SANDIA REPORT
SAND2021-11996
Printed September 2021

Sandia National Laboratories

Local limits of detection for anthropogenic aerosol-cloud interactions

Lyndsay Shand, Andrea Staid, Kelsie M. Larson, Lekha Patel, Erika L. Roesler, Don Lyons, Skyler Gray, James Hickey, Katherine M. Simonson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods

ABSTRACT

Ship tracks are quasi-linear cloud patterns produced from the interaction of ship emissions with low boundary layer clouds. They are visible throughout the diurnal cycle in satellite images from space-borne assets like the Advanced Baseline Imagers (ABI) aboard the National Oceanic and Atmospheric Administration Geostationary Operational Environmental Satellites (GOES-R). However, complex atmospheric dynamics often make it difficult to identify and characterize the formation and evolution of tracks. Ship tracks have the potential to increase a cloud's albedo and reduce the impact of global warming. Thus, it is important to study these patterns to better understand the complex atmospheric interactions between aerosols and clouds to improve our climate models, and examine the efficacy of climate interventions, such as marine cloud brightening. Over the course of this 3-year project, we have developed novel data-driven techniques that advance our ability to assess the effects of ship emissions on marine environments and the risks of future marine cloud brightening efforts. The three main innovative technical contributions we will document here are a method to track aerosol injections using optical flow, a stochastic simulation model for track formations and an automated detection algorithm for efficient identification of ship tracks in large datasets.

ACKNOWLEDGMENT

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND...

We would like to acknowledge Rob Wood and the atmospheric sciences department at University of Washington, conversations with whom helped propel this work forward. We would also like to thank the Earth Science investment area team for their support and encouragement of this work and Lori Dotson for contributing to the writing and editing processes.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

SUMMARY

We have developed novel data-driven techniques that advance our ability to assess the effects of ship emissions on marine environments and the risks of future marine cloud brightening efforts. We have developed a method to track aerosol injections using optical flow, a feature tracking algorithm adapted from the image processing literature. Using this approach, we can efficiently track local features in low-lying clouds, which are often difficult to isolate from regional high-cloud movement. This algorithm is also being made publicly available through Sandia's copyright assertion process. We proposed an approach to stochastically simulate the behavior of ship induced aerosols parcels within naturally generated clouds. Our method can use both real and simulated wind fields to determine the approximate movement of the cloud and simulated tracks, and uses a stochastic differential equation (SDE) to model the persistence behavior of cloud-aerosol paths. This SDE incorporates both a drift and diffusion term which describes the movement of aerosol parcels via wind and their diffusivity through the atmosphere, respectively. Lastly, we developed a novel statistical approach to automated ship track detection using a directional spatial filter applied to multi-spectral data combined with the probabilistic feature fusion technique developed by [48] to account for the known physical features of ship tracks.We show that accounting for the observed atmospheric conditions that lead to track formation and key track features, we achieve high detection accuracy and low false positive classification rates.

NOMENCLATURE

Table 0-1.

Abbreviation	Definition
ACI	Aerosol Cloud Interaction
CCN	Cloud Condensation Nuclei
DOE	Department of Energy
GHG	Green House Gas
MCB	Marine Cloud Brightening
PFF	Probabilistic Feature Fusion
SCI	Solar Climate Intervention

1. INTRODUCTION

It is well documented that aerosols from anthropogenic sources can apply direct radiative forcing by reflecting or absorbing sunlight, as well as apply indirect radiative forcing by altering the radiative properties of low-lying clouds [52, 1, 46, 9]. Natural aerosols can increase the concentration of cloud condensation nuclei (CCN) and lead to more, but smaller, cloud droplets for a fixed liquid water content. This, in turn, can cause changes in cloud albedo resulting in changes to a cloud's radiative properties [35]. The magnitude of the impact of aerosols on a cloud's radiative properties can vary greatly depending on the properties of the aerosol and the surrounding atmosphere [e.g., 7]. Most often, anthropogenic aerosols increase the amount of radiation reflected by clouds, but in some cases, they have been known to reduce a cloud's albedo [8].

Aerosols can directly and indirectly exert positive and negative radiative forcing on local and global climates [46, and references therein]. Direct radiative forcing occurs when aerosols scatter or absorb solar radiation to the surrounding air, either locally warming or cooling the atmosphere. Indirect radiative forcing by aerosols occurs when the aerosols interact with clouds and precipitation, impacting the clouds' lifetime, albedo, precipitation, and micro- and macro-physical properties [41, 28]. Currently, indirect radiative forcing is the largest documented source of uncertainty when it comes to overall radiative forcing in climate modeling [6, 36, 55]. This large uncertainty is due in part to the complexity of cloud dynamics, making it difficult to clearly separate the aerosol's radiative effect from that of the surrounding clouds [51]. Improving our understanding of aerosol-cloud interactions is necessary to reduce this uncertainty in climate models.

Increasing the reflectivity of all clouds has the potential to reduce positive radiative heating through targeted "climate cooling," the central concept in solar climate intervention (SCI). "SCI could potentially offer an additional strategy for responding to climate change but is not a substitute for reducing GHG emissions" [39]. Marine cloud brightening (MCB) and other intervention approaches have been proposed to intentionally increase the reflectivity of low altitude, boundary-layer clouds [e.g., 32, 37, 36] through the intentional increase of CCN via targeted aerosol injections in marine stratocumulus clouds. Low-lying marine clouds are most predisposed to albedo changes due to low densities of CCN, making this a promising approach to reducing global warming [3].

Ship tracks have been unintentional, natural examples of MCB for decades. For more than fifty years, satellite imagery has detected these bright linear features produced when the engine exhaust from large ocean-traversing ships mixes with low-lying marine clouds within 2 km of the earth's surface. Ship emissions have provided researchers with observable and traceable examples of aerosol-cloud interactions, which have been the focus of many studies to better understand the potential impacts of MCB [25, 19]. The Monterey Area Ship Track experiment off the coast of California [14] was one of the largest aircraft campaigns to study the formation of ship tracks. [12] and [53] were the first to document their observations of this phenomenon in visible-wavelength

images taken from the Television Infrared Observational Satellites (TIROSs). [42] found that ship tracks can impact the radiative properties of clouds long after they are visible.

This report details the work done under LDRD "Local limits of detection for anthropogenic aerocol-cloud interactions" to advance our understanding of aerosol-cloud interactions using observations in the form of ship tracks. Our focus is on developing data-driven methods to answer key questions such as: *How long can ship tracks persist?, Under what atmospheric conditions do ship tracks form? What are the limitations of local detection of ship tracks?* In an effort to answer these questions, this report is divided into three main chapters. Chapter 2 details a novel approach to following the formation and behavior of ship track formation and persistence. Chapter 3 proposes a stochastic hidden Markov modeling framework to simulate ship track formation from known ship locations and chapter 4 presents a statistical approach to the automatic detection of ship track using known physical features that contribute to track formation.

2. AN EFFICIENT APPROACH FOR TRACKING THE AEROSOL-CLOUD INTERACTIONS FORMED BY SHIP EMISSIONS USING GOES-R SATELLITE IMAGERY AND AIS SHIP TRACKING INFORMATION

2.1. Background

Although not all atmospheric conditions and ship exhaust have the potential to produce ship tracks [38], these features are more abundant than previously thought [59]. However, key questions remain regarding how long these tracks persist and what local impact they have on cloud radiative properties after the source of emissions has passed.

A lack of high resolution data, as well as difficulties isolating and tracing observed aerosol-cloud interactions over time, have been limiting factors in studying the longevity and long-term effects of ship tracks. Costly air campaigns have been the most reliable method of tracking the behavior of aerosols from a known source. Recently, due to the vast improvements in satellite imaging technology, more has become possible. For example, [60] used the Advanced Very High Resolution Radiometer (AVHRR) from the NOAA satellites to study the global long-term indirect effects of aerosols. [59] used machine learning to automatically label ship tracks in images from the MODerate resolution Imaging Spectroradiometer (MODIS) aboard both the Aqua and Terra satellites. [22] combined MODIS imagery and retrievals of cloud droplet number concentration (N_d) with known ship positions and properties to demonstrate a positive effect of emission sulfate concentration on the likelihood of ship track formation and a decrease in ship track observations due to fuel sulfur content restrictions set by the International Maritime Organization (IMO). More recently, [13] applied spatial kriging methods to cloud property data retrieved from MODIS and reanalysis from the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2), to infer negative impacts on radiative forcing from emissions along a major shipping lane in the southeast Atlantic. New satellite observations of aerosol-cloud interactions have been a large source of untapped information since it is very difficult to infer the radiative impact of ship emissions from observational data collected from earlier generations of satellites [16, 30, 23].

Historically, our understanding of the interactions between anthropogenic aerosols and clouds has been primarily limited to simulations. In atmospheric computational fluid dynamic numerical models, aerosol injections are initiated in the model at known, precise locations in fully defined environments that are easily traceable [e.g., 56, 2, 42, 4]. Unlike simulated case studies, satellite observations of ship tracks have many uncertainties including, for example, the size, composition, and concentration of the emitted aerosols and properties of the atmosphere, including temperature, wind, pressure, water content, and previous aerosol concentration. Real, observed ship tracks are initiated from an unknown source and form in a dynamic and only partially known environment, making it challenging to trace and fully characterize their behavior. In the research we present

herein, we show how high-resolution spatial and temporal observations from the National Oceanic and Atmospheric Administration's (NOAA) GOES-17 Advanced Baseline Imager (ABI) data can resolve ship tracks year-round.

In this chapter, we present two methods of effectively following the behavior and persistence of ship tracks and demonstrate these tools in determining how long ship tracks persist in a maritime environment. These methods could be used as a means to record ship track impact on climate. The first method relies on the NOAA HYSPLIT trajectory model [50], an atmospheric transport and dispersion model widely used by atmospheric scientists to estimate and study the trajectories of air parcels forward and backward in time. The second method relies on the image registration technique of [34], which is a purely image-based approach to local feature tracking. Both methods rely on radiance spectra collected from the GOES-17 ABI sensor [20] and provide an efficient approach to accurately and systematically characterize ship track persistence.

This research sets the stage for a more thorough exploration of the atmospheric conditions and exhaust compositions that produce ship tracks and factors that determine whether a track persists for 3, 9, or more than 24 hr. Many of these tracks have persisted as detectable linear cloud features for as long as 12 to 24 hr, much longer than the 6–8 hr typically assumed in climate simulation studies in pristine environments [e.g., 2].

2.2. Data

This research uses ship location information combined with L1b radiances measured from the Advanced Baseline Imager (ABI) instrument on the GOES-17 geostationary weather satellite [20], which provides four times higher spatial and three times higher temporal resolutions than previous generations of GOES imagers. Higher resolutions in both space and time allow us to study fast-changing cloud features such as ship tracks with greater precision. We rely on the near-infrared "cloud particle size" band (C06) and the infrared "shortwave window" band (C07) with central wavelengths of 2.24 and 3.90 μm, respectively, to visualize ship tracks throughout the diurnal cycle. To seamlessly visualize ship tracks during day-night transitions, we transform the data for each time stamp by subtracting spectral band C07 from band C06 and apply the image processing technique of histogram equalization to systematically control the contrast of each image. Histogram equalization enhances the contrast of each image, making it easier to visualize and identify key ship track features that may otherwise be invisible to the naked eye. The spatial resolution of both bands (C06 and C07) is 2 km, and the temporal resolution is every five minutes for the GOES-17 CONUS imager.

For the first half of 2019, we relied upon both satellite and terrestrial-based Automatic Identification System (AIS) data provided by the SeaVision database [54], which has near real-time resolution as frequent as every 15 minutes. In accordance with "Regulation 19 of SOLAS Chapter V *Carriage requirements for shipborne navigational systems and equipment*" of the International Maritime Organization (IMO), all ships must have on-board transponders capable of automatically transmitting ship information to other ships and coastal authorities [IMO, 2003]. The AIS provides positional data (latitude and longitude) and attributes, such as vessel name, size, type, and speed. AIS data was used to (1) cross-reference observed track locations with nearby ship positions to

confirm which ship produced the track emissions and (2) determine a precise location to initialize track trajectories via HYSPLIT, as described in Section 2.3.

We began by processing and plotting GOES-17 CONUS satellite imagery for a selection of dates in the first half of 2019 that coincided with available ship location information, which allowed us to identify a large number of dates with clearly visible ship tracks. We then selected three examples for further case study. These three examples consist of ship tracks in the Northern Pacific Ocean, often some distance off the western coast of the U.S. in February, April, and June 2019. Each example had a different composition of high and low clouds and thus exhibited different ship track behaviors (i.e., different feature formations and movements). While not a comprehensive study, these examples demonstrate the robustness of our approaches under different ship track conditions.

In the February 2019 example, we follow an intersection of two tracks, which is often easier to distinguish, and begin tracking on February 20, 2019, at 17:00 UTC at approximately 44.25° latitude, −139.45° longitude. We follow this intersection in a rotating cloud field until it dissipates and is no longer a distinguishable feature. In the April 2019 example, we follow a single track when it becomes clearly distinguishable and begin tracking on April 23, 2019, at 19:00 UTC at approximately 37.76° latitude, −131.37° longitude. This track quickly begins mixing with other nearby ship tracks until they all become indistinguishable in a larger cloud field. In the June 2019 example, we again follow an intersection, which becomes visible just as one ship passes into a large cloud bank. We begin tracking on June 17, 2019, at 05:00 UTC at approximately 36.4° latitude, -135.65° longitude. This is likely the best example of the three because it remains clearly visible throughout the 24-hr period even as it begins to dissipate and dim. The following sections demonstrate our two proposed tracing methods across these three examples.

2.3. Validating track persistence with HYSPLIT

We demonstrate the feasibility of using NOAA's HYSPLIT model [50] to predict ship track trajectories and assess track persistence, noting the model's advantages as well as its limitations. The HYSPLIT trajectory model [50] estimates air parcel locations using forward and backward trajectory analyses to determine either the future locations or historical origins of air masses or sources. Our simulations used analysis data from the Global Data Assimilation System (GDAS) provided by the National Weather Service's National Centers for Environmental Prediction. GDAS data is often used when gridded observational data is required, and the Air Resources Laboratory processes this data into HYSPLIT-useable formats. Until June 12, 2019, GDAS data had a spatial resolution of 0.5° grid cells (≈ 50 km), but it now has a resolution of 0.25° grid cells (≈ 25 km), and our simulations rely on data both before and after that date (both resolutions). For all trajectory analyses, the temporal resolution is hourly.

We used HYSPLIT to project the movements of pre-formed ship tracks forward in time. The "head" of a ship track is the position at which a new, visible cloud track is forming and is also referred to as the "initialization point" in this paper. We then constructed the forward trajectory of an air parcel using this position as the initial location for HYSPLIT and let the simulation run 24 hr forward in time. To quantify persistence, we visually verified whether the portion of the track we were following was still visible as a linear feature at the predicted locations.

We manually inspected satellite imagery to identify a ship track initialization point (e.g., intersection of ship tracks for easier feature matching), then matched the AIS location data to the feature of interest, and used the latitude and longitude values of the ship at the initialization point. We ran it forward in time for 24 hr, collecting the position data where the air parcels would have moved each hour. We then overlaid the HYSPLIT trajectory points onto the satellite imagery for these forward time points and visually assessed a) how well the air parcel projections estimated the observed ship track movement and b) whether there was still a clear remnant of the feature at each forward time step. Fig. 2-1 shows what this looks like for an initialization point (here, at the head of a track) from June 17, 2019.

Overall, the predicted HYSPLIT trajectories estimated the track feature movement reasonably well within the first 8 to 12 hr, although the model's performance was noticeably sensitive to the height initialization values required by HYSPLIT. The height initialization value identifies the correct air parcel to follow, is the starting point of a forward trajectory, and is used to interpolate pressure level data that largely drives the trajectory estimates. Initial exploration of cloud top height data, collected from the GOES-17 ABI as an L2 product, revealed significant variations in cloud top height measurements. While certain time periods had stable height values in a region of interest, others had values ranging from hundreds to thousands of meters within a small grid cell. In the latter case, due to the low spatial resolution of the cloud top height product (10 km) compared to L1b radiance data (2 km for bands > $2\mu m$), it was difficult to identify an appropriate initialization height for an observed track. Given that ship-induced cloud tracks are likely to form at the boundary layer, height values greater than 1 km are unlikely and greater than 2 km are unreasonable [33]. However, there is still uncertainty about the best choice of initialization height for these HYSPLIT trajectory runs. Therefore, we initialized HYSPLIT at multiple height values (0, 200, 400, and 600 m), which provided some information about the sensitivity of trajectory analysis to height and allowed us to qualitatively assess the uncertainty in predicted trajectories.

Fig. 2-1 shows our best example of the HYSPLIT trajectory model accurately estimating the predicted trajectory of the ship track feature that was clearly visible for more than 20 hr in June 2019 (also shown in Fig. 2-3). For this case, we initialized simulations at a discernible head of a ship track. Initializing the height at zero meters (sea level) seemed to provide the best match, and trajectories initialized at 200, 400, and 600 m are not unreasonable up to 9 hr. HYSPLIT's predictions of the underlying air parcel trajectory reasonably align with the observed movement of this cloud feature with minimal variability due to height initialization until about the 9-hr mark when the trajectory variability becomes more prominent. Although the track is still clearly visible at 12 hr, the HYSPLIT trajectories no longer seem to be centered on the track feature of interest and have become much more variable for different height initializations. HYSPLIT projections no longer match the movement of this feature past 12 hr as shown in Fig. 2-1. Two more examples of trajectory analysis for dates in February and April 2019 revealed similar trajectory variability due to height initialization, and misalignment to the observed track movements was much greater than for the June 2019 case. These trajectories are shown in Figs. 2-9 and 2-10.

In summary, HYSPLIT can be useful in predicting the trajectories of visible cloud-aerosol features up to approximately 8–12 hr depending on the case. However, due to the hourly time resolution, it can only provide a rough estimate of the length of time the track persists as a linear feature. For ship tracks that are less distinguishable from the surrounding clouds, it becomes difficult to confirm

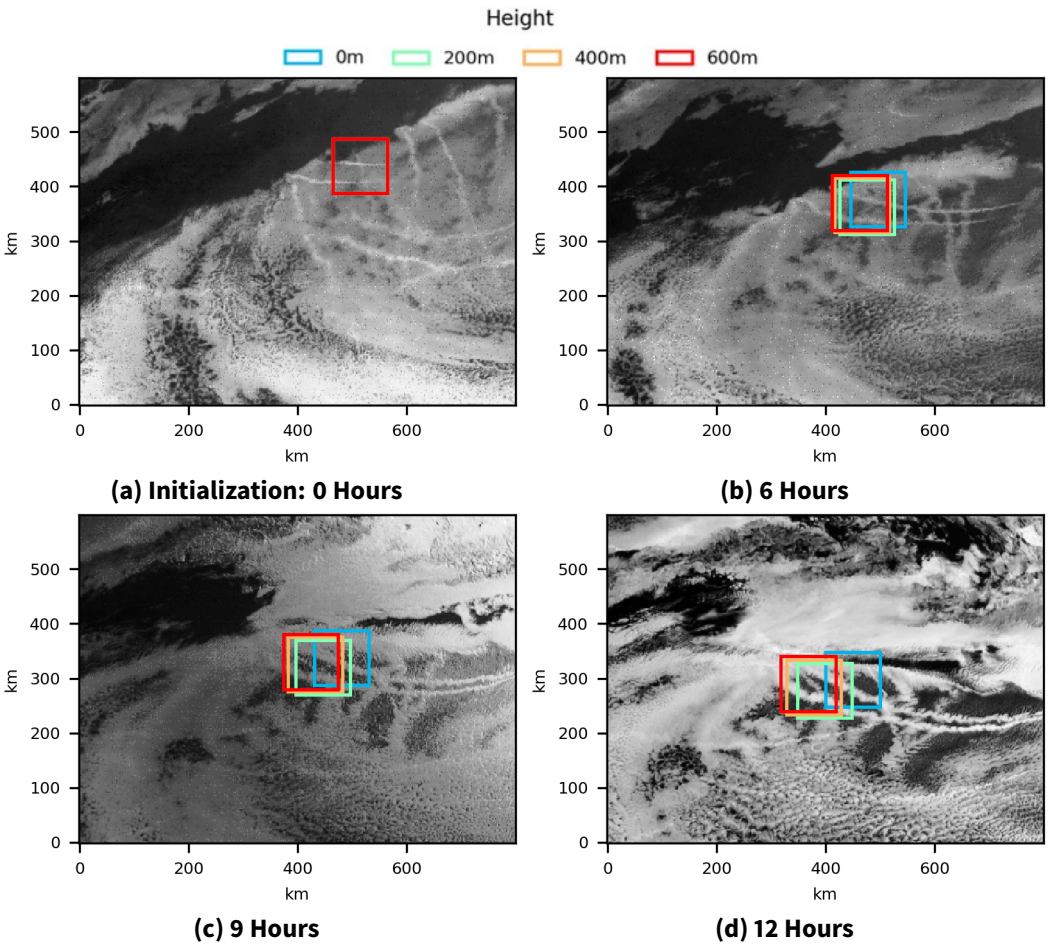

Fig. 2-1. These figures show the forward trajectory of air parcel movement, starting with the head of a track on June 17, 2019, at 07:02 UTC (a) and stepping forward in time, with snapshots shown 6 (b), 9 (c), and 12 (d) hr later. HYSPLIT is able to track the movement of this cloud feature well, and the track is still clearly visible at 12 hr

that HYSPLIT projections align with the same portion of the track where we initialized the forward trajectory. This is likely due to the fact that 1) HYSPLIT is initialized at a single location and then runs forward with no intermediate checks to see if it is still tracking the same (initial) feature, 2) HYSPLIT only projects at 25 km spatial and 1-hr increments (compared to the 0.5–2 km and 5–15 minute resolutions the GOES-17 imager collects) and might miss nonlinear movement within the hour, and 3) HYSPLIT predicts the locations of the emitted aerosols rather than the aerosol-cloud interactions that form ship tracks. However, an interesting application of this proposed HYSPLIT approach is to use it to infer the height at which the aerosol track resides. For example, if we initialize at several more height increments, the initialization height that best predicts the trajectory of the feature might be a more accurate estimate of the aerosol track height than the cloud top heights retrieved from the GOES-R L2+ products. Additionally, for improved predictions, we could re-initialize the starting point for trajectory analysis at each time step as long as the feature can still be clearly identified.

