Predicting Covid-19 outbreaks using social media images

Mentors: Marianthi-Anna Kioumourtzoglou, Carl Vondrick

Members: Ashutosh Kanitkar, Chris Russo, Gideon Teitel, Luke Beasley, Shahen Mirzoyan

Project Purpose: Enrich geospatial COVID data using images from social media

High-level Goal: Using computer vision, leverage image data from social media to enrich existing geospatial datasets

COVID-19 Use Case: Pipeline to scrape social media images, build useful COVID-related features to enhance predictive features in existing models

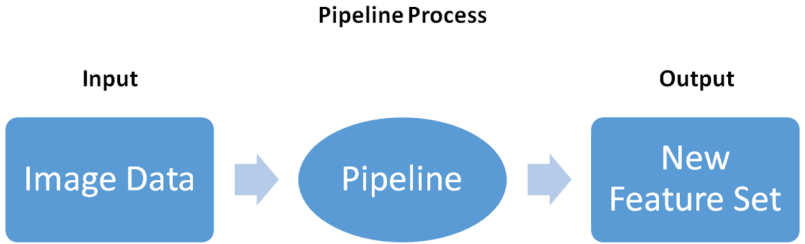

Pipeline Overview

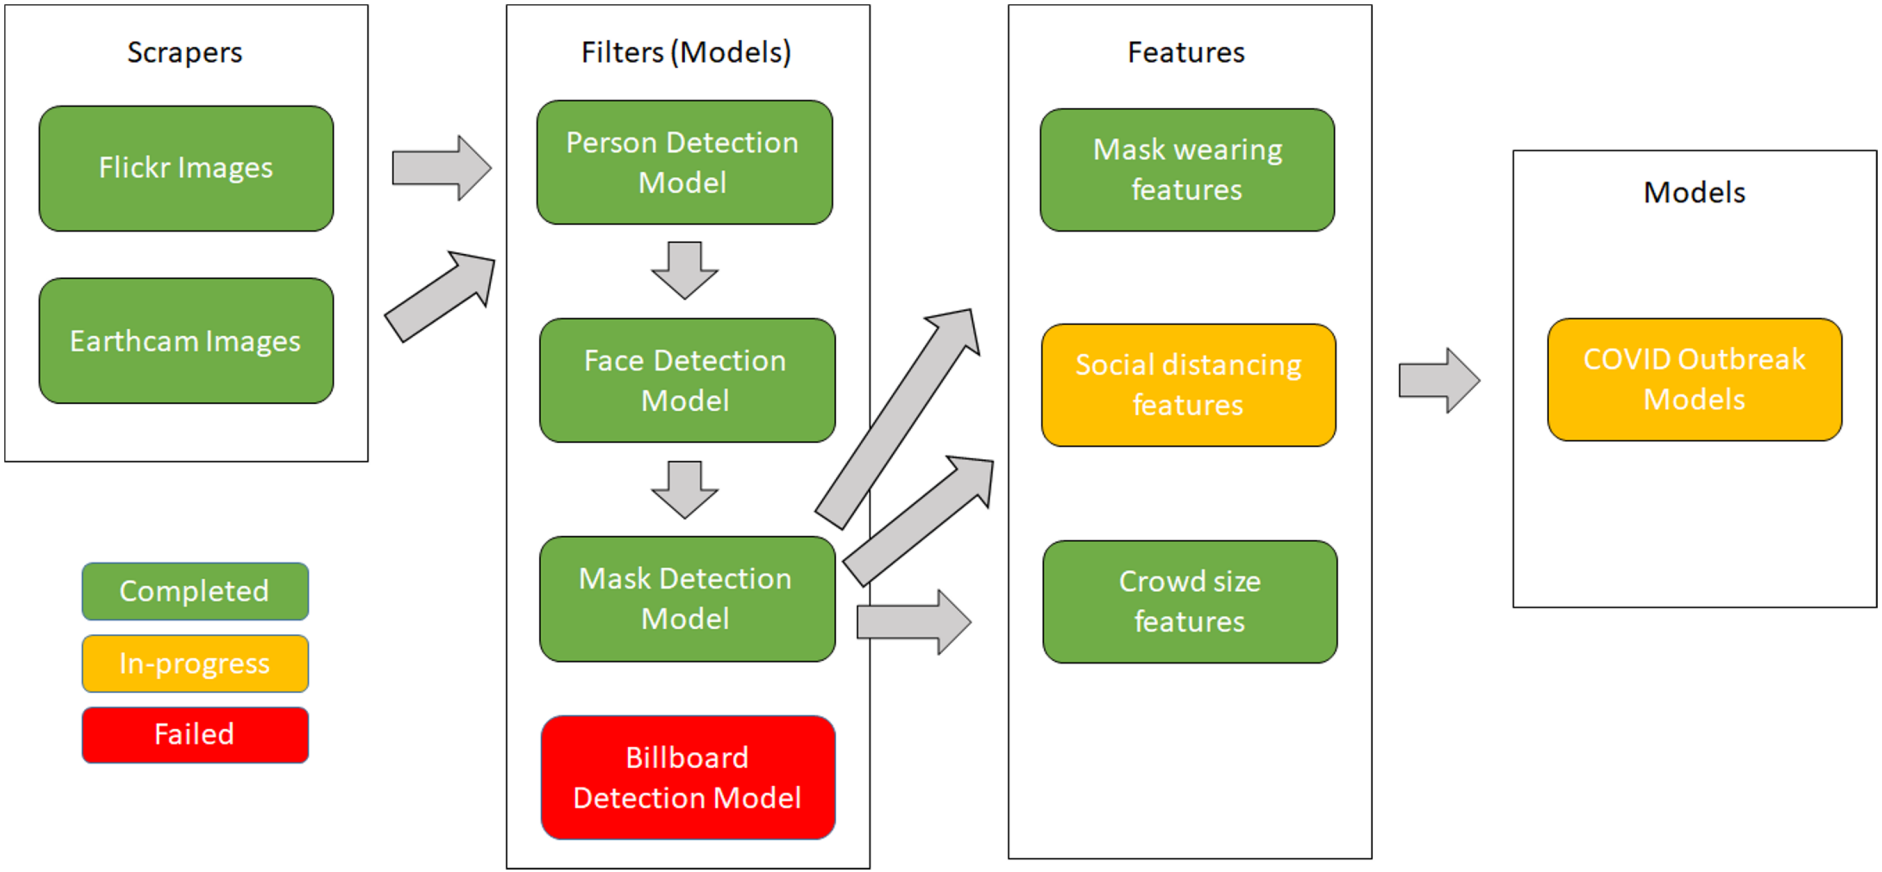

Image data

EarthCam

- Street view cameras
- 62,673 images (2,079 from Little Italy, rest from Times Square)

