

A Comparative Study of Contrail Frequency Indices and GOES-16 Contrail Data Set

Ryan O'Hara¹

Universities Space Research Association (USRA), Moffett Field, CA, 94035 USA

Hok K. Ng², Swati Saxena³

NASA Ames Research Center, Mountain View, CA, 94035, USA

Contrail formations have been shown to contribute to the greenhouse effect: they are practically transparent to incoming solar radiation and do little to reflect heat away from Earth but are highly effective at trapping heat within Earth's atmosphere. To understand the impact contrails have on climate change, contrail frequency indices (CFIs) can be used as a method to quantify aircraft-induced persistent contrails. These indices are capable of tracking long-term contrail formation and identify regions of airspace with the highest contrail formation rates. In this paper, an algorithm is proposed which is capable of using NASA Sherlock and Global Forecast System (GFS) datasets and computing CFIs over large geographic regions and long temporal intervals using NASA Ames' High-End Computing Capability (HECC) supercomputing system. CFIs are computed using nowcast weather data and previously flown flight tracks. This paper calculated the CFIs of all twenty Air Route Traffic Control Centers in the National Airspace System on October 28th, 2019 and compared the distribution of non-zero CFIs with observed contrail data collected from GOES-16 Satellite data in order to assess the accuracy of the CFI system as a contrail prediction model. It was ultimately determined that the computed CFIs were broadly distributed in the same way as the GOES-16 contrail data and that the individual CFIs computed at the latitude/longitude points at which GOES-16 contrail masks were available had high precision and recall (at 0.75 and 0.86 respectively). While these validation results bode well for the accuracy of the CFI method, the number of provided GOES-16 masks was quite small. Future work should aim to increase the size of the GOES-16 dataset in order to perform a more comprehensive comparison between these two datasets.

I. Introduction

Condensation trails (contrails) are formed when water vapor expelled by aircraft engines condenses and freezes around engine exhaust particles. Under particularly low temperatures, any water added to the atmosphere will immediately freeze; this creates persistent contrails that can linger for hours. Such formations have been shown to contribute to the greenhouse effect: they are practically transparent to incoming solar radiation and do little to reflect heat away from Earth but are highly effective at trapping heat within Earth's atmosphere. As a whole, contrails have been found to increase nighttime atmospheric warming by 33%, while persistent contrails specifically have been shown to have an up to ten times greater contribution to the greenhouse effect than other aircraft emissions [1]. Furthermore, the warming effect of contrails is, unlike carbon dioxide emissions, immediate, making them an urgent driver of climate change [2]. For both this reason as well as concerns that, as the commercial airline industry grows, contrails will form at more frequent rates, contrail mitigation is a primary objective of modern aviation industries [3].

As the aviation industry looks to tackle the problems posed by contrails, steps have been taken to better understand the ambient conditions that are conducive to persistent contrail formation so that aircraft can avoid these regions, as

¹ Research Assistant, Aviation Systems Division, USRA.

² Aerospace Engineer, Aviation Systems Division, NASA Ames Research Center, AIAA Member.

³ Aerospace Engineer, Aviation Systems Division, NASA Ames Research Center, AIAA Senior Member.

previous research indicates that small flight diversions can mitigate most contrail-induced warming [4]. Most of this contrail research has been achieved with observational satellite data, which can visually obtain insights on regions and times of contrail formation. However, different satellites have different drawbacks: low-Earth orbit (LEO) satellites have been used to quantify contrail coverage, but such satellites are unable to shed much light on the temporal shifts in contrail formation due to their infrequent passes of the same region of Earth’s surface, while geostationary satellites, despite bypassing this issue by collecting continuous data, are often plagued with low-fidelity images due to coarse spatial resolutions [5] [6]. As such, it has been necessary to develop algorithms capable of predicting contrail formation regions. Such models are also of value to the aviation industry as they can be used to make short-term predictions of potential contrail formation regions, allowing flights to be rerouted to avoid these airspace regions.

Historically, contrail prediction models have been developed around hand engineered, bottom-up algorithms which use physical properties of the atmosphere to determine the locations of contrails; however, these algorithms lack fidelity and long-term accuracy [7]. As a result, the aviation field is currently working to develop novel models capable of making more accurate predictions about contrails.