2.4. Tracking ship track features with optical flow

As discussed in Section 2.3, following a ship track feature using HYSPLIT can be challenging and has its limitations. In this section, we present a more precise "optical flow method" that relies on the Shi-Tomasi feature selection technique [47] and the Lucas-Kanade feature tracking technique [34]. These techniques are widely used in video processing to find and estimate motion of features between video frames. This approach relies on visual features alone and is thus unaffected by the meteorological errors and uncertainties to which HYSPLIT is subjected. By augmenting the two techniques to fit a simple motion prediction framework, we are able to reliably follow ship track features well past the point where HYSPLIT projections do not agree with observations.

Previously, the Lucas-Kanade algorithm was successfully used to estimate cloud motion in ground-based video feeds to forecast solar irradiance [58] and in satellite image sequences to track individual cloud banks [26]. [47] showed that Lucas-Kanade is most successful with high-contrast textural features (specifically those with high contrast in both x and y directions), making it an appropriate choice for tracking features within the textured cloud regions where we observe ship tracks. The Shi-Tomasi algorithm was designed to detect the most distinguishable features within an image. Our method combines the Shi-Tomasi algorithm for feature detection with the Lucas-Kanade algorithm for tracking. This approach is strictly computational, relying only on image pixel values with no consideration of other atmospheric or meteorological data that might govern aerosol movement. Thus, it is an attractive method for tracking cloud features observed at any altitude, but it is sensitive to image data corruption and intensity variation between frames. The latter can be common between GOES-17 CONUS frames, especially when transitioning between nighttime and daytime images. As described later in this section, we make modifications to the Lucas-Kanade algorithm to allow continuous feature tracking for more than 24 hr.

Optical Flow Overview

The optical flow approach used here is a combination of the Shi-Tomasi and Lucas-Kanade algorithms. The Shi-Tomasi algorithm is used to determine the center pixel locations of high-contrast features within an image by assigning a quality value to each pixel in the image, where higher quality is associated with higher-contrast features. The quality q assigned to a pixel at location (u_x, u_y) in the image I is the minimum eigenvalue of the associated structure tensor M, which is computed over an $n \times n$ neighborhood of pixels, i.e.,

$$M(u_x, u_y) = \sum_{u_x-n}^{u_x+n} \sum_{u_y-n}^{u_y+n} \begin{bmatrix} \left(\frac{\partial I}{\partial x}\right)^2 & \frac{\partial I}{\partial x}\frac{\partial I}{\partial y} \\ \frac{\partial I}{\partial x}\frac{\partial I}{\partial y} & \left(\frac{\partial I}{\partial y}\right)^2 \end{bmatrix} \tag{2.1}$$

If λ_1 and λ_2 are the eigenvalues of M, then $q(u_x, u_y) = \min(\lambda_1, \lambda_2)$. When implementing the Shi-Tomasi algorithm in practice, all pixel locations with a quality less than a selected threshold t_q are first rejected. Further refinement is then performed by non-maximum suppression within an $m \times m$ sliding window, and then all remaining pixel locations are selected as feature centers.

The Lucas-Kanade algorithm is then applied iteratively for each pair of frames in a sequence to estimate where features have moved from one frame to the next. A feature $I_{W(u_x,u_y)}$ is the collection of pixel intensities in the neighborhood W, about a feature center at location (u_x, u_y) in frame I. The width and height of neighborhood W, are given by $2\omega_x + 1$ and $2\omega_y + 1$, respectively, where ω_x and ω_y are non-negative integers chosen by the user.

Assuming that a given feature $I_{W(u_x,u_y)}$ has moved between two consecutive frames, that same feature will be given by $J_{W(u_x+d_x,u_y+d_y)}$ in consecutive frame J, where $d_x, d_y \in \mathbb{R}$ denote the distances (in pixels) the feature has moved between frames in the x and y directions, respectively. To identify $J_{W(u_x,u_y)}$, the distance $\boldsymbol{d}^* = [d_x, d_y]^T$, also known as the *optical flow*, is estimated between $I_{W(u_x,u_y)}$ and $J_{W(u_x,u_y)}$ by minimizing the sum of squares between the feature neighborhood in frame I and the shifted neighborhoods in frame J, i.e.,

$$\boldsymbol{d}^* = \underset{(d_x,d_y)}{\operatorname{argmin}}\, \epsilon(d_x, d_y) \text{ and} \tag{2.2}$$

$$\epsilon(d_x, d_y) = \sum_{x=u_x-\omega_x}^{u_x+\omega_x} \sum_{y=u_y-\omega_y}^{u_y+\omega_y} [I(x,y) - J(x+d_x, y+d_y)]^2 \tag{2.3}$$

The Lucas-Kanade algorithm uses a Taylor approximation about $(d_x, d_y) = (0,0)$ to estimate $\boldsymbol{d}^*$. This approach assumes minimal motion between consecutive frames, which is not a reasonable assumption for cloud features. To account for larger motion in practice, we apply Lucas-Kanade iteratively, such that $\boldsymbol{d}^*$ is estimated several times for each feature with each estimate updating the previous one, and use the pyramidal implementation described in [5]. This implementation constructs a "pyramid" of image copies of various resolutions with each copy having half the resolution of the previous one.

Application to cloud feature tracking

We applied the optical flow tracking method to three ship track case studies in February 20, April 23, and June 17 of 2019 and compared its performance to that of the HYSPLIT approach described previously. Using the combination of L1b radiances from GOES-17 ABI bands C06 and C07 described in Section 2.2, we implemented the optical flow method to identify and track features within a user-defined region of clouds surrounding a recently-formed ship track. This was accomplished using the OpenCV implementations of both the Shi-Tomasi and Lucas-Kanade algorithms [40]. The default values provided by OpenCV were used for any parameters not described here.

First, we manually selected a cloud region of interest (approximately 50×50 pixels, or 100 km $\times$ 100 km) immediately surrounding an initialization point, i.e., the head of a ship track or the intersection of two ship tracks, as described in Section 2.3. The size of the region of interest was chosen such that at least five high-contrast features, as selected by the Shi-Tomasi algorithm, could be found within the region. The robustness of our approach for tracking a cloud region relies on the selection of more than one feature because the Lucas-Kanade algorithm may fail for some features over time. We applied the Shi-Tomasi feature detection algorithm within the cloud region to

identify the locations of features for tracking with parameters $n = 3$, $m = 3$, and t_q as 20% of the maximum $q(u_x, u_y)$ within the region. Next, these features were tracked in all the following frames using the Lucas-Kanade technique. We found that a neighborhood size of 15 pixels ($\omega_x, \omega_y = 7$) was appropriate for creating a unique set of features for effective tracking. We chose the neighborhood size such that some visible cloud texture was encompassed within each feature, though we expect the results to be robust for neighborhood sizes between 10 and 20 pixels for our studies. For each frame, we iterated using 3-level pyramids until either the estimated displacement was less than 0.03 pixels or 10 iterations were completed, whichever occurred first. The features were tracked in consecutive frames until the ship track was no longer recognizable in the frame.

Although best assessed in video format, we demonstrate tracking performance in this paper as a tracking box overlaid on the satellite imagery at 4 to 6-hr intervals. The tracking box is parameterized by the user-defined cloud region of the first frame and is updated in each frame to be roughly centered around the ship track cloud feature by adding it to the mean displacement of the features between the current and previous frames. Fig. 2-2 shows the performance of the optical flow approach on a cloud region on June 17, 2019. The tracking follows two distinct ship tracks over 18 hr, throughout which the tracks clearly persist. Fig. 2-3 shows the isolated user-defined cloud region of the same tracking result in 4-hr increments over a 28-hr period. We clearly observe ship tracks from the beginning of tracking at 05:02 UTC until at least 01:02 UTC the following day, accurately quantifying the duration of persistence as 20 hr. Using this approach, we can observe the persistence and dispersion of ship emissions in a cloud layer by tracking that region over several hours in increments of 5 min, the temporal resolution of the GOES-17 ABI CONUS scan.

Figs. 2-4 and 2-5 demonstrate the optical flow tracking technique for the two case studies in February and April 2019. We compare optical flow performance with HYSPLIT trajectory predictions in the next section and Figs. 2-8, 2-9, and 2-10. Based on these comparisons and due to the high sensitivity of HYSPLIT to height inializations over time, we prefer the optical flow method over the HYSPLIT method to reliably predict the path of ship track features that persist beyond 8 hr.

Limitations of optical flow

Although we demonstrated the efficacy of the optical flow approach for these case studies, it is important to point out some of its limitations, as well as the modifications we made to track cloud features. A well-known source of error is abrupt changes in image brightness due to the algorithm's dependence on pixel intensity values. These changes may occur if the data is corrupted, if there are large thermal changes in the atmosphere such as at sunrise and sunset, or if other artifacts cause large intensity disparities in one frame compared to others in the sequence. Errors due to data corruption are avoided by considering the data quality flag (DQF) field in the ABI data. For example, we chose to omit a frame of data from the optical flow computations if the percent of corrupt pixels, as specified by the DQF, surpassed 2%. This method was sufficient for the data we used because there were very few frames with corrupt pixels, but it could be improved by considering the DQF for nearby pixels.

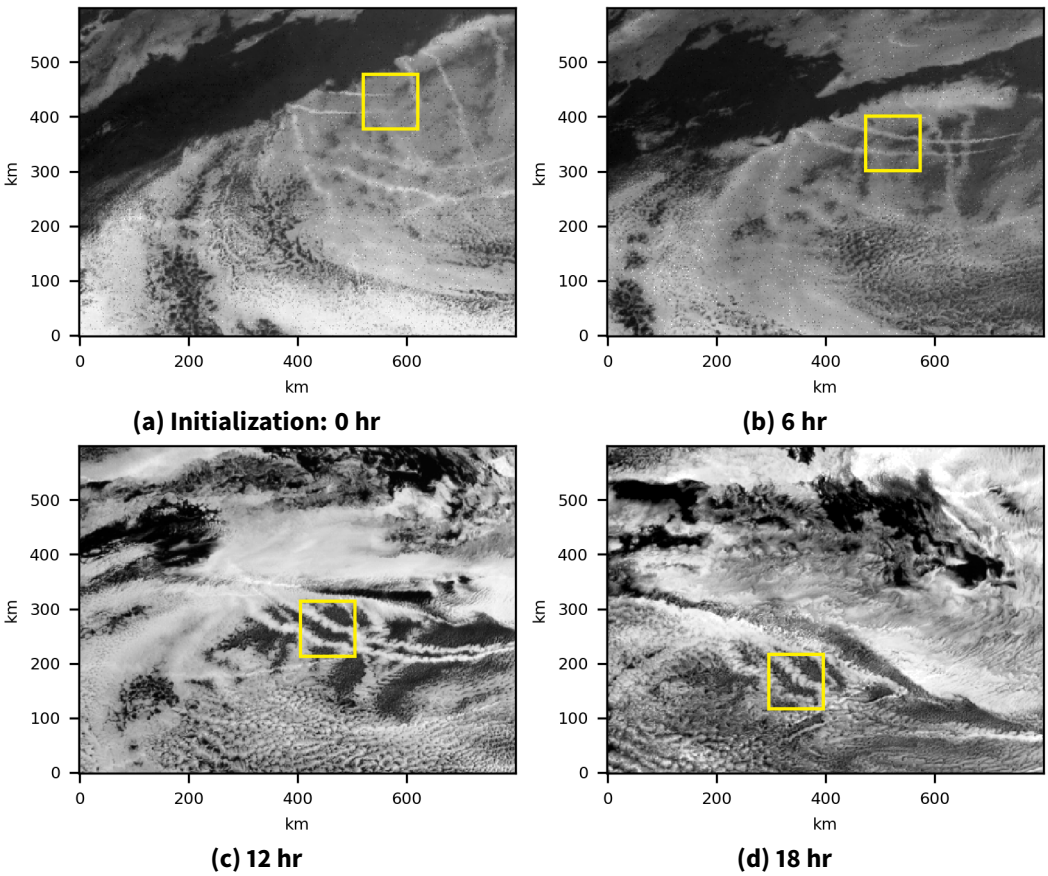

Fig. 2-2. These figures show the result of the optical flow method applied to a manually-selected local cloud region, starting with the heads of two ship tracks on June 17, 2019, at 05:02 UTC (a) and stepping forward in time, with snapshots shown at 6 (b), 12 (c), and 18 (d) hr later. The tracking algorithm is able to follow the movement of the cloud region well, and the tracks are still clearly visible 18 hr later

Intensity changes in satellite images captured during sunrise and sunset nearly always cause the Lucas-Kanade technique to incorrectly identify features from one frame to the next; this is unavoidable when following features such as ship tracks because they frequently persist for more than 8 hr. We were able to continuously follow a cloud region of interest across these boundaries by predicting the start and duration of these transitions and using observed trajectories rather than the Lucas-Kanade method to predict its trajectory.

We determined the starting and ending frames of a diurnal transition using calculated solar zenith angles at the right and left perimeter of our tracking box, respectively. These angles were calculated using the National Renewable Energy Laboratory's Solar Position Algorithm [43], which calculates the sun's apparent altitude with a precision of about 0.0003 degrees given the date, time, and location. For a given frame, the minimum and maximum angles are used for both the right and left edges of the tracking box to determine the start and stop of a diurnal transition. Thresholds for the start and end of each transition were chosen empirically and conservatively to ensure the transition periods are estimated accurately. Using the solar zenith angle, we found diurnal transitions to be quite predictable. More details for these computations and our threshold choices are discussed in

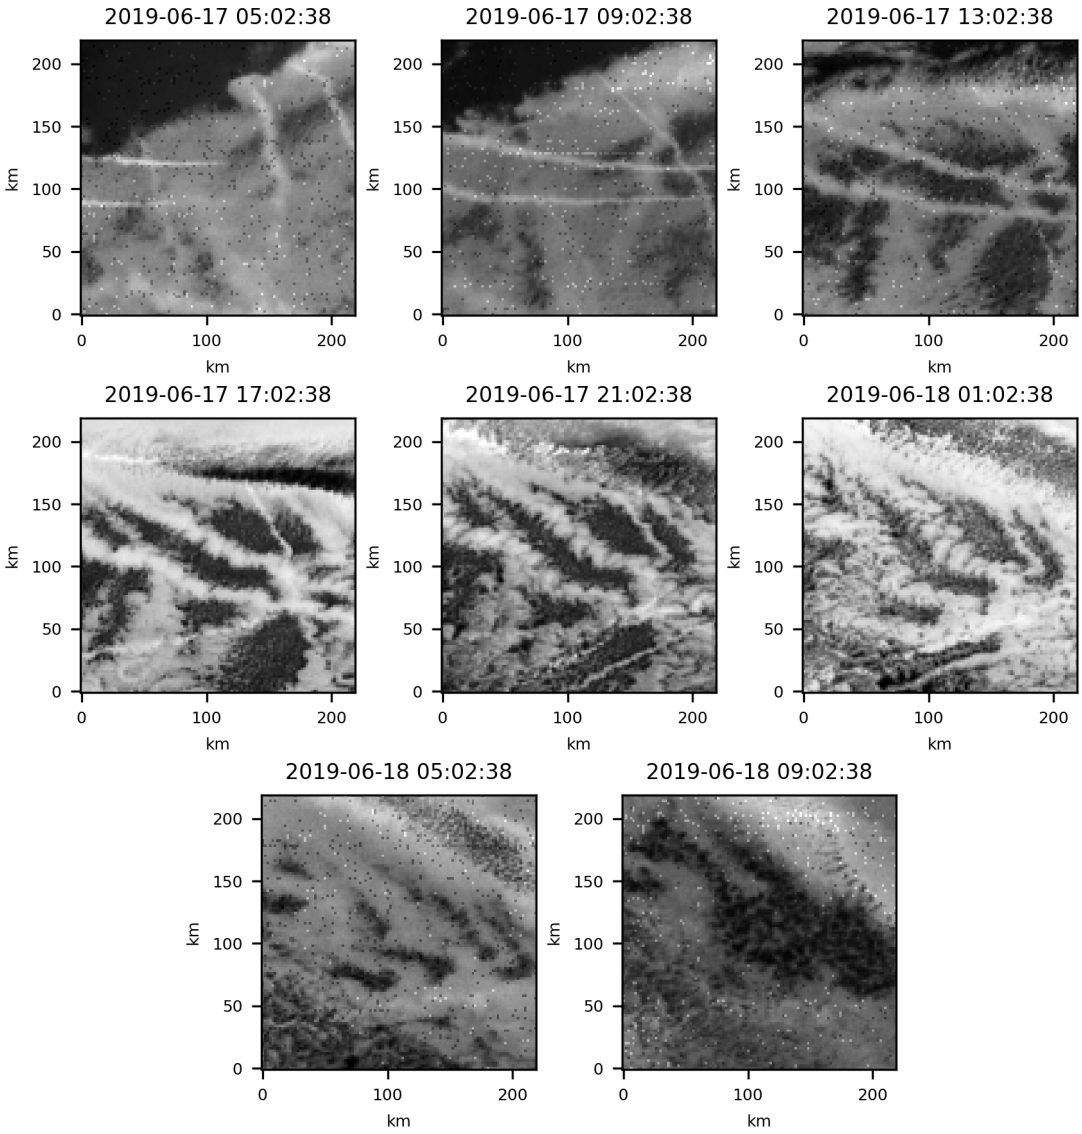

Fig. 2-3. These figures show the 110 px by 110 px image region centered on the tracking box of the optical flow method at each frame in 4-hr time intervals beginning on June 17, 2019, at 05:02 UTC. The box is removed from these images for better visualization. The remnants of the ship tracks are still visible up to 20 hr after they first appeared

section 2.4.0.1.

At the start of a boundary, feature tracking is stopped, and the average motion of the region of interest is computed using the motion of the features across the six prior frames. The velocity of the region of interest is then assumed to be constant over the transition period, and the location of the region of interest at any point during this transition is predicted using this assumption. Over this short transition period, this augmentation has proved to work well for our cases. At the end of the transition boundary, new features are selected within the region of interest using the Shi-Tomasi algorithm and followed until the ship track is no longer visible in the region of interest.

