

Contrail Detection on GOES-16 ABI With the OpenContrails Dataset

Joe Yue-Hei Ng¹, Kevin McCloskey¹, Jian Cui¹, Vincent R. Meijer¹, Erica Brand¹, Aaron Sarna¹, Nita Goyal¹, Christopher Van Arsdale, and Scott Geraedts

Abstract—Contrails (condensation trails) are line-shaped ice clouds caused by aircraft and are a substantial contributor to aviation-induced climate change. Contrail avoidance is potentially an inexpensive way to significantly reduce the climate impact of aviation. An automated contrail detection system is an essential tool to develop and evaluate contrail avoidance systems. In this article, we present a human-labeled dataset named *OpenContrails* to train and evaluate contrail detection models based on GOES-16 Advanced Baseline Imager (ABI) data. We propose and evaluate a contrail detection model that incorporates temporal context for improved detection accuracy. The human labeled dataset and the contrail detection outputs are publicly available on Google Cloud Storage at gs://goes_contrails_dataset.

Index Terms—Advanced Baseline Imager (ABI), benchmark datasets, contrails, geostationary satellites, GOES-16, image segmentation, machine learning, neural network.

I. INTRODUCTION

PERSISTENT contrails are cirrus clouds created by aircraft when flying through cold and ice supersaturated regions. Contrails both warm the atmosphere by trapping infrared radiation which would have otherwise escaped into space, and cool it by reflecting incoming solar radiation. Recent works have found that the net effect is warming and comparable to the historical effect of CO₂ emissions from aviation [1], [2]. Studies have suggested that the majority of contrail warming can be attributed to a small fraction of flights, and their warming impact can be significantly reduced by using a different route [3], [4]. However, considerable uncertainties exist both around the overall impact of contrails [1], and how well their formation and persistence can be predicted in advance [5]. More complete knowledge of where contrails are actually forming is key to reducing this uncertainty, and this can be obtained by creating an automated contrail detector with high spatio-temporal coverage.

Detecting contrails from satellite images is challenging because of their visual similarity to natural cirrus. Contrails are formed as line-shaped ice clouds and then slowly deform

over time, becoming difficult to distinguish from natural cirrus. In this work, we focus on detecting young contrails that are still linear in shape. While computer vision has made substantial progress thanks to the rise of large-scale datasets, like ImageNet [6], and the use of neural networks, recognition on domain specific problems like infrared satellite images is still challenging due to the lack of large in-domain datasets for pre-training. Apart from that, unlike traditional object recognition tasks where textures and colors are informative, contrails often look similar to or even indistinguishable from natural cirrus. In addition to color and texture, human experts must take into account the overall shape of the contrails (linear) and their temporal evolution (unlike natural cirrus, contrails appear very quickly and spread out rapidly over time).

Challenging computer vision tasks have often benefited from the existence of high-quality datasets, such as the MNIST dataset [7] for handwriting recognition. Such datasets lower the barrier to entry for different research groups to try to solve the problem, and provide a standard way to compare new models to prior work to assess whether a given new technique has led to improvement. In this work, we provide a public dataset, *OpenContrails*, to facilitate model comparison and reproducible research. Examples of images and labels are shown in Fig. 1.

While previous works have trained contrail detection models with deep convolutional neural networks [8], [9], the performance of those models is still significantly behind that of a human expert, and therefore improvement in model performance is possible. One possible reason why humans can do better is that during labeling, humans can evaluate the temporal behavior of the contrails, whereas existing computer models perform detection based on only one frame. We develop a model that takes multiple input frames into account for temporal context. We show that such a model improves the detection performance when compared to models that only process a single frame.

In this work, we present a dataset of human-labeled contrail detections on the geostationary satellite GOES-16, which provides high spatio-temporal coverage, and we make it publicly available. The proposed *OpenContrails* dataset contains per-pixel human labeled contrail masks for each image in the dataset. We train a neural network for detecting contrails and show that the resulting model produces more accurate contrail detection outputs than the existing models. We run

Manuscript received 24 February 2023; revised 12 August 2023 and 10 November 2023; accepted 18 November 2023. Date of publication 19 December 2023; date of current version 17 January 2024. (Corresponding author: Joe Yue-Hei Ng.)

Joe Yue-Hei Ng, Kevin McCloskey, Jian Cui, Erica Brand, Aaron Sarna, Nita Goyal, Christopher Van Arsdale, and Scott Geraedts are with the Research Department, Google Inc., Mountain View, CA 94043 USA (e-mail: yhng@google.com).

Vincent R. Meijer is with the Laboratory for Aviation and the Environment, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Digital Object Identifier 10.1109/TGRS.2023.3345226




Fig. 1. Example images and labels in the dataset. (Left) False color images from GOES-16 satellite. (Right) Human labeled binary contrail mask.

the contrail detection model on multiple years of available GOES-16 images and confirm previous findings on contrail research, such as contrail coverage patterns and diurnal effects. We hosted a Kaggle contest using an earlier version of the dataset, which attracted over 900 teams participated [10]. We believe that the *OpenContrails* dataset will be a valuable resource for the contrail research community by making large-scale contrail detection results widely available.

A. Related Work

The effect of contrails on climate change has been widely studied in previous literature [1], [2], [11], [12]. Simulations indicate that a small set of flights are responsible for the majority of contrail warming, which suggests that it is possible to adjust a small number of flight routes to substantially reduce aviation's climate impact [3], [4].

Empirical analyses of contrails from satellite imagery in the last two decades [13], [14], [15], [16] were typically accomplished by the algorithm proposed by Mannstein et al. [13]. Mannstein et al. proposed to detect contrails by applying a set of filters and thresholds on brightness temperature imagery.

Zhang et al. [9] trained convolutional neural networks to detect contrails on Himawari-8 satellite images for evaluating contrail coverage. Meijer et al. [8] also trained a neural network for contrail detection on GOES-16 images to examine what effects the reduced flight traffic during the COVID-19 pandemic had on contrails. None of these works made the dataset or contrail detections publicly available, and while Meijer et al. only labeled the contiguous US (CONUS) region, we sample scenes to be labeled from the majority of the GOES-16 full-disk viewing extent.

McCloskey et al. [17] released a contrail dataset on Landsat-8 data. With 100-m pixels in its infrared channels, Landsat-8 images have sufficient spatial resolution to identify even very young contrails. However, due to its low-earth sun-synchronous orbit and mission goals, imagery is only available for scenes that are mostly land, mostly in the daytime, at fixed local times of day and with a 16-day repeat cycle. Thus, Landsat-8 imagery has clear limits for conducting large-scale analysis for contrail research. We construct our dataset on GOES-16 ABI images [18] taken from geostationary orbit, which nominally have 2-km pixels in the infrared channels but have much higher spatio-temporal coverage. Consequently, our dataset and contrail detection model can be used as a foundation for contrail warming impact assessment and validating contrail avoidance experiments in the western hemisphere.

Deep neural networks have been extensively used for remote sensing applications [19], [20]. The task of contrail detection can be cast as image segmentation, on which deep neural networks have shown great success. Here, we employ promising image segmentation techniques including using the ResNet backbone [21] and Deeplabv3+ [22], [23] architectures.

II. DATASET

We build our dataset using GOES-16 Advanced Baseline Imager (ABI) imagery [18], specifically brightness temperatures calculated from the Level 1B radiances using the Planck constants provided by [24]. GOES-16 views the North and South American region with a full-disk image taken every 10 min since April 2019 (15 min interval before April 2019). The nominal pixel size of GOES-16 ABI is 2×2 km at nadir. This relatively coarse resolution means that contrails cannot be seen when they initially form, instead we must wait some time for them to spread out enough to be visible. This is not necessarily a disadvantage: the warming effects of contrails are dominated by contrails which persist for hours [4], so the inability to detect shorter-lived contrails may not hinder our ability to assess and prevent the vast majority of contrail warming.