II. Methods

The current work proposes one such contrail prediction model. In previous work, a quantification method was proposed to define the contrail frequency index, or CFI [1]. The CFI of a particular region can be calculated by determining [1]:

- i. The supersaturation of ice with respect to water in the surrounding atmosphere. When ice is supersaturated, the atmosphere is cold enough to maintain persistent contrails. This can be calculated using temperature and pressure readings from NOAA’s National Centers for Environmental Information (NCEI)’s Global Forecast System (GFS) dataset [13] using Eq. (1) [8], where T is the temperature in Celsius and RH_w is the relative humidity with respect to water (both of which are stored within the GFS dataset). The NCEI collects new readings globally on a 361×720 degree grid (that is, the data has a fidelity of 0.5 degrees) every six hours (at 0:00, 6:00, 12:00, and 18:00 UTC) and minimizes uncertainty by generating 21 different forecasts, all of which are used when producing the final GFS dataset [8].
- ii. The number of aircraft flying through the chosen geospatial point over the course of one hour.

$$RH_i = RH_w \times \frac{6.0612e^{18.102T/(249.52+T)}}{6.1162e^{22.577T/(237.78+T)}} \tag{1}$$

Previous work indicates that knowledge of these two datapoints will allow for predictions about whether a persistent contrail will form. When the atmosphere is in a state of supersaturation, the relative humidity with respect to ice (RH_i) will exceed 100%, and the long-lived ice crystals that are characteristic of persistent contrails could form. Aircraft must then fly through this section of airspace while RH_i remains above 100% to form a persistent contrail [9]. The CFI method makes use of this information to quantify the number of aircraft flying through a specific region over the course of one hour that could create persistent contrails [1].

The CFI of a particular Air Route Traffic Control Centers (ARTCC) for a particular point in time is thus calculated by summing the product of supersaturation from a persistent contrail formation matrix (a binary value of 1 when supersaturation occurs and 0 otherwise) and aircraft count across all geospatial points within that air traffic region’s airspace as in Eq. 1. ARTCCs are defined by the FAA as per Fig. 1.

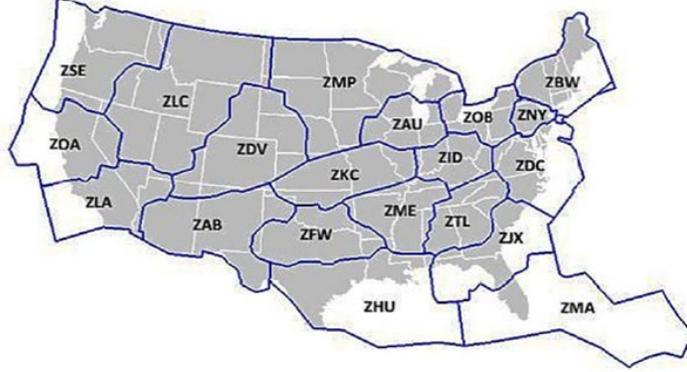


Fig 1. A map of the FAA's 20 ARTCCs over the continental US [10].

$$CFI = \sum_{i=0}^m \sum_{j=0}^n r_{i,j} c_{i,j} a_{i,j}$$

(2)

Using Eq. (2), the CFI of a particular ARTCC composed of an $m \times n$ grid, where each point on the grid is separated by 0.5 degrees on a traditional latitude and longitude plot, can be computed. The higher the CFI value, the greater the chance that a plane flying through that airspace will produce a persistent contrail, which is defined as a contrail that exists for at least 2 minutes. Eq. (2) sums over the product of each term in the persistent contrail formation matrix $r_{i,j}$ the localized ARTCC matrix ($c_{i,j}$, which is 1 if the given spatial data point is within the ARTCC region of interest and 0 otherwise), and $a_{i,j}$, the number of aircraft track points within that point over the hour of interest. Note that this grid size can be generalized to a region of any arbitrary size so long as the product of $r_{i,j}$, $c_{i,j}$, and $a_{i,j}$ is summed over a grid consisting of regions 0.5 degrees by 0.5 degrees in size. This study only produced CFIs over the continental United States due to the fact that only domestic air traffic and flight data was available for use.

Previous work on CFIs has indicated that the CFI predictions made using this quantification model are highly correlated with the true CFI [1]. The predicted CFIs were computed using Eq. (2), predicted future aircraft locations, and NOAA's Rapid Updated Cycle (RUC) atmospheric forecast dataset, while the true CFIs made use of known flight tracks and nowcast dataset. CFIs were computed at the 1-hour mark with an average correlation coefficient of 0.85 and had decreased correlation over longer prediction intervals (with three- and six-hour average correlation coefficients of 0.64 and 0.52 respectively). Results indicate that CFIs are stable in relation to small variations in air traffic and, instead, are largely dependent on shifting atmospheric conditions inherent in a contrail prediction model. Thus, with further analysis on larger datasets, CFIs show great promise as a new state-of-the-art model for making contrail predictions.