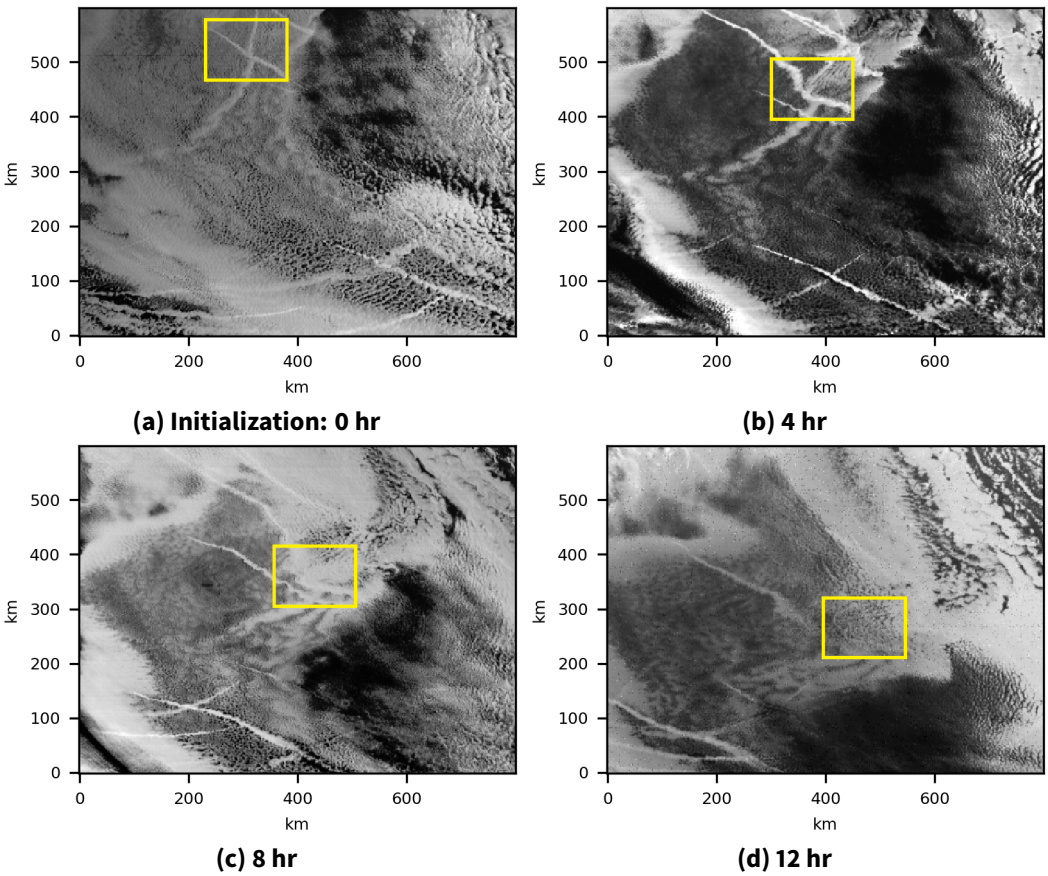

Fig. 2-4. Shown are the results of the optical flow approach for the February case study. This case was initialized on February 20, 2019, at 17:02 UTC, shown in (a), and snapshots are shown (b) 4, (c) 8, and (d) 12 hr later. The ship tracks remain visible 8 hr after initialization and remnants remain after 12 hr

An equally appealing approach to overcoming large changes in intensities between images might be to correct for significant pixel intensity changes between images so that the Lucas-Kande method could be leveraged for the full 24 hr period. For example, [17] show a promising approach for correcting pixel inhomogeneities in magnetic resonance brain images. However, we found our approach to be sufficient and did not explore this route

In addition to changes in pixel intensity, the optical flow method is also sensitive to non-affine changes in the shapes of features. Common cloud motions can introduce warping, but the warping between two consecutive GOES-17 CONUS frames, measured 5 min apart, is minimal and does not affect tracking success. A larger temporal gap between frames may allow for a greater change in features and thus introduce notable tracking errors. Large temporal differences between frames can occur when data is absent from the database or the frames are rejected because they contain corrupt data. In our experience, a temporal gap of up to one hr between two frames is generally reasonable for successful feature tracking.

Other natural phenomena that can cause the Lucas-Kanade technique to fail include interference from high-altitude clouds passing over the region of interest, dispersion of the boundary cloud layer, and disappearance of texture from the cloud layer. In the first case, the features may be obscured or

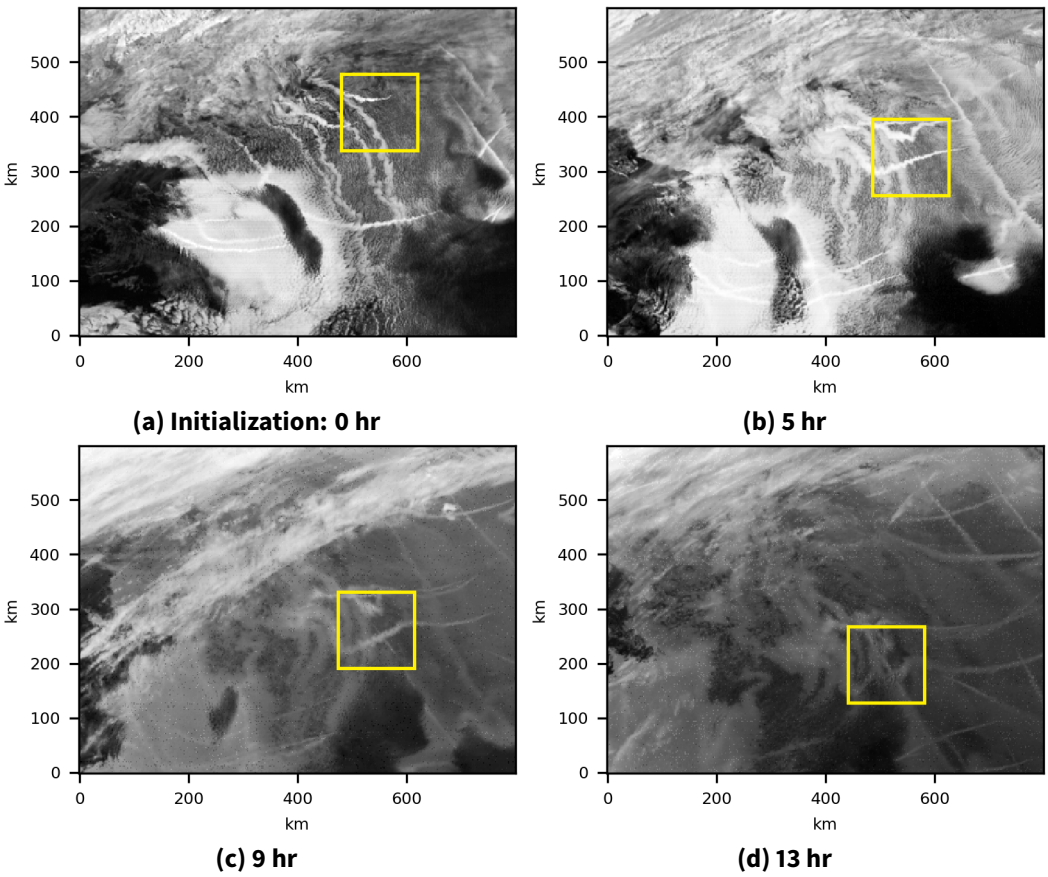

Fig. 2-5. Shown are the results of the optical flow approach for the April case study. This case was initialized on April 24, 2019, at 19:02 UTC, shown in (a). GOES-R ABI data from 22:58 to midnight on April 24, 2019, is unavailable, so snapshots are shown (b) 5, (c) 9, and (d) 13 hr later. The ship tracks remain clearly visible until after 13 hr after initialization when remnants still remain

confused with similar features in the higher-altitude clouds; in the latter two cases, the features disappear entirely and cannot be tracked further. Although the local feature tracker presented here has its disadvantages, we believe its advantages outweigh these challenges. Many of the obstacles described here could be circumvented by integrating known physical and/or meteorological factors that contribute to cloud feature movement. There is also the potential to integrate the HYSPLIT tracking approach described in Section 2.3 over short periods of time (<6 hr) when we expect the optical flow approach to fall short.

2.4.0.1. Details on sunset/sunrise transitions using solar zenith angle

We determined the starting and ending frames of a diurnal transition using calculated solar zenith angles at the right and left perimeter of our tracking box, respectively. These angles were calculated using the National Renewable Energy Laboratory's Solar Position Algorithm [43], which calculates the sun's apparent altitude with a precision of about 0.0003 degrees given the date, time, and location. For a given frame, the minimum and maximum angles are used for both the right and left

edges of the tracking box to determine the beginning and end of a diurnal transition. Thresholds for the start and end of each transition were chosen empirically and conservatively to ensure the transition periods are estimated accurately.

Specifically, let α_r and α_l be vectors of pixel solar zenith angles for the right and left half of the tracking box perimeter, respectively, and let $c, d \in \mathbb{R}$ such that $0 < c < d$ be the thresholds for transition periods. Then a diurnal transition is occurring if the following condition is true:

$$(\min(\alpha_r) < d \cap \max(\alpha_l) > c) \cup (\max(\alpha_r) > c \cap \min(\alpha_l) < d).$$

The first set in the union describes a sunrise and the second describes a sunset. A demonstration of this condition is displayed in figure 2-7 , which shows the diurnal pattern of rising and falling slopes of $\max(\alpha_l)$, $\min(\alpha_r)$, $\max(\alpha_r)$, and $\min(\alpha_l)$ for the February case study. The sunrise boundary begins at the frame in which $\min(\alpha_r) < d$ on a falling slope of $\min(\alpha_r)$ values and ends at the frame where $\max(\alpha_l) \leq c$. The sunset boundary begins at the frame in which $\max(\alpha_r) > c$ on a rising slope of $\max(\alpha_r)$ values and ends at the frame where $\min(\alpha_l) \geq d$.

To determine empirical values for c and d, we sampled a total of eight sunrise and nine sunset images from multiple dates. Using ten points manually selected from each image where sunrise or sunset significantly impacted cloud radiance, we derived a total of 160 solar zenith angles. Figures 2-6 (a) and (b) show an example of selected points along the start of a sunrise and sunset transition, respectively.

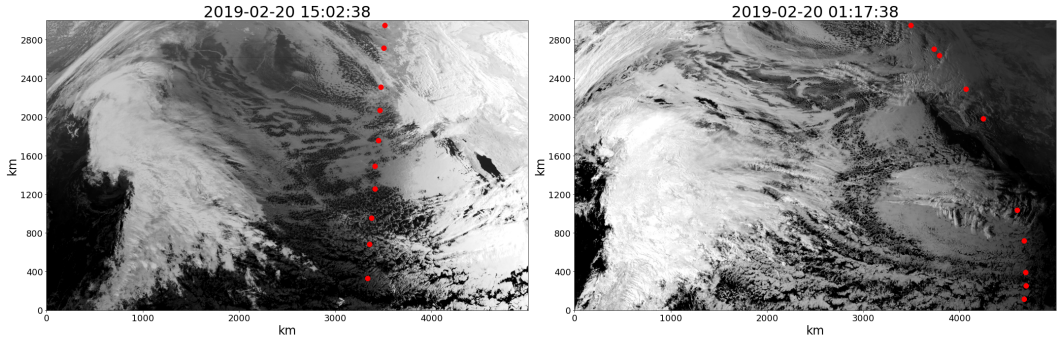

(a) Sunrise **(b) Sunset**

Fig. 2-6. Examples of selected observations (red points) for a sunrise (a) and sunset (b) transition on February 20^{th}

We initially chose a sunrise threshold of $d = 91°$ or the 95^{th} percentile of sampled angles during a sunrise event. Although in application, we found a better estimate of the start of sunrise to be when the minimum angle at the right hand side of the box falls below 96°. Similarly, we found that when the maximum angle on the right side of the box exceeds 84°, it is a good indicator of the start of a sunset transition. The last frame of a sunrise transition was determined when the maximum angle falls below 84° and sunset transitions completed when the minimum angle is above 96°. Using the solar zenith angle, we found diurnal transitions to be reliably predictable.

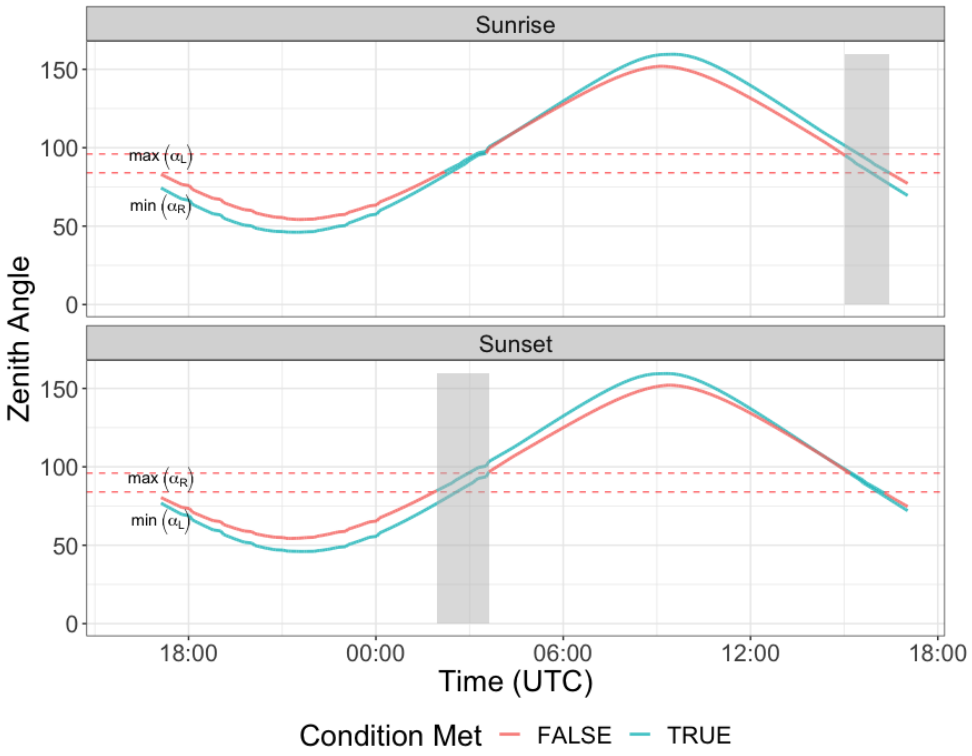

Fig. 2-7. This figure demonstrates the diurnal cycle of the minimum and maximum solar zenith angles for the right and left side of the tracking box over a 25 hour period starting on February 20^{th} at 17:00 UTC and ending on February 21^{th} at 18:00 UTC. The top figure shows the behavior of max(α_L) min(α_R), which indicate the start and stop of a sunrise transition while the bottom figure shows the behavior of max(α_R) min(α_L), which indicate the start and stop of a sunset transition. The vertical grey boxes indicate the time period where the transition conditions are met and thus the time periods for a sunrise or sunset event

2.5. Comparison of HYSPLIT to optical flow

Figures 2-8, 2-9, and 2-10 show the results of both the HYSPLIT predictions and the optical flow for each of the three case studies. In each case, HYSPLIT is initialized at four different altitudes, showing the sensitivity of the trajectory to height. The total image area shown is 500 px wide by 400 px tall, covering approximately 1000 km by 800 km. For February 20^{th}, initialization heights below 600 meters seem to work best for the first 6 hours. In this case, the HYSPLIT trajectory does not follow the track feature exactly, although the human eye can still clearly distinguish the feature after 6 hours. At 12 hours, the feature is still visible, but the HYSPLIT projection is quite far off. Beyond 12 hours, it appears that the track feature has dissipated into surrounding clouds. For April 23^{rd}, there is minimal variation in trajectory paths for different height initializations up to 6 hours, although an initialization height of zero meters seems most reasonable past 6 hours. This is again a case where it is hard for the human eye to distinguish the track after 12 hours.

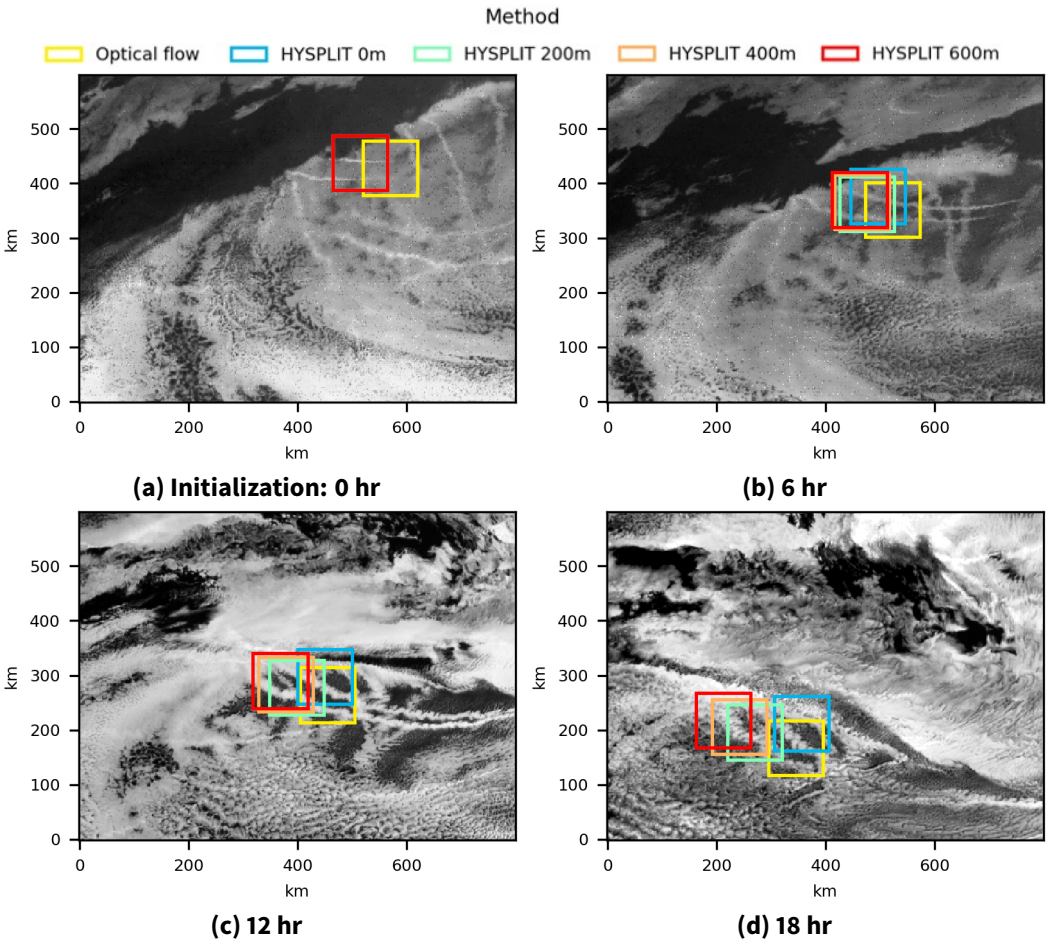

Fig. 2-8. Shown are the combined results of HYSPLIT trajectory analysis and optical flow approach for the June case study. Both HYSPLIT and optical flow were initialized on June 17, 2019, at 05:02 UTC, shown in (a), and snapshots are shown (b) 6, (c) 12, and (d) 18 hr later. The HYSPLIT analyses were initialized at four different altitudes for comparison. Though both methods agree with the motion of the ship tracks, some of the HYSPLIT projections appear to accumulate error over time. The ship tracks remain clearly visible 18 hr after initialization

2.6. Discussion

We have presented two methods of systematically following ship track behavior observed in imagery from the GOES-17 geostationary weather satellite. We have demonstrated the cloud feature tracking capabilities of each and shown how they can be used to quantify track persistence using three case studies. HYSPLIT and optical flow have complementary advantages and disadvantages, and together they make a very efficient tracking method for aerosol-cloud interactions. This capability can be leveraged in the future to track and assess long-term local impacts of proposed solar climate intervention efforts such as MCB.