Flickr

- Images uploaded by users/photographers
- Results of search by location
- 4,325 total images from Times Square, Harlem, Greenwich Village, Columbus Circle, Williamsburg, Prospect Park, Flushing Meadows

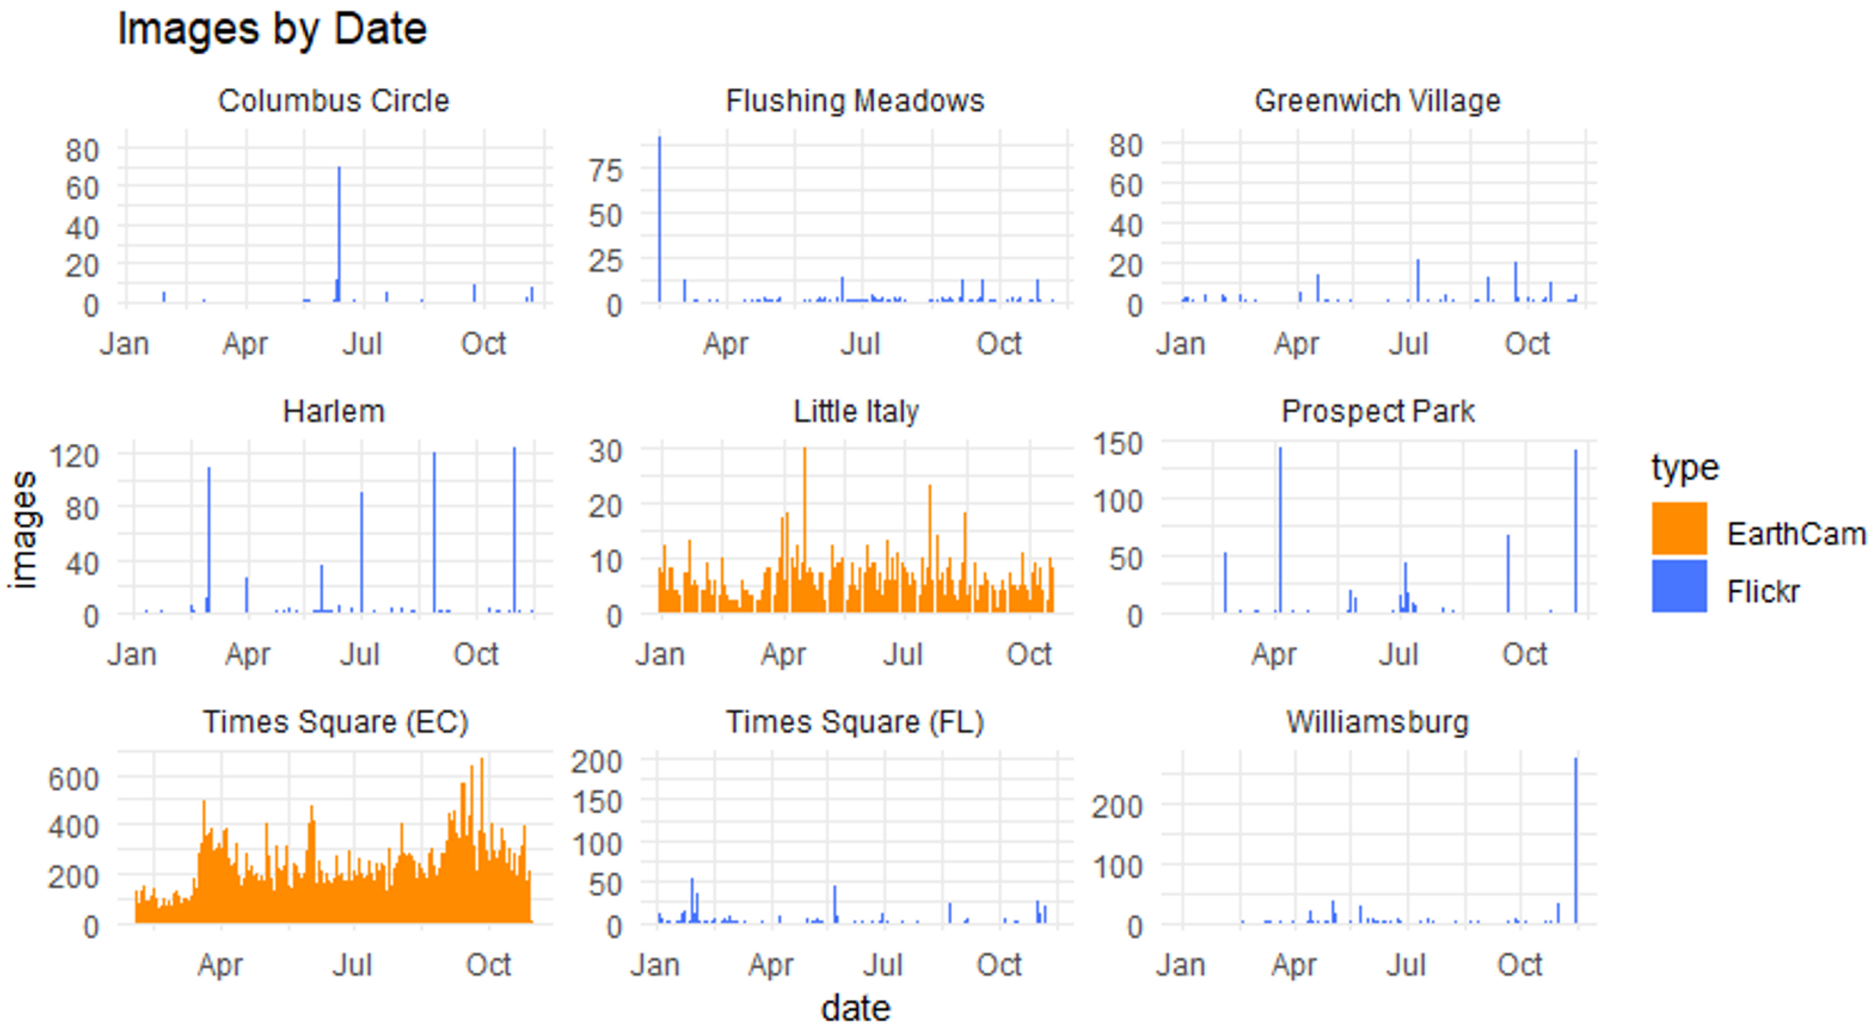
Images by Date
Columbus Circle
Flushing Meadows
Greenwich Village
Harlem
Little Italy
Prospect Park
Times Square (EC)
Times Square (FL)
Williamsburg
images
date
Jan
Apr
Jul
Oct
type
EarthCam
Flickr