Thermal infrared bands are more useful than visible bands for recognizing contrails, as visible bands do not provide much information at night. Mannstein et al. [13] suggested to use 10–12 μm bands for contrail detection, as young contrails tend to show higher transmissivity at 10–11 μm than 11–12 μm band due to smaller crystal sizes. Following Kulik [25], we show imagery to human labelers in an “ash” false color scheme that combines three longwave GOES-16 brightness temperatures. The red, green, and blue channels are represented by the difference between 12 and 11 μm , and

difference between 11 and 8 μm , and 12 μm , respectively. This color scheme is chosen to help identify contrails by highlighting ice clouds as darker colors in a lighter background. We have qualitatively found that the ash color scheme helps with spotting more contrails compared to grayscale brightness temperature difference. An example of the ash color scheme and the true color RGB image is shown in Fig. 2. Each image patch in the dataset corresponds to an approximately 500×500 km region and is reprojected to the Universal Transverse Mercator (UTM) coordinate system (with the zone decided by the northwest corner of the region) using bilinear resampling. Labelers are shown image patches with a size of 281×281 . As the labels near the boundary are often noisy due to the lack of spatial context, we center crop to 256×256 pixels on the images and contrail masks for training and evaluation.

We generate the examples by randomly sampling image patches from the GOES-16 viewable extent from April 2019 to April 2020.¹ We aim to increase the spatial coverage of automatic contrail detection to a larger GOES-16 region, so we sample images from both the northern and southern hemispheres. We restrict the sampling of image patches to be between -50° to 50° latitude and -135° to -30° longitude to avoid large viewing angles. We first generate candidate image patches using a non-overlapping sliding window to scan the bounding region of a GOES image every 40 min. Uniform sampling of GOES-16 images would result in very few positive examples, as contrails occur infrequently. To increase the number of positive examples in the dataset, we downsampled regions unlikely to contain contrails. We did this in three ways. First, we obtained aircraft flight tracks obtained from terrestrial ADS-B data licensed for publication from FlightAware, LLC (<https://flightaware.com>). We advected these flight tracks according to European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 [26] wind data, using the Runge–Kutta method [27], for 4 h. We kept only 5% of the images which contained no advected flight tracks, and only 20% of the images which contained less than ten tracks. Second, we calculated the relative humidity over ice from ECMWF ERA5 temperature, pressure, and specific humidity values. Since contrails can only persist if this humidity is above 100% [28], we kept only 5% of images for which no location in the image had relative humidity $>90\%$ for any of the altitudes where contrails might form (7000 – 12000 m). Finally, we applied the Mannstein et al. [13] contrail detection algorithm tuned as in McCloskey et al. [17] to achieve very high recall, as opposed to high precision ([17] showed that a similar method applied to images of contrails in Landsat-8 detected every contrail, but for every contrail it detected it also detected around 20 false positives.) We kept all the images in which the high recall algorithm detected a contrail, as well as 5% of the examples in which no contrails were detected. Note that the FlightAware data does not contain all flights and the ECMWF wind and humidity data are not always accurate. Therefore, the filters are not expected to be perfectly accurate, and it is still necessary




Fig. 2. (Left) True color image. (Right) False color image in ash color scheme.

to have human labelers annotate the images. We randomly subsampled the remaining images for human labeling.

To further boost the number of positives in the dataset, we also included some GOES-16 ABI imagery (in the training set only) at locations in the US where contrails had been identified in collocated Google Street View images of the sky, using an approach similar to Juan et al. [29]. These additional labeled GOES-16 images are only in the training set because Street View cars operate in the daytime and mostly on days with sunnier weather, so it may be easier than usual to identify contrails in the GOES-16 imagery of those locations. We found in our experiment that these additional training data slightly improve the contrail detection model performance on the validation set. Note that there are no Street View images in the dataset or used in models reported here: they were only used to identify contrail-rich GOES-16 scenes for human labeling.

Temporal context is important for recognizing contrails. Similar to the labeling procedure in Meijer et al. [8], we show labelers five images (50 min) before and two images (20 min) after the image being labeled. An example image sequence of contrail evolution is shown in Fig. 3, which shows how contrails appear spontaneously and spread out over time, while natural cirrus often occur for a longer time and may get sharper over time instead of spreading out. The advected flight density is also shown to the labelers to assist labeling. It is generated from FlightAware flight data and is advected using the same method as for screening (ECMWF ERA5 wind vector data, with the Runge–Kutta method). We spread out the advected flight density of flights which are older (at the time the image was taken) to reflect how uncertainty in the wind vector data makes us less certain about where a contrail will ultimately advect to. We ask the labelers to still label the contrails if flights are not shown in the flight density only when they are very confident. An example GOES-16 “ash” image and the corresponding advected flight density is shown in Fig. 4.

We ask human labelers to draw polygons on the images to enclose the contrail pixels in a web-based labeling interface. The labelers are instructed on how to label contrails through a set of slides containing guidelines, including examples of contrails and non-contrail cirrus.

We provide detailed instructions for labelers to identify and label contrails, for example: contrails are typically seen as dark linear objects in the satellite images using the ash color scheme; contrails appear as sharp lines and then spread out over time; and static structures across multiple frames

¹We mistakenly sampled about 9000 of the 20088 examples from October 2019 instead of the full year.




Fig. 3. Example sequence showing evolution of contrails. Contrails appear quickly and start as straight lines. They deform and fade over time slowly from satellite observations. The images correspond to an area centered at coordinates (48.75, -66.50). The first image was taken on July 21, 2018 at 14:10 UTC. The images were taken at 15-min intervals.




Fig. 4. Example GOES-16 false color image with advected flight density shown to labelers.




Fig. 5. Distribution of solar zenith angle in the dataset. Orange and blue bars correspond to daytime and nighttime examples, respectively.

are likely to be ground-based objects such as roads, rivers, or coastlines. We also provide the following additional labeling guidelines.

- 1) Contrail must contain at least 10 pixels to be labeled.
- 2) At some time in their life, contrails must be at least $3 \times$ longer than their width.
- 3) Contrails must either appear suddenly or enter from the sides of the image.
- 4) Contrails should be visible in at least two images.

Labelers were asked to initially label a small set of images, and then authors of this work gave feedback on their initial label quality. The process iterates until a labeler generates labels with satisfactory accuracy. As we observe variance across individual labelers, each GOES-16 scene is labeled by four human labelers, and pixels are considered positive if at least three labelers considered the pixel a contrail.

Even when instructions and trainings are provided, identifying and labeling contrails on satellite images are still challenging. We compared the labeler performance compared to the groundtruth (labeled by authors of this work) on about 200 examples. Individual labelers give about 64% precision




Fig. 6. Spatial distribution according to the image center in the dataset. The top and the right histograms correspond to the marginal distributions along the latitude and longitude axes.

and 76% recall measured per-pixel. The aggregated labels from 4 labelers with majority vote give about 80% precision and 78% recall.

The full dataset contains 20088 examples in the training set and 1879 examples in the validation set. The examples are partitioned by pre-defined time ranges so that the time difference between any training and validation examples is at least three hours. Satellite scenes that were identified as likely to have contrails by Google Street View are only included in the training set. 9103 of the training examples contain at least one annotated contrail. About 1.2% of the pixels in the training set are labeled as contrails. The dataset contains a wide variety of times and locations, as shown in Figs. 5 and 6. The examples are not uniformly distributed in space and time as the images are sampled to include more contrail examples as described above.

