Development of a Method to Detect High Ice Water Content Environments Using Machine Learning

JULIE A. HAGGERTY,^a ALLYSON RUGG,^a RODNEY POTTS,^b ALAIN PROTAT,^b J. WALTER STRAPP,^c THOMAS RATVASKY,^d KRISTOPHER BEDKA,^e AND ALICE GRANDIN^f

^a National Center for Atmospheric Research, Boulder, Colorado

^b Bureau of Meteorology, Melbourne, Victoria, Australia

^c MetAnalytics, Aurora, Ontario, Canada

^d NASA Glenn Research Center, Cleveland, Ohio

^e NASA Langley Research Center, Hampton, Virginia

^f Airbus SAS, Toulouse, France

(Manuscript received 6 November 2019, in final form 6 February 2020)

ABSTRACT

This paper describes development of a method for discriminating high ice water content (HIWC) conditions that can disrupt jet-engine performance in commuter and large transport aircraft. Using input data from satellites, numerical weather prediction models, and ground-based radar, this effort employs machine learning to determine optimal combinations of available information using fuzzy logic. Airborne in situ measurements of ice water content (IWC) from a series of field experiments that sampled HIWC conditions serve as training data in the machine-learning process. The resulting method, known as the Algorithm for Prediction of HIWC Areas (ALPHA), estimates the likelihood of HIWC conditions over a three-dimensional domain. Performance statistics calculated from an independent subset of data reserved for verification indicate that the ALPHA has skill for detecting HIWC conditions, albeit with significant false alarm rates. Probability of detection (POD), probability of false detection (POFD), and false alarm ratio (FAR) are 86%, 29% (60% when IWC below 0.1 g m^{-3} are omitted), and 51%, respectively, for one set of detection thresholds using in situ measurements. Corresponding receiver operating characteristic (ROC) curves give an area under the curve of 0.85 when considering all data and 0.69 for only points with IWC of at least 0.1 g m^{-3} . Monte Carlo simulations suggest that aircraft sampling biases resulted in a positive POD bias and the actual probability of detection is between 78.5% and 83.1% (95% confidence interval). Analysis of individual case studies shows that the ALPHA output product generally tracks variation in the measured IWC.

1. Introduction

Since the mid-1990s, observations of jet-engine power loss in areas with high ice particle concentrations have generated interest from the aviation industry, regulatory agencies, and the meteorological research community. Over 200 ice crystal icing (ICI) events causing power loss or engine rollback (loss of engine control) in commuter and large transport aircraft have been documented as of 2019 (Lawson et al. 1998; Bravin et al. 2015; Bravin and Strapp 2019). Air data system performance disruptions, such as anomalies in total air temperature measurements, have also been noted in these conditions (Lawson et al. 1998; Rogers et al. 2002; Mason et al. 2006;

Beswick et al. 2015). The majority of these events occurred in flight conditions where ice particles dominated surrounding clouds and icing by accretion of supercooled liquid to the airframe was not observed. In 2006, the Engine Harmonization Working Group, an international committee of regulators, engine manufacturers, and airframers concluded these events were likely caused by ingest of high mass concentrations of ice particles (Mason et al. 2006). Investigation of these events has provided a basic understanding of the meteorological conditions that support areas of high ice water content (HIWC).

Analyses of in-service ICI events (Mason et al. 2006; Grzych and Mason 2010; Bravin et al. 2015; Bravin and Strapp 2019) indicate that potentially hazardous HIWC conditions most often occur near convective systems in

Corresponding author: Julie Haggerty, haggerty@ucar.edu

DOI: 10.1175/JTECH-D-19-0179.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

warm air masses at high altitudes. The majority of events analyzed occurred outside of vigorous updraft regions. Reports of turbulence were light to moderate. Radar reflectivity at flight level from pilots' weather radar (X band) was generally below 20–30 dBZ, and many events occurred without flight-level radar echoes (<20 dBZ), suggesting that small ice crystals, which have a low radar reflectivity as compared with large ice crystals or supercooled liquid drops, constitute the bulk of the ice mass encountered in ICI events.

As aircraft continue to experience ICI events (e.g., [General Civil Aviation Authority 2015](#); [Aviation Safety Network 2016](#); [Bravin and Strapp 2019](#)), the need to develop methods for detecting and eventually forecasting locations of hazardous HIWC conditions has become apparent. Details of the cloud micro and macrophysical processes leading to HIWC conditions, however, are the subject of ongoing research. The most recent inventories of confirmed ICI events ([Bravin and Strapp 2019](#)) found that they cluster in regions of deep convection with high air traffic density. Events are more common in local summer months, but show no general diurnal pattern. While events can usually be linked to a convective surge, they do not necessarily occur at the time of peak convective activity.

Supplementing conclusions drawn from in-service icing events, analysis of research data from a series of field campaigns seeks to better understand the atmospheric processes responsible for HIWC. For example, in situ microphysics measurements collected in HIWC conditions associated with large mesoscale convective systems (MCS) showed typical mass median diameters of ice crystals ranging from 320 to 690 μm over a temperature range from -50° to -10°C ([Leroy et al. 2017a](#); [J. W. Strapp et al. 2019](#), unpublished FAA report). [Protat et al. \(2016\)](#) used bulk IWC measurements from research aircraft to develop a statistical relationship between IWC and 95 GHz radar reflectivity in MCS convective areas. Cloud-resolving model simulations by [Fridlind et al. \(2015\)](#) investigated the microphysical pathways that could produce substantial ice water content in combination with modest radar reflectivity. [Proctor et al. \(2017\)](#) simulated the three-dimensional evolution of an MCS and showed that regions of IWC exceeding 2 g m^{-3} persisted beyond the most intense stages of the storm.

Efforts to develop detection methods that might inform aviation operations are also an area of active research ([Haggerty et al. 2019a](#)). [Grzych et al. \(2015\)](#) described attempts to track local features of storms, including low-level convergence and overshooting cloud tops, associated with in-service ICI events. Other researchers have developed methods to discover HIWC

signatures in geostationary satellite data ([de Laat et al. 2017](#); [Yost et al. 2018](#)). [Defer et al. \(2015\)](#) described the use of concurrent multispectral multiple-technique observations from the low-orbit A-Train mission to identify specific signatures in HIWC regions. A modeling-observation hybrid approach identified areas of developing convection that may contain HIWC in the rapidly developing thunderstorm (RDT) system ([Gounou et al. 2015](#)).

In this paper we present a method for estimating the likelihood for HIWC conditions, based on techniques that optimally blend relevant information from multiple data sources. The Algorithm for Prediction of HIWC Areas (ALPHA) obtains input data from numerical weather prediction (NWP) models, satellite cloud retrievals, and ground-based radar products. An artificial intelligence approach based on fuzzy-logic technology is implemented to combine information derived from each data source. Research aircraft measurements acquired from field experiments are used to objectively train the algorithm using a machine-learning technique. [Section 2](#) explains the approach upon which ALPHA is based. Observations used to train and verify the method are presented in [section 3](#). Implementation of a machine-learning technique is summarized in [section 4](#). [Section 5](#) discusses results, performance statistics, and application to case studies.

2. Approach

Atmospheric conditions associated with ICI events span a range of temperatures, flight altitudes, and convective storm types (e.g., [Grzych and Mason 2010](#)). Nevertheless, there are enough common attributes to suggest that identification of HIWC areas from operationally available meteorological data is feasible. Development of ALPHA is intended to demonstrate the feasibility and evaluate the skill of a product that would offer guidance to aviation operators for avoiding encounters with high concentrations of ice crystals.

In the absence of a full and rigorous physical understanding of processes that lead to hazardous concentrations of ice crystals in anvil clouds, an artificial intelligence (AI) approach provides a means for estimating the potential for HIWC conditions using readily available data products. ALPHA employs a set of meteorological predictors of HIWC extracted from satellite, radar, and NWP model output. Each of these data sources offers advantages and disadvantages for solution of the HIWC detection problem. Geostationary satellite data are available over most of the globe with good spatial and temporal resolution. Products derived from visible and infrared (IR) channels

TABLE 1. Input variables considered as input to ALPHA derived from satellite, radar, and NWP model products. Variables implemented into the final algorithm are indicated with italics.