This paper aims to address this need by comparing CFI matrices produced over large timescales with an observed, human-labeled contrail dataset, thus shedding light on the accuracy of this method. The work first outlines a novel Python algorithm developed to compute CFI matrices using input atmospheric and flight track data, then discusses the use of Ames Research Center's HECC supercomputing facilities to generate large quantities of CFI matrices for long-term comparisons, and finally discusses a method for treating and comparing these matrices to the ground truth, human-labeled dataset.

For this study, a human-labeled contrail dataset collected by NASA Ames colleagues from the GOES-16 geostationary weather satellite was used. This dataset was collected from generated contrail masks at various latitude/longitude pairs across a period of months over 2019 and 2020 containing information on the specific spatial location of observed contrail pixels. This dataset functioned as an effective validation method with which to test the CFI method, as described below.

A. Contrail Frequency Index Matrix Formation

As discussed in the previous section, a matrix containing the CFIs of a particular ARTCC can be built using atmospheric and flight track data. This work developed a Python language algorithm capable of reading in data from GFS and from NASA's Sherlock Data Warehouse [11]. GFS can be used to determine the RH_i at a given spatial point,

and Sherlock contains flight track information of all domestic commercial airlines over a given hour. Our algorithm parses the data from both GFS and Sherlock for all flights through the region of interest over the hour of interest and builds a flight matrix. The flight matrix contains information on the elevation and changing geospatial locations of each flight. The flight matrix data is then converted into the binary data needed to calculate the CFI for each ARTCC within the region of interest.

Next, the data from the GFS is used to build a binary supersaturation matrix, a graphical representation of which can be found in Fig 2, which contains information on whether each geospatial point in the region of interest has a saturation index of above or below 100%. Finally, these matrices are multiplied together and summed over the region of the ARTCC to compute the CFI of each ARTCC within the region of interest; these CFIs are output in an easy-to-read table (Table 1).

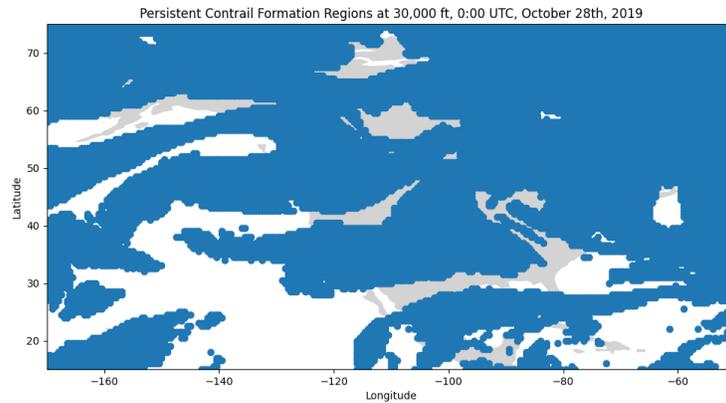


Fig 2. A graphical representation of the RHi matrix computed at an elevation of 30,000 ft from 0:00-1:00 UTC on October 28th, 2019. Regions in blue indicate an RHi above 100% (which is associated with a value of 1 on the binary RHi matrix), regions on grey indicate the presence of landmasses, and regions in white indicate water.