Contrails occur more often in cloudy scenes. We estimate the cloud cover fraction using the GOES-16 ABI L2 Cloud Top Phase product and regard all non-clear-sky pixels as clouds. We show the example counts in our dataset by the cloud coverage fraction in Fig. 7. Our dataset contains more cloudy scenes than clear sky images, and the positive examples are

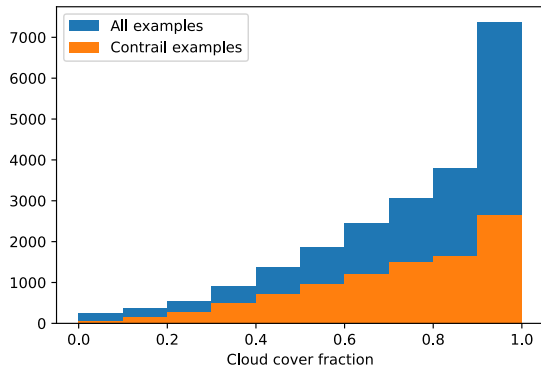


Fig. 7. Statistics of cloud coverage in the dataset.

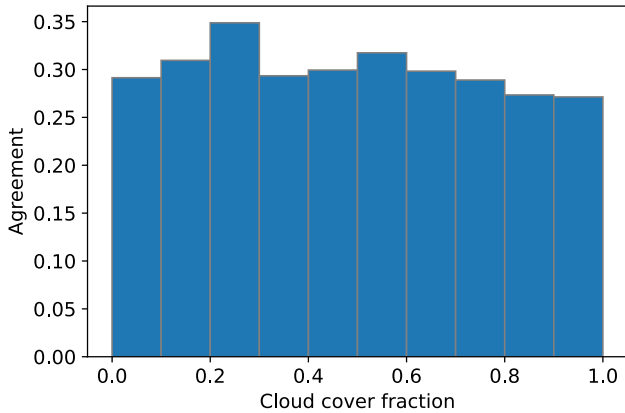


Fig. 8. Labeler agreement at different cloud cover fractions.

roughly proportionally distributed at different cloud coverage fractions. It is possible that contraails are harder to be identified in cloudy scenes. Fig. 8 shows the labeler agreement at different cloud cover fractions in the scenes, where agreement is defined as the number of contrail pixels labeled by the majority of the labelers divided by the number of contrail pixels labeled by at least one labelers. Labelers agree slightly less in scenes with higher cloud cover fractions.

The dataset is provided in TFRecord format. Each example includes an image with a size of 256×256 . Each example is provided with multiple GOES-16 ABI brightness temperatures and brightness temperature differences. The dataset is uploaded to Google Cloud bucket and made publicly available at `gs://goes_contraails_dataset`.

A. Evaluation Metrics

The best evaluation metric for contrail detection depends on the downstream application. A per-pixel metric would be useful for assessing the energy forcing of contraails as it captures the area covered by contraails, and a per-contrail metric would be more useful for associating the contraails to flights and computing the contrail lengths. Here, we suggest two evaluation metrics for contrail detection.

Contrail detection can be considered as an image segmentation task which is traditionally evaluated with per-pixel metrics. Typical semantic segmentation metrics like intersection over union (IoU) and DICE coefficient are sensitive to the choice of a single binarization threshold. Since the detection

outputs could be useful at different operating points for different applications, we evaluate the per-pixel performance using area under the precision recall curve (AUC-PR). We use 10 000 thresholds sampled uniformly between 0 and 1 to compute the precision and recall for each pixel to obtain the AUC-PR.

We also evaluate contrail detections as a line segment detection task. The contrail detection system should output a set of line segments for each image. Similar to [17], we evaluate the resulting linear contraails with a precision recall curve. The contrail detected is considered a correct detection if its orientation angle is within 10° from the groundtruth and within 10 km (approximately 5 pixels) mean distance to the contrail [17].

III. MODELS

We use a convolutional neural network model for contrail detection. The model outputs a score (between 0 to 1) for each pixel indicating the confidence that pixel is part of a contrail. The per-pixel contrail detection outputs are then postprocessed to produce line segments representing the contrail instances.

To consistently identify contraails including at night, we train on infrared channels represented as brightness temperatures as input. Our contrail detection model takes six channels as input: 8, 10, 11, 12 μm , difference between 12 and 11 μm , and difference between 11 and 8 μm . Each input channel is standardized by subtracting the per-channel global mean and dividing by the per-channel global standard deviation of the entire training set before feeding it into the network.

A. Single-Frame Model

We first train a binary image segmentation model for classifying each pixel as contrail or background. We employ the semantic segmentation architecture Deeplabv3+ [23] and experimented with different image backbones. Following [23], we use a dilated ResNet [22] by modifying the last residual block with dilated convolutions, resulting in an output stride of 16 in the output feature maps. An Atrous Spatial Pyramid Pooling (ASPP) module [23] follows to incorporate information at multiple scales. Low-level feature maps are then combined with upsampled higher level feature maps to preserve the resolution of the feature maps. The final segmentation head consists of three convolutional layers and a two-channel output representing “background” or “contrail.” The model is trained with a per-pixel cross-entropy loss.

The performance of the original ResNet, which has been used previously for contrail detection model [8], can be improved by various recent techniques [30]. We follow the best practice documented in [30] and use an improved architecture with ResNet-D stem [31] and Squeeze-and-Excitation blocks [32]. We set the squeeze-and-excitation ratio to 0.25. Stochastic depth [33] was used to regularize the network with 0.2 drop rate.

The output resolution of the segmentation model can be smaller than the input images by a factor of the output stride of the backbone. We apply bilinear upsampling to the network prediction to resize the outputs to the groundtruth image resolution during both training and evaluation.




Fig. 9. Network architecture for multiframe detection model. The model takes a sequence of frames as inputs, and take a slice of the feature maps that correspond to a single timestep, and then run the decoder.

B. Multiframe Model

To incorporate temporal context into our model, we extend the Deeplabv3+ model to incorporate multiple input frames. We use a spatio-temporal encoder (based on 3-D convolutions) to incorporate temporal information.

Inspired by the wide success of convolutional networks in video classification, we use the inflated 3-D (I3D) convnet as the backbone of our detection model. Our inflated convnet is based on the ResNet architecture [21]. Since 3-D convolution is computationally expensive, we implement the inflated ResNet block by factoring spatial and temporal convolutions similar to [34], [35]. A regular ResNet block consists of three convolution layers with 1×1 , $k \times k$, and 1×1 kernel sizes. We inflate each ResNet block by expanding the first convolutional layer with a temporal component, while keeping the second and third convolutional layers with a temporal kernel size of 1. The inflated ResNet block thus has kernel sizes of $t \times 1 \times 1$, $1 \times k \times k$, and $1 \times 1 \times 1$. In our implementation, we apply the temporal convolution with a kernel size of 3, at every ResNet block in stages 3 and 4.

Differently from video classification (where only a single label is returned for each example), fine-grained spatio-temporal information is required for precisely predicting the contrail detection masks. We augment the backbone network to have a temporal stride of 1 in the first max pooling and convolution layers in the stem. Thus, the temporal resolution is fixed throughout the backbone.

Since the dataset only has one labeled frame available for each example, we slice one frame of features to pass to the decoder. Similar to the single-frame model, we use the ASPP module for the decoder and stacked convolutional layers for classifier. The overall model architecture is illustrated in Fig. 9.