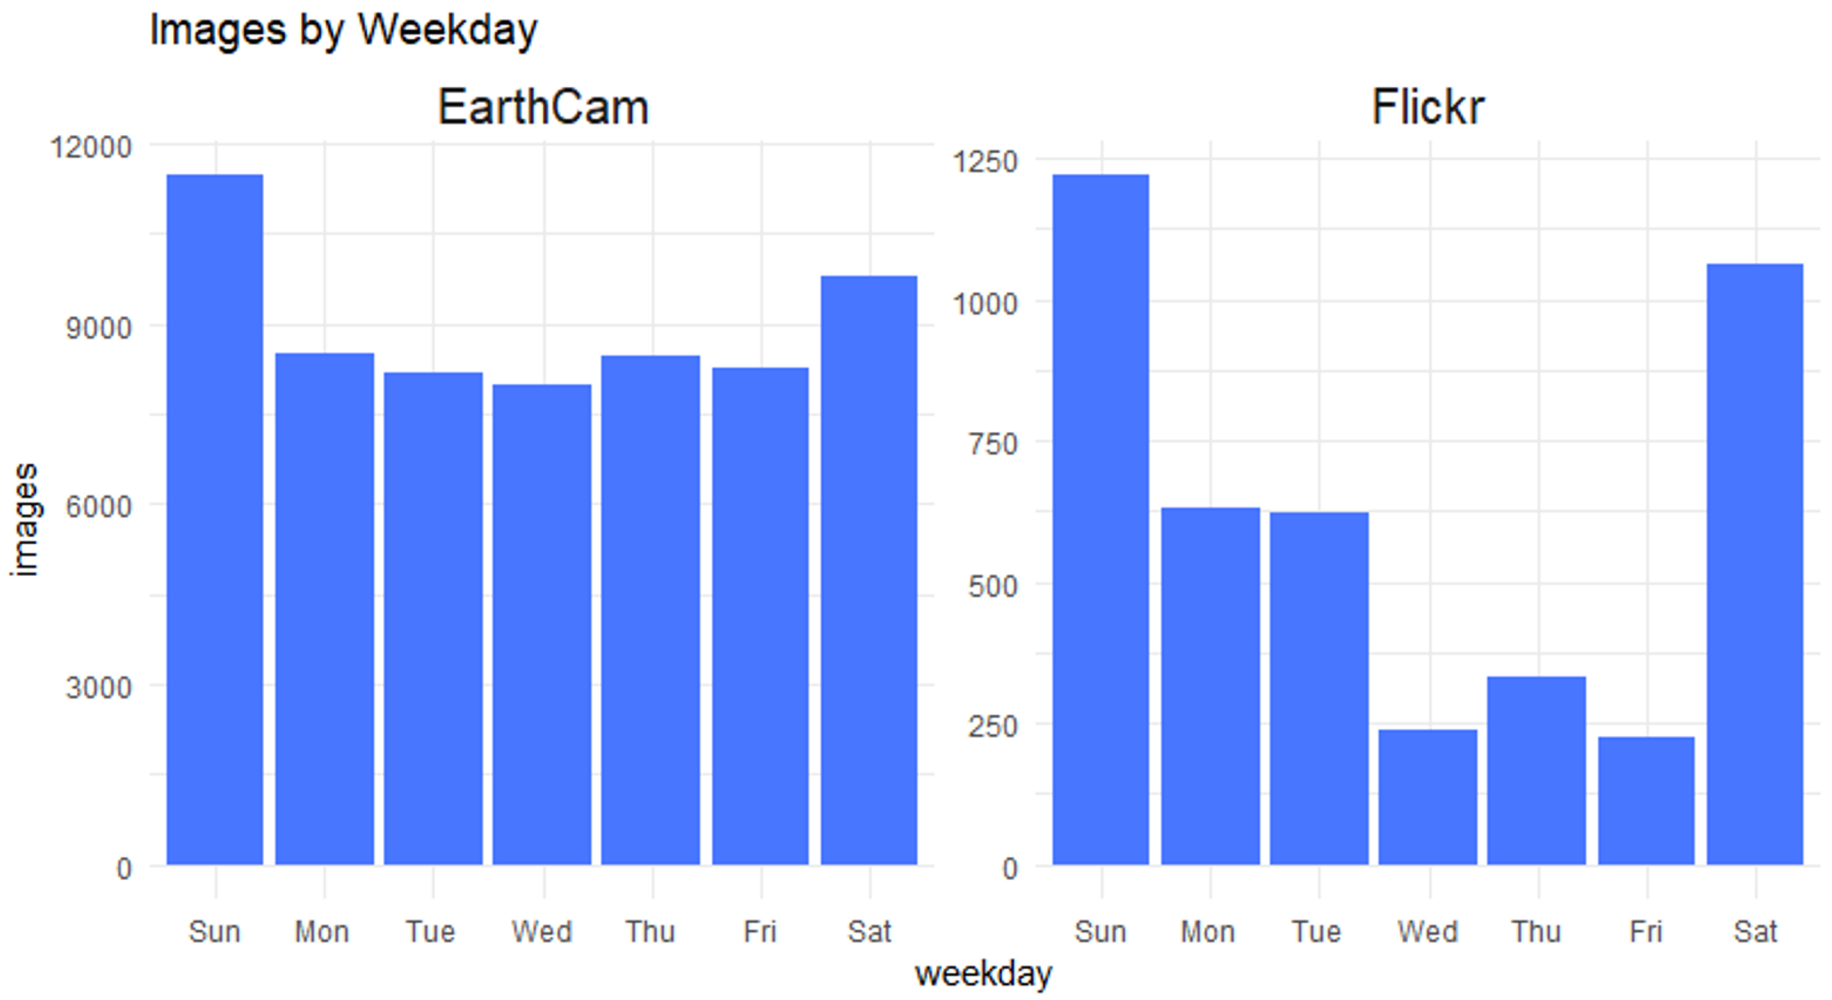
Images by Weekday
EarthCam
Flickr
images
12000
9000
6000
3000
0
1250
1000
750
500
250
0
Sun
Mon
Tue
Wed
Thu
Fri
Sat
weekday