Satellite	Radar	WRF Model	ACCESS model
<i>Cloud-top temperature</i>	<i>Composite reflectivity</i>	Ice water content	Ice water content
<i>Cloud-top height</i>	<i>Max height with at least 10 dBZ</i>	<i>Relative humidity</i>	<i>Relative humidity</i>
<i>Temperature difference</i> (6.8 – 10.8 μm)	Max height with at least 30 dBZ	Vertical velocity	Vertical velocity
<i>Optical depth</i>	Echo top (18-dBZ threshold)	1-h convective precipitation	<i>1-h convective precipitation</i>
<i>Brightness temperature</i> (6.8 μm)	<i>Cloud depth (18-dBZ threshold)</i>	1-h total precipitation	1-h total precipitation
<i>Brightness temperature</i> (10.8 μm)	<i>Ice mass</i>	<i>Convective available potential energy (CAPE)</i>	<i>High cloud cover</i>
Tropopause height – cloud-top height	<i>Volume-averaged height-integrated radar reflectivity (VAHIRR)</i>	Convective initiation (CIN)	<i>Middle cloud cover</i>
	Volume integrated liquid (VIL)	<i>Temperature</i>	<i>Low cloud cover</i> <i>Temperature</i>

provide useful information about cloud properties, although retrievals largely represent conditions at cloud top. Radar products give a three-dimensional view of convective systems at high resolution, but spatial coverage is limited globally, especially over oceans where many ICI events occur. Operational NWP models provide a three-dimensional picture of the atmosphere and forecasts of its evolution, but at coarser resolution and with limited skill for simulating convective processes. Proper implementation of an AI approach exploits the strengths and minimizes the weaknesses of each data source.

Statistics compiled from research aircraft measurements in tropical convective conditions guide the selection of specific meteorological predictors from each dataset and the determination of their range of values that suggest HIWC conditions. Fuzzy-logic method is used to blend information derived from each predictor. Implementation of this method follows techniques applied in the “current icing potential,” which has been successfully employed for operational detection of supercooled liquid water (Bernstein et al. 2005). Advantages inherent in this approach include the ability to exploit empirical relationships between basic meteorological variables and the desired product and the capability to optimally blend information from a variety of sources. Fuzzy logic also accounts for uncertainties evident in the datasets and can allow for a gradual transition from non-HIWC to HIWC environments.

ALPHA is composed of multiple fuzzy-logic membership functions that quantify the level of interest in individual variables. These relationships approximate empirical relationships between a single variable and the occurrence of HIWC conditions. Membership functions are defined for all variables related to the presence of HIWC based on available data. Applied to each input

variable, membership functions create an interest field for that variable. Resulting interest fields are then blended together using optimized weighting factors to derive an estimate of the likelihood of HIWC at a given location. In this way, ALPHA produces a three-dimensional diagnosis of HIWC potential (a parameter scaled from 0 to 1, where 0 represents minimum likelihood of HIWC and 1 is maximum likelihood).

Table 1 lists ALPHA input variables grouped by data source. Input variables were initially selected based on analysis of in-service ICI events by Grzych and Mason (2010) and fundamental knowledge of oceanic convective systems provided in the literature (e.g., Zipser et al. 2006). Satellite cloud products from visible and IR channels on geostationary satellites are used to identify cold, high cloud tops associated with deep convection. NWP model fields indicate the potential for convection, the presence of various condensate types, the presence and type of precipitation, and updrafts associated with convective activity. Model-derived temperature profiles are also used to limit the vertical extent of HIWC potential to layers with sufficiently cold temperatures. Ground-based radar products are used to identify areas with active updrafts and characterize the vertical extent of high reflectivity values. Spatial and temporal resolution of the ALPHA output products mirror those of the satellite products used as input. Typically, pixel sizes of 5 km are available from the various geostationary satellite products, while model resolutions of approximately 13 km and radar product resolutions on the order of 1 km require interpolation and subsampling, respectively, to match the three datasets.

Initially, the input variable set, associated membership functions, and blending weights were defined using findings from ICI in-service events and qualitative information derived from procedures used by human forecasters to identify convective activity (Haggerty et al. 2012).

FIG. 1. Flight tracks from HIWC experiments: (top) flights based in Cayenne (C), Fort Lauderdale (F), Palmdale (P), and Kona (K) and (bottom) flights based in Darwin (D).

As research-quality data from dedicated field experiments became available, the ALPHA parameters were optimized using objective machine-learning techniques as described in [section 4](#).

3. Datasets

a. Airborne observations for algorithm training and verification

Engine events on commercial airliners provide insufficient data for thorough analysis of the meteorological processes that control formation of HIWC conditions and for description of the in situ properties of HIWC clouds ([Strapp et al. 2016a](#)). Hence, the need for targeted field campaigns was identified by the North American HIWC project partners ([Strapp et al. 2016a](#)) and the European High Altitude Ice Crystal (HAIC) team ([Dezitter et al. 2013](#)), who conducted two collaborative flight campaigns in 2014 and 2015. Subsequently, two additional flight campaigns focusing on radar detection of HIWC conditions were conducted in 2015 and 2018 ([Harrah et al. 2016](#); [Ratvasky et al. 2019](#)). The resulting observations provide an extensive set of measurements in areas of large tropical oceanic deep convection, as well as some additional measurements in stratiform regions of moderate-sized tropical continental convection. These datasets were collected to address multiple research objectives including

development of nowcasting tools for the avoidance of regions conducive to jet-engine power loss and air data probe failures due to encounters with ice crystals.

Two tropical locations (Darwin, Australia, and Cayenne, French Guiana) and a subtropical site (Fort Lauderdale, Florida) served as bases for the campaigns ([Fig. 1](#); note that a subset of the 2018 flights were based in Palmdale, California, and Kona, Hawaii). These locations provided access to large oceanic MCS, the primary target for measurements given their frequent association with ICI events ([Mason et al. 2006](#); [Grzych and Mason 2010](#); [Bravin et al. 2015](#)). Tropical storms and moderate-sized continental convection were also sampled on several occasions. Airborne sampling along straight and level in-cloud segments averaging approximately 100 km in length was performed at specified temperature levels (-10° , -30° , and -50°C , with less frequent sampling at -20° and -40°C) with the intent of addressing aviation industry objectives to characterize the IWC and particle size distribution associated with HIWC conditions ([Strapp et al. 2016a](#)).

The first of the HAIC–HIWC field campaigns was conducted in Darwin (12.46°S , 130.85°E) from 16 January to 18 February 2014 ([Leroy et al. 2017b](#)). The Service des Avions Français Instrumentés pour la Recherche en Environnement (SAFIRE) Falcon-20 aircraft performed 17 flights that operated mainly in large MCS but also sampled a tropical cyclone in the region ([Fig. 1](#), lower

panel). The payload included a new isokinetic probe (IKP2) (Davison et al. 2010), which is a total water evaporator probe with a near-unit efficiency for water droplets and ice crystals that was developed by the HIWC project to measure high values of total water content (TWC) during these campaigns. The accuracy of the IKP2 has been estimated using system accuracy calculations and from wind tunnel performance testing and calibrations. The system accuracy estimates and tunnel comparisons established that the IKP2 accuracy is within the 20% target set for high-IWC conditions at the outset of the probe development, and likely was within about 10% in the temperature range of most of the campaign measurements (Davison et al. 2016; Strapp et al. 2016b). TWC values up to a maximum of 4.1 g m^{-3} were observed during the flight campaign over a distance scale of about 1 km, with the vast majority of observations conducted in glaciated conditions. Cloud particle size distributions were measured by a Droplet Measurement Technologies (DMT) Cloud Droplet Probe (CDP2), a DMT Precipitation Imaging Probe (PIP), and a SPEC, Inc., 2D-S imaging probe. The Radar System Airborne (RASTA) W-band cloud radar provided vertical profiles of reflectivity, winds and retrieved TWCs (Protat et al. 2016).

A second HAIC–HIWC field campaign occurred in Cayenne (4.92°N , 52.31°W), from 5 to 29 May 2015 (Fig. 1, upper panel). The SAFIRE Falcon-20 aircraft flew 16 flights in oceanic and continental MCS carrying the same payload as described for the Darwin experiment.

Last, the HIWC RADAR I flight campaign used the NASA DC-8 aircraft for 10 flights over the Gulf of Mexico, Caribbean Sea, and western Atlantic Ocean from 12 to 23 August 2015 (Fig. 1, upper panel). Based in Fort Lauderdale, (26.12°N , 80.14°W), the aircraft sampled clouds associated with oceanic MCS plus four flights in two tropical storms. The payload included the IKP2 sensor along with CDP2, PIP, and 2D-S cloud probes. The HIWC RADAR II flight campaign occurred from 27 July to 20 August 2018 with seven flights based from Fort Lauderdale, Palmdale, and Kona.

b. Satellite products as algorithm input

Retrieval of cloud properties from *Multifunction Transport Satellite 1R* (MTSAT-1R) (Darwin), *GOES-13* (Cayenne and Florida during the HIWC RADAR I flight campaign), and *GOES-16* (Florida during HIWC RADAR II) were based on methods described by Minnis et al. (2011). Cloud properties retrieved from *GOES-13* Imager, *MTSAT-1R* Japanese Advanced Meteorological Imager (JAMI), and *GOES-16* Advanced Baseline Imager radiances using the Satellite Cloud and Radiation Property

Retrieval System (SatCORPS) were provided by the NASA Langley Research Center Cloud and Radiation Research Group. As explained by Yost et al. (2018), the Visible Infrared Shortwave-Infrared Split-Window Technique (VISST) is used to retrieve daytime cloud properties such as cloud optical depth, cloud-top height and temperature, and other variables. The Solar Infrared Split-Window Technique (SIST) calculates these parameters at night, albeit with significant limitations on the ability to detect optical depth variations in deep convective clouds using only IR channels (Hong et al. 2007). Because of nighttime limitations, the satellite component of ALPHA is dependent on solar zenith angle (explained in section 4). Furthermore, to reduce the impact of shadowing, optical depth at all times of day is smoothed as in Yost et al. (2018) using a weighted average over a 5-by-5-pixel neighborhood with weight of 9 given to the center pixel; weights of 3 given to the first, 3-by-3, ring; and weights of 1 given to the second, 5-by-5, ring. Parallax corrections are also applied to retrievals in the Pacific Ocean because of high viewing angle by the *GOES-16* satellite. Note that most of the aircraft data were collected during daylight hours.