C. Training

We train the network using stochastic gradient descent with 10 000 iterations, using a batch size of 256 and a momentum of 0.9. All models are implemented in TensorFlow and trained on 16 Tensor Processing Unit V3 (TPUv3) chips. The learning rate is set to 0.2 and gradually decays to zero with half-cycle cosine learning rate decay. To improve training stability, we employ a linear warm-up learning rate schedule, increasing the learning rate linearly from 0 for the first 500 iterations. L2 weight decay is used with a factor of 0.0001. Input images

are resized to a resolution of 512×512 before being fed into the network for both training and evaluation, unless otherwise specified. Data augmentation is applied during training to reduce overfitting: a scale factor is chosen uniformly randomly at each iteration, and then the image is scaled according to the scale factor, and a crop of the scaled image is used for training. We found this significantly helps overall performance.

As our detection dataset is relatively small compared to natural image datasets like ImageNet [6], pre-training the neural network on larger datasets can often help—even pre-training on images which are very different from satellite images. We initialize the network with ImageNet pre-trained weights for the single-frame model. We have also experimented with weights pre-trained on image segmentation datasets like COCO but did not see improvement. For the multiframe model, we pre-train the backbone model on the Kinetics-400 dataset [36] which consists of about 240 000 videos collected from YouTube: each video clip is 10 s long and is annotated with an action from 400 classes. Because our model takes six input channels (instead of three used by the pre-trained network), we average the convolutional filters along the channel dimension in the first layer, and then replicate them to the six channels during initialization.

D. Converting Binary Masks to Line Segments

The network outputs corresponding to pixelwise contrail probabilities are further postprocessed to produce the line segment outputs. We first choose a threshold to convert the pixel probabilities to binary masks, and then use off-the-shelf OpenCV’s LineSegmentDetector, which is an implementation of [37], to detect line segments from the binary masks. We then iteratively merge line segments if one segment is part of (or a continuation of) another line segment. We found that this procedure works similarly well compared to a more complicated algorithm which convolves multiple line kernels at multiple scales as described in [13] and [17], which requires tuning about ten hyperparameters for the best performance.

IV. RESULTS

A. Main Results

We compare the single-frame models and multiframe model with different backbones. The results are summarized in

TABLE I
DIFFERENT ARCHITECTURES FOR CONTRAIL DETECTION

Backbone	Pixel AUC-PR
ResNet-50	68.0
ResNet-101	68.3
Dilated ResNet-101	69.6
ResNet3D-101	71.2




Fig. 10. Per-contrail precision/recall curves, showing better precision and recall for the multiframe model compared to the best single-frame model we explored.

TABLE II
CONTRAIL DETECTION (MULTIFRAME MODEL) WITH DIFFERENT INPUT RESOLUTIONS

Input resolution	Single-frame Model Pixel AUC-PR	Multi-frame Model Pixel AUC-PR
256x256	61.2	63.4
384x384	67.2	69.4
512x512	69.6	71.2
784x784	69.9	-
1024x1024	70.3	-

TABLE III
DETECTION PERFORMANCE WITH DIFFERENT NUMBERS OF INPUT FRAMES

Size of temporal context	Pixel AUC-PR
$T_{\text{before}} = 0, T_{\text{after}} = 0$	69.8
$T_{\text{before}} = 0, T_{\text{after}} = 1$	70.7
$T_{\text{before}} = 1, T_{\text{after}} = 0$	70.3
$T_{\text{before}} = 0, T_{\text{after}} = 3$	70.6
$T_{\text{before}} = 2, T_{\text{after}} = 1$	70.6
$T_{\text{before}} = 3, T_{\text{after}} = 0$	71.2
$T_{\text{before}} = 4, T_{\text{after}} = 3$	70.5

Table I. All models achieve reasonable detection performance, showing that our dataset is sufficiently large for training contrail detection models. The multiframe model slightly outperforms the single-frame-based models, showing the model is able to use temporal context to improve detection.

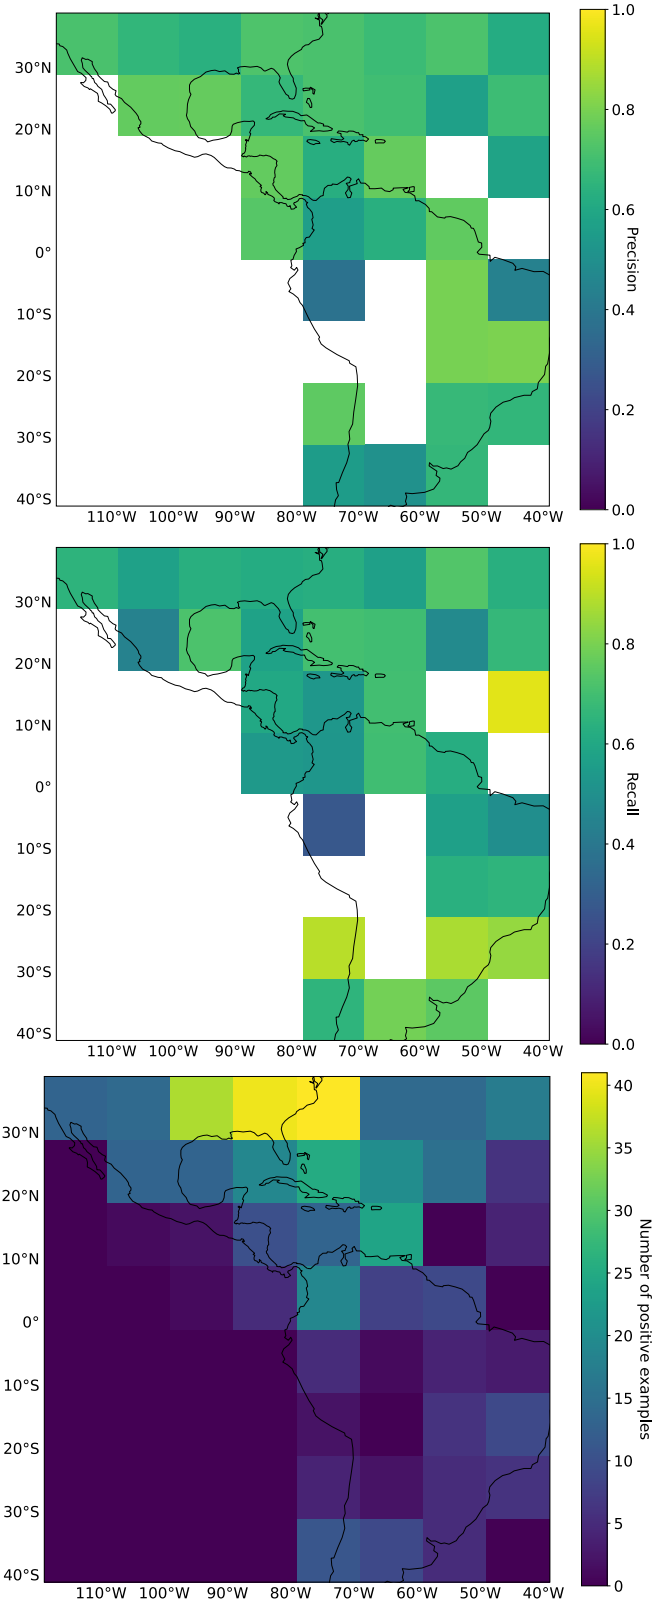


Fig. 11. (Top) Detector precision, (middle) recall, and (bottom) number of positive examples at different spatial locations.

Linear Contrail Detection: To assess model accuracy on linear contrail objects, we linearize the output prediction mask as discussed in Section III-D. Before each linearization, we vary




Fig. 12. (Left) Detector precision, (middle) recall, and (right) number of examples at different solar zenith angles on the validation set.




Fig. 13. (Left) Detector precision, (middle) recall, and (right) number of examples at different cloud cover fractions on the validation set.

the thresholding value for the model output mask to generate multiple binary contrail masks at different confidence levels. Fig. 10 shows the precision recall curves from this process. Once again our multiframe model shows better performance compared to the single-frame model.