Table 1. CFIs at each altitude (in feet) from 20000 ft to 46000 ft in 1000 ft increments for each US ARTCC on October 28th, 2019 from 0:00-1:00 UTC.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
0 Elevation	ZBW	ZTL	ZSE	ZOB	ZOA	ZNY	ZMP	ZME	ZMA	ZLC	ZLA	ZKC	ZJX	ZHU	ZID	ZFW	ZDV	ZDC	ZAB	ZAU	
1 20000	34	47	0	232	0	24	41	19	0	60	161	13	214	418	100	135	101	0	68	68	
2 21000	41	99	0	235	0	39	15	50	0	88	129	148	145	386	91	254	87	0	123	186	
3 22000	56	142	0	340	7	27	23	17	3	252	331	14	218	453	138	197	89	0	271	83	
4 23000	67	306	0	256	6	56	37	44	58	82	463	48	489	376	169	251	80	0	325	198	
5 24000	157	327	0	441	17	111	40	73	146	206	693	74	538	509	174	423	90	5	306	221	
6 25000	329	264	0	384	0	187	252	122	131	120	479	103	404	471	277	414	185	15	432	438	
7 26000	84	301	186	553	0	240	591	169	96	495	599	226	469	422	631	438	819	20	474	798	
8 27000	288	684	3	725	0	400	71	292	122	131	544	363	572	541	418	454	539	42	564	527	
9 28000	119	496	8	1802	0	695	389	646	71	379	910	1146	362	619	934	531	1026	83	1001	985	
10 29000	149	553	33	891	0	373	73	336	123	60	1445	505	576	727	453	865	656	166	1149	441	
11 30000	45	981	53	1747	0	821	1511	987	252	837	1780	1284	606	1113	1216	710	2948	136	1842	1376	
12 31000	218	875	290	1319	0	889	447	547	246	571	838	1340	1023	427	612	463	1121	228	1747	864	
13 32000	4	2030	1071	3176	0	867	1398	2754	261	1522	1246	2654	1099	1730	2309	1312	1877	177	1309	2910	
14 33000	148	2663	1116	1192	158	828	787	2089	147	2566	375	1414	1238	302	2054	692	3151	769	659	1522	
15 34000	3	3995	2490	1273	406	685	3471	5165	267	6333	1005	5036	2611	3160	1844	3711	7873	780	1640	1789	
16 35000	0	4552	2310	21	315	310	2717	3619	229	4132	66	5297	3851	1443	985	1895	5884	2270	1740	1497	
17 36000	0	3529	2376	0	131	350	2018	5960	581	4421	180	4791	3745	4468	639	3500	6409	1504	1126	699	
18 37000	0	2198	652	0	6	231	766	4598	0	779	0	4825	1552	345	918	2059	4338	1520	484	48	
19 38000	0	1324	66	0	0	57	0	1469	0	254	0	741	708	35	340	401	911	505	138	0	
20 39000	0	553	0	0	0	413	0	1134	0	0	0	199	382	0	339	219	38	533	0	0	
21 40000	0	21	0	0	0	70	0	85	0	0	0	18	25	0	6	33	0	25	0	0	
22 41000	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
23 42000	0	0	0	0	0	0	0	0	1	0	0	0	0	2	0	0	0	0	0	0	
24 43000	0	0	0	0	0	0	0	0	6	0	0	0	6	0	0	0	0	0	0	0	
25 44000	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
26 45000	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
27 46000	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

CFI matrices are rapidly produced over long timescales (generally a period of several months, for the purposes of this project) by running this algorithm on the HECC, NASA Ames’ supercomputing facility. As each CFI matrix contains the CFIs for all continental United States ARTCCs for one hour, a run of several months produces approximately 1,000 CFI matrices. This offers the advantage of allowing the authors to test the efficacy of this algorithm on large quantities of aviation data, providing in-depth information on the robustness of this approach.

For this particular study, CFI tables were generated for 0:00-1:00, 6:00-7:00, 12:00-13:00, and 18:00-19:00 UTC on October 28th, 2019. This day was chosen due to the high volume of contrails observed over its 24-hour period in the GOES-16 human-labeled dataset, making it an ideal time period with which to compare the generated CFIs and ground truth data owing to the relatively large number of contrail pixels in the latter dataset. Meanwhile, this set of hours was selected due to limitations in the ways the GFS’ nowcast data was stored. The GFS only collected atmospheric data at 0:00, 6:00, 12:00, and 18:00 UTC every day; any CFIs generated outside of the hour following this data’s collection would be generated using out-of-date atmospheric data, which could result in inaccurate CFIs that didn’t reflect the current state of the atmosphere.

Notably, the generated CFI matrices and GOES-16 pixel count have different dimensions (the matrices are 4-dimensional in terms of latitude, longitude, elevation, and time while the GOES-16 pixel counts are 3-dimensional in terms of latitude, longitude, and time). Therefore, to produce CFI matrices with the same dimensions as the GOES-16 dataset and allow for accurate comparisons, the CFI matrices were next flattened along the elevation axis. This flattening process was done by summing the CFIs of all elevations of a particular ARTCC together, producing a singular CFI for each ARTCC for each studied hour. This flattening process can be justified by noting that, for the purposes of this study, no information is lost by merging the different elevation levels together. The GOES-16 satellite collected its information by treating the atmosphere of the Earth as a two-dimensional plane; that is, contrails that existed at different elevations were flattened onto a single plane when the satellite images were captured. The CFI matrix flattening procedure replicates this process. Furthermore, when flattening the CFI matrices, all information is conserved. If a specific latitude/longitude point had a nonzero CFI at a specific elevation point, then it will still have a nonzero CFI at that same latitude/longitude point after the flattening procedure takes

place. For these reasons, the flattening procedure was the most effective and accurate way to compare the two datasets.