The first method we demonstrated relies on NOAA's HYSPLIT trajectory model, and it accurately followed ship track behavior with some degree of uncertainty within the first ~8 hr before its

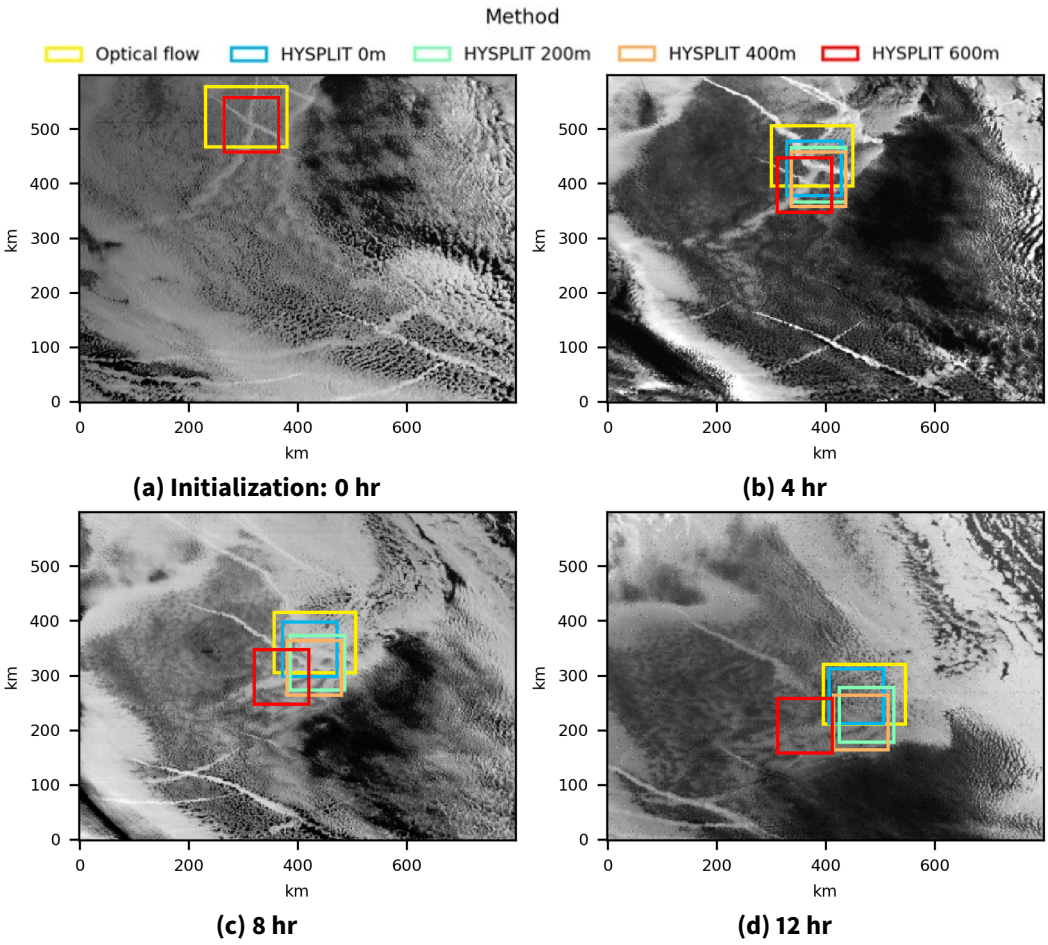

Fig. 2-9. Shown are the combined results of HYSPLIT trajectory analysis and optical flow approach for the February case study. Both methods were initialized on February 20, 2019, at 17:02 UTC, shown in (a), and snapshots are shown (b) 4, (c) 8, and (d) 12 hr later. The HYSPLIT analyses were initialized at four different altitudes for comparison. Though both methods agree with the motion of the ship tracks, some of the HYSPLIT projections appear to accumulate error over time. The ship tracks remain visible 8 hr after initialization

tracking accuracy declined, which could be due a number of factors. Most notably, HYSPLIT results are sensitive to the reanalysis data it depends upon. It is likely that we could see improved results using reanalysis data from ERA-5 rather than the GDAS data we used here. [29], for example, find the meteorology data of the ECMWF atmospheric reanalysis [24, ERA5] work well to follow boundary layer trajectories within their Lagrangian large-eddy simulations over a 2-day period. Assuming ship tracks are a delayed response of the cloud to ship emissions, this time delay may also need to be taken into account to improve the accuracy. Additionally, the continual forward trajectory estimates of HYSPLIT have no intermediate checks to see if it is still tracking the initial feature, and the low hourly temporal resolution of HYSPLIT could also cause issues as any nonlinear movement that occurs within the hour will not be taken into account. Lastly, HYSPLIT is heavily dependent on atmospheric measurement retrievals and reanalysis products, which can have large associated measurement errors. An interesting observation in all three case studies is that

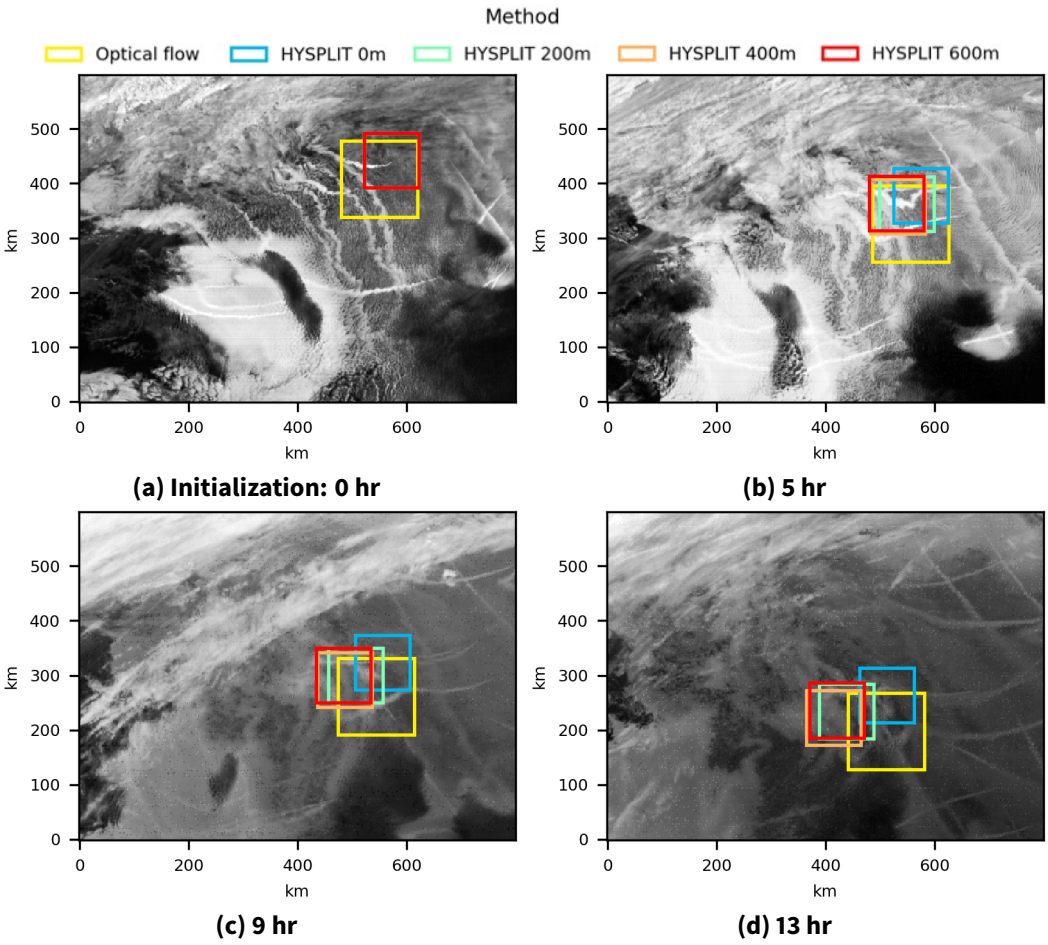

Fig. 2-10. Shown are the combined results of HYSPLIT trajectory analysis and optical flow approach for the April case study. Both methods were initialized on April 24, 2019, at 19:02 UTC, shown in (a). GOES-R ABI data from 22:58 to midnight on April 24, 2019, is unavailable, so snapshots are shown (b) 5, (c) 9, and (d) 13 hr later. The HYSPLIT predictions were initialized at four different altitudes for comparison. The ship tracks remain clearly visible 9 hr after initialization

HYSPLIT seems to follow the ship track reasonably well during the first 5 to 6 hr when initialized at a height of zero meters (sea level). Height is used by HYSPLIT to infer surface pressure for that location, which at sea level would be the highest vertically. Some of the NOAA higher resolution data files have five or more vertical levels in the boundary layer (<850 hPa), which could explain why we see so much variation between initialization height increments of only 200 m. Another potential use of the proposed HYSPLIT forward trajectory approach is to infer the height at which the aerosol track resides, which is difficult to infer from satellite measurements alone. Since we expect to see the first signs of track formation in boundary layer clouds, we can initiate HYSPLIT at varying altitudes and assess its tracking capability over the first few hours to potentially extract a boundary layer height estimate.

The second method we examined is a local feature tracking, or optical flow, approach that adapts the Lucas-Kanade algorithm for tracking ship track features and only requires radiance data from the GOES-17 ABI sensor. Note that we rely on spectral bands C06 and C07 to seamlessly observe

the tracks throughout day-night transitions off the coast of California, but any method of feature visualization would work for this approach. We have shown that our method can follow a ship track with high accuracy well past 12 hr and throughout day-night transitions, allowing more precise characterization of ship track feature persistence. However, due to its reliance on high-contrast features within the tracking window, it will not satisfactorily follow weak ship track features without further image preprocessing techniques, which can be tedious to perform over the many images a track can appear. In this case, using both the local feature tracking method (optical flow) combined with the HYSPLIT approach can provide a complete picture of the ship track's trajectory. Aside from following ship track features, there is the potential to use the optical flow approach to estimate the motion of the low-lying cloud bank where the ship track was last visible in an attempt to track its effect well beyond being a visible linear cloud feature. The algorithms used to produce the GOES-R ABI L2 product derived motion winds (DMW) [21] are promising approaches for following large high-cloud regions but are not suitable for local movement tracking needed to estimate motion at varying altitudes. This capability is currently being explored and is a topic for future discussion.

3. A STOCHASTIC SIMULATION MODEL FOR CLOUD-AEROSOL INTERACTIONS

3.1. Background

Most current knowledge on specific conditions under which tracks form have come from physical simulation studies under pristine conditions, which do not necessarily represent reality. In climate simulation studies of aerosol-cloud interactions [56, 2, 42, 4], aerosol injections are initiated by the user at a known location in fully defined environments. Satellite-observed tracks, however, are instead “initiated" by an unknown source and form in a dynamic and only partially known environment that is difficult or near impossible to replicate in a physical simulation study, which can also be quite computationally expensive.

In this work, we present a computationally efficient, mathematical simulation approach to emulating the observed formation and behavior of ship tracks. Existing methods focus on modeling the chemical evolution of aerosol composition [44, 49] and are not applicable to the physical modeling of cloud-aerosol paths through the atmosphere. Our approach aims to do so by flexibly accounting for the effect of weather and atmospheric conditions.

Our method is different from physical simulation approaches in that it attempts to emulate what is observed via satellite, rather than generate the full 3-D micro-physics environment. Ultimately, we would like to infer from imagery and atmospheric data under what conditions do track form or not form to incorporate in our emulation approach. For now, without more information on to what degree atmospheric conditions effect the visibility or behavior, we present the general framework accounting for the cloud movement and point out where atmospheric effects can be incorporated.

For a given ship, we consider modeling each aerosol *emission burst* as a single target. Each target is transported vertically from the ship through the atmosphere until it reaches a specific altitude near the cloud top height at which the target can become visible to orbital satellites and form a linear tracks in a cloud. figure 3-1 outlines the general behaviors of the aerosols that are observed or unobserved via satellite. The green box in figure 3-1 represents the portion of the track formation process that is visible via satellite. A ship track is the visible effect of the exhaust aerosols mixing with the low-lying clouds. The vertical transport of the aerosols between the ship’s smoke stack and the boundary clouds is largely unknown and unobserved. The exact altitude of the boundary clouds in which the ship track forms and the time lag between an aerosol burst released from a ship and reaching the visibility height largely depends on the complex weather and cloud dynamics. The visibility height can be approximated using cloud top height measurements obtained from satellite retrievals but the time lag is likely impossible to infer from satellite images with spatial resolution greater than a kilometer such as those retrieved from the GOES-R imager. Aerosol transport from

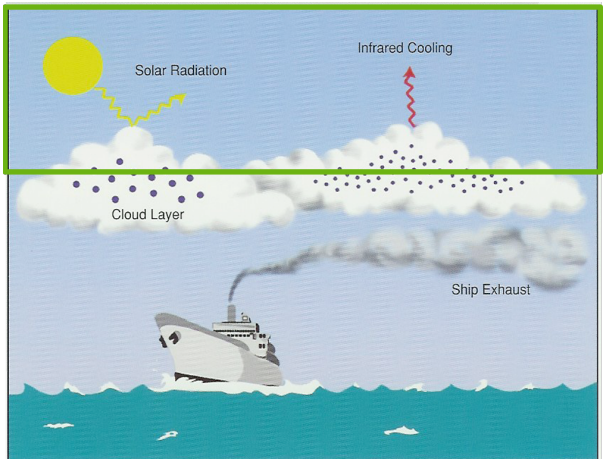

Observed via Satellite	Not Observed Via Satellite
Effects after aerosols from exhaust mixes with the boundary layer clouds	Vertical transport from ship to cloud layer
Evolution of formed ship track	Ship location (although can be inferred from AIS)
Disappearance of part or all of the ship track	Ship exhaust which does not lead to ship track formation

Fig. 3-1. Left: A diagram of track formation from ship exhaust (https://ral.ucar.edu/staff/jwolff/aerosols.html/intro.html). The green box highlights the radiative effects that are observed from satellite sensors after track formation. Right: A list of phenomenon associated with track formation broken down into what is and is not measured by remote sensing instruments.

ship to boundary layer (height at which cloud formation starts) should be fairly vertical without much resistance but tracking it vertically through the clouds is not trivial.

Due to variations in of fuel types and quantities emitted and complex atmosphere dynamics, not all *emission bursts* will produce visible tracks. Thus, we only observe ship tracks under the appropriate conditions. This not only means that not all ship emissions will produce a ship track, but also that interruptions in the visibility of an existing ship track can occur when ships pass under different atmospheric conditions.

To the naked eye, new ship track observations appear in imagery directly above known ship locations due to the resolution of the imaging so it is reasonable think of the entire vertical transport path from ship to boundary layers nearly “instantaneous" with some epsilon error. For this reason, in this paper, we implicitly impose a known but random time lag between ship emissions and their first detection at the cloud top layer in our simulations. Existing ship track formations will then move with wind dynamics, a variable that is straight-forward to simulate and is independent of actual ship movement. The visible tracks then persist in the clouds for an unknown time as *ship tracks* until the aerosols are fully diffused into the atmosphere and are no longer distinguishable from the surrounding clouds.

3.2. Modeling aerosols using a Hidden Markov Model (HMM)

To model the formation and behavior of the aerosol tracks, in this section, we construct a state-space point process representation relating imaging observations of emission tracks to partially observed, known locations of aerosol emission bursts from ships. A constructed Hidden Markov Model (HMM) is outlined in this section to characterize the relationship between between

the image observations and partially observed truth. We are interested in building a computational model that can emulate the persisting behavior the ship tracks to understand how this behavior changes with changing atmospheric dynamics.

3.2.1. State-space representation

The true emission path is generated by the continuously emitted aerosol emission packets by a single ship over the spatial window $\mathcal{X} \subset \mathbb{R}^2$ up to time $T \in \left[0, \sum_{n=1}^{N-1} \Delta_{n,n+1}\right]$ where N is the number of frames and $\Delta_{n,n+1} > 0$ is the time between frames n and $n+1$ (typically between 5 and 15 minutes). For simplicity, we assume in this article that $\Delta_{n,n+1} \equiv \Delta$, so that $t_{n+1} - t_n = \Delta$ for all n.

We first define the unobserved spatio-temporal point process $\{X_n : (x, y, t_n) \in \mathbb{R}^2 \times \mathbb{R}\}$ which characterizes the true behavior of the aerosol emission bursts, continuously released prior to (and still visible at) time t_n. Second, we define the observed spatio-temporal point process $\{Y_n : (x, y, t) \in \mathbb{R}^2 \times \mathbb{R}\}$ which characterizes the patterns of the partially observed ship tracks in image frame n, generated by X_n. Using this state-space representation, we formulate a Hidden Markov Model relating the two processes. T can also be defined in terms of number of image frames such that $T \in \left[0, \sum_{n=1}^{N-1} \Delta_{n,n+1}\right]$ where N is the number of frames and $\Delta_{n,n+1} > 0$ is the time between frames n and $n+1$. For simplicity, and since many imagers tend to collect data at regular intervals, we assume that $\Delta_{n,n+1} \equiv \Delta$, so that $t_{n+1} - t_n = \Delta$ for all n.

The true emission path is generated by the continuously emitted aerosol emission packets by a single ship over the spatial window $\mathcal{X} \subset \mathbb{R}^2$ up to time $T > 0, T \in \mathbb{R}$, with $\mathcal{X}$ and time T typically defined by the imager or the user. Although in practice, the observed satellite imagery and our partially observed data Y_{t_n} is observed discreetly, we will treat time as continuous in our simulation model. For ship $k = 1 \ldots K$ which produces a track, we assume that its entire emission path is comprised of $P_k > 0$ aerosol bursts (packets) which may or may not become visible. Assuming that only k_{t_n} of K ships that are expected to be observed prior to T, have entered the window $\mathcal{X}$ by time $t_n < T$, for an arbitrary single track $k = 1 \ldots k_{t_n}$, only $p_{k_{t_n}} \leq P_k$ packets are expected to become visible. To show proof of concept, for now we will ignore the complex cloud dynamics and assume all emission packets reach reach the boundary layer clouds and become visible with time lag $< \epsilon$. This will allow us to start with a general simulation framework and build in more atmospheric conditions when needed at a later time.

In the region of interest $\mathcal{X}$, we denote the set of true positions or *states* of each packet as $\{\mathbf{x}_{i,n}\}_{i=1}^{p_{k_{t_n}}}$, where $\mathbf{x}_{i,n} \in \mathcal{X}$ denotes the state of the ith packet of emission track k at time t_n.

Existing ship tracks are only modified at the next time step t_{n+1} in three possible ways:

- the oldest aerosol emission packets at the end of the track diffuse completely and mix back into the atmosphere (leaving no detectable trace), or
- surviving packets diffuse and become less distinguishable as part of the track (but are still visible), according to cloud dynamics and wind motion, or
- new packets appear at the front of the track in the direction the ship movement.

These situations result in $p_{k_{n+1}}$ new states (locations) $\{\mathbf{x}_{i,t_{n+1}}\}_{i=1}^{p_{k_{n+1}}}$ in each of the new and existing emission tracks present at time t_{n+1}.

In practice, however, the full lifespan (from first appearance to permanent disappearance) of each emission packet is *unknown*. Instead, at each observed image frame n, the GOES-R ABI sensor captures a snapshot in time of all estimated packet locations without information on age of the packet, i.e. how long the observations have visibly persisted in the atmosphere. It is also the case that the locations of the emission packets over their lifespan are not unique and can share a location with another emission packet. Specifically, for a track k_{t_n}, a set of $o_{k_{t_n}} \leq p_{k_{t_n}}$ observations $\{\mathbf{y}_{i,t_n}\}_{i=1}^{o_{k_{t_n}}}$, is recorded, where $\mathbf{y}_{i,t_n} \in \mathcal{Y}$ denotes the state of the ith observation at time t_n. We may assume that $\mathcal{Y} = \mathcal{X}$.

At time $t_n \in \mathbb{R}$, a newly observed track can be generated from newly released emission packets into the atmosphere. Due to the complex dynamics of the atmosphere, it is not often possible to link new observations to their true source. An observed aerosol track from GOES-R is not always visible directly above the known ship location. Thus, we will assume that there is no information about which emission packet generates which observation. Since there is no ordering on the respective collections of emission packet states and measurements at time t_n, they can be naturally represented as finite spatial-temporal point processes. Specifically, for $n = 1, \ldots N$, we denote

$$X_{t_n} = \{\{\underbrace{\mathbf{x}_{1,t_n}, \ldots, \mathbf{x}_{p_{1,t_n}}}_{\substack{p_1 \text{ packets} \\ \text{from emission } 1}}\}, \ldots, \{\underbrace{\mathbf{x}_{k_{t_n},t_n}, \ldots, \mathbf{x}_{p_{k_{t_n},t_n}}}_{\substack{p_{k_{t_n}} \text{ packets} \\ \text{from emission } k_{t_n}}}\}\} \in \mathcal{F}(\mathcal{X}) \quad k_{t_n} \leq K$$

$$Y_{t_n} = \{\{\underbrace{\mathbf{y}_{1,t_n}, \ldots, \mathbf{y}_{o_{1,t_n}}}_{\substack{o_1 \text{ packets} \\ \text{from emission } 1}}\}, \ldots, \{\underbrace{\mathbf{y}_{k_{t_n},t_n}, \ldots, \mathbf{y}_{o_{k_{t_n},t_n}}}_{\substack{o_{k_{t_n}} \text{ packets} \\ \text{from emission } k_{t_n}}}\}\} \in \mathcal{F}(\mathcal{Y}) \quad 0 \leq o_n \leq p_{k_{t_n}}$$

where $\mathcal{F}(\mathcal{X})$ and $\mathcal{F}(\mathcal{Y})$ denote the collections of all finite subsets of $\mathcal{X}$ and $\mathcal{Y}$ respectively. The target point process X_{t_n} is referred to as the *multi-target state* and the measurement set Y_{t_n} is referred to as the *multi-target observation*. With this model specification, the objective is to recover the true states of emission packet point processes $X_{t_1}, X_{t_2}, \ldots, X_{t_N}$ from their measurement sets $Y_{t_1}, Y_{t_2}, \ldots, Y_{t_N}$.

3.2.1.1. Multi-target state model

In this section, we describe a finite point process model for the time evolution of the multiple-target state $X_{t_n}, n = 1, \ldots, N$, which incorporates emission packet motion, birth and death. Specifically, we mathematically define the processes of aerosol packets first being conceived in boundary layer clouds, their motion and diffusion through the atmosphere until their permanent disappearance.