McCloskey et al. [17] provided a similar precision recall curve on the contrail dataset based on the Landsat-8 satellite with the baseline Mannstein et al. algorithm [13]. It is worth noting that our precision presented here is significantly higher. At 60% recall, The Mannstein et al. algorithm gave 15% precision on the Landsat-8 dataset, while our model gives more than 70% precision on our GOES-16-based dataset. While the evaluation datasets and models are different, this suggests that our model detects contrails more accurately for contrail research.

B. Detailed Analyses

1) *Input Image Size*: Using a larger input image size often improves performance on visual recognition tasks [38], [39]. This is because the network can use more computation on smaller patches, and the output resolution is increased. We perform bilinear resampling to enlarge the input image before feeding it to the network, during both training and evaluation. The output predictions are resampled to match the groundtruth resolution when computing the AUC-PR. Table II shows the results with different resolutions. We see that larger input image sizes give slightly better results, at the cost of slower inference and more memory usage during training.

2) *Number of Input Frames*: Using multiple input frames improves the performance by providing more temporal context. The results are summarized in Table III. In general, we observe that the temporal context before the labeled frames is more important than after the labeled frames. This

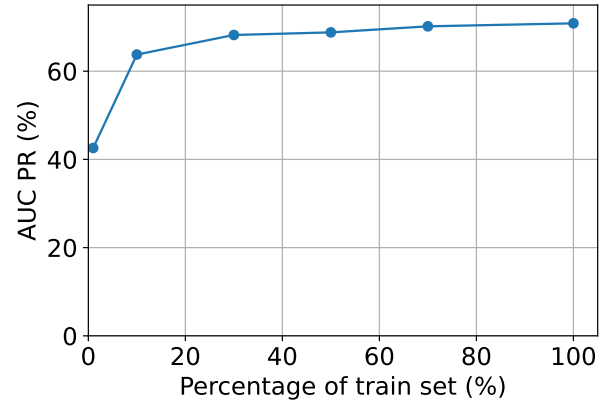


Fig. 14. Validation performance (pixelwise AUC-PR) when the multiframe model is trained with different numbers of labels.

TABLE IV
MODEL COMPARISON WITH MEIJER ET AL.'S DETECTION [8] ON DIFFERENT DATASETS. THE RELAXED PRECISION AND RECALL ($\rho = 2$) OF A FIXED THRESHOLD ARE REPORTED

	Relaxed Precision	Relaxed Recall
Performance on Meijer et al's dataset [8]		
Meijer et al [8]	71.6	57.2
Ours	74.5	59.7
Performance on our dataset		
Meijer et al [8]	64.7	50.7
Ours	92.1	74.6

matches our intuition as contrails dissipate over time and become harder to recognize. Using frames after the labeled frames shows diminished model performance. This could be an artifact of our current model architecture; we plan to further experiment with other architectures to validate this in the future.




Fig. 15. Example detection outputs from our model. Columns from left to right: 1) false color images from GOES; 2) groundtruth contrail labels; 3) pixel outputs from our model; and 4) binarized detection outputs (with threshold 0.35) and the red line segments correspond to the linearization outputs from the line segment detector.

3) *Performance Across Spatial Locations:* We visualize the precision and recall (with 0.4 detection threshold) at different spatial locations in Fig. 11 for each 10×10 degree gridbox that contains at least three positive examples in our validation set. The precision and recall estimates are noisier in South America due to smaller number of positive examples. In general, we observe no substantial spatial bias in different

spatial locations, except with slightly worse performance in the northwest part of South America.

4) *Performance at Different Solar Zenith Angles:* The solar zenith angle is how many degrees away from “directly overhead” the sun is at any given moment; greater than 90° indicates night scenes. Fig. 12 shows the precision and recall of the detector at different solar zenith angles.