Person Detector Model Used

Faster RCNN with Inception V2 has an 76 map across all categories of MS COCO dataset and has map of 85 in particular for person detection

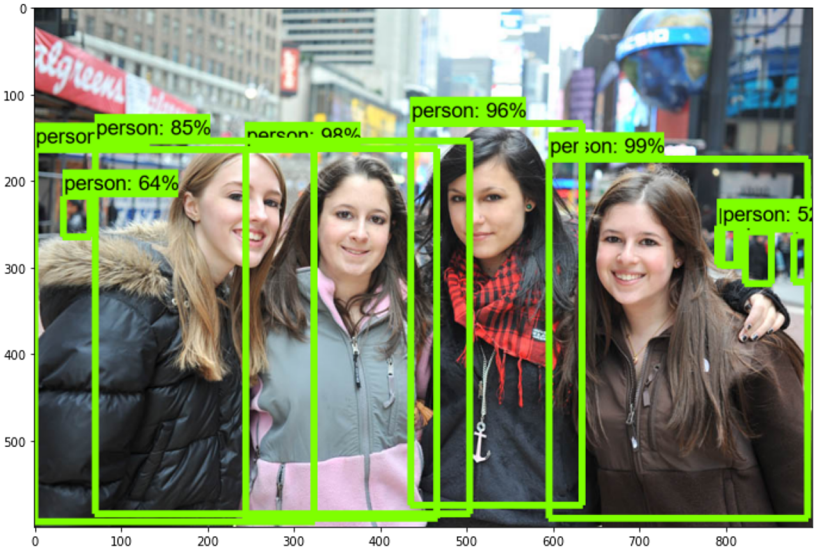

Face Detection using Multi-task Cascaded Convolutional Networks

Mask Detection

We have decided to use a mask dataset that was created by Chandrika Deb given that the images used were real images of faces wearing masks. This dataset consists of 3835 images belonging to two classes:with_mask: 1916 images
without_mask: 1919 images.
Example of dataset