Spatial and temporal resolutions vary by dataset: *MTSAT-1R* (4 km and 10 min), *GOES-13* (4 km and 30 min during the HAIC–HIWC Cayenne experiment; 4 km and 15 min during the HIWC RADAR Study, Phase I), *GOES-16* (4 km and 10–30 min during the HIWC RADAR Study, Phase II). Note also that there are slight differences in central wavelengths of channels available on each satellite (e.g., $10.7 \mu\text{m}$ on the *GOES-13* Imager vs $10.8 \mu\text{m}$ on *MTSAT-1R*); it is assumed that the algorithm development described below will not be significantly affected by such small differences.

c. NWP models

Regional NWP models were implemented during each of the three field campaigns. The Australian Community Climate and Earth-System Simulator (ACCESS) covered the Darwin experimental domain (Fig. 1). This regional version (ACCESS-R) ran four times daily providing hourly forecasts out to 72 h and horizontal resolution of 12 km at standard pressure levels and the surface (Australian Bureau of Meteorology 2013; Puri et al. 2013). For the domains around French Guiana, Florida, California, and Hawaii, the Weather Research and Forecasting (WRF) Model (version 3 for the 2015 projects and version 4 in 2018; Skamarock et al. 2008) was run four times daily at 13-km resolution with forecasts out to 48 h; a separate 3-km WRF domain was run in the 2018 experiment. The Thompson aerosol-aware microphysics and Tiedke cumulus parameterizations were used in both versions.

d. Radar observations

Data from operational ground-based radars were obtained in three of the four experiments. Parts of the Darwin flight domain were covered by the Northern Australia 3D Radar Mosaic products (Australian Bureau of Meteorology 2015). National Weather Service Radar Mosaic products from the U.S. Southeast sector covered the domain for the 2015 Florida-based experiment and part of the 2018 experiment (excluding the flights based in Hawaii). Radar coverage was available for portions of 21 of 23 Darwin-based flights, 7 of 10 Florida-based flights in 2015, and 2 of 7 flights conducted from Florida, California, and Hawaii in 2018. Products extracted from these datasets included reflectivity, heights of specified reflectivity thresholds, and derived products such as ice mass and Volume-Averaged Height-Integrated Radar Reflectivity (VAHRR; Dye et al. 2007).

4. Objective determination of ALPHA parameters

The HAIC–HIWC and the HIWC RADAR field campaigns yielded a substantial dataset for optimizing the set of input variables used in ALPHA, formulating representative membership functions for each variable, and determining optimal weighting factors for blending information derived from the input variables. The parameter set can be objectively optimized using in situ measurements of TWC from the airborne IKP2 (since the vast majority of data used herein were collected in glaciated clouds, IKP2 measurements will henceforth be referred to as IWC). By comparing ALPHA input data, interest values calculated from proposed membership functions, and resulting HIWC potential along the aircraft flight tracks, an extensive set of matched in situ IWC measurements and corresponding ALPHA parameters was obtained. This matched dataset enables the development and training of ALPHA.

Because of the high sampling rate of the IKP2, in situ data are available at far higher resolution than ALPHA input products. IWC measurements are provided as 5-s moving averages, which covers a distance of about 1 km at typical airspeeds, while ALPHA spatial resolution is about 5 km. Hence multiple IWC measurements are associated with each ALPHA grid cell. To obtain a mean IWC across a grid cell, the 5-s IKP2 measurements were averaged over 25-s intervals before being matched to ALPHA input data. Some uncertainty in the collocation of IWC measurements with ALPHA input data was introduced by satellite geolocation errors, but *GOES-16* geolocation error is less than 0.5 km according to Kalluri et al. (2018), smaller than the distance represented by a 5-s average IWC measurement. Only

FIG. 2. Diagram showing the development of ALPHA components.

points for which the NWP model temperature was below 0°C were used in training.

Given the multiple degrees of freedom inherent in the fuzzy-logic method used in ALPHA, adjusting the various parameters individually is difficult to accomplish manually. Membership functions and weights are not independent, so changing one could lead to a different optimal solution for another. Additionally, the best set of functions and weights might not be intuitively recognized by a human. For these reasons, an objective means of optimizing the parameter set was employed.

Two machine-learning techniques were applied to optimize the various components of ALPHA: particle swarm optimization (PSO) and canonical correlation analysis (CCA). The ALPHA high-level structure and optimization process is described in Fig. 2. At the top level, input variables are grouped by data source, and PSO is used to derive the optimal fuzzy-logic membership functions and blending weights for each input variable. Variables assigned near-zero weights by the PSO process were eliminated. Individual membership functions applied to each variable produce interest fields that are then combined within each data category using weighting factors derived by PSO. The result is a set of three intermediate interest fields from the satellite, model, and radar input variable sets. Note that because different numerical models were used in the Darwin experiment and the Cayenne and Florida experiments, the model component of ALPHA is trained once on the

FIG. 3. Autocorrelation of IWC 25-s averages measured by the IKP2 sensor during four field campaigns.

ACCESS model and once on the WRF Model. As indicated in the second level of Fig. 2, the interest fields from each data source are blended using CCA. Because radar data are not always available, there are two options at this level: a two-input version that excludes radar data, and a three-input version that includes all three data sources. The weights for the two- and three-input products were optimized separately using CCA and then combined using the three-input version whenever possible or the two-input product where radar data are unavailable.

a. Training dataset

In the context of machine learning applied to ALPHA development, training consists of discovering the set of parameter values that result in optimal agreement between the data and the algorithm's output through the use of PSO and CCA. A subset of IWC measurements from the IKP2 (25-s mean values) was compiled for training purposes, and another independent subset was reserved for subsequent verification of the resulting algorithm.

Proper partitioning of the training and verification datasets is imperative to ensure independence while maintaining similar IWC distributions in each subset. Since consecutive data points are adjacent in space and time and share meteorological properties, they cannot be considered independent. Hence random assignment of 25-s IWC values would not produce independent training and verification datasets. As shown in Fig. 3, consecutive 25-s average IWC data points have a correlation of 0.92. The autocorrelation between data points levels off to 0.3 at a separation of about 500s (20 data points). To achieve independence between the

FIG. 4. The distribution of IWC observations assigned to the training, verification, and buffer datasets.

training and verification datasets, a 20-point segment of buffer data was excluded between segments assigned for training and verification. The remaining data points were grouped into 40-point (1000s) segments, which balances the desire to retain as many data points as possible with the need to keep segments short enough that the training and verification sets have the same IWC distributions. Figure 4 shows that training and verification datasets obtained via this process have roughly identical distributions.

b. Particle swarm optimization

PSO was developed to model the use of cognition and cooperation by animal groups to find food or habitats (Kennedy 1997), but it can be applied to any optimization problem. The entities defined as “particles” move about a search space with a velocity composed of three terms: an inertial term pointing in the direction in which a particle was moving in the previous iteration, a memory term that points toward the best location that the particle has personally experienced, and a cooperation term that points toward the best location that any particle in the simulation has experienced. The addition of these three vectors is illustrated in Fig. 5. The “best” outcome (whether personal or group best) is determined by a user-specified cost function that the particles aim to minimize.

Applied to ALPHA, a particle's position represents a set of fuzzy-logic membership functions and blending weights. Each membership function is constrained to be piecewise linear with two inflection points. The functions vary from 0 to 1 (or vice versa) between inflection points, and have a constant value of either 0 or 1 outside that region. The sign of each membership function (i.e., whether interest increases or decreases between

FIG. 5. Illustration of particle movement for each iteration of a particle swarm optimization. A weighted sum of three vectors is added together to obtain the new position of each particle ($\mathbf{x}_{i+1}$): the inertial term points in the direction in which the particle was last moving, the memory term points toward the best position that particle has experienced ($\mathbf{x}_p$), and the cooperation term points toward the best position found by any particle ($\mathbf{x}_g$). The diagram is for illustration purposes only and is not to scale.

inflection points as the variable increases) is predetermined using the empirical relation of an input variable to IWC. Therefore, for each input variable, the PSO space has three dimensions: the location of the first inflection point, the distance between inflection points, and the weight of the input variable in the blending step. Figure 6 illustrates the search space for a simplified algorithm using only cloud optical depth. The particles move through the search space until eventually arriving at the point marked by a star where the cost function is minimized. In this way the optimized membership function and weighting factor are determined for each variable.