B. Ground Truthing Via GOES-16 Contrail Dataset Comparisons

Next, metadata from the human-labeled GOES-16 dataset was processed. This dataset contained the latitude and longitude of a contrail mask (each of which spanned approximately 0.5 x 0.5 degrees), the time and date at which the contrails within the mask were observed, and the number of contrail pixels within the mask (with a higher contrail pixel count correlating with a larger observed contrail). Note that each mask has pixel dimensions of 256 by 256, so the maximum possible number of contrail pixels in a single mask is 65,536. The contrail information from 0:00-1:00, 6:00-7:00, 12:00-13:00, and 18:00-19:00 GMT on October 28th, 2019 was extracted using a Python script and directly compared to the latitude and longitude points at which non-zero CFIs were generated. This methodology allowed for the validation of the CFI method by treating the GOES-16 images as a ground truth dataset. Each of these comparison points could be classified as one of four possible values (TP, FP, TN, FN): a true positive (TP; a non-zero CFI was generated and a non-zero number of contrail pixels were observed), a false positive (FP; a non-zero CFI was generated, but zero contrail pixels were observed), a true negative (TN; a zero CFI was generated and zero contrail pixels were observed), or a false negative (FN; a zero CFI was generated, but a non-zero number of contrail pixels were observed). This data could then be used to compute the precision and recall of each studied hour, where:

$$\text{Precision} = \frac{TP}{TP+FP} \tag{2}$$

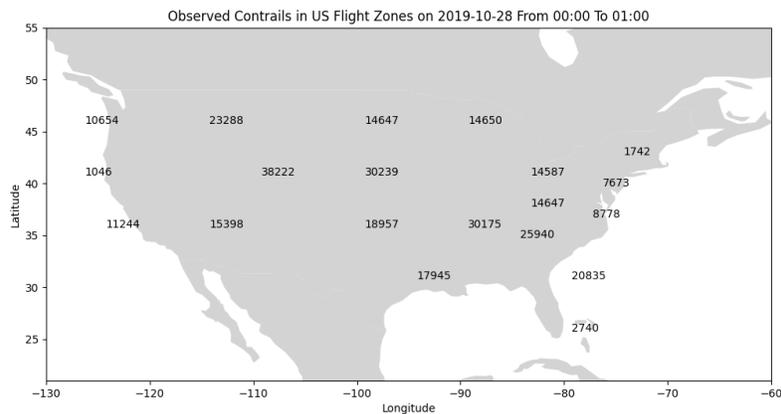
and

$$\text{Recall} = \frac{TP}{TP+FN} \tag{3}$$

with a high precision and recall indicating a high level of correlation between the generated CFIs and ground truth GOES-16 dataset [12].

III. Results

After running the above algorithms for October 28th, 2019, the flattened CFI matrices visualized in Fig. 3 were produced.