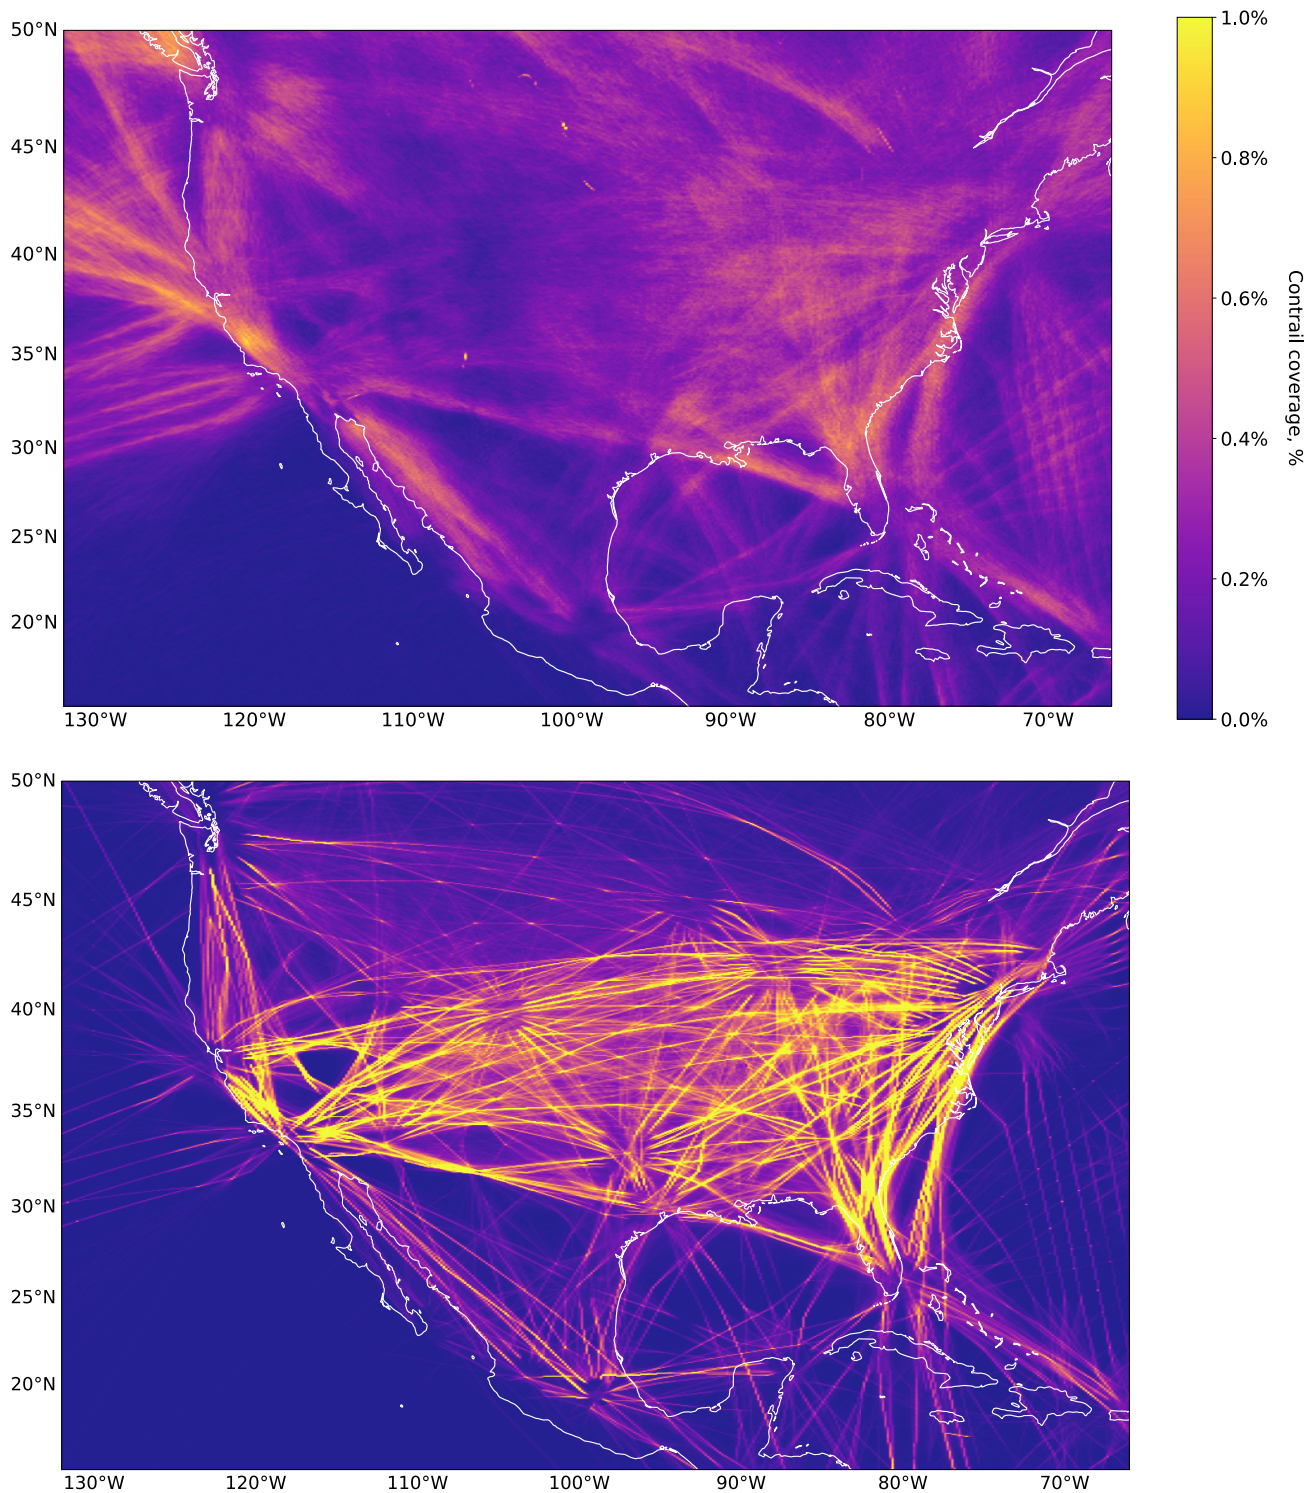


Fig. 16. (Top) Averaged detected contrail mask, 2018–2019. (Bottom) Flight waypoint density above 7000 m, 2018–2019.

We observe relatively uniform performance across different solar zenith angles, except in the zenith angle bucket with small number of examples which may lead to noisy results. The example distribution is different from Fig. 5 because the training set includes GOES-16 images that were sourced by virtue of having a contrail in a collocated Street View picture of the sky, and those are only daytime images.

5) *Performance at Different Cloud Cover Fractions:* Contrails in cloudy scenes could be more difficult to recognize. We show the precision and recall (with 0.4 detection threshold) in Fig. 13. The precision is relatively similar across different cloud cover levels, while the recall slightly drops in higher cloud cover levels.

6) *Dataset Size:* Larger training datasets often improve the quality of machine learning models. Fig. 14 shows the effect of

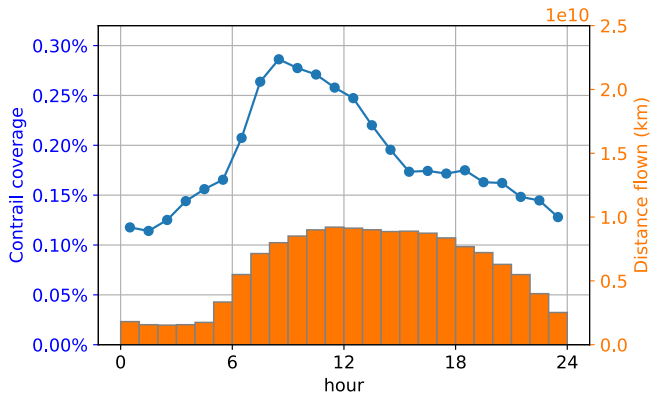


Fig. 17. Linear contrail coverage and flight density by the time of day. The blue curve shows the linear contrail coverage derived from our contrail detector and the orange bar shows the flight densities.

training with different amounts of training data. More training data improves the detection performance, but diminishing returns are observed at the current scale. This suggests that it may be more important at this point to improve the model architectures, input features, or learning algorithms than it is to collect more labeled data.

C. Comparison of Single-Frame Model With Meijer et al. [8]

Meijer et al. [8] have also trained a contrail detection model on GOES-16 images. Their model is based on a U-Net [40] architecture with ResNet-18 backbone, and they trained the model on their dataset which consists of 103 labeled images covering the contiguous US (CONUS) region, with each image having 3000×2000 pixels. We obtained the evaluation set and detection results from [8] and here we provide an inter-comparison of both models on both datasets. We report on the images from the year 2018 in the evaluation set from [8], which contains human labels on the CONUS region for 12 different GOES-16 scans spread throughout the year. We have verified that there is no overlap between the evaluation set and OpenContrail’s training set. Since their dataset uses GOES-16 images from 2018 where full-disk images are taken every 15 min instead of 10 min, we evaluate our single-frame model on this dataset instead of the multiframe model. When evaluating on OpenContraails, we only use examples located in the CONUS region because Meijer et al.’s detection only covers the CONUS region.

Due to a discrepancy in our respective implementations of image projection, during initial comparisons, we observed a slight misalignment between the Meijer et al. human labels when compared to our projections of the GOES-16 images. Therefore, we use a “relaxed” version of precision and recall as the metric, similar to [41]. We consider that the relaxed precision represents the fraction of the predicted contrails that are within ρ pixels of the true contrails, and the relaxed recall represents the fraction of true contrails that are within ρ pixels of the predicted contrails. We use $\rho = 2$ in our evaluation.

We compare the detection results without model retraining. We select a threshold value of 0.35 to generate binary detection masks. The results are summarized in Table IV. Our model gives better accuracy than [8] on our dataset,

which is unsurprising given our model is trained on data with a similar distribution. Despite the distribution shift of labels from our training set to the labels from [8], our model slightly outperforms their model trained on their dataset. This highlights the good quality and the degree of generalization of our contrail detection model trained on our new dataset.

D. Visualization

Fig. 15 shows some qualitative results of our detections. We observe good visual agreement between the model detections and human label masks, and the line segment detector appears able to convert the detection mask to linear segments reasonably well. The model sometimes misses some detections when there are many contrails in the scene. It is also clear from qualitative inspections of multitemporal satellite imagery (data not shown) that the model generally does not assign much (if any) probability to nonlinear anthropogenic contrail cirrus clouds that evolved from linear contrails. This is unsurprising given the instructions we have provided to labelers (to only mark linear contrails), but it has an important implication for using the detection model as the basis of contrail warming assessments: we expect the warming from contrail pixels detected by the model to comprise a lower bound on the actual contrail warming impact.

V. CONTRAIL COVERAGE

To compute contrail coverage visuals and statistics, we run our single-frame contrail detection model as an overlapping sliding window, and then stitch the detection outputs to obtain the combined detection masks for the CONUS region in the year 2018–2019. To give an unbiased contrail coverage estimate, following Meijer et al. [8] we select a per-pixel detection threshold (for our model, 0.4) which gives a precision rate and recall rate on our validation set that are approximately equal. We show the average contrail coverage in 2018–2019 in the CONUS region and the corresponding flight density (computed from FlightAware data) in Fig. 16. The contrails cluster along the high traffic flight routes, which confirms the pattern from [8]. The average linear contrail coverage in the considered domain is 0.19% which is similar to 0.17% reported in [8].