Pipeline Illustration

Input Image

Person Detector Output

Face Detector Output

Mask O/P

Another example

Output of person detector for input image

Person detector info

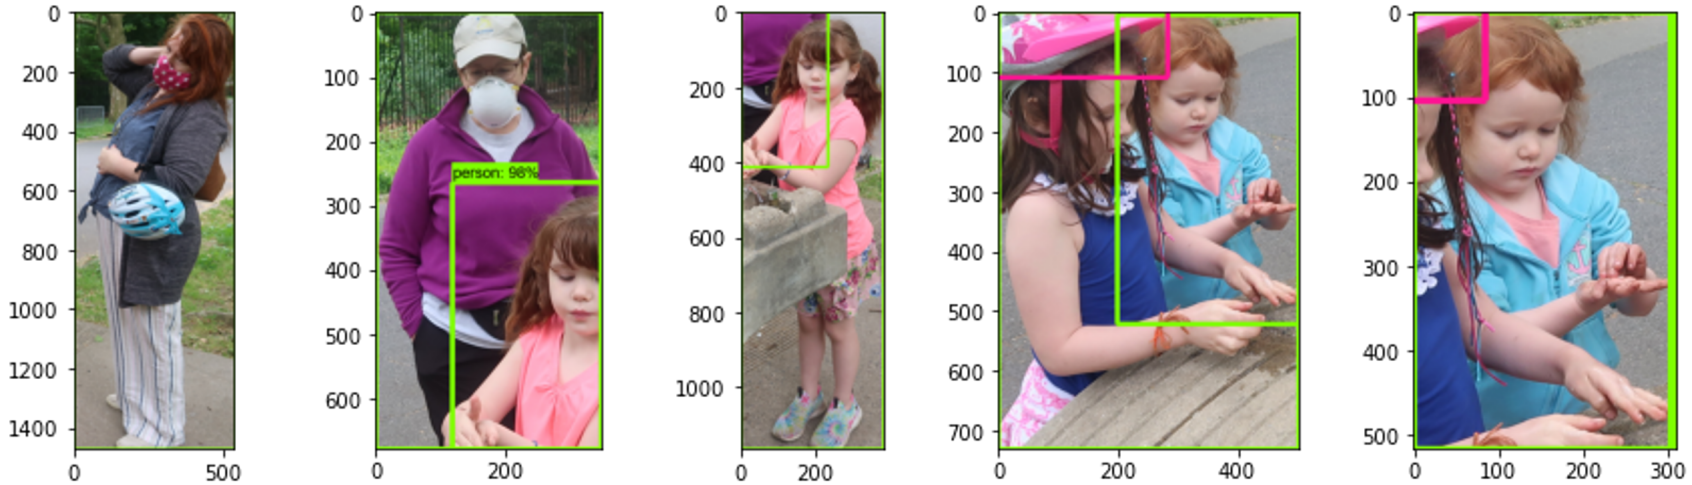

Output of face detector for selected people

Face detector info

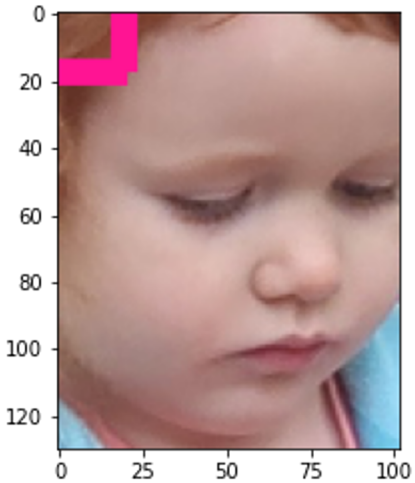

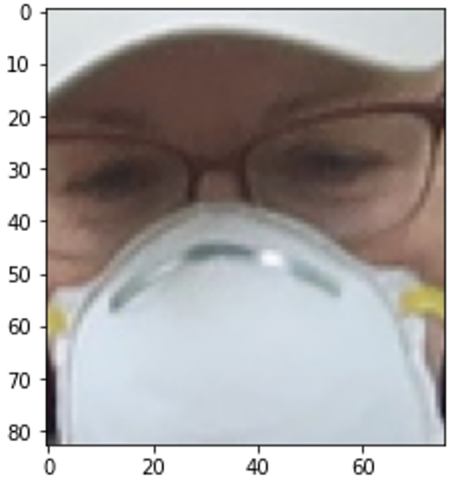

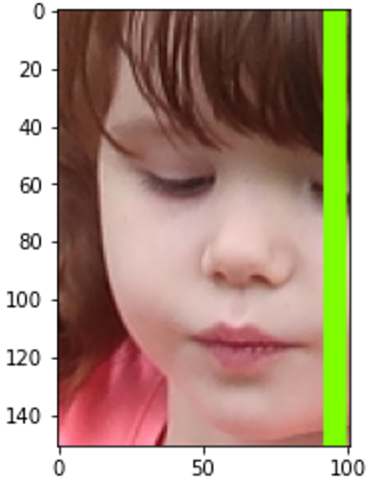

Output of mask detector

Mask detector info

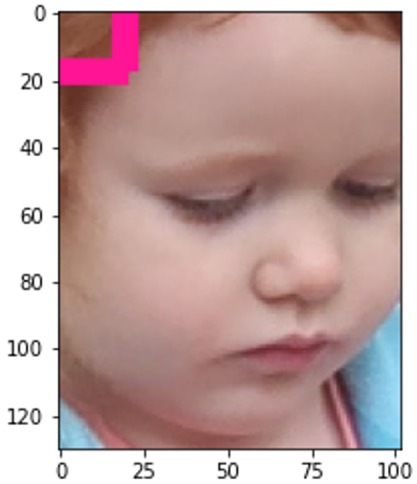

No

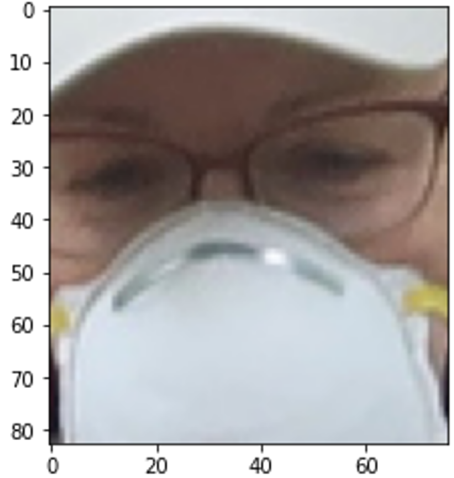

Yes

Pipeline output

Final Pipeline output shows total number of people , total number of people wearing masks and total people of people not wearing masks for each date starting from 1 january 2020 to current date for each of different locations (along with their source) present in dataset