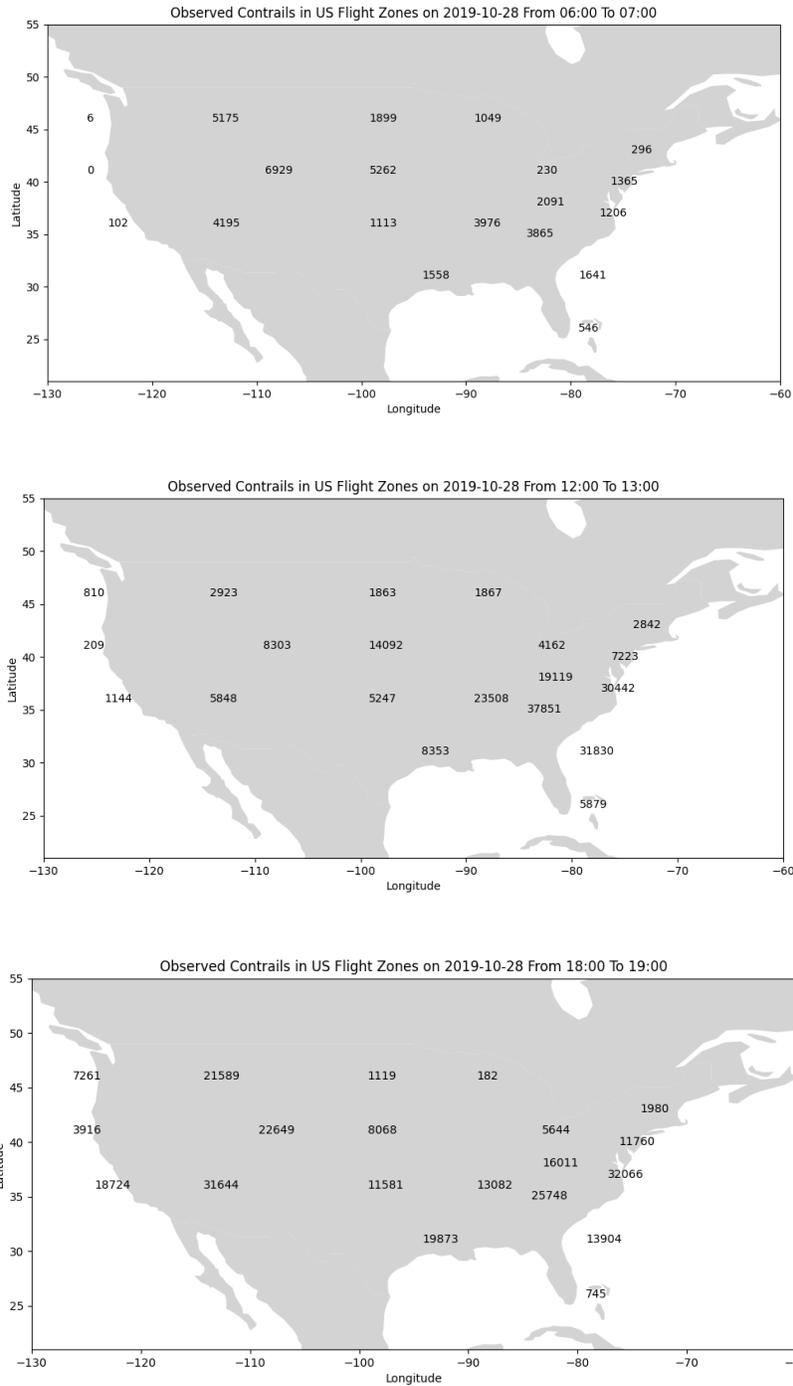


Fig 3. A visual representation of the flattened CFIs generated for all twenty ARTCCs at 0:00-1:00, 6:00-7:00, 12:00-13:00, and 18:00-19:00 on October 28th, 2019. Over the course of the entire day, contrails are most likely to form across the midwestern and southeastern regions of the US.

Similarly, the GOES-16 human-labeled dataset at the same four times on October 28th, 2019 contained contrail masks at the locations visualized in Fig. 4.

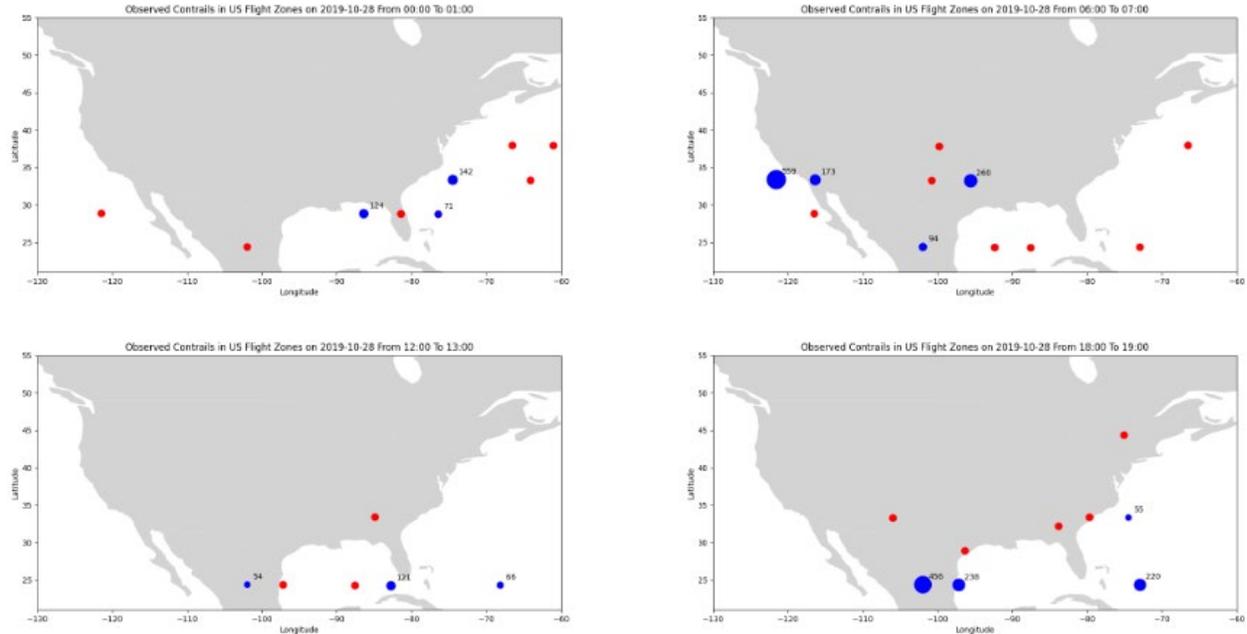


Fig 4. A visual representation of the GOES-16 contrail masks produced during 0:00-1:00, 6:00-7:00, 12:00-13:00, and 18:00-19:00 on October 28th, 2019. A red circle indicates a mask with 0 contrail pixels, while blue circles indicate masks with a non-zero number of contrail pixels. The blue circles are sized according to the magnitude of the observed number of contrail pixels (and the total number of pixels is inscribed within the circle). All circles are centered around the latitude/longitude pair at which the associated mask was produced.