Fig. 17 shows our estimated linear contrail coverage and flight density by the time of day in 2019. The local hour of the contrails is approximated by $\text{hour}_{\text{UTC}} + \text{longitude}/360 \times 24$. We observe a similar diurnal pattern as in the previous work [8]: the contrail coverage is lowest around midnight with about 0.1% coverage. It peaks at around 8 A.M. with about 0.28% contrail coverage (as the distance flown goes up) but then drops in the afternoon despite the flight distance remaining high.

Fig. 18 shows the contrail coverage versus flight distance by the day of year from 2018 to 2021. We observe similar trends to those observed by Meijer et al. [8]. Contrail coverage drops as the flight distance drops significantly in April 2020, due to COVID-19. We also observe a seasonal effect of contrail coverage, where fewer contrails are observed consistently in the summer months across multiple years, regardless of the flight distance flown.




Fig. 18. Linear contrail coverage from our single-frame model with flight distance by the day of year from 2018 to 2021. 15-day rolling means are displayed.

VI. DISCUSSION

While our dataset provides high-quality contrail labels on GOES-16 satellite images, some contrails may not be observable from geostationary satellites and thus cannot be labeled by human raters. For example, young contrails may already be deformed (or dissipated) before they are visible from geostationary orbit. Also, contrails forming below existing optically thick cirrus are unlikely to be recognizable by human labelers; however, it has been suggested that those contrails do not generate a significant amount of warming [42]. Further research is needed to empirically determine the fraction of contrails and their corresponding warming which can be detected in geostationary satellite images.

Our current system detects contrails at each target frame independently, and does not provide correspondence of contrails between difference frames. Such correspondence could be useful for identifying the contrail age, and improving flight attribution (i.e., which flight caused the contrail). It is also currently estimated that a large fraction of contrail warming happens after contrails are no longer linear [42]. Our proposed models and labels only include linear-shaped contrails, and do not identify deformed contrail-cirrus clouds which are no longer distinguishable from natural cirrus without tracing back to when they are formed. To fully assess the warming impact of contrails, it may be crucial to track contrail pixels for an extended period of time once they are detected. Such a contrail tracking system can be built based on the linear contrail models demonstrated in this work.

Due to the high cost of pixelwise labeling of contrails, our proposed dataset is relatively small compared to large-scale natural image datasets like ImageNet (20k versus 1M examples). Large-scale labeled datasets for geostationary satellite images are unfortunately not currently available for model pre-training. Therefore, our proposed models are initialized on checkpoints trained on natural image or video datasets such

as ImageNet or Kinetics, and their benefits may be limited by domain mismatch between internet images and satellite images. Self-supervised and semi-supervised machine learning techniques that leverage unlabeled data have been shown to improve machine learning model performance for remote sensing segmentation tasks [43], and should be explored here in the future.

Our work here can be extended to cover other geostationary satellites. We plan to collect labels on other geostationary satellites (e.g., Himawari-8 [44] and Meteosat-11 [45]) to cover other regions including Europe and Asia-Pacific. We believe that models trained on the proposed OpenContrails dataset based on GOES-16 can be used for transfer learning to obtain high-quality contrail detectors on other satellites with fewer human labels.

VII. CONCLUSION

We present here a large human-labeled contrail detection benchmark dataset based on GOES-16 ABI imagery. Using it, we explore several contrail detection model architectures, and achieve a new state of the art in linear contrail detection on GOES-16 ABI imagery. We show that leveraging temporal context improves contrail detection, and confirm previously reported seasonal and diurnal contrail coverage trends. The contrail detection dataset and the multiframe detection model outputs are made publicly available on Google Cloud Storage at gs://goes_contrails_dataset. We believe that these contributions will advance contrail empirical studies including contrail detection, climate impact assessment, and mitigation.

ACKNOWLEDGMENT

The authors gratefully acknowledge: Rachel Soh, Jessica Ferguson, Jeanie Pearson, Ladislav Honsa, and Rob von Behren for assistance acquiring and processing flight path data;

Nathan Kiner for assistance setting up the contrail annotation pipeline; the Google Research data operations team for labeling the GOES-16 imagery; the Kaggle team for hosting the Kaggle contest; and the Kaggle contest participants for identifying issues and providing valuable feedback on the dataset.

DECLARATION OF CONFLICTS OF INTEREST

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Some authors are employees of Google Inc. as noted in their author affiliations. Google is a technology company that sells machine learning services as part of its business.

