murphy, and Heather Baier.

Finally, I would like to thank my family for their incredible support in this endeavor. I could not have done this without them.

Bibliography

- Adnan, Myasar Mundher et al. (2021). "Automatic Image Annotation Based on Deep Learning Models: A Systematic Review and Future Challenges". In: *IEEE Access* 9, pp. 50253–50264. DOI: [10.1109/ACCESS.2021.3068897](https://doi.org/10.1109/ACCESS.2021.3068897).
- Barman, Rahul et al. (2019). "Transfer learning for small dataset". In: *Proceedings of the National Conference on Machine Learning, Mumbai, India*. Vol. 26.
- Basiri, Mohammad Ehsan et al. (2021). "ABCDM: An Attention-based Bidirectional CNN-RNN Deep Model for sentiment analysis". In: *Future Generation Computer Systems* 115. ISSN: 0167-739X. DOI: <https://doi.org/10.1016/j.future.2020.08.005>. URL: <https://www.sciencedirect.com/science/article/pii/S0167739X20309195>.
- Brewer et al. (2021). *Predicting road quality using high resolution satellite imagery: A transfer learning approach*. URL: <https://doi.org/10.1371/journal.pone.0253370>.
- Deng, Jia et al. (2009). "ImageNet: A large-scale hierarchical image database". In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. DOI: [10.1109/CVPR.2009.5206848](https://doi.org/10.1109/CVPR.2009.5206848).
- Dong Du, Gardner (2020). *An interactive web-based dashboard to track COVID-19 in real time*. DOI: [https://doi.org/10.1016/S1473-3099\(20\)30120-1](https://doi.org/10.1016/S1473-3099(20)30120-1).
- Dua, Mohit et al. (2020). "An improved RNN-LSTM based novel approach for sheet music generation". In: *Procedia Computer Science* 171, pp. 465–474.
- Ensafi Winter, Mueen and Crandall (2015). *Analyzing the Great Firewall of China Over Space and Time*. DOI: <https://doi.org/10.1515/popets-2015-0005>.
- Gavrikov, Paul (2020). *visualker*. <https://github.com/paulgavrikov/visualker>.
- Gharavi, Erfaneh, Neda Nazemi, and Faraz Dadgostari (2020). "Early outbreak detection for proactive crisis management using twitter data: Covid-19 a case study in the us". In: *arXiv preprint arXiv:2005.00475*.

- Gorelick, Noel et al. (2017). "Google Earth Engine: Planetary-scale geospatial analysis for everyone". In: *Remote Sensing of Environment*. DOI: 10.1016/j.rse.2017.06.031. URL: <https://doi.org/10.1016/j.rse.2017.06.031>.
- He, Kaiming et al. (2016). *Identity Mappings in Deep Residual Networks*. DOI: 10.48550/ARXIV.1603.05027. URL: <https://doi.org/10.48550/arxiv.1603.05027>.
- Ienco, Dino et al. (2017). "Land Cover Classification via Multitemporal Spatial Data by Deep Recurrent Neural Networks". In: *IEEE Geoscience and Remote Sensing Letters* 14.10, pp. 1685–1689. DOI: 10.1109/LGRS.2017.2728698.
- Jean et al. (2016). "Combining satellite imagery and machine learning to predict poverty". In: *Science* 353.6301, pp. 790–794. DOI: 10.1126/science.aaf7894. URL: <https://www.science.org/doi/abs/10.1126/science.aaf7894>.
- Kemp, Simon (2022). *DIGITAL 2022: ZIMBABWE*. URL: <https://datareportal.com/reports/digital-2022-zimbabwe>.
- Komenda Bulhart, Karolyi Jarkovský Mužík Májek Šnajdrová Růžicková Rážová Prymula Macková Březovský Marouněk Černý Dušek (2020). "Complex Reporting of the COVID-19 Epidemic in the Czech Republic: Use of an Interactive Web-Based App in Practice". In: *Journal of Medical Internet Research* 22.5. URL: <https://doi.org/10.2196/19367>.
- Liu, Haoran (2021). *The Long-Term Effects of Urban Cover Changes on Thermal Environment Using Dense Landsat Time Series*. URL: <https://proxy.wm.edu/login?url=https://www.proquest.com/dissertations-theses/long-term-effects-urban-cover-changes-on-thermal/docview/2563495297/se-2?accountid=15053>.
- Milan and Treré (2020). "The Rise of the Data Poor: The COVID-19 Pandemic Seen From the Margins". In: *Sage Journals*. URL: <https://doi.org/10.1177/2056305120948233>.
- Monitoring, Copernicus Land (2020, 2021).
- Nindrea Sari, Lazuardi Aryandono (2020). *Validation: The Use of Google Trends as an Alternative Data Source for COVID-19 Surveillance in Indonesia*. URL: <https://doi.org/10.1177/1010539520940896>.
- Papadomanolaki, Maria et al. (2019). "Detecting Urban Changes with Recurrent Neural Networks from Multitemporal Sentinel-2 Data". In: pp. 214–217. DOI: 10.1109/IGARSS.2019.8900330.

- Rahnemoonfar, Maryam et al. (2018). "Flooded Area Detection from Uav Images Based on Densely Connected Recurrent Neural Networks". In: pp. 1788–1791. DOI: [10.1109/IGARSS.2018.8517946](https://doi.org/10.1109/IGARSS.2018.8517946).
- Runfola Stefanidis, Baier (2021). *Using satellite data and deep learning to estimate educational outcomes in data-sparse environments*. URL: <https://doi.org/10.1080/2150704X.2021.1987575>.
- Science, NASA Landsat (2013). URL: <https://landsat.gsfc.nasa.gov/satellites/landsat-8/>.
- (2021). URL: <https://landsat.gsfc.nasa.gov/satellites/landsat-9/>.
- Skiera, Bernd et al. (2020). *How to Best Predict the Daily Number of New Infections of Covid-19*. DOI: <https://doi.org/10.48550/arXiv.2004.03937>. URL: <https://arxiv.org/abs/2004.03937>.
- Sulyok, Mihály, Tamás Ferenci, and Mark Walker (2021). "Google Trends Data and COVID-19 in Europe: Correlations and model enhancement are European wide". In: *Transboundary and Emerging Diseases* 68.4, pp. 2610–2615.
- The Centers for Disease Control and Prevention (2021). *How We Conduct Case Surveillance*. URL: <https://www.cdc.gov/nndss/about/conduct.html>.
- The New York Times (2021). URL: <https://github.com/nytimes/covid-19-data>.
- Timoneda, Joan C. and Sebastián Vallejo Vera (2021). *Will I die of coronavirus? Google Trends data reveal that politics determine virus fears*. URL: <https://doi.org/10.1371/journal.pone.0258189>.
- Vincent, François and Olivier Besson (2019). *Subpixel Target Detection in Hyperspectral Imaging*. DOI: [10.1109/CAMSAP45676.2019.9022454](https://doi.org/10.1109/CAMSAP45676.2019.9022454).
- Williams, Graham et al. (2002). "A comparative study of RNN for outlier detection in data mining". In: *2002 IEEE International Conference on Data Mining, 2002. Proceedings*. IEEE, pp. 709–712.
- Yuan, Ye and You Zhou (2015). "Twitter sentiment analysis with recursive neural networks". In: *CS224D course projects*.