After an aerosol track has already formed at time t_{n-1}, if an emission packet $\mathbf{x}_{t_{n-1}} \in X_{t_{n-1}}$ which makes up part of that track survives to time $t_{n+1} > t_n$, its subsequent state is determined by a *drift* term which is described by the wind motion at $\mathbf{x}_{t_{n-1}}$, and a *diffusion* term which describes the diffusion of the emission packet within the clouds it is situated in. This type of process is known as

a **Markov diffusion process** and is described by the following (continuous time) stochastic differential equation:

$$d\mathbf{x}_t = \underbrace{\mu(\mathbf{x}_t, t)}_{\text{drift}}\, dt + \underbrace{\sigma(\mathbf{x}_t, t)}_{\text{diffusion}}\, dB_t, \tag{3.1}$$

where $B_t \sim \mathcal{N}_2(\mathbf{0}, tI_2)$ denotes a standard Brownian motion in two dimensions, with I_2 denoting the 2-dimensional identity matrix. The drift function $\mu(\mathbf{x}_t)$ denotes the **wind velocity** at point $\mathbf{x}_t$, at time t and is in general known. For this problem, we choose the diffusion function $\sigma(\mathbf{x}_t) \equiv \sigma_x$ to be a **constant** that describes the diffusivity of an aerosol parcel within the atmospheric boundary layer. The solution to (3.1) with a changing wind velocity in space and time is in general unknown and requires numerical solvers which may be computationally cumbersome and time consuming. For simulation purposes, we therefore propose the following approximation.

Given discrete time intervals of the form $I_n = (t_n, t_{n+1}] \equiv (n\Delta, (n+1)\Delta]$, with $n \in \mathbb{Z}^+$, we assume that the simulation interval time $t_{n+1} - t_n = \Delta$ is taken small enough so that the wind velocity within the interval is approximately constant. That is to say, for a continuous time point $t \in I_n$, we use the approximate SDE

$$d\mathbf{x}_t = \mu(\mathbf{x}_t)dt + \sigma_x dB_t, \tag{3.2}$$

with $\mu(\mathbf{x}_t)$ denoting the wind velocity field for a parcel with state $\mathbf{x}_t$ at time $t_n = n\Delta$. Given a previous state $\mathbf{x}_s$ at time $s \in I_n, t > s$, equation (3.1) can be solved directly

$$\begin{aligned} d\mathbf{x}_t &= \mu(\mathbf{x}_t)dt + \sigma_x dB_t \\ \implies \mathbf{x}_t - \mathbf{x}_s &= \int_s^t \mu(\mathbf{x}_t)\, \mathrm{d}w + \sigma_x(B_t - B_s) \\ &= \mu(\mathbf{x}_t)(t-s) + \sigma_x B_{t-s}, \end{aligned}$$

where $B_t - B_s \overset{D}{\equiv} {}^1 B_{t-s} \sim \mathcal{N}_2(\mathbf{0}, (t-s)I_2)$. This implies the corresponding transition density is

$$\mathbf{x}_t | \mathbf{x}_s \sim \mathcal{N}_2(\mathbf{x}_s + \mu(\mathbf{x}_s)(t-s), \sigma_x^2(t-s)I_2).$$

In particular, the probability density of the parcel's transition to state $\mathbf{x}_{t_n} \in X_{t_n}$ from state $\mathbf{x}_{t_{n-1}}$ is given by the Markovian density $f^M_{t_n|t_{n-1}}(\mathbf{x}_{t_n}|\mathbf{x}_{t_{n-1}}) \sim \mathcal{N}_2(\mathbf{x}_{t_{n-1}} + \mu(\mathbf{x}_{t_{n-1}})\Delta, \sigma_x^2 \Delta I_2)$. Its behavior at this time is therefore modeled by the point process $S_{t_n|t_{n-1}}(\mathbf{x}_{t_{n-1}})$, where

$$S_{t_n|t_{n-1}}(\mathbf{x}_{t_{n-1}}) = \begin{cases} \mathbf{x}_{t_n} & \text{where } \mathbf{x}_{t_n} \sim f^M_{t_n|t_{n-1}}(\cdot|\mathbf{x}_{t_{n-1}}) \text{ with probability } p_{S,t_n}(b_{\mathbf{x}_{t_{n-1}}}) \\ \emptyset & \text{otherwise.} \end{cases} \tag{3.3}$$

Here, $p_{S,t_n}(b_{\mathbf{x}_{t_{n-1}}})$ denotes the *survival probability* of packet $\mathbf{x}_{t_{n-1}}$ at time t_n (described in more detail below) and where the motion diffusion coefficient σ_x is unknown and requires estimation from the model.

[1] $\overset{D}{\equiv}$ denotes equivalence in distribution.

A new emission packet at time $t_n \in \mathbb{R}$ can arise in two ways. The first is as a spontaneous birth (of a newly risen emission track), which is independent of any existing emission track. The second is by *spawning* from an existing emission source, resulting in a newly visible emission packet. We denote the *birth time* of packet $\mathbf{x}_{t_n}$ (partially) observed at time t_n as $b_{\mathbf{x}_{t_n}}$.

Spontaneous births of new emission tracks at time t are denoted by the finite point process Γ_t. We model Γ_t as a finite **Poisson point process** with *intensity function* $\gamma_t(\mathbf{x}) = \lambda_{\gamma_t} f_{b,t}(\mathbf{x}))$, for $\mathbf{x} \in \mathcal{X}$:

$$\Gamma_t \sim \text{Poisson}(\lambda_{\gamma_t} f_{b,t}(\mathbf{x})). \tag{3.4}$$

- Here, $N_{b,t} \sim \text{Poisson}(\int_{\mathcal{X}} \lambda_{\gamma_t} f_{b,t}(\mathbf{x})\text{d}\mathbf{x})$ denotes the number of births occurring in $\mathcal{X}$ at time t.
- $f_{b_t}(\mathbf{x})$ denotes their spatial distribution.

Assuming we have (simulated or real) of the boat positions/path that produce these new emissions, we may let this inform $f_{b_t}(\mathbf{x})$. Specifically, if $\mathbf{x}_{b,t_n}$ is the position of a new boat at time t_n, then $f_{b_{t_n+\epsilon}}(\mathbf{x}) = \mathcal{N}_2(\mathbf{x}_{b,t_n}, \sigma_b^2 I_2)$ where ϵ denotes the time lag between ship emission and aerosol observation at the cloud boundary layer.

Spawned births occurring within the time interval I_{n-1} denote newly visible emission packets from existing emission tracks that reach the cloud top layer at time t_n. Newly spawned targets can only be spawned by packets that were birthed in the previous time interval I_{n-2}, as this models the continuous emission of aerosol packets in a single stream.

We model the set of spawned births $B_{t_n|t_{n-1}}(\mathbf{x}_{t_{n-1}})$ at time t_n from a packet $\mathbf{x}_{t_{n-1}}$ at time t_{n-1} as a finite point process. An example used in this paper is **Bernoulli point process** with spawning probability p_{β,t_n}:

$$B_{t_n|t_{n-1}}(\mathbf{x}_{t_{n-1}}) = \begin{cases} \{\mathbf{x}\};\ \mathbf{x} \sim f^{\beta}_{t_n|t_{n-1}}(\mathbf{x}|\mathbf{x}_{t_{n-1}}) \text{ with probability } p_{\beta,t_n} & t_{n-2} < b_{\mathbf{x}_{t_{n-1}}} \leq t_{n-1} \\ \emptyset & \text{otherwise.} \end{cases}$$

1. The number of spawned targets N_{s,t_n} from $\mathbf{x}_{t_{n-1}}$ follows $N_{s,t_n} \sim \text{Bernoulli}(p_{\beta,t_n})$.
2. Therefore <u>at most one</u> packet can be spawned from a target $\mathbf{x}_{t_{n-1}}$ born in the previous time step.
3. If $N_{s,t_n} = 1$, then the spatial distribution of the spawned target $\mathbf{x}$ follows $f^{\beta}_{t_n|t_{n-1}}(\cdot|\mathbf{x}_{t_{n-1}})$ from $\mathbf{x}_{t_{n-1}}$.

In this paper, we assume knowledge of ship positions that continuously emit aerosols whilst moving, thereby corresponding to this spawning process. For simulation purposes, we therefore use the spawning density

$$f^{\beta}_{t_n|t_{n-1}}(\mathbf{x}|\mathbf{x}_{t_{n-1}}) = \mathcal{N}_2(\mathbf{x}_{b,t_n-\epsilon} + \epsilon\sigma_\beta^2 I_2).$$

The spawning probability p_{β,t_n} is directly related to the number of aerosol packets each ship emits during the observed time window. For simulation purposes, we assume that each boat continuously emits aerosols up to the simulation time $T = N\Delta$ and that they exist when within the observed window $\mathcal{X}$. This enables $p_{\beta,t_n} = 1$ when $t_n \leq T$ and is zero otherwise.

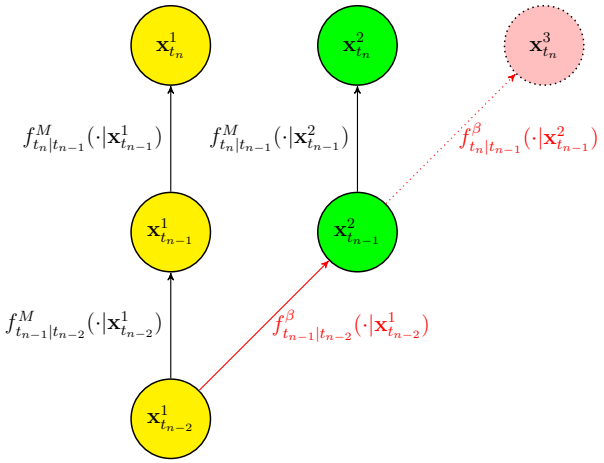

Fig. 3-2. The motion and spawning process. The first packet $\mathbf{x}^1_{t_{n-2}}$ (indicated in yellow) born at time t_{n-2} undergoes Markovian motion (black) to reach state $\mathbf{x}^1_{t_{n-1}}$ and spawns a target (indicated in green) $\mathbf{x}^2_{t_{n-1}}$ at time t_{n-1}. This packet undergoes Markovian motion (black) to reach state $\mathbf{x}^2_{t_n}$ and is to spawn a packet (indicated in pink) $\mathbf{x}^3_{t_n}$ at time t_n.

Figure 3-2 illustrates the motion and spawning process of an aerosol packet described by our procedure.

For a given multi-target state $X_{t_{n-1}}$ at time t_{n-1}, each packet $\mathbf{x} \in X_{t_{n-1}}$ either continues to exist (survives) at time $t_n > t_{n-1}$ with probability $p_{S,t_n}(\mathbf{x}, b_\mathbf{x})$, or "dies" with probability $1 - p_{S,t_n}(\mathbf{x}, b_\mathbf{x})$. Here, a "death" of an emission packet occurs when it sinks back through the atmosphere and ceases to be visible. Furthermore, the survival probability of each emission packet is written as a function of the time t_n, its spatial location $\mathbf{x}$, and the packet's "birth" time $b_\mathbf{x}$ in the region. However, since the effects of up and downward drafts in the atmosphere on each packet can be considered negligible, this enables the survival probability to only be a function of t and $b_\mathbf{x}$, i.e. that $p_{S,t}(\mathbf{x}, b_\mathbf{x}) \equiv p_{S,t}(b_\mathbf{x})$.

In this simulation model, we assume that each boat produces a cloud-aeorosol track that has an *average* lifetime $T_d \sim \text{Exp}(\lambda_T)$ from birth. Given T_d, the individual aerosol packets that are contained in its emission then each have an independently and identically distributed (i.i.d.) death time

$$d \sim \text{Log-normal}\left(\mu_d = \log\left(\frac{T_d}{\sqrt{\sigma_{p_d}^2 + T_d^2}}\right), \sigma_d^2 = \log\left(\frac{\sigma_{p_d}^2 + T_d^2}{T_d^2}\right)\right),$$

where $\sigma_{p_d}^2$ is the variance of the packet death time and is a fixed simulation input.

Altogether, we have a multi-target state X_{t_n} of the following form

$$X_{t_n} = \underbrace{\left[\bigcup_{\mathbf{x} \in X_{t_{n-1}}} S_{t_n|t_{n-1}}(\mathbf{x})\right]}_{\text{Survived packets}} \cup \underbrace{\left[\bigcup_{\mathbf{x} \in X_{t_{n-1}}} B_{t_n|t_{n-1}}(\mathbf{x})\right]}_{\text{Spawned packets}} \cup \underbrace{\Gamma_{t_n}}_{\text{New emissions}}. \tag{3.5}$$

An important modeling characteristic is that the unions in (3.5) are *independent*.

3.2.1.2. Multi-target observational model

In this section, we describe a finite point process model for the time evolution of the multiple-target observation, $Y_{t_n}, n = 1, \ldots, N$, which incorporates observations generated from emission tracks. and those falsely generated by changes in cloud or atmospheric humidity.

When a packet $\mathbf{x}_{t_n} \in X_{t_n}$ is generated according to the process described above, an observation of it $\mathbf{y}_{t_n} \in Y_{t_n}$ is generated from an observational model $g_{t_n}(\cdot|\mathbf{x}_{t_n})$. This function is typically chosen to take the form $\mathbf{y}_{t_n}|\mathbf{x}_{t_n} \sim \mathcal{N}_2(\mathbf{x}_{t_n}, \Sigma_{\mathbf{x}_{t_n}})$, where $\Sigma_{\mathbf{x}_{t_n}}$ can be taken to be the *marginal* covariance of $\mathbf{x}_{t_n}$. Specifically, for packet $\mathbf{x}_{t_n}$ birthed at time $b_{\mathbf{x}_{t_n}}$, its marginal density can be calculated via

$$f(\mathbf{x}_{t_n}) = \int_{\mathcal{X}} f^M_{t_n|b_{\mathbf{x}_{t_n}}}(\mathbf{x}_{t_n}|\mathbf{x}_{b_{\mathbf{x}_{t_n}}})\pi(\mathbf{x}_{b_{\mathbf{x}_{t_n}}})\, \mathrm{d}\mathbf{x}_{b_{\mathbf{x}_{t_n}}},$$

with $\pi(\mathbf{x}_{b_{\mathbf{x}_{t_n}}})$ being the initial probability density of packet $\mathbf{x}_{t_n}$ in $\mathcal{X}$ at the time of its birth. For this paper, we take $\pi(\mathbf{x}_{b_{\mathbf{x}_{t_n}}}) = \delta_{\mathbf{x}_{b_{\mathbf{x}_{t_n}}}}(\mathbf{x}_{b_{\mathbf{x}_{t_n}}})$, the dirac delta function centered at $\mathbf{x}_{b_{\mathbf{x}_{t_n}}}$, yielding $\Sigma_{\mathbf{x}_{t_n}} = \sigma_x^2(t_n - \mathbf{x}_{b_{\mathbf{x}_{t_n}}})I_2$ and

$$\mathbf{y}_{t_n}|\mathbf{x}_{t_n} \sim \mathcal{N}_2(\mathbf{x}_{t_n}, \sigma_x^2(t_n - \mathbf{x}_{b_{\mathbf{x}_{t_n}}})I_2).$$

When simulating across pixelated grids, we discretize the above equation such that the pixel intensity of a pixel P at time t_n denoted $\mathcal{I}_{t_n}(P)$ follows

$$\mathcal{I}_{t_n}(P) \propto \sum_{\mathbf{y} \in Y_{t_n}} \int_P f(\mathbf{y}|\mathbf{x}_{t_n})\, \mathrm{d}\mathbf{y}$$

with the normalization constant given by the highest pixel intensity simulated across the video.

For observations generated by true emission packets, we note that a packet $\mathbf{x} \in X_t$, at time t is only detected by satellites with probability $p_{D,t}(\mathbf{x})$. This *detection probability* has a spatio-temporal dependence structure which is needed to first, model the spatial randomness of cloud humidity/density and second, to account for cloud movement across the observation time window. In the field of view $\mathcal{X}$, if the cloud humidity is too low or too high, emission packets cannot be detected. In the former case, packets cannot be observed since clouds cannot form to produce the necessary observations. In the latter, the cloud density may be too high, or may already be contaminated with existing aerosols which would subsequently not produce observations of new packets.

To deal with this, we may choose to model $p_{D,t}(\mathbf{x})$ as a function of the existing cloud humidity/density. This may be formulated by modeling pixel intensities measured by the GOES-R ABI sensor, and utilizing a lower and upper threshold ι_L, ι_U. For example, setting

$$p_{D,t_n}(\mathbf{x}) = \begin{cases} 1 & \text{if } \iota_L < \mathcal{I}_{t_n}(\mathbf{x}) < \iota_U \\ 0 & \text{otherwise,} \end{cases} \tag{3.6}$$

enables a packet to be observed with probability one if its true location $\mathbf{x}$ lies within a pixel of the nth frame, with an intensity $\bar{I}_{t_n}(\mathbf{x}) \in (\iota_L, \iota_U)$, sufficient for it to be observed via satellite.

Subsequently, the **observational point process** $\Theta_{t_n}(\mathbf{x}_{t_n})$ from an emission packet $\mathbf{x}_{t_n} \in \mathcal{X}$ follows

$$\Theta_{t_n}(\mathbf{x}_{t_n}) = \begin{cases} \{\mathbf{y}\} \text{ where } \mathbf{y} \sim \mathbf{y}_{t_n}|\mathbf{x}_{t_n} & \text{with probability } p_{D,t_n}(\mathbf{x}_{t_n}) \\ \emptyset & \text{with probability } 1 - p_{D,t_n}(\mathbf{x}_{t_n}). \end{cases} \tag{3.7}$$

Altogether, we have a multi-target observation Y_{t_n} of the following form

$$Y_{t_n} = \bigcup_{\mathbf{x} \in X_{t_n}} \Theta_{t_n}(\mathbf{x}). \tag{3.8}$$

In this paper, we have described a computational method to simulate cloud-aerosol tracks given wind and boat simulated fields, using a stochastic differential equation that incorporates aerosol packet birth, motion, diffusion and death. A simulation example has been provided to highlight each step of our algorithm.

Using the presented methodology, a next step would be to verify that this surrogate model is accurate in representing cloud-aerosol paths that are observed in satellite imagery. This is challenging as real cloud-aerosol tracks have an unknown or unidentifiable source and the relationship between observed atmospheric properties and track behavior is not trivial to infer from imagery alone. Incorporating observed wind data from ERA-5 reanalysis and available atmospheric information that are well-documented to contribute to cloud track formation such cloud condensation nuclei (CCN) and liquid water paths (LWP) also available from reanalysis products such as MERRA-2, into an improved simulation algorithm would aid in the simulation of realistic cloud-aerosol behaviors. This would thereby allow us to focus on developing methodology and inference mechanisms to estimate atmospheric conditions under which ship tracks form or not form, which can then directly inform simulation inputs and unknown aerosol properties with particular reference to their longevity in the atmosphere. Additionally, this level of inference would require labeling image pixels and extracting observation points of cloud tracks, which may necessitate feature extraction algorithms such as Convolution Neural Networks (CNN) such as the one developed in [59] or the detection algorithm we have developed and demonstrate in chapter 4.

4. A STATISTICAL AND PHYSICS-BASED APPROACH TO AUTOMATICALLY LABELING SHIP TRACKS

4.1. Background

Identification of ship tracks is also important to regulators seeking to identify noncompliance with the 2020 International Maritime Organization (IMO) regulations, which limits the burning of low-grade, high-sulphur content fuels without any type of exhaust treatment [31]. [18] suggested that regulators might use satellite imagery to detect noncompliance with the IMO regulations out on the open ocean where it is more difficult to monitor ship activity, and they successfully demonstrated that NO2 emission plumes from ships can be detected and attributed to individual ships using satellite measurements.

Researchers have discovered that many ship tracks go undetected using current identification methods [e.g., 22]. Ship tracks may not be readily apparent because they do not typically follow straight lines or other low-order polynomial curves [57]. Sometimes tracks become hidden behind dense clouds and only reappear again under the appropriate atmospheric conditions. On the other hand, many linear cloud patterns may look like ship tracks when they are not. Manually identifying ship tracks in satellite data is a time-consuming and exhaustive process; thus, researchers have been developing advanced tools and algorithms to automatically label ship tracks in satellite images. Although, none aim to limit false detections leaving researchers to still resort to manual labeling ([e.g. 45, 10, 22]. An effective method for automatic ship track detection will enable more efficient data processing of satellite imagery and lead to the curation of large ship track datasets [59] to support stronger scientific conclusions than those made from small case studies [e.g. 14, 25, 11, 10] and modeling improvements.