A high-level description of the PSO implementation is provided here, with additional details included in the appendix. Figure 7 depicts the PSO process as applied to ALPHA development. The process is *initialized* by assigning starting velocities and positions to 50 particles. For each input variable, the initial three dimensions of the particles' positions are specified at random, and all particles start "at rest" with an inertia term of zero.

The next step consists of *correlation computation* (Fig. 7). The goal of optimization is to maximize the correlation between observed IWC and HIWC potential resulting from the set of parameters represented by each particle's location. Thus, the cost function to minimize during PSO is defined as one minus this correlation. When a position resulting in a lower cost is found, personal and group best positions are *updated*. As the process continues, the number of iterations is monitored

FIG. 6. A simplified diagram of the space in which particles move during the PSO application. For each input variable to ALPHA there are three degrees of freedom: the location of the first membership function inflection point, the distance between inflection points, and the weight of the variable. The particles (10 shown here) move according to Fig. 5, until they reach some final position, denoted by the blue star. For this example, the satellite input field cloud optical depth is used; the star symbol represents the optimal membership function and weight used in ALPHA.

in the *check iterations* step and convergence criteria are tested. If the convergence criteria are not met, computations continue.

Following the check for convergence, the *velocity* vector for each particle is computed using the inertia, memory, and cooperation terms (Fig. 7). Computed velocities are applied to *change* the particle positions by adding the velocity and position vectors. At this point bounds on the domain are *enforced* and particles are corrected for any violations. Constraints imposed here are explained in the appendix. Last, all steps are repeated until the convergence criteria are met (see the appendix). In total, about 1.4 million sets of possible ALPHA parameters were tested during this process.

Because of major differences in the day and night optical depth products, two versions of the satellite PSO were implemented. The daytime version included optical depth as an input but restricted solar zenith angles (SZA) to between 10° and 58° to eliminate glare and terminator effects. A second PSO was implemented using only IR-based products (i.e., excluding optical depth), for all SZA. These thresholds were determined empirically by considering the correlation of optical depth with IWC for different ranges of SZA (not shown). While the daytime algorithm was trained using only the best optical depth data, the IR version used input data from all solar zenith angles. Blending the daytime and IR-based versions for terminator and nighttime hours increases overall performance. This is demonstrated by

FIG. 7. Steps within the particle swarm optimization process used to develop ALPHA.

using CCA, which provides a closed-form solution for the weights producing the highest correlation between IWC and satellite interest (Hotelling 1936). Figure 8 (top panel) shows that, for SZA bins 58° – 68° and larger, a weighted average of the daytime and IR-based versions increases correlation relative to either version alone. The weight for the daytime version as computed by CCA is also shown in Fig. 8 (middle panel). This curve is simplified slightly and then used to blend the daytime and IR-based versions of the satellite algorithm. As a result, when SZA is less than 58° , the daytime version is assigned a weight of 1 and the IR-based version is not used. During terminator hours ($58^{\circ} < \text{SZA} < 83^{\circ}$), the daytime-version weight decreases linearly to 0.3 as SZA approaches 82.5° , the point at which the optical depth retrieval provided by SatCORPS changes to the nighttime algorithm. After this point the optical depth becomes more beneficial again, although a correction must be made to account for differences in the day and nighttime retrievals. This is done with a lookup table that converts the nighttime values to the daytime equivalent by matching percentiles of their observed distributions, allowing the same daytime membership function to be applied to nighttime optical depth. Nevertheless, the nighttime optical depth estimate is hindered by lack of visible imagery, so the daytime-version

weight is only 0.6. Blending the two satellite algorithms in this manner ensures 1) better performance during terminator and nighttime hours and 2) a smoother transition across the domain because the blending weights change gradually. Once the daytime and IR versions are blended into a single satellite interest, that result is combined with the model and radar interests using CCA once again (Fig. 2).

The final component of the algorithm is a simple mask (Fig. 2) that prevents areas that are either too warm for the existence of ice crystals and/or are above cloud top from being classified as HIWC. The model-derived 0°C isotherm and the satellite-derived cloud-top height are used to identify regions with temperatures above freezing and above cloud top, respectively.

c. Resulting algorithm

The optimal set of ALPHA input variables derived from the machine-learning process is listed in Table 1. Although the initial selection of variables is controlled by the user, the PSO process objectively eliminates variables that contribute minimally to the best solution or supply information redundant with other variables. Individual membership functions and weighting factors are given in Table 2. While the selection of final input variables is accomplished objectively with the assignment of weights

FIG. 8. Effect of SZA on the use of satellite data in ALPHA: (top) correlation of ice water content with the daytime version (blue) and IR-based version (red) of the satellite interests as a function of SZA and correlation of blended daytime and IR-based versions with IWC (black circles), (middle) weight of the daytime-version satellite interest when blended with the IR-based version (blue) using CCA, along with the weight used in the actual algorithm (red), and (bottom) the sample size of each SZA bin. To obtain a sample size of at least 500 points in each bin, all data (training, verification, and buffer) were used.

by the PSO, a more comprehensive list of variables thought to be indicative of convective conditions and/or associated with in-service ICI events was considered during the ALPHA development process (Haggerty et al. 2012; Black et al. 2014).

The relationship of select variables retained in the PSO process to IWC measurements are shown in Fig. 9. Some variables have an obvious correlation to IWC. For example, increasing satellite-derived optical depth is clearly associated with increasing IWC, and the membership function derived in the PSO process monotonically increases as optical depth varies from about 20 to 130 where it reaches its maximum value of 1. Similarly, the maximum height of at least 10 dBZ radar reflectivity generally increases with IWC, resulting in a membership function that maximizes interest above 14 km. In some cases, there is little separation between the membership function's inflection points, so the result is basically a step function. Correlations of some variables with IWC are weaker, particularly for the WRF and ACCESS model variables. This result is consistent with anecdotal observations of relatively poor spatial and temporal placement of convective conditions by models during the campaigns.

Weighting factors indicate the relative importance of each variable to the detection of HIWC conditions (Table 2). Among the (daytime) satellite variables, cloud optical depth and cloud-top temperature have the highest correlations with IWC and hence have the largest weights. This result is consistent with observations that engine icing events occur near the highest, coldest, and deepest convective clouds (Bravin and Strapp 2019). It should be noted that the weights assigned to a given variable are not always directly proportional to its correlation with IWC measurements. These results can be explained by the fact that information contained within a single variable may be partially redundant with that provided by other variables used in the algorithm. Radar variables with the largest weights are height of the highest 10 dBZ echo and VAHRR. NWP model variable sets are different for the WRF and ACCESS models; the most significant WRF variable is relative humidity while fractional cloud coverages at low- and midlevel altitudes have the highest weights among the ACCESS fields, potentially providing information about the vertical structure of cloud layers not available from satellite data. Weighting factors associated with individual variables in Table 2 are

TABLE 2. Complete set of membership function parameters and weights used in ALPHA.

Variable	Sign	Inflection points		Weight (%)
Daytime satellite module				
Cloud-top temperature	Negative	−76.6°C	−76.6°C	31.9
Optical depth	Positive	16.2	129.2	61.2
10.8-μm brightness temperature	Negative	−58.2°C	−47.9°C	6.9
IR satellite module				
Cloud-top height	Positive	12.9 km	14.7 km	5.5
Brightness temperature difference (6.8 − 10.8 μm)	Positive	−5.8°C	−1.7°C	33.1
6.8-μm brightness temperature	Negative	−79.1°C	−72.7°C	41.2
10.8-μm brightness temperature	Negative	−57.4°C	−57.4°C	20.2
Radar module				
Composite reflectivity	Positive	20.0 dBZ	31.1 dBZ	2.3
Height of 10 dBZ	Positive	9.8 km	13.6 km	38.4
Cloud depth	Positive	8.5 km	8.5 km	5.1
Ice mass	Positive	0.76 g m ^{−2}	0.76 g m ^{−2}	5.0
VAHRR	Positive	2.7 dBZ km	65.9 dBZ km	49.2
WRF module				
Relative humidity	Positive	76.9%	77.1%	66.2
CAPE	Positive	2.3 kJ kg ^{−1}	2.3 kJ kg ^{−1}	33.8
ACCESS module				
Relative humidity	Positive	82.4%	82.6%	1.6
Convective precipitation	Positive	0.1 mm h ^{−1}	0.1 mm h ^{−1}	15.7
High cloud coverage	Positive	1.00	1.00	12.5
Midlevel cloud coverage	Positive	0.05	0.39	38.5
Low cloud coverage	Positive	0.01	0.57	31.9

applied to obtain intermediate interests from the satellite, model, and radar variables.