Snippet of Pipeline Output

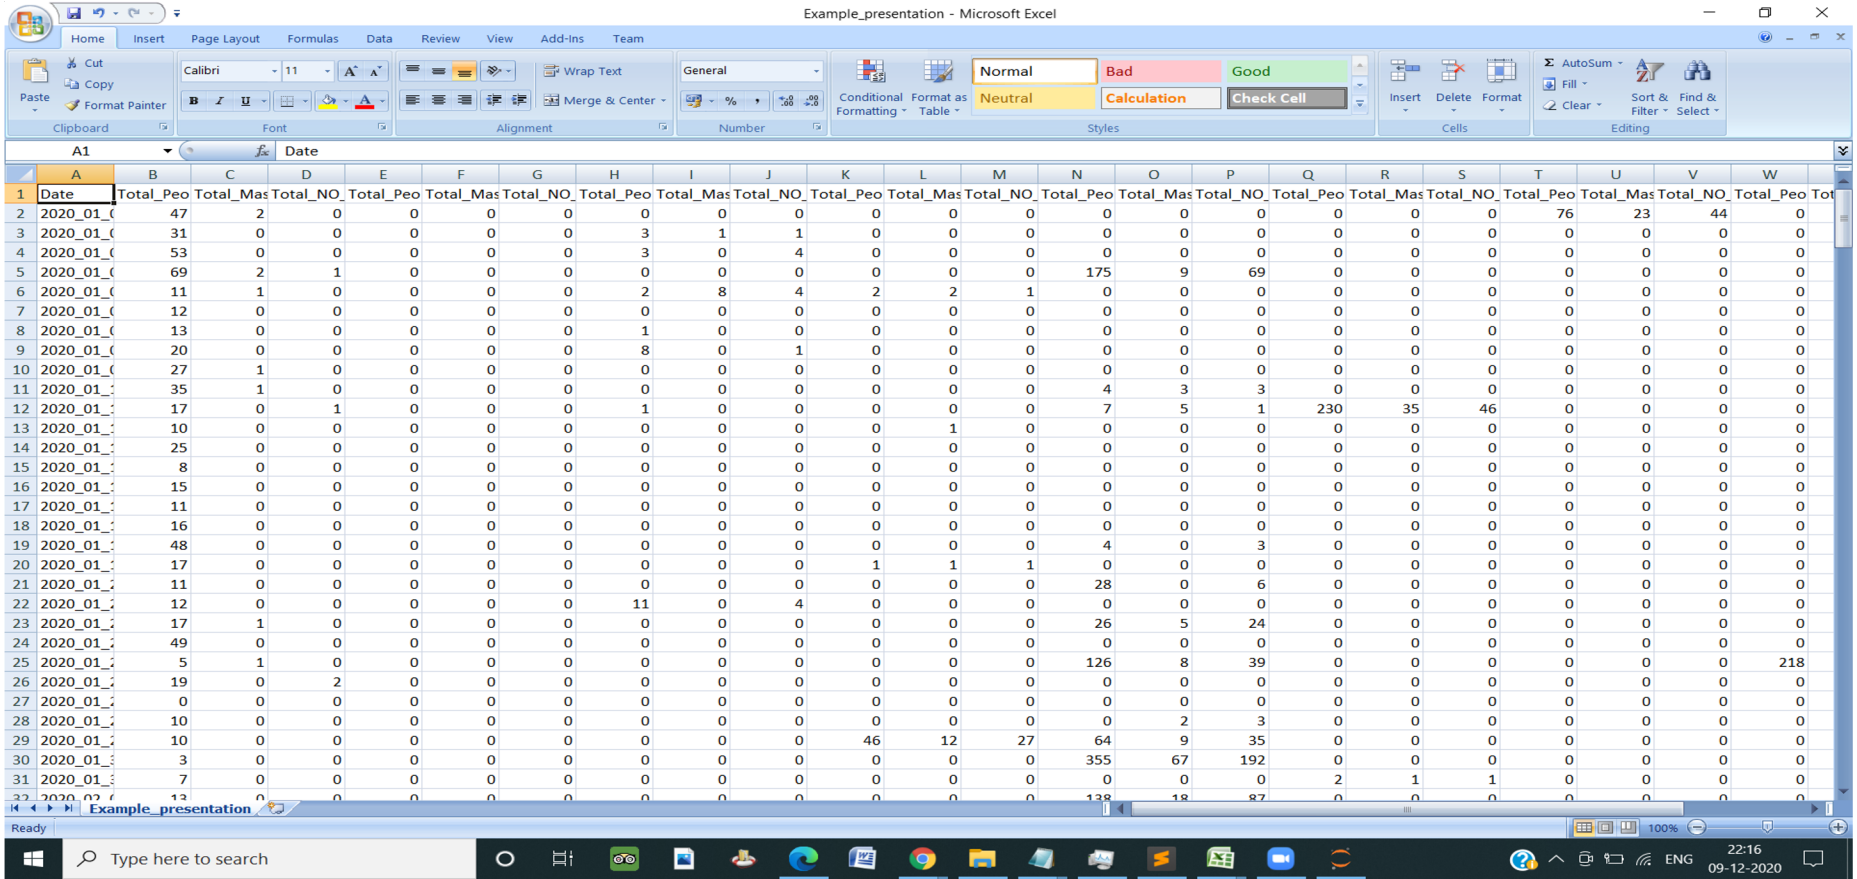

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
1	Date	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Total_Mas	Total_NO_	Total_Peo	Tot
2	2020_01_(	47	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	76	23	44	0	
3	2020_01_(	31	0	0	0	0	0	3	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
4	2020_01_(	53	0	0	0	0	0	3	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	
5	2020_01_(	69	2	1	0	0	0	0	0	0	0	0	0	175	9	69	0	0	0	0	0	0	0	
6	2020_01_(	11	1	0	0	0	0	2	8	4	2	2	1	0	0	0	0	0	0	0	0	0	0	
7	2020_01_(	12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
8	2020_01_(	13	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
9	2020_01_(	20	0	0	0	0	0	8	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
10	2020_01_(	27	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
11	2020_01_1	35	1	0	0	0	0	0	0	0	0	0	0	4	3	3	0	0	0	0	0	0	0	
12	2020_01_1	17	0	1	0	0	0	1	0	0	0	0	0	7	5	1	230	35	46	0	0	0	0	
13	2020_01_1	10	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	
14	2020_01_1	25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
15	2020_01_1	8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
16	2020_01_1	15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
17	2020_01_1	11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
18	2020_01_1	16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
19	2020_01_1	48	0	0	0	0	0	0	0	0	0	0	0	4	0	3	0	0	0	0	0	0	0	
20	2020_01_1	17	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	
21	2020_01_2	11	0	0	0	0	0	0	0	0	0	0	0	28	0	6	0	0	0	0	0	0	0	
22	2020_01_2	12	0	0	0	0	0	11	0	4	0	0	0	0	0	0	0	0	0	0	0	0	0	
23	2020_01_2	17	1	0	0	0	0	0	0	0	0	0	0	26	5	24	0	0	0	0	0	0	0	
24	2020_01_2	49	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
25	2020_01_2	5	1	0	0	0	0	0	0	0	0	0	0	126	8	39	0	0	0	0	0	0	218	
26	2020_01_2	19	0	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
27	2020_01_2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
28	2020_01_2	10	0	0	0	0	0	0	0	0	0	0	0	0	2	3	0	0	0	0	0	0	0	
29	2020_01_2	10	0	0	0	0	0	0	0	0	46	12	27	64	9	35	0	0	0	0	0	0	0	
30	2020_01_3	3	0	0	0	0	0	0	0	0	0	0	0	355	67	192	0	0	0	0	0	0	0	
31	2020_01_3	7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	1	1	0	0	0	0	
32	2020_02_(	13	0	0	0	0	0	0	0	0	0	0	0	138	18	87	0	0	0	0	0	0	0	