First, to understand whether the broad distribution of CFIs agrees with the GOES-16 data for October 28th, this study considers the day’s total contrail pixel counts collected by colleagues at NASA Ames. These pixel counts, as illustrated in Table 2, show the number of contrail pixels counted over the entire day in 10 degree delineations. This data indicates that, on October 28th, the GOES-16 dataset observed the highest number of contrail pixels between a latitude of 30 and 40 and longitude of -80 and -90 (a region which covers the US’ southeastern states). Meanwhile, the number of contrail pixels drops off when progressing south along Florida or west toward California. As indicated in Fig. 3, the overall magnitude of the day’s CFIs match this pattern; across all four computed CFI maps, the flattened CFI count for the ZMA, ZOA, ZSE, and, to a lesser degree, ZLA ARTCCs tend to have the lowest magnitudes while the ARTCCs ZJX, ZTL, and ZDC tend to have some of the largest magnitudes.

Table 2. Total contrail pixel counts from the GOES-16 dataset for October 28th, 2019. Pixel counts are grouped into bins with a size of 10 degrees. The table’s x-axis represents each pixel count’s longitude while the table’s y-axis represents the pixel count’s latitude.

	-120	-110	-100	-90	-80	-70	-60
30	559	221	0	654	106,522	335	130
20	0	4,551	775	937	1,002	464	816

In order to better understand the accuracy of the CFI method within smaller 0.5 x 0.5 degree regions, rather than only considering the large-scale CFI trends, the study examined the precision and recall of the computed CFIs as per the method outlined above. Comparisons were made at a fidelity of 0.5 degree by 0.5 degree grids (that is, considering all GOES-16 masks cover an area less than this grid, the total number of analyzed grids for any one hour is equal to the number of points on each hour’s GOES-16 plot in Fig. 4). The number of TP, FP, TN, and FN grids as defined above can be found in Table 3.

Using the values in Table 3, the precision and recall of the CFI algorithm during each hourly period on October 28th, 2019 can be determined using (2) and (3) respectively. This information can be used to determine the algorithm’s

F1-score, which measures the overall predictive performance of the CFI method during the associated hourly period. Note that, due to the usage of different latitude/longitude points during each period, an overall F1-score for this dataset on October 28th, 2019 cannot be determined. The F1-score is calculated as per (4):

$$F1 = \frac{2}{\left(\frac{1}{precision}\right) + \left(\frac{1}{recall}\right)} \tag{4}$$

The precision, recall, and F1-score values for each hour are listed in Table 3. Previous literature suggests that, to be considered a viable candidate for real-world contrail prediction applications, the CFI method should generate an F1-score of at least 0.80. While the size of this study’s dataset prevents sweeping conclusions about the accuracy of the CFI method to be made, the F1-scores for each hour are either just slightly lower than or, in the case of the 0:00-1:00 dataset, above, this baseline.

Table 3. Number of true positive, false positive, true negative, and false negative 0.5 degree by 0.5 degree grids across the four analyzed hours. For ease of comparison, any GOES-16 contrail masks that crossed into multiple grids were only compared with the CFIs in the 0.5 degree by 0.5 degree region holding the mask’s minimum latitude and longitude.

	Number of TP	Number of FP	Number of TN	Number of FN	Precision	Recall	F1-Score
00:00 – 01:00	3	1	5	0	0.75	1.00	0.85
06:00 – 07:00	4	2	5	0	0.66	1.00	0.79
12:00 – 13:00	2	0	3	1	1.00	0.66	0.79
18:00 – 19:00	3	1	4	1	0.75	0.75	0.75

IV. Conclusion

As outlined above, this ground truth analysis suggests that the CFI method produces results that largely agree with the observed contrails. The computed CFIs follow the same general magnitude pattern as the GOES-16 contrail counts and, while the number of contrail masks may be small in number, the individual CFIs and GOES-16 contrail masks had high precision and recall.

One notable drawback of this study is that the GOES-16 human-labeled dataset used was not comprehensive (that is, this study only examined a small handful of contrail points across the country). As such, while broad trends could be observed using the day’s total contrail pixel counts and the individual comparisons made between masks and CFIs indicate the general accuracy of the CFI method, a larger human-labeled dataset would allow for more rigorous comparisons to be made. As is, it is possible that the positive results observed in this study are not indicative of the overall performance of the CFI method. Future work could augment the GOES-16 dataset used in this study (or, alternatively, make use of a larger human-labeled contrail dataset) and then repeat the study’s methodology to develop more ironclad conclusions about the robustness of the CFI method. Another direction for future work could involve comparing the contrails predicted using CFIs with observed contrails over the elevation dimension. As discussed above, one of the drawbacks of the observed contrails dataset used in this study was its 3-dimensional, rather than 4-dimensional, structure, which prevented any comparisons to be made between elevations. There would be value in a future observed contrail dataset being built to differentiate between contrails at different elevation levels to allow for these comparisons to be made. This would provide valuable insight into the accuracy of the CFI method.