All intermediate interests are used at locations where radar data are available (ALPHA three-input version). Satellite and model sources alone are blended when radar data are not available (ALPHA two-input version; note that the NWP model interest has a near-zero weight in the two-input version, but model temperature is still used as a mask at a later stage in the processing). Table 3 gives the weights applied to intermediate interest fields from each data source. Weighting factors applied in the blending process reflect confidence in the ability of each dataset to contribute to the identification of HIWC.

The intermediate satellite interest is given the largest weight in the two-input version and the radar interest is highest in the three-input versions of ALPHA, demonstrating that these products show good correlation with measured IWC. The intermediate satellite interest has the next highest weight in the three-input version. The intermediate model interest is assigned the smallest weight, reflecting the challenges of using a 13-km NWP model for resolving convective processes. From the weights in Table 3, it appears that model variables are able to improve detection relative to satellite information only when radar data are available. Although the PSO scheme assigns far smaller weights to model variables than to satellite and

radar variables, the vertical temperature structure provided by the model is a significant piece of information that is not available from either of the other two sources.

The output from ALPHA is a three-dimensional likelihood of HIWC conditions at each grid cell. This interest field, which varies from 0 to 1, can be considered an uncalibrated probability of the occurrence of HIWC. The calibration process partitions the uncalibrated interest field into 20 discrete bins. A linear fit relates the bin center to the fraction of observations in the bin with IWC of at least 0.5 g m^{-3} (Fig. 10). This equation transforms uncalibrated interest $\mathbf{X}$ to HIWC potential $\mathbf{Y}$, which is the final output product from ALPHA. Any HIWC potential values that fall below 0 or above 1 are corrected to 0 and 1, respectively, after the transformation. The final HIWC potential field product has spatial and temporal resolutions roughly corresponding to the satellite products used (about 5 km). An example of the HIWC potential product during the 2015 Florida flight campaign is shown in Fig. 11.

d. ALPHA performance

Statistically, the performance of HIWC potential is evaluated by comparing it with IWC measurements from the independent verification dataset described above. Multiple metrics are utilized to quantify the performance of ALPHA for detecting HIWC conditions. Some of

FIG. 9. Four of the membership functions used in ALPHA (optical depth, maximum height with at least 10 dBZ from radar, midlevel cloud coverage from ACCESS, and $10.8\text{-}\mu\text{m}$ brightness temperature from the IR-only satellite module). Training data are shown as boxplots, where the short red lines indicate the medians, the rectangle edges indicate the 25th and 75th percentiles, the dashed lines extend to the minimum and maximum nonoutlier values, and outliers are shown as black dots. Although the y axis extends from 0 to 3 g m^{-3} IWC, the membership functions were scaled so that minimum interest (0) appears at the bottom of the axis (0 g m^{-3}) and maximum interest (1) appears at the top (3 g m^{-3}).

these comparisons require definition of an IWC threshold that constitutes “high ice water content” conditions. While such a threshold is still under discussion by the research community and aviation regulators, values of 0.5 g m^{-3} and 1.0 g m^{-3} have been used previously to define HIWC conditions (e.g., de Laat et al. 2017; Yost et al. 2018). Studies of air data system anomalies on research aircraft have demonstrated that disruptions occur for IWC values of 0.5 g m^{-3} when averaged over a 10-min period preceding an event (Haggerty et al. 2019b). Thus, for the analysis presented here, an HIWC threshold of 0.5 g m^{-3} was assumed unless otherwise noted. Conclusions are similar when the analysis is repeated with a higher threshold, albeit with a smaller sample size.

Figure 12 indicates the skill of HIWC potential field by showing the fraction of measured IWC data points that exceed the HIWC threshold for each HIWC potential bin. In general, this histogram, which includes data from the verification dataset only, demonstrates that the fraction of moderate or greater IWC measurements

increases with increasing HIWC potential. The trend is inconsistent at the highest values of HIWC potential where sample sizes are very small.

In selecting an appropriate HIWC potential threshold for use in a future operational product, the trade-off between maximizing probability of detection (POD) and minimizing the probability of false detection (POFD) must be balanced. A contingency table (Table 4) summarizes the categories used for calculation of these statistics. POD is defined as the fraction of HIWC observations for which the HIWC potential is above the specified threshold $[H/(H + M)]$, using the notation from Table 4. Similarly, POFD is defined as the fraction

TABLE 3. Optimized weighting factors used to blend intermediate interest fields in ALPHA.

ALPHA version	Satellite	Radar	Model
Three input	24.0	58.3	17.7
Two input	97.6	—	2.4

FIG. 10. Calibration of (top) two-input ALPHA and (bottom) three-input ALPHA. Equations for the resulting line fits, where X is the ALPHA interest bin center and Y is the fraction of measurements qualifying as HIWC (assuming a threshold of 0.5 g m^{-3}), are shown.

of non-HIWC observations for which the HIWC potential is above the threshold $[F/(F + N)]$. Conversely, the false alarm ratio (FAR) is the fraction of cases with an HIWC potential above the threshold but the measured IWC is below 0.5 g m^{-3} $[F/(F + H)]$. As indicated in Fig. 13, setting an HIWC potential threshold of 0.25 would yield POD, POFD, and FAR of 86%, 29%, and 51%, respectively, for the verification dataset using all IWC values (omitting points where IWC was less than

0.1 g m^{-3} , thus forcing the algorithm to distinguish between moderate and high IWC, increases POD to 60% while POD and FAR remain unchanged). A lower threshold would give a higher POD, but would also produce a higher POFD.

FIG. 11. HIWC potential as estimated by ALPHA in the Gulf of Mexico on 16 Aug 2015. The red-outlined box surrounds an MCS sampled by the NASA DC-8 aircraft.

FIG. 12. Fraction of IWC measurements from the airborne IKP2 that exceed the high-IWC threshold of 0.5 g m^{-3} (blue bars) and 1.0 g m^{-3} (green bars). Sample sizes for each bin are shown in blue text above the bars or white text within them.

TABLE 4. Contingency table defining categories used for POD statistics.

	$0.1 < \text{IWC} < 0.5 \text{ g m}^{-3}$	$\text{IWC} > 0.5 \text{ g m}^{-3}$
HIWC potential < threshold	Null (<i>N</i>)	Miss (<i>M</i>)
HIWC potential > threshold	False alarm (<i>F</i>)	Hit (<i>H</i>)

Receiver operating characteristic (ROC) curves compare the POD with the POFD, providing a single performance metric with the “area under the curve” (AUC) statistic. The curves shown in Fig. 14 are derived by assuming an HIWC threshold of 0.5 g m^{-3} , and then varying the ALPHA HIWC potential threshold to obtain different combinations of POD versus POFD. A perfect algorithm would have a curve reaching the upper-left corner of the plot (POD of 1; POFD of 0). In this plot, the blue curve represents the performance using all verification observations whereas the red curve uses only in-cloud observations ($\text{IWC} > 0.1 \text{ g m}^{-3}$). The

yellow and black curves use the same in-cloud definition as the red one but are separated to represent daytime ($\text{SZA} \leq 58^\circ$) and night/terminator ($\text{SZA} > 58^\circ$) algorithms, respectively. Overall performance is related to the AUC; for the verification dataset the AUC is 0.85 when considering all data and 0.69 for $\text{IWC} > 0.1 \text{ g m}^{-3}$. When stratified by SZA, in-cloud AUC values are 0.74 for daytime and 0.65 for night and terminator hours.

Given biases inherent in the aircraft dataset used for training the algorithm, an objective assessment of confidence in the POD, POFD, and FAR statistics is useful. For example, the aircraft data are biased toward high values of IWC because sampling such conditions was the stated objective of the field campaigns. Also, time series data such as IWC measurements have some degree of autocorrelation as demonstrated in Fig. 6. For this reason, a stationary bootstrap (Politis and Romano 2012) was applied to determine the representativeness of the data used. In contrast to the traditional bootstrap (Efron 1979), stationary bootstrap works by sampling

FIG. 13. ALPHA performance statistics showing probability of detection, probability of false detection (for data in and out of cloud), and false alarm ratio for various combinations of IWC and HIWC potential thresholds. Statistics are computed using only the independent verification dataset.