Cases/Masks by location

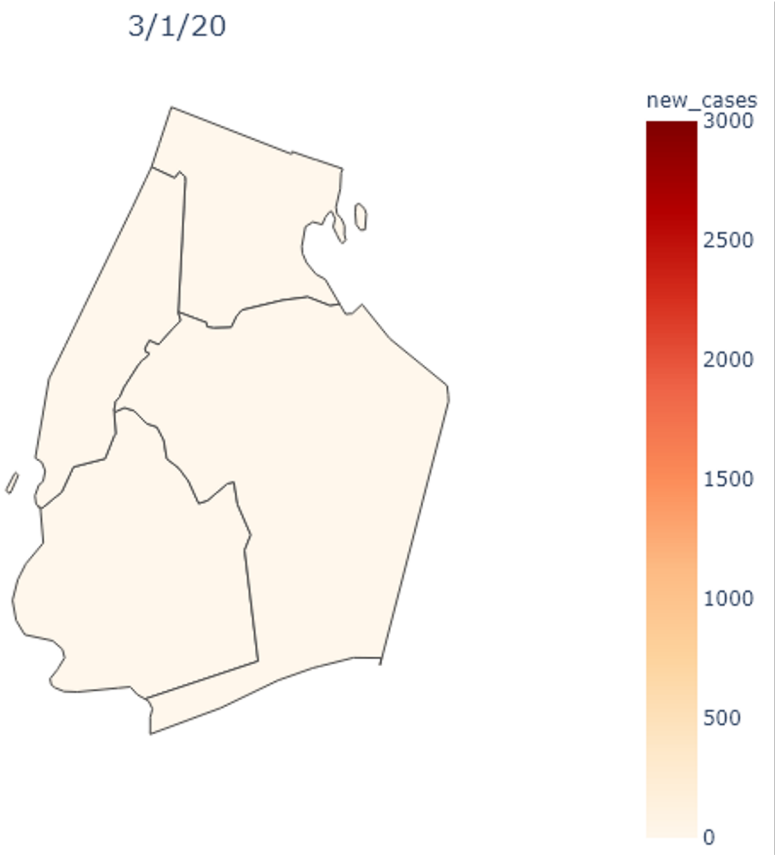

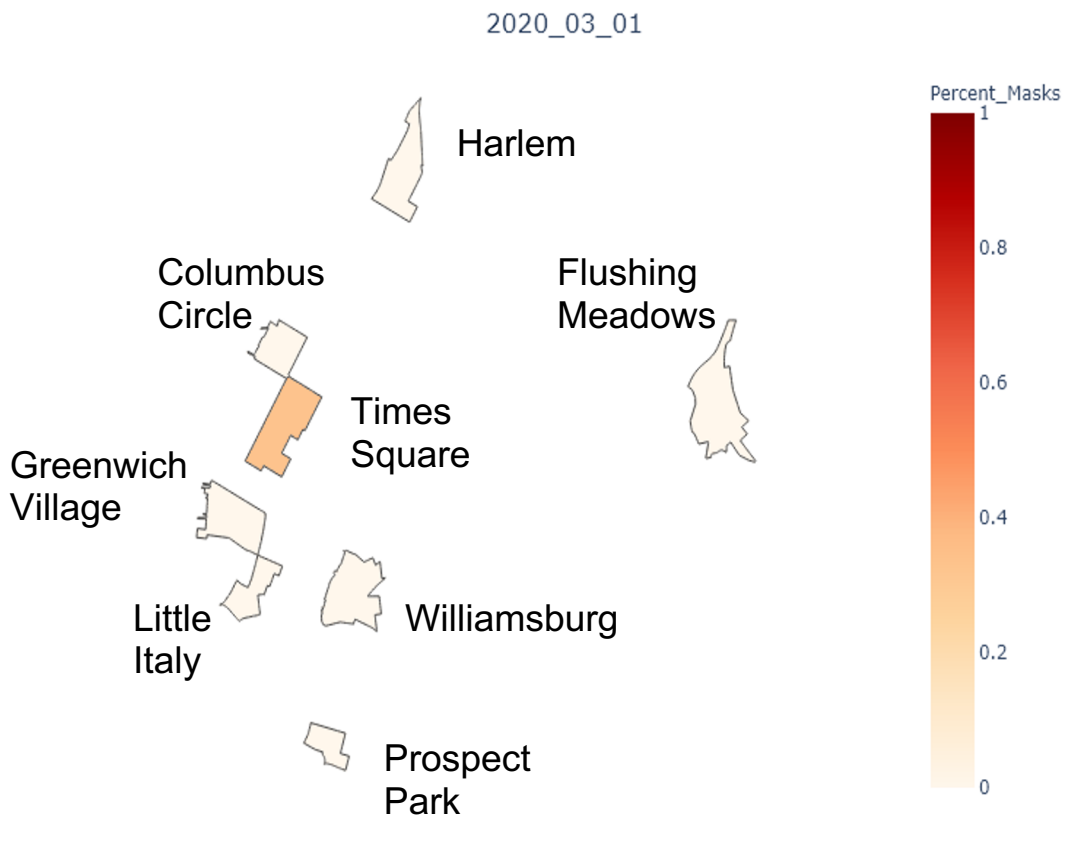

Covid-19 and Google Mobility Data

- Google Mobility data accessed through smart phone signal
- Shows changes in traffic to stores, parks, transit stations
- Daily per county Covid data from usafacts.org

Models

- Response variable:
  - Covid-19 cases (New York County)

- Explanatory variables:
  - 6 Google Mobility indicators
  - Percent of people wearing masks / Percent of people not wearing masks (percent of people that we are unsure of is the third category implied by the first two)

- Methods:
  - Linear Regression, Random Forest, Support Vector Machine, Gradient Boosting

Time component

- The actions of people on a given day are not likely to be reflected in Covid-19 cases immediately

- We averaged the values of the explanatory variables over a certain number of past days to align with the current day's case number
  - Did so for 5, 10, and 14 previous days

- Models evaluated by RMSE on validation sets, and best performing model applied to a separate test set