Despite this shortcoming, this study’s novel comparison of the accuracy of CFI-based algorithms that make use of large amounts of nowcast weather data, recorded flight tracks, and ground truth, human-labeled GOES-16 contrail masks further strengthens the case for the broader adoption of the CFI method. This study’s results showing the accuracy of CFIs lays the groundwork for future research, which can focus on determining the accuracy of the

method's near future forecasting abilities. If further work indicates that CFIs are consistently capable of making accurate predictions, then this algorithm can be applied to predict the contrails formed by upcoming commercial flights and reroute planes as necessary to mitigate each flight's environmental impact.

With the impacts of climate change growing more apparent and the speed at which the aviation industry is expanding, it is essential to develop methods that can decrease the significant environmental impact of contrails. By continuing to develop the CFI method – which, as evidenced by this study, produces results that align with ground truth data – the aviation field has the ability to propel itself into a global leader in climate change mitigation and help preserve the skies for years to come.

V. Acknowledgements

The authors would like to thank Sebastian Gutierrez-Nolasco at USRA for providing the GOES-16 contrail masks used in this study. This material is based upon work supported by NASA under Contract Number NNA16BD14C, managed by the USRA.

References

- [1] N. Chen, B. Sridhar and H. Ng, "Prediction and Use of Contrail Frequency Index for Contrail Reduction Strategies," *AIAA Guidance, Navigation, and Control Conference*, 2012.
- [2] D. Avila, L. Sherry and T. Thompson, "Reducing global warming by airline contrail avoidance: A case study of annual benefits for the contiguous United States," *Transportation Research Interdisciplinary Perspectives*, vol. 2, 2019.
- [3] B. Karcher, "Formation and radiative forcing of contrail cirrus," *Nature Communications*, vol. 9, no. 1824, 2018.
- [4] R. Teoh, U. Schumann, A. Majumdar and M. Stettler, "Mitigating the Climate Forcing of Aircraft Contrails by Small-Scale Diversions and Technology Adoption," *Environmental Science Technology*, vol. 54, no. 5, pp. 2941-2950, 2020.
- [5] V. Meijer, L. Kulik, S. Eastham, F. Allroggen, R. Speth, S. Karaman and S. Barrett, "Contrail coverage over the United States before and during the COVID-19 pandemic," *Environ. Res. Lett.*, vol. 17, pp. 34-39, 2022.
- [6] D. Spangenberg, P. Minnis, S. Bedka, R. Palikonda, D. Duda and F. Rose, "Contrail radiative forcing over the Northern Hemisphere from 2006 Aqua MODIS data," *Advancing Earth and Space Sciences*, vol. 40, no. 3, pp. 595-600, 2013.
- [7] L. Kulik, "Satellite-based detection of contrails using deep learning," *Massachusetts Institute of Technology, Department of Aeronautics and Astronautics*, p. thesis, 2019.
- [8] O. Alduchov and R. Eskridge, "Improved Magnus Form Approximation of Saturation Vapor Pressure," *Journal of Applied Meteorology*, vol. 35, pp. 601-609.
- [9] K. Gierens, S. Matthes and S. Rohs, "How Well Can Persistent Contrails Be Predicted?," *MDPI Aerospace*, vol. 7, no. 12, 2020.
- [10] Committee for Study of Federal Aviation Administration Air Traffic Controller Staffing, "Federal Aviation Administration's Approach for Determining Future Air Traffic Controller Staffing Needs," *Transportation Research Board*, 2014.
- [11] H. Arneson, P. Hegde, M. La Scola, A. Evans, R. Keller and J. Schade, "NASA New Technology Report for "Sherlock Data Warehouse", Report No. ARC-E-DAA-TN67994," NASA Ames Research Center, Moffett Field, 2019.
- [12] V. Meijer, S. Eastham and S. Barrett, "Using satellite-based observations of contrails to inform contrail avoidance strategies," in *MIT Laboratory for Aviation and the Environment*, 2022.
- [13] National Center for Environmental Information, "Global Ensemble Forecast System (GEFS)," NOAA, [Online]. Available: <https://www.neci.noaa.gov/products/weather-climate-models/global-ensemble-forecast>.