FIG. 14. ROC curve for HIWC potential. The blue curve represents the performance using all verification observations, and the red curve uses only in-cloud observations ($IWC > 0.1 \text{ g m}^{-3}$). The yellow and black curves use the same in-cloud definition as the red one but are separated to represent only daytime (solar zenith $\leq 58^\circ$) and night/terminator (solar zenith $> 58^\circ$) algorithms, respectively. AUC values are shown in the legend. An IWC threshold of 0.5 g m^{-3} was assumed to define HIWC.

segments, or blocks of data, which mimics the manner in which aircraft observations occur in a sequence. The number of 25-s IWC observations in a block N is varied according to a geometric distribution, so the probability of N being some nonnegative integer n is given by

$$P(N = n) = p(1 - p)^n,$$

where p is a value between 0 and 1. The parameter p is computed so that the mean length of samples is the same length as the buffer regions used in separating the independent training and verification datasets. The mean of a geometric distribution is

$$(1 - p)/p.$$

So the geometric parameter p is set to $1/21 = 0.0476$, resulting in a mean of 20 points (500 s).

In addition to accounting for autocorrelations in the IWC data, the resampling procedure must consider the disproportionate number of extreme data points resulting from actively seeking HIWC conditions. To mimic this bias in the sampled data, the probability of choosing a point as the start of a sample block was weighted by $(1 + 2 \times \text{HIWC potential})$. This weighting assumes that the flight tracks were 3 times as likely to pass through an area of HIWC potential equal to 1 as an area equal to 0. Methods would be more rigorous if the sampling bias were quantified by comparing the distribution of sampled HIWC potential to the

FIG. 15. Distributions of the probability of detection of 10 000 bootstrapped samples using the stationary bootstrap technique (see text). Blue bars show the number of resampled sets with a given range of values; the red lines indicate the mean of resampled values and the value of the original dataset, which are also written in the top left.

distribution of HIWC potential over the larger geographic domain of the field campaigns. However, it would be too time consuming to run ALPHA for all the field campaign domains to obtain the actual distribution. Note that these resampling statistics were produced using all 18 734 data points where temperature was below freezing ($T < 0$) because the training dataset alone is too small considering that the resulting confidence intervals are already wide. Blocks of data are thus taken until the total number of points exceeds 18 734. The data are then truncated to give a sample size of exactly 18 734. This was done 10 000 different times to provide a distribution of performance statistics.

Figure 15 shows the distribution of POD from the 10 000 resulting samples. The mean of the distribution is indicated on the plot. Because the resampled mean POD is higher than the observed statistic, the true statistic is probably lower than observed. Confidence intervals derived with thresholds of IWC at 0.5 g m^{-3} and HIWC potential at 0.25 are shown in Table 5. The confidence intervals indicate a 95% probability that the actual value of the statistic falls within the interval listed. For example, these results show that there is a 95% probability that the ALPHA POD is between 78.5% and 83.1%, with an expected (mean) value of 80.7%. Because both the estimated POD and POFD are lower than observed while FAR remains roughly unchanged (Fig. 13), it would be prudent for users to lean toward using a lower HIWC potential threshold.

e. ALPHA case studies

The HIWC potential field for 16 August 2015 shows a large MCS positioned over the northern Gulf of Mexico

TABLE 5. ALPHA performance statistic 95% confidence intervals derived from bootstrap resampling method. All statistics use an observed IWC threshold of 0.5 g m^{-3} and an HIWC potential threshold of 0.25.

Statistic	Median (%)	Lower bound (%)	Upper bound (%)
Probability of detection	80.7	78.5	83.1
Probability of false detection for $\text{IWC} > 0.1$	53.3	49.7	57.0
Probability of false detection for all IWC measurements	19.8	16.2	23.2
False alarm ratio	50.3	47.1	53.6

(Fig. 11). The NASA DC-8 aircraft operated in this system south of the Florida Panhandle, conducting a series of straight and level flight segments through the storm at altitudes ranging from about 9 to 11 km (ambient temperatures from -30° to -50°C). Cloud-top heights exceeded 14 km, with occasional overshooting cloud tops that penetrated the tropopause. Sustained IWC values of $1.0\text{--}2.0 \text{ g m}^{-3}$ were observed by the IKP2

sensor during these segments, with peaks exceeding 2.5 g m^{-3} . The HIWC potential estimated by ALPHA at 1500 UTC gives moderate-to-high values of 0.3–0.6 in most of the flight domain.

A sector of the domain in Fig. 16 shows a 1.5-h flight segment overlaid on the ALPHA HIWC potential field. Measurements of IWC from the aircraft are color-coded along the flight track. Figure 17 compares time series

FIG. 16. HIWC potential (gray shading) estimated from ALPHA on 16 Aug 2015. Valid times correspond to the initial time on each panel; altitudes correspond to a flight level of 29 000 ft (8839 m) in the first five panels and 30 000 ft (9144 m) in the last panel. The flight segment across the middle of each image occurred at the valid time of the HIWC potential background image.

FIG. 17. For the flight segment shown in Fig. 16, (top) the observed IWC, (top middle) the HIWC potential along the flight path, and the (middle) satellite, (bottom middle) radar, and (bottom) model intermediate interests. Areas where the HIWC potential threshold of 0.25 did not capture observed HIWC ($IWC > 0.5 \text{ g m}^{-3}$) are highlighted in red, and areas where HIWC potential gave a false alarm are indicated in gray.

of the HIWC potential with the in situ IWC along the flight track. Intermediate interest values from the satellite, NWP model, and radar modules of ALPHA are also plotted. Missed events where the HIWC potential threshold did not exceed the 0.25 threshold during observed HIWC ($IWC > 0.5 \text{ g m}^{-3}$) are highlighted as red bands. False alarms (HIWC potential > 0.25 and $IWC < 0.5 \text{ g m}^{-3}$) are indicated by gray bands. In general, misses and false alarms were short in duration and occurred near edges of clouds. For example, the missed HIWC around 1505 UTC occurred in the small pocket of HIWC potential less than 0.25 between two larger storm cells (Fig. 16). Satellite and radar intermediate interests both track the IWC time series data closely, and the HIWC potential time series shows a high similarity to the satellite interest with adjustments in magnitude due to the radar interest and calibration. The model interest indicates an awareness of the storm, although the break in high model interest from 1500 to 1550 UTC suggests that placement and/or horizontal extent of the convective activity was inaccurate. This

result is consistent with findings of Qu et al. (2018), who found differences between simulated and observed cloud amounts using the Canadian Global Environmental Multiscale model to analyze an MCS near French Guiana on 16 May 2015 during the Cayenne field campaign.

Similar comparisons for a flight based in Darwin on 23 January 2014 are shown in Figs. 18 and 19. The target was an area of convection associated with the monsoon trough in the Bonaparte Gulf southwest of Darwin. Eight flight segments were conducted at altitudes ranging from about 9.5 to 12 km corresponding to ambient temperatures from -30° to -50°C . Sustained IWC values of 2.0 g m^{-3} and peak values of 2.5 g m^{-3} were observed, although the system decayed in the later part of the flight.

Figure 19 shows brief periods during which false alarms and misses occur on the edges of clouds as in Fig. 17. The aircraft was not within range of a ground-based radar for most of this flight, so HIWC potential estimates are composed of satellite and NWP model interest fields (two-input ALPHA version). Of note in

01/23/2014 Darwin Flight

FIG. 18. As in Fig. 16, but showing segments from a Darwin-based flight on 23 Jan 2014. Altitudes used are 39 000 ft (11 887 m) in the first three panels, 33 000 ft (10 058 m) in the next two panels, and 31 000 ft (9449 m) in the last panel.

this case is an extended period of false alarm (gray shading) centered at 2210 UTC where measured IWC was below the 0.5 g m^{-3} threshold, but HIWC potential was slightly above 0.25. A possible explanation may be found in the onboard cloud radar data shown in Fig. 20 (Protat et al. 2016). At the time of the false alarm, which occurred during the satellite terminator period, the RASTA cloud radar IWC retrieval suggests that the aircraft was in thin cloud, while the satellite-estimated cloud-top height is well above flight level. An over-estimated satellite cloud-top height results in a higher satellite interest field, and thus HIWC potential was slightly overestimated. The RASTA IWC retrieval does indicate HIWC in the column below the aircraft at these times. As seen in the bottom panels of Fig. 18, the areas responsible for the false alarms decayed shortly after.

5. Summary and conclusions

This paper describes development of a method for discerning HIWC conditions near deep convective clouds using input data from satellites, NWP models, and ground-based radar. Airborne in situ measurements of IWC from a series of field experiments were used to discover which fields from each data source have the strongest correlation to measurements. A machine-learning technique was applied to determine the optimal way to combine available information using fuzzy logic. Algorithm training used a subset of the airborne IWC measurements; an independent subset was reserved for verification. The resulting algorithm estimates the potential for HIWC conditions over a three-dimensional domain.