In the early years of ship track studies, [57] developed a multi-step automated approach for detecting ship tracks in images from the Advanced Very High-Resolution Radiometer (AVHRR) sensor aboard the National Oceanic and Atmospheric Administration polar orbiting satellites. The approach consists of preprocessing images to enhance contrast between ship tracks and the surrounding background; detecting the ship tracks using a technique called ridge detection by "iterated erosion" to account for varying track widths; and post-processing the enhanced images using a connected components analysis that effectively eliminates ridges not associated with ship tracks. The results show that it is possible to detect ship tracks using this method, but differences in ship track brightness and image quality can significantly affect the results. More than twenty years later, [59] automated the detection of ship tracks by training a deep neural network to detect ship tracks in night-time satellite images from the MODerate resolution Imaging Spectroradiometer (MODIS) on the Aqua satellites. Specifically, they trained a multi-layer convolutional neural network with skip connection architecture. Their results were checked against a database of manually labeled ship tracks and they show model performance as high as a 91% rate of detecting

real tracks. The research also revealed that ship tracks are much more prevalent than previously thought over ocean areas with high ship traffic. Both [57] and [59]show progress towards automated ship track detection but both methods rely solely on multispectral imager and are thus equally likely to detect linear non-persistent cloud features leading to high false positive detection rates and still leaving the user to manually remove falsely identified tracks. With both approaches, it is also challenging to back out any inference on the detection results, e.g., understand why the algorithm detected a particular track and with what uncertainty the identified feature is a real ship track. This is critical to reducing the false positive detection rate and understanding how ship tracks form.

In this chapter, we introduce a novel statistical approach to ship track identification composed of three steps. We first implement a clustering algorithm using a directional spatial filter designed to detect the quasi-linear shapes of ship tracks. This filter is designed to be rather permissive so that it identifies a high percentage of true ship tracks in our image as well as possible false tracks. Second, we apply simple threshold filters, or masks, to remove false tracks that are either fleeting or appear over a significant portion of land where tracks are not expected. Third, we use the probabilistic feature fusion technique of [48] to remove many remaining false tracks using distinguishing physical features of ship tracks. With this probabilistic framework, we can weight physical features such as length and width of a track and background conditions under which tracks most likely form or do not form, e.g., cloud top height and cloud optical depth, according to their contribution to true track identification. We demonstrate our approach on three case studies using multi-spectral radiance and derived product data collected from the Geostationary Operational Environmental Satellite-R series (GOES-R) Advanced Baseline Imager (ABI). Using our directional filter combined with informative background conditions, track persistence information, and physical features of ship tracks, we have developed a tailored approach to accurately and systematically identify ship tracks that might otherwise be indistinguishable from similar cloud features. The remainder of this chapter is organized as follows: Section 4.2 outlines our data sources and Section 4.3 outlines our methods and presents some examples;

4.2. Data

To demonstrate our approach, we use remotely sensed data, both raw multi-spectral and derived, from the GOES-R ABI sensor (GOES-R Calibration Working Group and GOES-R Series Program, 2017). The GOES-R ABI is a multi-spectral channel, two-axis scanning radiometer designed to provide radiometrically calibrated and geolocated observations of the Earth (GOES-R Series Product Definition and Users' Guide, 2019). From the L1b radiance data collected from GOES-17, we rely on the near-infrared "cloud particle size" band (C06) and the infrared "shortwave window" band (C07) with central wavelengths of 2.24 and 3.90 μm, respectively. To visualize ship tracks day or night, we follow section 2.2 and transform the data for a given image by subtracting spectral band C07 from band C06. The spatial resolution of both bands (C06 and C07) is 2 km, and the temporal resolution is every five to fifteen minutes depending on the image scan. Similar to previous sections, we focused on the CONUS image scan which is focused over the North Pacific Ocean and has a temporal resolution of five minutes.

4.2.1. GOES-R products

We additionally explored GOES-R meteorological and solar derived products to understand key observed environments that are susceptible to cloud brightening. The details of 30+ products can be found at https://www.goes-r.gov/products/overview.html. Note that many reanalysis products and products from polar orbiting satellites in the A-train such as the Aqua, Terra, CloudSAT and CALIPSO satellites but we did not have time to explore these sources and most have low temporal resolution.

Specifically, we narrowed our focus aerosol and cloud optical depth (AOD and COD), total precipitable water (TPW), cloud top height (CTH), cloud top pressure (CTP) and Convective Available Potential Energy (CAPE) , a derived stability index, to help discriminate true ship tracks from falsely labeled tracks and show proof of concept that this additional information improves classification results. These variables were chosen after two screening studies were performed with the goal of identifying key atmospheric measurements that define background environments that are most susceptible to aerosol perturbation and thus environments in which we are most likely to observe ship tracks.

In the first study we manually labeled observed ship tracks from three 2–4 day time periods in July, September and December of 2020 and identified the positions of the ships creating these tracks from the AIS data base described in section 2.2. We also identified a number of ships which were not producing tracks in similar regions to those actively producing tracks as a control group. We then collected GOES-R product data from a small box of pixels around the ship locations 30-60 minutes prior to track formation. This data was used to fit a logistic regression across each time period accounting for the temporal lag. Although the fit of this model was not great, a common theme arose from the significant covariates found from this study. From the three different studies, we found that the most significant covariates owing to ship track formation were:

1. **Aerosol Optical Depth (AOD)**, which measures the amount of light lost due to the presence of aerosols on a vertical path through the atmosphere.

2. **Cloud Height Top (HT)**, which identifies the Geopotential height at the top of a cloud layer.

3. **Total Precipitable Water (TPW)** which is the amount of water that can be obtained from the Earth's surface to the top of the atmosphere from condensation.

4. **Derived Stability Index Convective Available Potential Energy (CAPE)** which is a measure of atmospheric stability, in particular for the development of cumulus clouds contributing to severe weather hazards.

In the second study, we looked at slightly larger subregions within an image across a set of days and manually labeled whether ship tracks appeared in these subregions or not. We explored a number of days and different box sizes for data collection on the presence of ship tracks. A more extensive exploration is likely needed, but this was beyond the scope of our current research effort. For discussion here, we present results for the three-day period from August 30, 2020 through September 1, 2020. Figure 4-1 shows a sample day and hour with AIS ship locations and a set of manually-identified boxes that either contain tracks or do not. Areas that do not contain tracks were

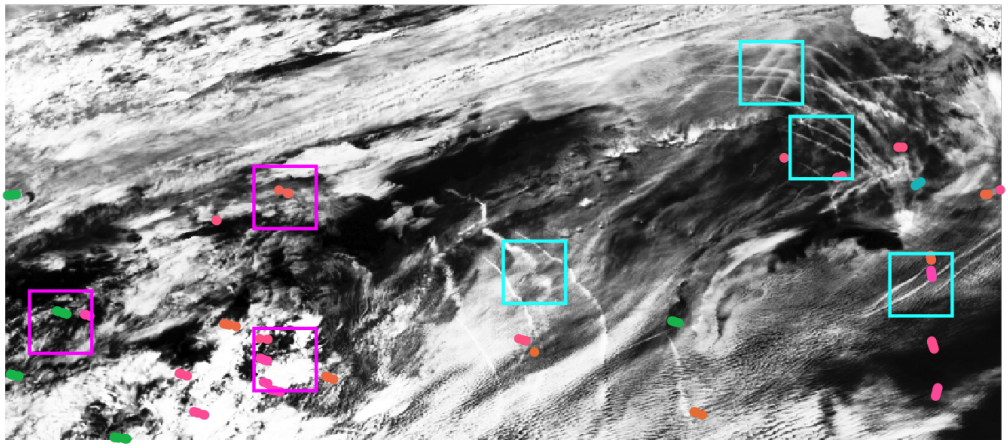

Fig. 4-1. Plot of the North Atlantic Ocean for September 1, 2020 at 22:02:00. The boxes show a sample of manually-identified regions that either do (cyan) or do not (magenta) contain ship tracks. The points represent ship locations at the time of the image.

chosen based on the location of ships, as any area without a ship cannot be determined to be incapable of producing tracks.

For the set of three days, we collected atmospheric data that corresponds with the manually-identified boxes, as we seek to identify differences between the subregions that either do or do not contain ship tracks. For this example, we identified 29 subregions that do contain tracks and 27 subregions that do not at four time periods per day across the three days (averaging 4.6 subregions per image). Each subregion is 50 x 50 pixels. We experimented with different size subregions, but for simplicity will only show one set of results here. We compared values for Total Precipitable Water (TPW), Cloud Top Pressure (CTP), Cloud Optical Depth (COD), and Cloud Top Height (HT). Figure 4-2 shows box plots comparing the values of these values for the subregions with and without tracks. Largely, the values are comparable on average, despite often stark visual differences in the plotted imagery. Of greatest promise is the Cloud Optical Depth, which shows a slight difference in mean value as well as a very long tail for regions without tracks. While much more effort is needed to collect more data, determine the optical subregion size, and expand our manual identification process, this small set of initial results shows some promise in improving our understanding of key atmospheric conditions for track formation and persistence.

In summary, many variables from the available GOES-R L2 products that we have identified to correspond with track formation conditions agree with existing results from modeling studies. Although, as expected, the measurement noise present in this observational data reduces the signal we are looking for and limits our ability to clearly distinguish conditions that are more susceptible to perturbations from ship exhaust. More advanced methods are needed to account for this noise. Atmospheric measurements from instruments with potentially more relevant atmospheric data and higher spatial resolutions, such as those in the polar-orbiting https://atrain.nasa.gov/, would provide additional insights. Although the low temporal resolution of polar-orbiting instruments would

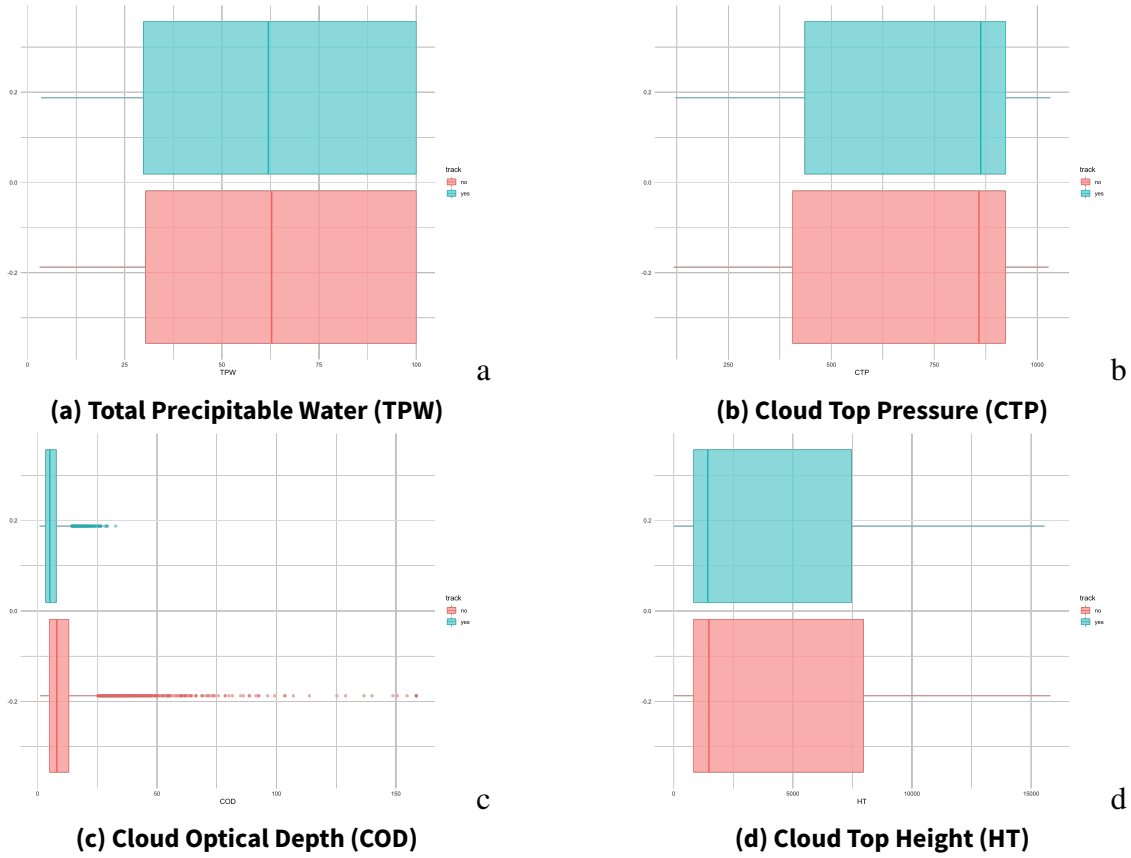

a

(a) Total Precipitable Water (TPW)

b

(b) Cloud Top Pressure (CTP)

c

(c) Cloud Optical Depth (COD)

d

(d) Cloud Top Height (HT)

Fig. 4-2. Boxplots showing the distribution of four atmospheric variables for small subregions that either do (blue) or do not (red) contain ship tracks.

inhibit our understanding of track formation across a 24-hr period.

4.3. Methods

Our technique for automated identification of ship tracks incorporates three distinct algorithmic stages: *pre-processing*|of GOES imagery to enhance the desired signatures, *detection* of candidate ship tracks in the preprocessed imagery, and *classification* of candidate ship tracks by measuring their consistency with a trained statistical model. This general approach is consistent with that of traditional automatic target recognition (ATR) systems [15]. Furthermore, we adopt ATR terminology where the "target" class represents valid ship tracks and the "non-target" class consists of all other patterns seen in marine-environment clouds. Here, we will discuss the pre-processing and detection stages in sequence.

4.3.1. Pre-processing

The preprocessing stage incorporates operations that both smooth and normalize input image data. These operations are controlled by the parameters introduced in Table 4-1.

Table 4-1. Parameters controlling the preprocessing stage.

Parameter	Description
BOX	Size of median box applied in first smoothing step
GB	Guard band applied in directional filtering
$BASE$	Baseline pixels used in directional filtering
$HALF$	Half-width of final median smoother

The first operation is the application of a median box filter over the entire image, to smooth pixels that are unusually bright or dim without blurring any scene edges. The parameter BOX represents the size of each side of the centered area over which the filter operates, and should be an odd number. With $BOX = 1$, this step becomes an identity mapping.

Ship tracks are elongated bands of pixels that are brighter than their neighboring pixels. They may be oriented in any direction relative to the axes of the image, and can gradually change in orientation with time, due to motion of the overlying clouds or adjustments to the motion trajectory of the vessel generating the track. The directional filters applied in the next operation highlight narrow bands of locally bright pixels using directional filters that compare each pixel's intensity to its local neighborhoods in the row and column directions. *Vertical* gradients are highlighted by operating across columns in a single row, while *horizontal* gradients are highlighted by operating across rows in a single column. Denote by x_{ij} the pixel located in row i and column j of the input image. Each filter implements a guard band of width GB on either side of the pixel being processed, and estimates the local background using $BASE$ pixels on either side of $(i, j) \pm GB$. The appropriate size for the guard band depends on the typical pixel width of ship tracks in the input image; for GOES-R imagery we have found that values in the range from 2 – 4 generally work well. The baseline size should be sufficient to allow for estimation of the mean and variance, without extending the neighborhood too far from the pixel being filtered. We set the default value at $BASE = 5$, which provides a total of ten pixels for background estimation.

The mean and variance of background pixels for the vertical gradient filter applied across columns in row i are calculated as:

$$\bar{x}_{ij(V)} = \frac{1}{2 \times BASE} \left(\sum_{k=j-GB-BASE}^{j-GB-1} x_{ik} + \sum_{k=j+GB+1}^{j+GB+BASE} x_{ik} \right)$$

and

$$s^2_{ij(V)} = \frac{1}{2 \times BASE - 1} \left(\sum_{k=j-GB-BASE}^{j-GB-1} (x_{ik} - \bar{x}_{ij(V)})^2 + \sum_{k=j+GB+1}^{j+GB+BASE} (x_{ik} - \bar{x}_{ij(V)})^2 \right).$$

The mean, denoted $\bar{x}_{ij(H)}$, and variance, $s^2_{ij(H)}$, for the horizontal gradient filter applied to column j is similarly defined, but with summation across rows. Guard band and background areas for vertical gradient filtering of a single pixel are depicted in Figure 4-3.

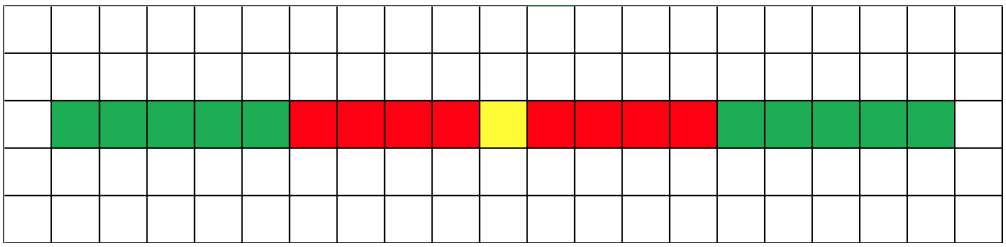

Fig. 4-3. Layout for filtering to highlight vertical gradients with $GB = 4$ and $BASE = 5$. Because ship tracks are more than one pixel wide, a guard band (red) is implemented to keep the pixels within a track outside of the background estimate (green) for a pixel being filtered (yellow).

The horizontally and vertically filtered values for pixel (i, j) are computed as

$$z_{ij(H)} = \frac{x_i j - \bar{x}_{ij(H)}}{s_{ij(H)}} \quad \text{and} \quad z_{ij(V)} = \frac{x_i j - \bar{x}_{ij(V)}}{s_{ij(V)}}, \quad \text{respectively.}$$

Figure 4-4 shows an input GOES-R image, along with outputs of the horizontal and vertical gradient filters.

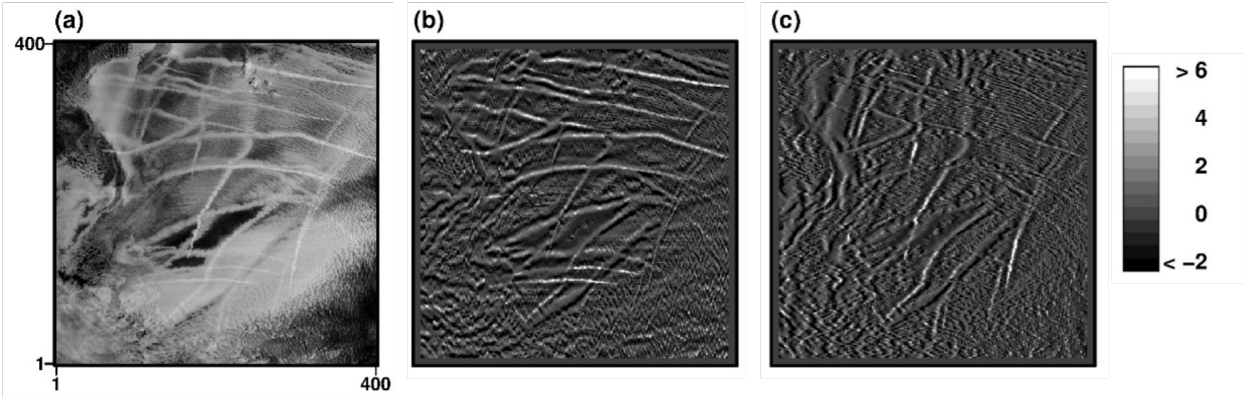

Fig. 4-4. Illustration of the directional filters for a data subset from the GOES-17 CONUS scan on April 24, 2019 at 17:02 UTC. (a) 400 × 400 px image input to the filters (GOES-R band C06 - band C07). (b) Output from horizontal gradient filter. (c) Output from vertical gradient filter. The legend at right represents the greyscale for panels (b) and (c), which are both expressed in units of standard deviations above or below the neighborhood mean, as in equations (3) and (4). Parameter values $BOX = 1$, $GB = 3$, $BASE = 5$, $HALF = 0$.

In the filtered outputs shown in panels (b) and (c) of Figure 4-4, the ship tracks are characterized by long streaks of bright pixels that maintain the same general orientation over a reasonable distance. By contrast, many of the bright outputs caused by other cloud formations (as in the lower left-hand corner of the image) are more spatially compact. To further highlight ship tracks relative to other cloud patterns, we may apply a second smoothing operation in the space of the filtered outputs. This is accomplished by running medians along three orientations, centered on each pixel in the filtered image. Because the goal is to smooth along candidate ship tracks, no guard band is applied

during these secondary smoothing operations. For the horizontally-filtered outputs, smoothing is conducted in the horizontal direction (to enhance horizontal features), along the main diagonal (to enhance features oriented in a NW-SE orientation), and along the off diagonal (to enhance features oriented in a SW-NE orientation). For vertically-filtered outputs, smoothing occurs in the vertical dimension along with the same two diagonals. The length of the secondary smoother is controlled by the parameter $HALF$, which specifies the number of pixels on each side of the center box that are included in the calculation. Choosing $HALF = 6$ gives a filter length of 13, which is small enough to prevent the cancellation of cloud features at intermediate orientations.