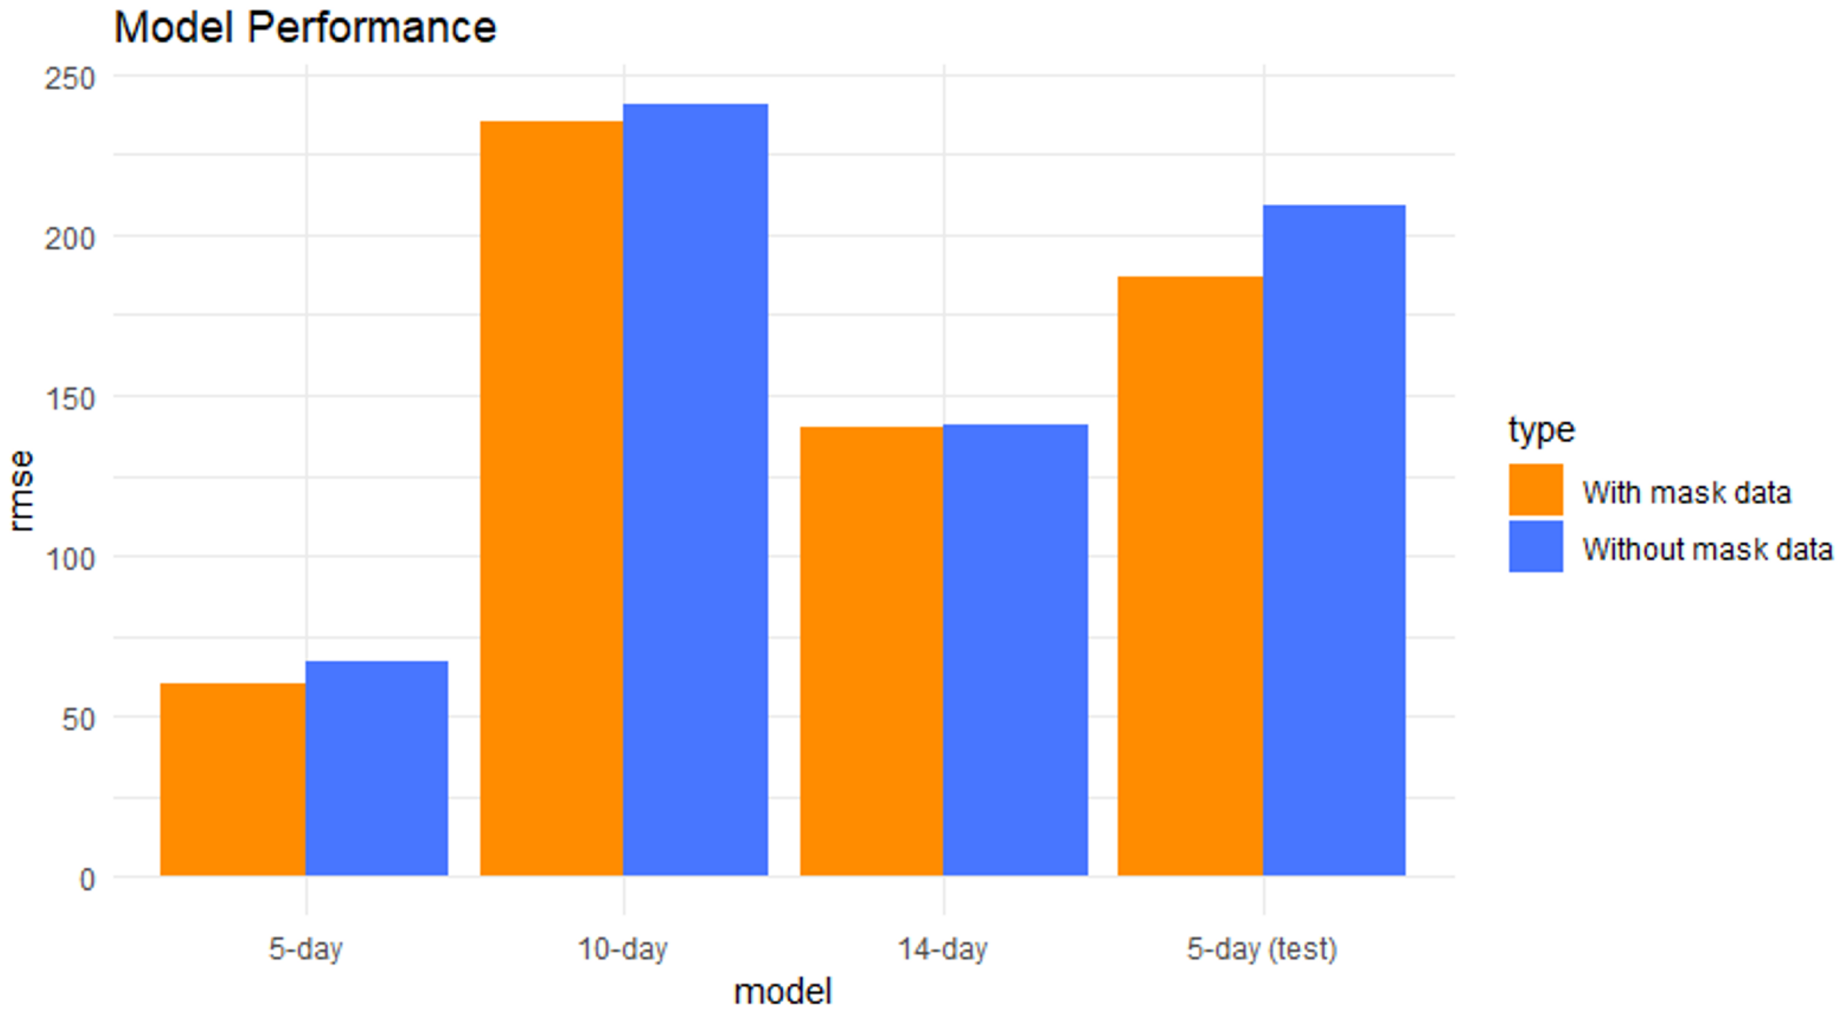
Model Performance
250
200
150
100
50
0
rmse
5-day
10-day
14-day
5-day (test)
model
type
With mask data
Without mask data

Discussion

- Not surprising that the models do not consistently predict cases well
  - Small sample size and limited information from Google Mobility data

- Encouraging that even with varying RMSE values, the slight difference between mask/no mask information holds

- Some evidence that our generated mask data provides useful information

Limitations

Misclassifying billboards as people

- Billboard Classifier not solved
- Next Step: build “real person” classifier on output of person detector

Imbalanced data: by date and by location

- Next Step: add latitude/longitude parameters to scrapers

Potential Bias

- Next step: Exploratory data analysis on image frequency across different demographics

Conclusion/Potential Next Steps

- Leverage pipeline to enhance other geospatial datasets that could improve performance of additional models

- Predictions can be made more robust with data from additional neighborhoods

- Incorporate social distancing analysis
  - Use pedestrian detection models to build classifier for social distancing
  - Requires inferring distances from a single view
  - Previous literature on the topic: People Watching: Human Actions as a Cue for Single View Geometry

References

Deb, Chandrika. "GitHub." n.d. *https://github.com/chandrikadeb7/Face-Mask-Detection#:~:text=Face%20Mask%20Detection%20system%20built%20with%20OpenCV%2C%20Keras%2FTensorFlow,images%20as%20well%20as%20in%20real-time%20video%20streams.*

Girshick, R. ""Fast R-CNN "." (2015).

Google, LLC. ""Google COVID-19 Community Mobility Reports"." n.d. 5 December 2020.

Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi. "You Only Look Once:Unified, Real-Time Object Detection." (n.d.).

MTA, Metropolitan Transportation Authority. "Turnstile Data." n.d. 21 October 2020.

Navneet Dalal, Bill Triggs. "“Histograms of Oriented Gradients for Human Detection“." (n.d.).

R. Girshick, J. Donahue, T. Darrell, and J. Malik. " "Rich feature hierarchies for accurate object detection and semantic segmentation."." *CVPR* (2014).

References

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,* (2017): 1137-1149.

Simonyan, Karen, Zisserman, and Andrew. "Very Deep Convolutional Networks for Large-Scale Image Recognition." *arXiv.org* (2015).

Telegraph. "https://www.telegraph.co.uk/technology/2018/11/25/chinese-businesswoman-accused-jaywalking-ai-camera-spots-face/." n.d. *Telegraph.*

USAFACTS. “US Coronavirus Cases and Deaths”. n.d. 5 December 2020.

Vanhoucke, Vincent, Sergey, Jonathon, and Zbigniew. "Rethinking the Inception Architecture for Computer Vision." *arXiv.org* (2015).

W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. "SSD: Single Shot MultiBox Detector." (n.d.).

Zhang, Ren, Sun, and Jian. "Deep Residual Learning for Image Recognition." *arXiv.org* (2015).