Performance statistics indicate varying skill level, depending on the ALPHA HIWC potential value and

FIG. 19. As in Fig. 17, but for the 2100–2300 UTC segment of the 23 Jan 2014 flight of the SAFIRE Falcon 20 from Darwin.

on the IWC threshold used to define HIWC. Assuming an IWC threshold of 0.5 g m^{-3} , POD, POFD, and FAR are 86%, 29%, and 51%, respectively, for an HIWC potential value of 0.25. These statistics are derived from the verification dataset using all points (POFD increases to 60% when only data points with $\text{IWC} \geq 0.1 \text{ g m}^{-3}$ are used). The corresponding ROC curves give an AUC of 0.85 when considering all data and 0.69 for data points with $\text{IWC} \geq 0.1 \text{ g m}^{-3}$. A statistical resampling procedure used to estimate confidence intervals suggests that there is a 95% probability that the ALPHA POD is between 78.5% and 83.1% with a median of 80.7%. Analysis of individual case studies shows that the ALPHA HIWC potential generally tracks variation in the measured IWC. Results are similar when the analysis is repeated with an IWC threshold of 1.0 g m^{-3} , but the sample size is much smaller.

These results demonstrate the ALPHA method has skill for detecting HIWC conditions, albeit with significant false alarm rates. HIWC conditions are associated with ICI events experienced by jet aircraft, so ALPHA may offer some guidance to the aviation community for avoiding this risk.

Note that the algorithm was trained using measurements of IWC from research aircraft, rather than data from actual ICI engine events. Point values of high IWC are just one factor associated with ICI events.

FIG. 20. Ice water content (color shading) from the RASTA radar retrievals and the in situ IKP2 sensor (between the white lines). The red curve indicates the satellite cloud-top height retrieval used in ALPHA. Aircraft altitude is midway between the white lines.

The horizontal scale of the HIWC feature (i.e., the duration of exposure of an aircraft to these conditions) and specific engine characteristics are also thought to contribute to the occurrence of events. Another limitation of the training dataset, and thus the ALPHA method, is that the research data may not encompass all of the environmental conditions associated with in-service ICI events. J. W. Strapp et al. (2019, unpublished FAA report) reported on the representativeness of the dataset gathered in the first three of the four experiments that compose the ALPHA training and verification datasets. Their analysis addresses geographic diversity of sampling locations, as well as possible biases in the research dataset when compared with the current inventory of ICI events. Biases may result from diurnal sampling limitations, impact of sampling in predominantly low aerosol environments, and the absence of data in large vigorous continental MCS associated with as many as 40% of ICI events according to [Bravin and Strapp \(2019\)](#).

Ongoing research will address the need for regional tuning of algorithms, better characterization of vertical variation of HIWC conditions, and short-term forecasting methods for predicting the ICI hazard. Improvements to ALPHA can likely be made with use of higher-resolution NWP models and more advanced satellite sensors than were available during the experiments that began in 2014. In areas where a network of ground-based radars are available, better use of three-dimensional radar products is warranted. In addition, derived satellite products, such as the location of overshooting cloud tops, may be useful for improving detection capability.

Feedback from potential users will be an important source of information for refining the capability and utility of products like ALPHA in a real-world setting. Outreach efforts are currently underway to introduce HIWC detection capabilities to weather forecasters and airline users. Under a joint effort by the Australian Bureau of Meteorology, NCAR, and the Federal Aviation Administration (FAA), an HIWC nowcasting trial is being conducted with real-time ALPHA products provided to industry users ([Haggerty and Potts 2017](#)). In parallel with research activities in this area, the International Civil Aviation Organization has acknowledged a requirement by the international aviation industry for HIWC guidance products, noting the need for feedback on the outcome of the trial, including how the information may change decision-making processes for pilots.

Acknowledgments. The authors are grateful to Gary Cunning and George McCabe of NCAR for assistance with software development. Frank McDonough and Jennifer Black contributed to early versions of

ALPHA. Cathy Kessinger and James Bresch provided valuable insight on the use of radar and satellite data and NWP model fields, respectively. Rabindra Palikonda and Christopher Yost provided processed satellite datasets. The authors also thank the organizers, collaborators, and sponsors of the HAIC and HIWC projects who enabled the collection of essential datasets for this work and provided support for development of these products. This research is in response to requirements and funding by the Federal Aviation Administration (FAA). The views expressed are those of the authors and do not necessarily represent the official policy or position of the FAA. The HAIC program received funding from the European Union's Seventh Framework Program in research, technological development and demonstration under Grant Agreement ACP2-GA-2012-314314. The HIWC project received funding from the FAA Aviation Weather Research Program. This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement 1852977.

APPENDIX

Implementation of the Particle Swarm Optimization Process

The flow of the PSO applied to ALPHA is depicted in [Fig. 7](#). The process is *initialized* by assigning starting velocities and positions to 50 particles. A swarm size of 50 is consistent with standards from [Bratton and Kennedy \(2007\)](#), who found that for swarm sizes between 20 and 100, there is usually little difference in PSO performance. For each input variable, the initial three dimensions of the particles' positions are specified. The first inflection point is selected at random from a uniform probability density (UPD) between the 5th and 50th percentiles of observed values. The width of the membership function (distance between inflection points) is then taken so that the second inflection point is larger than the first, but less than the 95th percentile of observations. Finally, the variable's initial weighting factor is determined from a UPD function between 0 and 1. The weights are then normalized so that the sum of all weights for a given particle is 1. All particles start "at rest" with an inertia term of zero.

The next step consists of *correlation computation* ([Fig. 7](#)). The goal of the optimization is to maximize the correlation between the observed IWC and the HIWC potential resulting from the set of parameters represented by each particle's location. Thus, the cost function to minimize during PSO is defined as one minus this correlation. When a position resulting in a lower cost is found,

personal and group best positions are *updated*. As the process continues, the number of iterations is monitored. In the *check iterations* step, if there has been no change in the global best set of parameters in the last 100 iterations and the total number of iterations is at least 200, the algorithm stops and returns the group best position as the solution. While 100 and 200 are somewhat arbitrary, this ensures that with 50 particles, 5000 different parameter sets are tested after the last update without improved performance, and that a minimum of 10 000 total parameter sets are tested for each PSO run. On average, about 350 total iterations were necessary to satisfy this convergence criteria. If the convergence criteria are not met, computations continue.

Following the check for convergence, the *velocity* vector for each particle is computed using the inertia, memory, and cooperation terms (Fig. 7). The weights for the memory and cooperation terms are taken from a UPD ranging from 0 to 1 (denoted $U[0, 1]$) with new random weights for each iteration. These random weights make the process stochastic and improve the searching capability of the system. The inertial weight begins at 1.5 and decreases over the first 200 iterations to a value of 0.25 to encourage early exploration followed by eventual convergence, as published in previous studies using PSO to tune fuzzy-logic membership functions (e.g., Esmin and Lambert-Torrez 2006; Bratton and Kennedy, 2007). In mathematical terms, the velocity is computed as follows:

$$\mathbf{V}_i = \left[1.5 - 1.25 \times \min \left(1, \frac{i}{200} \right) \right] \mathbf{V}_{i-1} + U[0, 1](\mathbf{X}_p - \mathbf{X}_{i-1}) + U[0, 1](\mathbf{X}_g - \mathbf{X}_{i-1}),$$

where $\mathbf{V}_i$ is the particle velocity at iteration step i , and $\mathbf{X}_g$ and $\mathbf{X}_p$ are the global best and personal best positions, respectively. Using the computed velocities, an update is applied to *change* the particle positions by adding the velocity and position vectors:

$$\mathbf{X}_i = \mathbf{X}_{i-1} + \mathbf{V}_i.$$

At this point bounds on the domain are *enforced* and particles are corrected for any violations. The following constraints are imposed:

- The first inflection point of a membership function is bounded by the 5th and 95th percentiles of the observed values of the input variable.
- Bounds on the width of the membership function depend on the first inflection point so that the resulting second inflection is larger than the first and also between the 5th and 95th percentiles of the observations.

- Weights are permitted to be any nonnegative value but are normalized to sum to 1 during every iteration.
- If any of the three conditions above is violated, the position is corrected to the nearest valid point. The appropriate velocity component is also changed to be -0.5 times the original, so the inertial term for the next iteration points back toward the valid domain.

All steps are repeated until the convergence criteria are met.

For each of three data sources (top level of Fig. 2), 10 PSO runs are executed as above with different random initializations. A second round of 10 PSO runs follows using the same parameters except that 10 of the starting positions are those discovered by the initial 10 runs; the other 40 are random as before. The best result from the second set of runs comprises the fuzzy-logic parameters used in ALPHA. In total, among all PSO runs and all particles, about 1.4 million sets of possible ALPHA parameters were tested.