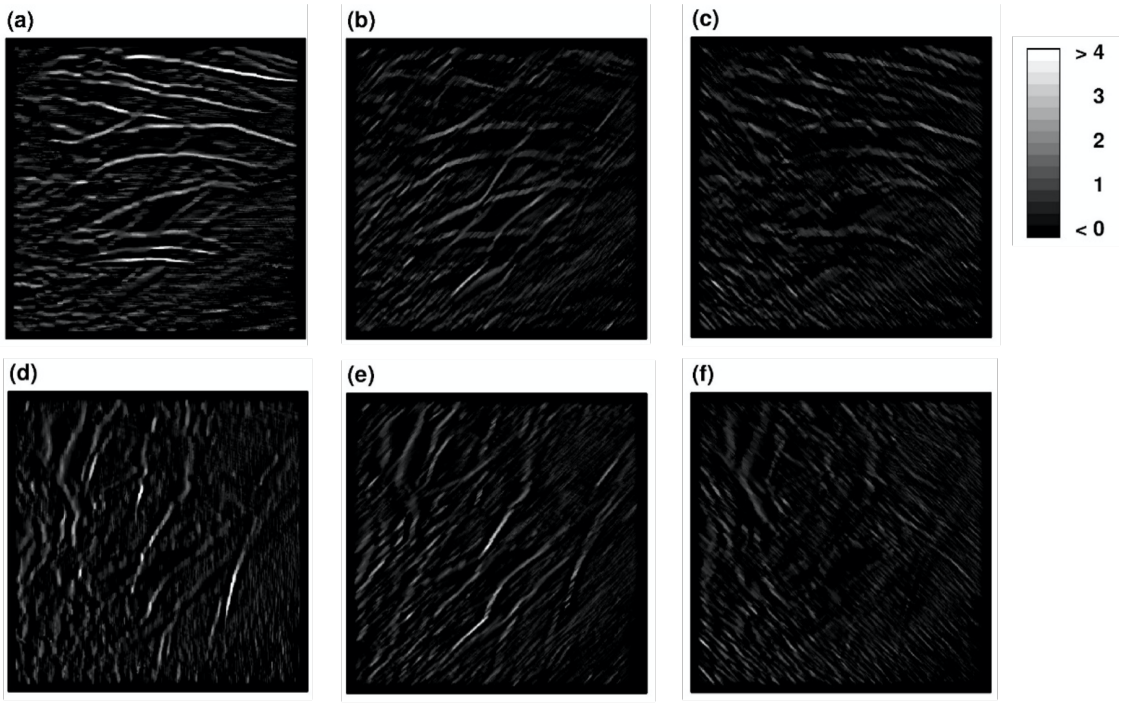

Fig. 4-5. Secondary smoothers applied to directionally filtered data shown in Figure 4-4. (a) Horizontally filtered image (Figure 4-4b), smoothed in the horizontal direction. (b) Horizontally filtered image smoothed along the off-diagonal (SW-NE). (c) Horizontally filtered image smoothed along the main diagonal (NW-SE). (d) Vertically filtered image (Figure 4-4c), smoothed in the vertical direction. (e) Vertically filtered image smoothed along the off-diagonal. (f) Vertically filtered image smoothed along the main diagonal. The greyscale shown in the legend is applied to all panels. Parameter values $BOX = 1$, $GB = 3$, $BASE = 5$, $HALF = 6$**.**

Smoothed outputs from the data of Figure 4-4 are shown in Figure 4-5. We note that the secondary smoothing operation enhances valid ship tracks along the smoothed orientation. Because the input image does not have clear ship tracks in the NW-SE orientation, both panels (c) and (f) show little signal. By contrast, numerous ship track signatures oriented in the horizontal, vertical, and off-diagonal orientations are clearly enhanced in the Figure 4-5.

Note that setting $HALF = 0$ turns off secondary median filtering; in this case the pre-processing stage produces the two filtered frames shown in panels (b) – (c) of Figure 4-4 rather than the six filtered and smoothed frames seen in Figure 4-5.

With the application of the secondary smoother, the pre-processing stage is complete, and attention turns to the algorithms used to detect candidate ship tracks.

4.3.2. Detection

The goal in the detection stage is to identify and delimit potential target patterns (ship tracks) in the pre-processed data. This is accomplished by thresholding pre-processed images, forming clusters of connected pixels lying above the specified threshold, and diluting to join clusters that are close to one another. Following these operations, all clusters with size above a specified number of pixels are considered to be "detected" and are passed to the final classification stage for quality assessment. The parameters that control the detection process are listed in Table 4-2.

Table 4-2. Parameters controlling the detection process.

Parameter	Description
$T1$	Threshold value for preliminary clustering
R	Dilation radius
$SZ1$	Minimum size for detected clusters, pre-dilation
$SZ2$	Minimum size for detected clusters, post-dilation

At the start of detection processing, a threshold denoted $T1$ is applied to each pre-processed frame. The threshold should be chosen so that most pixels on valid ship tracks will have a value above $T1$, while the vast majority of non-target pixels will have value below $T1$. Because the pre-processed data are expressed in units of standard deviations above or below the local mean, the appropriate threshold should not vary from one GOES-R image to the next. However, the directional median smoothing operation generally has the effect of reducing the filtered values, so that a threshold that is permissive when $HALF = 0$ may be stricter for a larger value of $HALF$.

Following the thresholding operation, clusters of connected pixels lying above the threshold are identified. Those with fewer than $SZ1$ pixels are eliminated. The surviving clusters are then dilated to allow clusters split by the thresholding process to re-join. Dilation is implemented by provisionally elongating each cluster in the principal filtered orientation (horizontal or vertical), out to at least R additional pixels. Only pixels with positive values (following directional filtering and median smoothing) are eligible for dilation. Following dilation, clusters are re-shrunk. Here, dilated pixels lying between two original clusters are retained (thus joining the two clusters), while those at the ends of a cluster are eliminated. The dilation, joining, and shrinkage processes are illustrated in Figure 4-6.

Following these steps, merged clusters with size below the post-dilation threshold, $SZ2$, are eliminated. The surviving clusters are now considered "detections" of potential targets (ship tracks). The overall goal in the detection stage is to pass the vast majority of targets through to the final classification stage. We have found benefit to running the detector with two different settings (one permissive, another strict), and merging the resulting clusters with a logical "AND" operator. This approach often provides merged solutions that contain fewer false alarms than are obtained with the

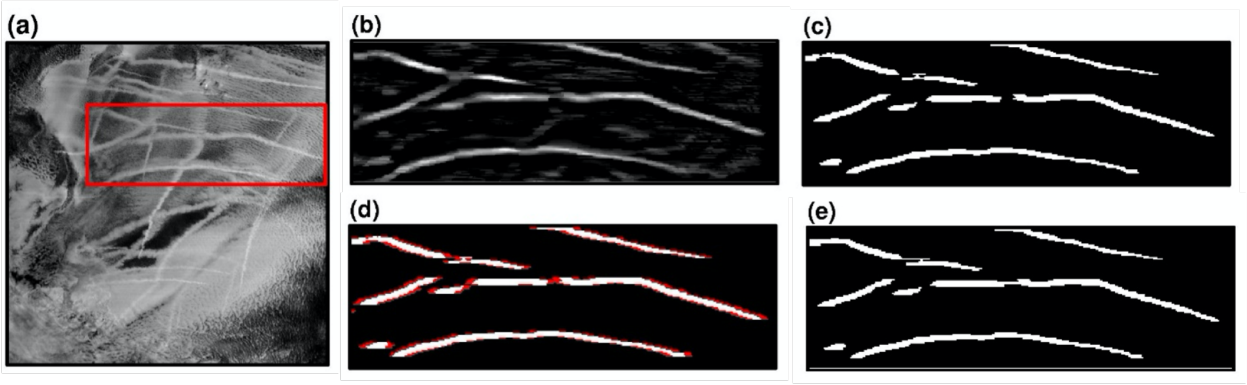

Fig. 4-6. Illustration of clustering, dilation, and shrinkage in the detection stage. (a) subset of Figure 4-4(a) outlined in red used to demonstrate (b) thresholding of pre-processed data, (c) surviving clusters, (d) dilation process of clusters (shown in red) and (e) resulting clusters after re-shrinking the dilated clusters. Parameter values $BOX = 1$, $GB = 3$, $BASE = 5$, $HALF = 6$, $T1 = 1.4$, $R = 5$, $SZ1 = 50$, $SZ2 = 100$.

permissive setting, and provide better coverage of valid ship tracks than is obtained with the strict settings alone. Figure 4-7 shows an example of this two-stage clustering approach applied to the sub image taken from the GOES-17 CONUS scan on June 18, 2019 at 17:00 UTC. Although not perfect, this approach is able to identify most of the obvious and even not so obvious ship tracks in the image.

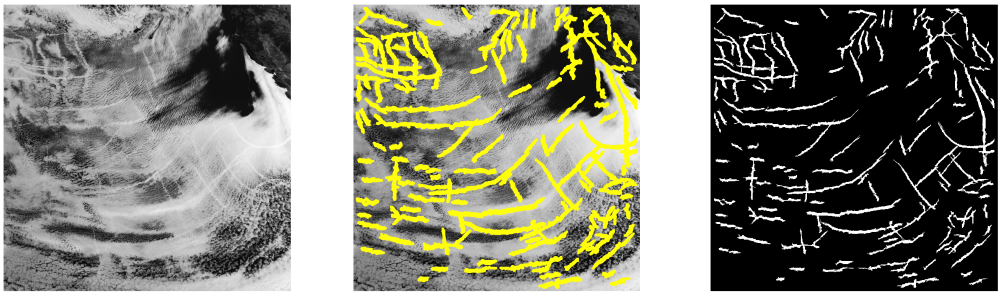

(a) (b) (c)

Fig. 4-7. Demonstration of ship track detection on June 18, 2019 17:00UTC using a two-stage clustering approach. (a) raw image of ship tracks using band transformation C06 - C07 (b) identified clusters in green superimposed on the raw image, and (c) identified clusters alone. Parameter values for permissive clustering $BOX = 3$, $GB = 4$, $BASE = 5$, $HALF = 0$, $T1 = 1$, $R = 3$, $SZ1 = 100$, $SZ2 = 200$. **Parameter values for second-stage restrictive clustering** $BOX = 1$, $GB = 3$, $BASE = 5$, $HALF = 6$, $T1 = 1$, $R = 3$, $SZ1 = 100$, $SZ2 = 200$.

4.3.3. Classification

Although we will not go into detail here, classification of ship tracks can be done in a number of ways. For example, the probabilistic feature fusion technique of [48] might be a promising

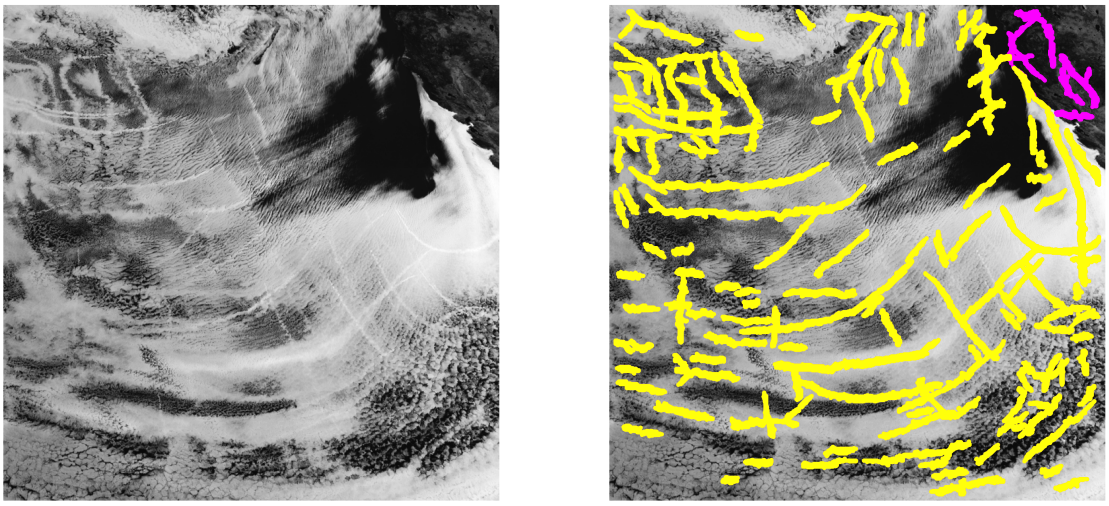

(a) (b)

Fig. 4-8. Applying land mask to clusters identified in Figure 4-7. (a) original image with visible ship tracks and (b) clusters remaining after applying land mask (yellow) and removed clusters(pink). $q_{land} = 50\%$.

approach to classify clusters leveraging cluster features (e.g. number of pixels in a cluster) and the atmospheric conditions mentioned in 4.2. In this work, however, we show some simple examples of how track persistence and the nature of tracks to be seen over the ocean can be used to classify tracks.

4.3.4. Land mask

An obvious indicator of a misclassified cluster is if that cluster is situated mostly over land rather than over the ocean where we expect tracks to form. To account for this, we implement a simple land mask where we consider an identified track mislabeled if q_{land} % of the track resides over land. The parameter q_{land} allows a more dynamic mask. In most cases, a detected track over land should be considered a false positive and removed from the pool of tracks. However, there are situations that legitimate tracks may occur partially over land, for example if the track drifts over land or is observed over islands which may be reasonable. Tuning q_{land} offers the flexibility to excise the tracks that are majority over land while maintaining tracks that are realistic, while still being fractionally over land. For our example here, we choose $q_{land} = 50\%$. Figure 4-8 demonstrates this land mask. The pink clusters were misidentified in the detection algorithm and our land mask algorithm appropriately does not classify these as tracks due more than 50% of the pixels in the cluster appearing over land.

Prior time to t (hours)	Fraction of persisting cluster
0.5	0.95
1	0.9
3	0.8
6	0.5
12	0

Table 4-3. Suggested choices for persistence-dependent fractions of clustered linear features that persist over 0.5, 1, 3, 6 and 12 hours. Clusters are classified as false ship tracks if less than the specified fractions of pixels remain at time t.

4.3.5. Persistence mask

Another distinguishing feature of ship tracks is that they typically persist for between 2 and 12 hours as documented in Section 2. This information provides the basis in formulating a mask that leverages key observations of ship-track persistence behavior.

Given clusters of linear features identified via the detection approach defined above, the goal at this stage is to use persistence behavior to filter out false positive, or misclassified, features that do not correspond to true tracks at the time the image was collected t. To do so, the detection algorithm are applied to images taken at $t - t_{lag}$. All detected clusters at time t are then searched for within the image at time $t - t_{lag}$ to identify tracks that persist for a duration of at least t_{lag}. Note that for practical applications as shown here, clusters identified at time $t - t_{lag}$ were identified using the detection algorithm with rather permissive settings rather than both permissive and restrictive as demonstrated in Figure 4-7.

Due to the impact of differing wind fields on track movement between between times t and $t - t_{lag}$, a time-varying fraction of the grouped pixels identified at time $t - t_{lag}$ that also appear at time t is calculated to determine the number of pixels from each cluster that has persisted in previous images. Linear features detected at time $t - t_{lag}$ that are comprised of more than the fraction of pixels of persisting tracks at time t are classified as false tracks and removed. From exploratory data analysis, fractions deemed suitable at each previous time are detailed in Table 4-3.

An example of applying this procedure to our example on June 18, 2019 is highlighted in Figure 4-9. A further breakdown of the persistence of the tracks in this image are shown in Figure 4-10. Here, the tracks that remain after the algorithm is run are largely determined by the filter applied at 3 hours, removing many of the persisting linear features which are not likely to be attributed as ship-tracks. This approach could be repeated for 12 hours or more but 3 hours was enough to demonstrate this persistence mask for our example here. Figure 4-11 shows the remaining tracks after applying the land mask and 3-hr persistence mask.

Note that this persistence mask is very preliminary. Variability inherent to ship-track formation over the most flexible wind fields render this analysis prone to volatile false positive rates, of which further analysis is deemed necessary. If persistence for longer than 3 hours is needed to reduce the false positive rate, the user will risk increasing the false negative rate. For example, the largest track identified to persist for between 30 minutes to 1 hour (Figure 4-10 (b) appears to be a true track is

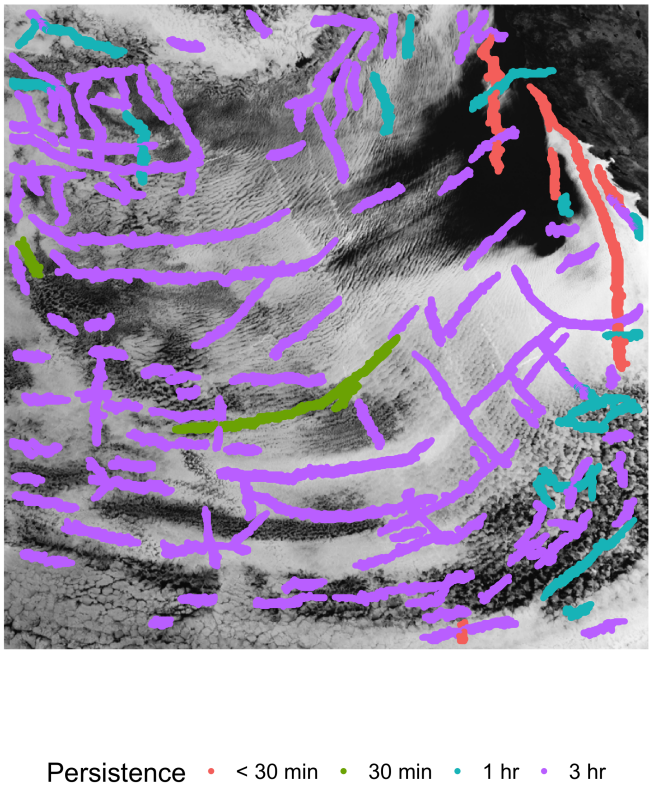

Fig. 4-9. Applying the persistence mask following the land mask, we can identify tracks that have persisted for less than 30 minutes (red), between 30 minutes and 1 hour(green), 1–2 hours (blue), and more than 3 hours (purple).