REFERENCES

- [1] D. S. Lee et al., “The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018,” *Atmos. Environ.*, vol. 244, Jan. 2021, Art. no. 117834.
- [2] M. Bickel, M. Ponater, L. Bock, U. Burkhardt, and S. Reineke, “Estimating the effective radiative forcing of contrail cirrus,” *J. Climate*, vol. 33, no. 5, pp. 1991–2005, Mar. 2020.
- [3] 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,” *Transp. Res. Interdiscipl. Perspect.*, vol. 2, Sep. 2019, Art. no. 100033.
- [4] R. Teoh, U. Schumann, A. Majumdar, and M. E. J. Stettler, “Mitigating the climate forcing of aircraft contrails by small-scale diversions and technology adoption,” *Environ. Sci. Technol.*, vol. 54, no. 5, pp. 2941–2950, Mar. 2020.
- [5] K. Gierens, S. Matthes, and S. Rohs, “How well can persistent contrails be predicted?” *Aerospace*, vol. 7, no. 12, p. 169, Dec. 2020. [Online]. Available: <https://www.mdpi.com/2226-4310/7/12/169>
- [6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [7] L. Deng, “The MNIST database of handwritten digit images for machine learning research [best of the web],” *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 141–142, Nov. 2012.
- [8] V. R. Meijer et al., “Contrail coverage over the United States before and during the COVID-19 pandemic,” *Environ. Res. Lett.*, vol. 17, no. 3, Mar. 2022, Art. no. 034039.
- [9] G. Zhang, J. Zhang, and J. Shang, “Contrail recognition with convolutional neural network and contrail parameterizations evaluation,” *SOLA*, vol. 14, pp. 132–137, Sep. 2018.
- [10] Kaggle. (2023). *Google Research—Identify Contrails to Reduce Global Warming*. [Online]. Available: <https://www.kaggle.com/competitions/google-research-identify-contrails-reduce-global-warming>
- [11] C.-C. Chen and A. Gettelman, “Simulated radiative forcing from contrails and contrail cirrus,” *Atmos. Chem. Phys.*, vol. 13, no. 24, pp. 12525–12536, Dec. 2013.
- [12] U. Schumann, “A contrail cirrus prediction model,” *Geoscientific Model Develop.*, vol. 5, no. 3, pp. 543–580, May 2012.
- [13] H. Mannstein, R. Meyer, and P. Wendling, “Operational detection of contrails from NOAA-AVHRR-data,” *Int. J. Remote Sens.*, vol. 20, no. 8, pp. 1641–1660, Jan. 1999.
- [14] R. Meyer, H. Mannstein, R. Meerkötter, U. Schumann, and P. Wendling, “Regional radiative forcing by line-shaped contrails derived from satellite data,” *J. Geophys. Res., Atmos.*, vol. 107, no. D10, p. ACL-17, 2002.
- [15] M. Vázquez-Navarro, H. Mannstein, and S. Kox, “Contrail life cycle and properties from 1 year of MSG/SEVIRI rapid-scan images,” *Atmos. Chem. Phys.*, vol. 15, no. 15, pp. 8739–8749, Aug. 2015.
- [16] D. P. Duda et al., “Northern hemisphere contrail properties derived from Terra and Aqua MODIS data for 2006 and 2012,” *Atmos. Chem. Phys.*, vol. 19, no. 8, pp. 5313–5330, Apr. 2019.
- [17] K. McCloskey et al., “A human-labeled Landsat-8 contrails dataset,” in *Proc. Int. Conf. Mach. Learn. Workshop Tackling Climate Change With Mach. Learn.*, 2021. [Online]. Available: <https://www.climatechange.ai/papers/icml2021/2>
- [18] T. J. Schmit, P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebar, “A closer look at the ABI on the GOES-R series,” *Bull. Amer. Meteorol. Soc.*, vol. 98, no. 4, pp. 681–698, Apr. 2017.
- [19] L. Ma, Y. Liu, X. Zhang, Y. Ye, G. Yin, and B. A. Johnson, “Deep learning in remote sensing applications: A meta-analysis and review,” *ISPRS J. Photogramm. Remote Sens.*, vol. 152, pp. 166–177, Jun. 2019.
- [20] X. X. Zhu et al., “Deep learning in remote sensing: A comprehensive review and list of resources,” *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [22] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, “Rethinking atrous convolution for semantic image segmentation,” 2017, *arXiv:1706.05587*.
- [23] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 801–818.
- [24] GOES-R. (2019). *GOES-R Series Product Definition and Users’ Guide*. [Online]. Available: <https://www.goes-r.gov/products/docs/PUG-L2+vol5.pdf>
- [25] L. Kulik, “Satellite-based detection of contrails using deep learning,” Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, USA, 2019.
- [26] H. Hersbach et al., “The ERA5 global reanalysis,” *Quart. J. Roy. Meteorol. Soc.*, vol. 146, no. 730, pp. 1999–2049, Jul. 2020.
- [27] P. Bogacki and L. F. Shampine, “A 3(2) pair of Runge–Kutta formulas,” *Appl. Math. Lett.*, vol. 2, no. 4, pp. 321–325, 1989.
- [28] K. Gierens, P. Spichtinger, and U. Schumann, *Atmospheric Physics: Background—Methods—Trends*. Heidelberg, Germany: Springer, 2012.
- [29] D.-C. Juan et al., “Graph-RISE: Graph-regularized image semantic embedding,” 2019, *arXiv:1902.10814*.
- [30] I. Bello et al., “Revisiting resnets: Improved training and scaling strategies,” in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 22614–22627.
- [31] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, and M. Li, “Bag of tricks for image classification with convolutional neural networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 558–567.
- [32] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7132–7141.
- [33] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger, “Deep networks with stochastic depth,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Cham, Switzerland: Springer, 2016, pp. 646–661.
- [34] X. Wang, R. Girshick, A. Gupta, and K. He, “Non-local neural networks,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7794–7803.
- [35] Z. Qiu, T. Yao, and T. Mei, “Learning spatio-temporal representation with pseudo-3D residual networks,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5534–5542.
- [36] W. Kay et al., “The kinetics human action video dataset,” 2017, *arXiv:1705.06950*.
- [37] R. G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall, “LSD: A line segment detector,” *Image Process. Line*, vol. 2, pp. 35–55, Mar. 2012.
- [38] M. Tan and Q. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6105–6114.
- [39] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2999–3007.
- [40] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Proc. 18th Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*, Munich, Germany. Cham, Switzerland: Springer, Oct. 2015, pp. 234–241.
- [41] V. Mnih and G. E. Hinton, “Learning to detect roads in high-resolution aerial images,” in *Proc. Eur. Conf. Comput. Vis. (ECCV)*. Berlin, Germany: Springer, 2010, pp. 210–223.
- [42] R. Teoh et al., “Aviation contrail climate effects in the North Atlantic from 2016–2021,” *Atmos. Chem. Phys. Discuss.*, vol. 22, no. 16, pp. 1–27, 2022.
- [43] C. Patel, S. Sharma, V. J. Pasquarella, and V. Gulshan, “Evaluating self and semi-supervised methods for remote sensing segmentation tasks,” 2021, *arXiv:2111.10079*.

- [44] K. Bessho et al., "An introduction to Himawari-8/9—Japan's new-generation geostationary meteorological satellites," *J. Meteorol. Soc. Japan. Ser. II*, vol. 94, no. 2, pp. 151–183, 2016.
- [45] J. Schmetz et al., "An introduction," *Bull. Amer. Meteorol. Soc.*, vol. 83, no. 7, pp. 977–992, 2002.

Joe Yue-Hei Ng received the B.Sc. degree in computer science from the Chinese University of Hong Kong, Hong Kong, in 2011, and the Ph.D. degree in computer science from the University of Maryland, College Park, MD, USA, in 2018.

He has been with Google Research, Mountain View, CA, USA, as a Software Engineer, since 2018. His research interests include computer vision research and applying machine learning for climate solutions.



Kevin McCloskey received the B.S. degree in computer science from Carnegie Mellon University, Pittsburgh, PA, USA, in 2006.

He has been a Software Engineer with Google Inc., Mountain View, CA, USA, since 2006. His research interests include applying machine learning for accelerating the natural sciences and climate change solutions.

Jian Cui received the Ph.D. degree in computer science from Purdue University, West Lafayette, IN, USA, in 2015.

He worked at Google Inc., Mountain View, CA, USA, as a Software Engineer, from 2015 to 2022, where he worked on building content-based indexes with computer vision models on Google Street View images.

Vincent R. Meijer received the B.Sc. degree in aerospace engineering from the Delft University of Technology, Delft, The Netherlands, in 2018, and the M.Sc. degree in mathematical modeling and scientific computing from the University of Oxford, Oxford, U.K., in 2019. He is currently pursuing the Ph.D. degree in aeronautics and astronautics with the Massachusetts Institute of Technology, Cambridge, MA, USA, with a focus on the climate impact of aviation.



Erica Brand received the S.B. degree in life sciences from MIT, Cambridge, MA, USA, in 1990, and the M.S. degree in developmental biology from Stanford University, Stanford, CA, USA, in 1996.

She has worked as a CAD Engineer and a Program Manager across many disciplines including semiconductors, medical applications of AI, software development, security, and climate change mitigation. She has been a Program Manager at Google Inc., Mountain View, CA, USA, since 2013, and handles strategy and operations for Google Research's Applied Science Division.

Aaron Sarna received the B.A. degree in computer science from Cornell University, Ithaca, NY, USA, in 2011.

He has been a Software Engineer with Google Inc., Cambridge, CA, USA, since 2011. He has worked in various fields, including machine learning, computer vision, user interface design, and climate and sustainability research.

Nita Goyal received the B.Tech. degree in CS from IIT Kanpur, Kanpur, India, in 1987, and the Ph.D. degree in computer science from Stanford University, Stanford, CA, USA, in 1993.

She works at Google Research applying AI to climate solutions. She has worked at Google Inc., Mountain View, CA, USA, since 2013 across research and products in natural language, search, and maps. Prior to Google, she worked at several tech startups. She is the first woman CS graduate from any of the IITs.



Christopher Van Arsdale received the B.S. degree in physics and computer science from the University of Washington, Seattle, WA, USA, in 2004.

He leads Google Research's Climate and Energy group. He has worked on numerous projects over two stints at Google Inc., Mountain View, CA, USA, and has a background in large-scale infrastructure, remote sensing, and information retrieval. He has previously worked in venture capital at Innovation Endeavors before focusing on Climate Change research and has served as an advisor to many startups in the Bay Area.

Scott Geraedts received the B.Sc. degree in physics from McMaster University, Hamilton, Canada, in 2010, and the Ph.D. degree in physics from the California Institute of Technology, Pasadena, CA, USA, in 2015.

He has been a Software Engineer and Researcher with Google Inc., Mountain View, CA, USA, since 2017. His research interests include software engineering and machine learning to solve impactful science problems.