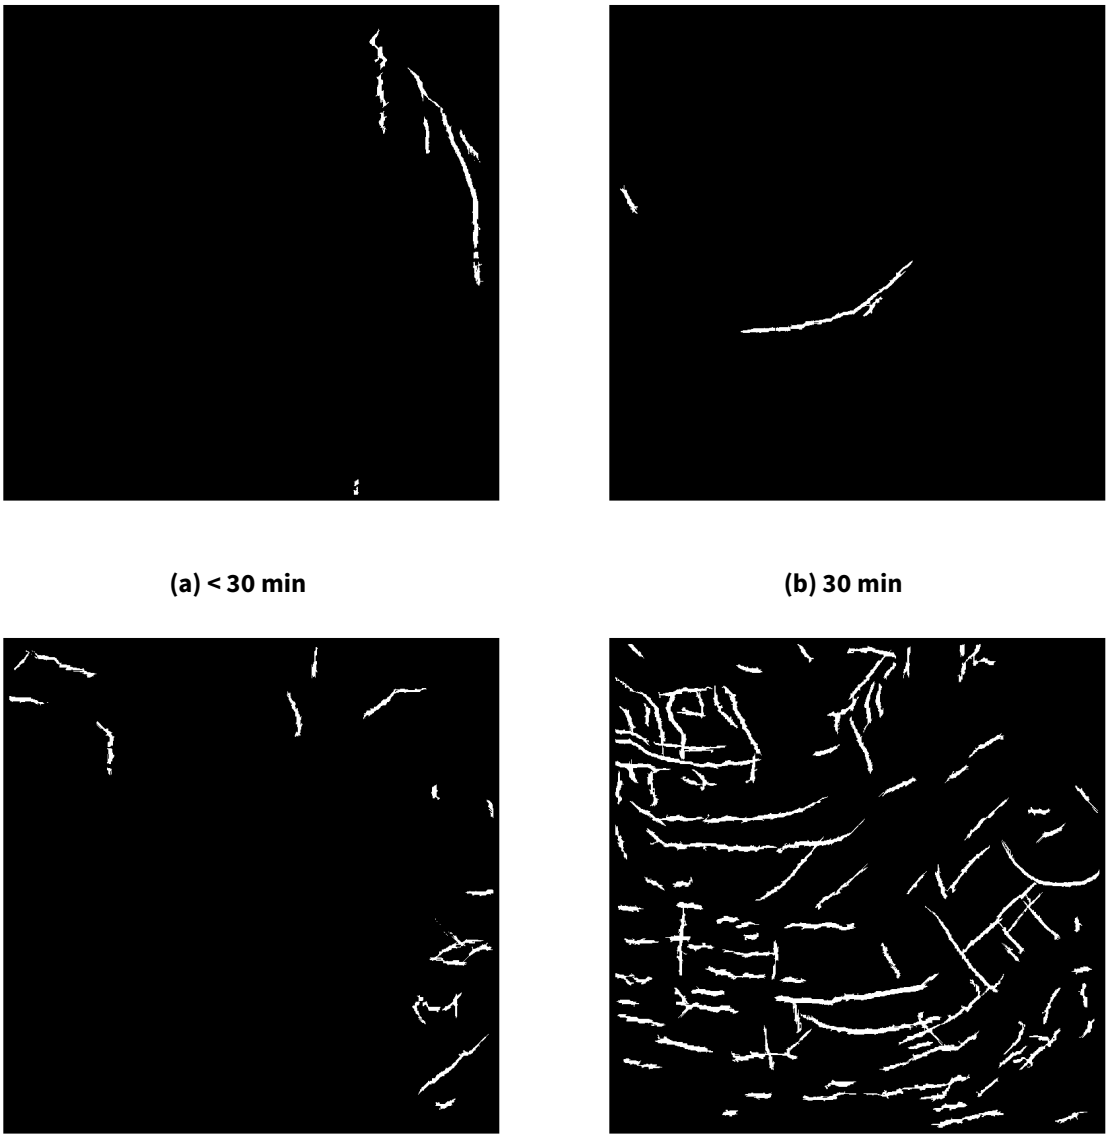

(a) < 30 min **(b) 30 min**

(c) 1 hr **(d) 3 hr**

Fig. 4-10. Applying the persistence mask following the land mask, we can identify tracks that have persisted for less than 30 minutes (a), between 30 minutes and 1 hour (b), 1–2 hours and 3 hours (c), and more than 3 hours (d).

falsely classified even though it was a clear track around 1 hour previously (see Fig. 4-12). This is likely just an artifact of a premature algorithm with insufficient testing. It could indicated that the more rigorous two-step detection algorithm is needed at previous time steps. Additionally, taking into consideration exact wind fields could significantly reduce this risk and has been preliminarily explored.

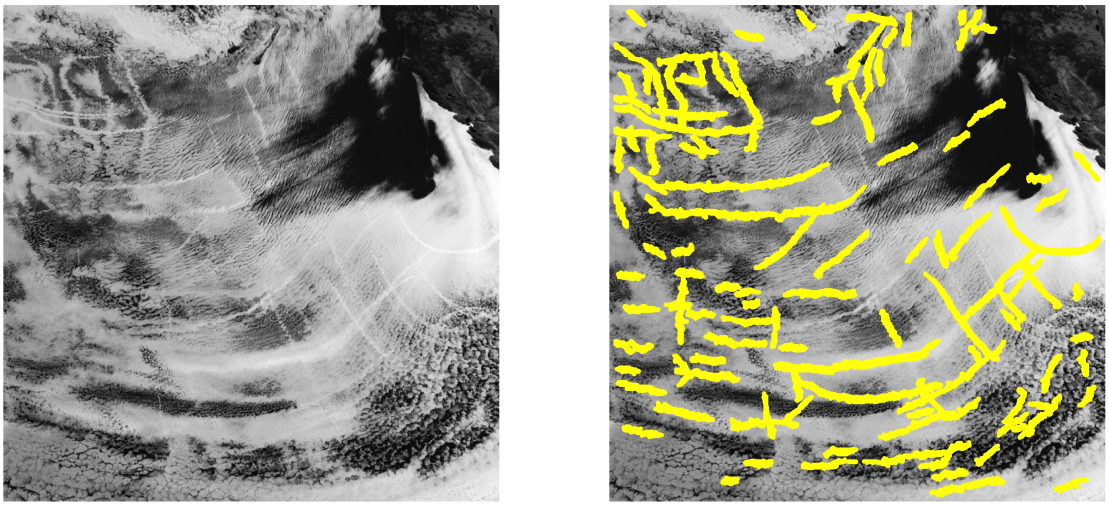

(a) (b)

Fig. 4-11. Remaining clusters classified as tracks after applying both land mask and 3-hr persistence mask.

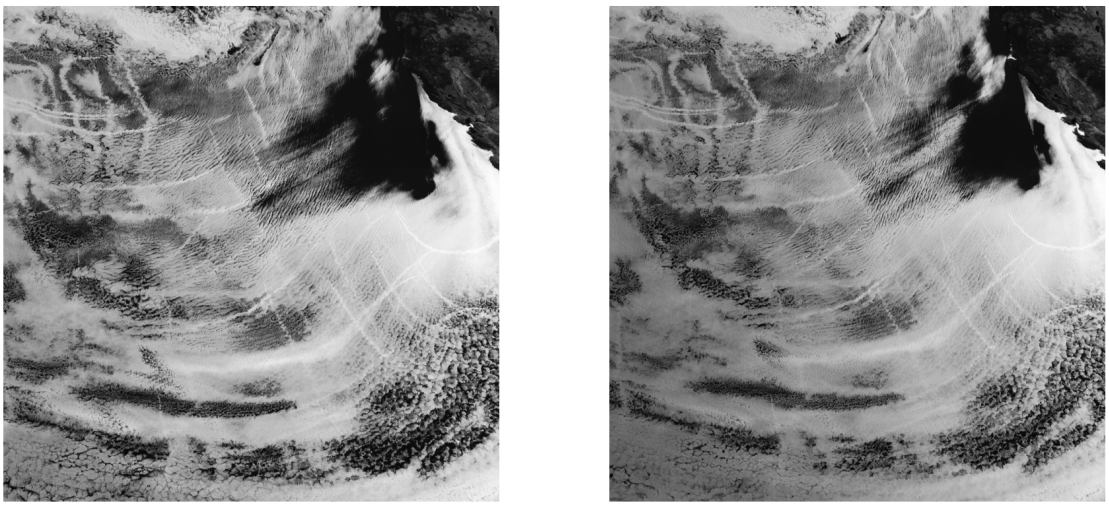

(a) (b)

Fig. 4-12. Original images at (a) 17:00 UTC and (b) 16:00 UTC on June 18, 2019

REFERENCES

B. A. Albrecht. Aerosols, Cloud microphysics, and Fractional Cloudiness. *Science*, 245(4923):1227–1230, 1989.

A. H. Berner, C. S. Bretherton, and R. Wood. Large eddy simulation of ship tracks in the collapsed marine boundary layer: a case study from the Monterey area ship track experiment. *Atmospheric Chemistry and Physics*, 15(10):5851–5871, 2015.

E. J. Bickel and L. Lane. Definition and description of srm solutions. In *An Analysis of Climate Engineering as a Response to Climate Change*, pages 14–28. Copenhagen Consensus Center, 2018.

P. N. Blossey, C. S. Bretherton, J. A. Thornton, and K. S. Virts. Locally Enhanced Aerosols Over a Shipping Lane Produce Convective Invigoration but Weak Overall Indirect Effects in Cloud-Resolving Simulations. *Geophysical Research Letters*, 45(17):9305–9313, August 2018.

J.-Y. Bouget. Pyramidal implementation of the Affine Lucas Kanade Feature Tracker, Description of the algorithm. Technical report, Intel Corporation Microprocessor Research Labs, 2001.

K. S. Carslaw, L. A. Lee, C. L. Reddington, K. J. Pringle, A. Rap, P. M. Forster, G. W. Mann, D. V. Spracklen, M. T. Woodhouse, L. A. Regayre, and J. R. Pierce. Large contribution of natural aerosols to uncertainty in indirect forcing. *Nature*, 503:67–71, 2013.

Y. C. Chen, M. W. Christensen, G. L. Stephens, and J. H. Seinfeld. Satellite-based estimate of global aerosol–cloud radiative forcing by marine warm clouds. *Nature Geoscience*, 7(6528):643–646, 2014.

Y. C. Chen, M. W. Christensen, L. Xue, A. Sorooshian, G. L. Stephens, R. M. Rasmussen, and J. H. Seinfeld. Occurrence of lower cloud albedo in ship tracks. *Atmospheric Chemistry and Physics*, 12(17):8223–8235, 2012.

M. W. Christensen, W. K. Jones, and P. Stier. Aerosols enhance cloud lifetime and brightness along the stratus-to-cumulus transition. *Proceedings of the National Academy of Sciences*, 117(30):17591–17598, 2020.

M. W. Christensen and G. L. Stephens. Microphysical and macrophysical responses of marine stratocumulus polluted by underlying ships: Evidence of cloud deepening. *Journal of Geophysical Research*, 116:D03201, 2011.

J. A. Coakley and C. D. Walsh. Limits to the aerosol indirect radiative effect derived from observations of ship tracks. *Journal of Atmospheric Sciences*, 59:668–680, 2002.

J. H. Conover. Anomalous Cloud Lines. *Journal of Atmospheric Science*, 23(6):778–785, 1966.

M. S. Diamond, H. M. Director, R. Eastman, A. Possner, and R. Wood. Substantial Cloud Brightening from Shipping in Subtropical Low Clouds. *AGU Advances*, 1(1), 2020.

P. A. Durkee, K. J. Noone, and R. T. Bluth. The Monterey Area Ship Track Experiment. *Journal of the Atmospheric Sciences*, 57(16):2523–2541, 2000.

K. El-Darymli, E. Gill, P. McGuire, D. Power, and C. Moloney. Automatic Target Recognition in Synthetic Aperture Radar Imagery: A State-of-the-Art Review. *IEEE Access*, 4:6014–6058, 10 2016.

P. G. Falkowski, Y. Kim, Z. Kolber, C. Wilson, C. Wirick, and R. Cess. Natural Versus Anthropogenic Factors Affecting Low-Level Cloud Albedo over the North Atlantic. *Science*, 256(5061):1311–1313, 1992.

M. Ganzetti, N. Wenderoth, and D. Mantini. Quantitative Evaluation of Intensity Inhomogeneity Correction Methods for Structural MR Brain Images. *Neuroinformatics*, 14:5–21, 2016.

A. K. Georgoulias, K. F. Boersma, J. van Vliet, X. Zhang, R. van der A, P. Zanis, and J. de Laat. Detection of NO2 pollution plumes from individual ships with the TROPOMI/s5p satellite sensor. *Environmental Research Letters*, 15(12):124037, dec 2020.

F. Glassmeier, F. Hoffmann, J. S. Johnson, T. Yamaguchi, K. S. Carslaw, and G. Feingold. Aerosol-cloud-climate cooling overestimated by ship-track data. *Science*, 371(6528):485–489, 2021.

GOES-R Calibration Working Group and GOES-R Series Program. NOAA GOES-R Series Advanced Baseline Imager (ABI) Level 1b Radiances, 2017.

GOES-R Calibration Working Group and GOES-R Series Program. NOAA GOES-R Series Advanced Baseline Imager (ABI) Level 2 Derived Motion Winds (DMW), 2018.

E. Gryspeerdt, T. W. P. Smith, E. O'Keeffe, M. W. Christensen, and F. W. Goldsworth. The impact of ship emission controls recorded by cloud properties. *Geophysical Research Letters*, 46(21):12547–12555, 2019.

Q. Han, W. B. Rossow, and A. A. Lacis. Near-Global Survey of Effective Droplet Radii in Liquid Water Clouds using ISCCP Data. *Journal of Climate*, 7(4):465–497, 1994.

H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.

P. V. Hobbs, T. J. Garret, R. J. Ferek, S. R. Strader, D. A. Hegg, G. M. Frick, W. A. Hoppel, R. F. Gasparovic, L. M. Russell, D. W. Johnson, C. O'Dowd, P. A. Durkee, K. E. Nielson, and G. Innis. Emissions from Ships with Respect to Their Effects on Clouds. *Journal of Atmospheric Sciences*, 57(16):2570–2590, 2000.

H. I. B. Idder and N. Laachfoubi. Cloud Motion Estimation in Satellite Image Sequences by Tracking Skeleton Critical Points using Lucas-Kanade method. In *13th International Conference Computer Graphics, Imaging, and Visualization*, pages 178–183, 2016.

IMO. Guidelines for the Installation of a Shipborne Automatic Identification System (AIS). Technical report, International Maritime Organization, 6 January 2003.

IPCC. Aerosols. In S. Solomon, Q. D., M. Manning, M. Marquis, K. Averyt, M. M. Tignor, H. LeRoy Miller, Jr., and Z. Chen, editors, *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter 2.4. Cambridge University Press, United Kingdom and New York, NY, USA, 2007.

J. Kazil, M. Christensen, S. Abel, T. Yamaguchi, and G. Feingold. Realism of lagrangian large eddy simulations: Tracking a pocket of open cells under a biomass burning aerosol layer. *Earth and Space Science Open Archive*, page 51, 2021.

Y. Kim and R. D. Cess. Effect of anthropogenic sulfate aerosols on low-level cloud albedo over oceans. *Journal of Geophysical Research*, 98(D8):14883–14885, 1993.

D. A. Lack and J. J. Corbett. Black carbon from ships: a review of the effects of ship speed, fuel quality and exhaust gas scrubbing. *Atmospheric Chemistry and Physics*, 12(9):3985–4000, 2012.

J. Latham. Control of global warming? *Nature*, 347:339–340, 1990.

W. Lin, M. Zhang, and N. G. Loeb. Seasonal Variation of the Physical Properties of Marine Boundary Layer Clouds off the California Coast. *Journal of Climate*, 22(10):2624–2638, 05 2009.

B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of DARPA*, pages 121–130, April 1981.

G. Myhre, C. E. L. Myhre, B. H. Samset, and T. Storelvmo. Aerosols and their relation to global climate and climate sensitivity. *Nature Education Knowledge*, 4(5):7, 2013.

G. Myhre, D. Shindell, and F. M. Bré on and W Collins and J Fuglestvedt and J Huang and D Koch and J F Lamarque and D Lee and B Mendoza and T Nakajima and A Robock and G Stephens and T Takemura and H Zhang. Anthropogenic and natural radiative forcing. In T. F. Stocker, D. Qin, G. K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley, editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chapter 8. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2013.

National Research Council. *Climate Intervention: Reflecting Sunlight to Cool Earth*. The National Academies Press, Washington, DC, 2015.

K. J. Noone, E. Ö strö m, R. J. Ferek, T. Garrett, P. V. Hobbs, D. W. Johnson, J. P. Taylor, L. M. Russell, R. C. Flagan, J. H. Seinfeld, C. D. O'Dowd, M. H. Smith, P. A. Durkee, K. Nielsen, J. G. Hudson, R. A. Pockalny, L. D. Bock, R. E. V. Grieken, R. F. Gasparovic, and I. Brookes. A case study of ships forming and not forming tracks in moderately polluted clouds. *Journal of the Atmospheric Sciences*, 57(16):2729–2747, 2000.

N. A. of Sciences Engineering Medicine (NASEM). Reflecting Sunlight: Recommendations for Solar Geoengineering Research and Research Governance, 2021.

OpenCV. Open Source Computer Vision Library, July 2018. Version 3.4.2.

J. Penner, M. Andreae, and M. Annegarn. Aerosol, their direct and indirect effects. In *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, chapter 5, pages 291–336. Cambridge University Press, 01 2001.

A. Possner, H. Wang, R. Wood, K. Caldeira, and T. P. Ackerman. The efficacy of aerosol-cloud radiative perturbations from near-surface emissions in deep open-cell stratocumuli. *Atmospheric Chemistry and Physics*, 18:17475–17488, 2018.

I. Reda and A. Andreas. Solar position algorithm for solar radiation applications. Technical report, National Renewable Energy Laboratory, 2008.

N. Riemer, M. West, R. Zaveri, and R. Easter. Simulating the evolution of soot mixing state with a particle-resolved aerosol model. *Journal of Geophysical Research: Atmospheres*, 114:D09202, 09 2008.

M. S. Segrin, J. A. Coakley, and W. R. Tahnk. Modis observations of ship tracks in summertime stratus off the west coast of the united states. *Journal of the Atmospheric Sciences*, 64(12):4330 – 4345, 2007.

J. H. Seinfeld, C. Bretherton, K. S. Carslaw, H. Coe, P. J. DeMott, E. J. Dunlea, G. Feingold, S. Ghan, A. B. Guenther, R. Kahn, I. Kraucunas, S. M. Kreidenweis, M. J. Molina, A. Nenes, J. E. Penner, K. A. Prather, V. Ramanathan, V. Ramaswamy, P. J. Rasch, A. R. Ravishankara, D. Rosenfeld, G. Stephens, and R. Wood. Improving our fundamental understanding of the role of aerosol-cloud interactions in the climate system. *Proceedings of the National Academy of Sciences*, 113(21):5781–5790, 2016.

J. Shi and C. Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR94)*, pages 593–600, 1994.

K. M. Simonson, R. D. West, R. L. Hansen, T. E. LaBruyere, and M. H. Van Benthem. A statistical approach to combining multisource information in one-class classifiers. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 10:199–210, 2017.

M. Sofiev, V. Sofieva, T. Elperin, N. Kleeorin, I. Rogachevskii, and S. Zilitinkevich. Turbulent diffusion and turbulent thermal diffusion of aerosols in stratified atmospheric flows. *Journal of Geophysical Research: Atmospheres*, 114:D18209, 09 2009.

A. Stein, R. R. Draxler, G. D. Rolph, B. J. Stunder, M. Cohen, and F. Ngan. NOAA's HYSPLIT Atmospheric Transport and Dispersion Modeling System. *Bulletin of the American Meteorological Society*, 96(12):2059–2077, 2015.

B. Stevens and G. Feingold. Untangling aerosol effects on clouds and precipitation in a buffered system. *Nature*, 461:607–613, 2009.

S. Twomey. Pollution and the planetary albedo. *Atmospheric Environment*, 8(12):1251–1256, 1974.

S. Twomey, H. B. Howell, and T. A. Wojciechowski. Comments on “Anomalous Cloud Lines". *Journal of Atmospheric Sciences*, 25(2):333–334, 1968.

U.S. Department of Transportation. Seavision, 2020.

H. Wang, R. C. Easter, R. Zhang, P.-L. Ma, B. Singh, K. Zhang, D. Ganguly, P. J. Rasch, S. M. Burrows, S. J. Ghan, S. Lou, Y. Qian, Y. Yang, Y. Feng, M. Flanner, L. R. Leung, X. Liu, M. Shrivastava, J. Sun, Q. Tang, S. Xie, and J.-H. Yoon. Aerosols in the e3sm version 1: New developments and their impacts on radiative forcing. *Journal of Advances in Modeling Earth Systems*, 12(1):e2019MS001851, 2020. e2019MS001851 2019MS001851.

H. Wang, P. J. Rasch, and G. Feingold. Manipulating marine stratocumulus cloud amount and albedo: a process-modelling study of aerosol-cloud-precipitation interactions in response to injection of cloud condensation nuclei. *Atmospheric Chemistry and Physics*, 11:4237–4249, 2011.

J. M. Weiss, R. Luo, and R. M. Welch. Automatic detection of ship tracks in satellite imagery. In *1997 IEEE International Geoscience and Remote Sensing Symposium Proceedings. Remote Sensing - A Scientific Vision for Sustainable Development*, volume 1, pages 160–162. IGARSS’97, 1997.

P. Wood-Bradley, J. Zapata, and J. Pye. Cloud tracking with optical flow for short-term solar forecasting. In *50th Annual Conference of the Australian Solar Energy Society*. Australian Solar Council, December 2012.

T. Yuan, C. Wang, H. Song, S. Platnick, K. Meyer, and L. Oreopoulos. Automatically finding ship tracks to enable large-scale analysis of aerosol-cloud interactions. *Geophysical Research Letters*, 46(13):7726–7733, 2019.

X. Zhao, Y. Liu, F. Yu, and A. K. Heidinger. Using long-term satellite observations to identify sensitive regimes and active regions of aerosol indirect effects for liquid clouds over global oceans. *Journal of Geophysical Research: Atmospheres*, 123(1):457–472, 2018.

DISTRIBUTION

Email—External [redacted]

Name	Company Email Address	Company Name
Rob Wood	robwood2@uw.edu	University of Washington

Email—Internal [redacted]

Name	Org.	Sandia Email Address
Benjamin K. Cook	08910	bkcook@sandia.gov
Adele Doser	05573	abdoser@sandia.gov
Andrew Glen	08913	aglen@sandia.gov
Lori K. Parrot	08913	lkparro@sandia.gov
Randy Smith	01464	ransmit@sandia.gov
Technical Library	1911	sanddocs@sandia.gov

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.