REFERENCES

- Australian Bureau of Meteorology, 2013: APS1 upgrade of the ACCESS-R numerical weather prediction system. National Meteorological and Oceanographic Centre Rep. 98, 28 pp., <http://www.bom.gov.au/australia/charts/bulletins/apob98.pdf>.
- , 2015: N. Australia 3D mosaic radar data, version 1.0. UCAR/NCAR Earth Observing Laboratory, accessed 20 November 2017, <https://data.eol.ucar.edu/dataset/384.017>.
- Aviation Safety Network, 2016: FAA orders engine icing fixes for GENx-powered Boeing 787 Dreamliners. Flight Safety Foundation, <http://news.aviation-safety.net/2016/04/23/faq-orders-engine-icing-fixes-for-genx-powered-boeing-787-dreamliners/>.
- Bernstein, B. C., F. McDonough, M. K. Politovich, B. G. Brown, T. P. Ratvasky, D. R. Miller, C. A. Wolff, and G. Cuning, 2005: Current icing potential: Algorithm description and comparison with aircraft observations. *J. Appl. Meteor.*, **44**, 969–986, <https://doi.org/10.1175/JAM2246.1>.
- Beswick, K., and Coauthors, 2015: Properties of small cirrus ice crystals from commercial aircraft measurements and implications for flight operations. *Tellus*, **67B**, 27876, <https://doi.org/10.3402/TELLUSB.V67.27876>.
- Black, J., J. Haggerty, G. McCabe, C. A. Wolff, G. Cuning, and A. Grandin, 2014: Application of a nowcasting tool for analysis of meteorological conditions associated with engine icing. *Sixth Atmospheric and Space Environments Conf.*, Atlanta, GA, American Institute of Aeronautics and Astronautics, 2014–2068, <https://doi.org/10.2514/6.2014-2068>.
- Bratton, D., and J. Kennedy, 2007: Defining a standard for particle swarm optimization. *Institute of Electrical and Electronics Engineers Swarm Intelligence Symp.*, Washington, DC, IEEE, 120–127, <https://doi.org/10.1109/SIS.2007.368035>.
- Bravin, M., and J. Strapp, 2019: A continuing investigation of diurnal and location trends in an ice crystal icing engine event database. *Int. Conf. on Icing of Aircraft, Engines, and Structures*, Minneapolis, MN, SAE International, 2019–01-1964, <https://doi.org/10.4271/2019-01-1964>.
- , —, and J. G. Mason, 2015: An investigation into location and convective lifecycle trends in an ice crystal icing engine

- event database. *Int. Conf. on Icing of Aircraft, Engines, and Structures*, Prague, Czech Republic, SAE International, 2015-01-2130, <https://doi.org/10.4271/2015-01-2130>.
- Davison, C. R., C. Landreville, and J. D. MacLeod, 2010: Initial development and testing of isokinetic probe to measure total water content during ground and airborne testing. National Research Council Canada Rep. LTR-GTL-2010-0002, 119 pp.
- , J. W. Strapp, L. Lilie, T. P. Ratvasky, and C. Dumont, 2016: Isokinetic TWC evaporator probe: Calculations and systemic error analysis. *Eighth Atmospheric and Space Environments Conf.*, Washington, DC, American Institute of Aeronautics and Astronautics, 2016-4060, <https://doi.org/10.2514/6.2016-4060>.
- Defer, E. J., and Coauthors, 2015: Overview of the HAIC “Space-borne Observation and Nowcasting of High Ice-Water Content Regions” sub-project and mid-term results. *2015 Int. Conf. on Icing of Aircraft, Engines, and Structures* Prague, Czech Republic, SAE International, 2015-01-2123, <https://doi.org/10.4271/2015-01-2123>.
- de Laat, A., and Coauthors, 2017: Analysis of geostationary satellite-derived cloud parameters associated with environments with high ice-water content. *Atmos. Meas. Tech.*, **10**, 1359–1371, <https://doi.org/10.5194/amt-10-1359-2017>.
- Dezitter, F., A. Grandin, J.-L. Brenguier, F. Hervy, H. Schlager, P. Villedieu, and G. Zalamansky, 2013: HAIC—High altitude ice crystals. *Fifth Atmospheric and Space Environments Conf.*, San Diego, CA, American Institute of Aeronautics and Astronautics, 2013-2674, <https://doi.org/10.2514/6.2013-2674>.
- Dye, J. E., and Coauthors, 2007: Electric fields, cloud microphysics, and reflectivity in anvils of Florida thunderstorms. *J. Geophys. Res.*, **112**, D11215, <https://doi.org/10.1029/2006JD007550>.
- Efron, B., 1979: Bootstrap methods: Another look at the jackknife. *Ann. Stat.*, **7**, 1–26, <https://doi.org/10.1214/aos/1176344552>.
- Esmin, A. A. A., and G. Lambert-Torrez, 2006: Fitting fuzzy membership functions using hybrid particle swarm optimization. *Int. Conf. on Fuzzy Systems*, Vancouver, BC, Canada, Institute of Electrical and Electronics Engineers, <https://doi.org/10.1109/FUZZY.2006.1681993>.
- Fridlind, A., A. Ackerman, A. Grandin, F. Dezitter, M. Weber, J. W. Strapp, A. V. Korolev, and C. R. Williams, 2015: High ice water content at low radar reflectivity near deep convection—Part I: Consistency of in situ and remote-sensing observations with stratiform rain column simulations, 2015. *Atmos. Chem. Phys.*, **15**, 11 713–11 728, <https://doi.org/10.5194/acp-15-11713-2015>.
- General Civil Aviation Authority, 2015: A346, en route, eastern Indian Ocean, 2013. General Civil Aviation Authority Rep., http://www.skybrary.aero/index.php/A346_en_route_eastern_Indian_Ocean_2013.
- Gounou, A., and Coauthors, 2015: The use of RDT nowcasting tool for detecting convective areas associated with high ice water content during HAIC/HIWC field campaign. *Int. Conf. on Icing of Aircraft, Engines, and Structures*, Prague, Czech Republic, SAE International, 2015-01-2124, <https://doi.org/10.4271/2015-01-2124>.
- Grzych, M., and J. Mason, 2010: Weather conditions associated with jet engine power loss and damage due to ingestion of ice particles: What we’ve learned through 2009. *14th Conf. on Aviation, Range and Aerospace Meteorology*, Atlanta, GA, Amer. Meteor. Soc., 6.8, <https://ams.confex.com/ams/pdfpapers/165923.pdf>.
- , T. Tritz, J. Mason, M. Bravin, and A. Sharpsten, 2015: Studies of cloud characteristics related to jet engine ice crystal icing utilizing infrared satellite imagery. *Int. Conf. on Icing of Aircraft, Engines, and Structures*, Prague, Czech Republic, SAE International, 2015-01-2086, <https://doi.org/10.4271/2015-01-2086>.
- Haggerty, J., and R. Potts, 2017: A high ice water content (HIWC) nowcasting trial exercise for enhanced situational awareness and decision-making support. *Aeronautical Meteorology Scientific Conf.*, Toulouse, France, WMO.
- , F. McDonough, J. Black, G. Cuning, G. McCabe, M. Politovich, and C. Wolff, 2012: A system for nowcasting atmospheric conditions associated with jet engine power loss and damage due to ingestion of ice particles. *Fourth Atmospheric and Space Environment Conf.*, New Orleans, LA, American Institute of Aeronautics and Astronautics, 2012-3234, <https://doi.org/10.2514/6.2012-3234>.
- , and Coauthors, 2019a: Detecting clouds associated with jet engine ice crystal icing. *Bull. Amer. Meteor. Soc.*, **100**, 31–40, <https://doi.org/10.1175/BAMS-D-17-0252.1>.
- , A. Rugg, T. Ratvasky, J. W. Strapp, L. Lilie, K. Bedka, J. Jensen, and A. Bansemmer, 2019b: Characterization of high ice water content conditions that impact air data system performance. *19th Conf. on Aviation, Range, and Aerospace Meteorology*, Phoenix, AZ, Amer. Meteor. Soc., 10.4, <https://ams.confex.com/ams/2019Annual/webprogram/Paper350460.html>.
- Harrah, S., T. Ratvasky, J. W. Strapp, L. Lilie, J. Strickland, and P. Hunt, 2016: Overview of NASA 2015 HIWC Radar flight campaign. *HAIC Final Forum*, Toulouse, France, Airbus, <https://www.haic.eu/media-2/events/conference-presentations/13/>.
- Hong, G., P. Yang, B. Gao, B. A. Baum, Y. X. Hu, M. D. King, and S. Platnick, 2007: High cloud properties from three years of MODIS Terra and Aqua collection-4 data over the tropics. *J. Appl. Meteor. Climatol.*, **46**, 1840–1856, <https://doi.org/10.1175/2007JAMC1583.1>.
- Hotelling, H., 1936: Relations between two sets of variates. *Biometrika*, **28**, 321–377, <https://doi.org/10.2307/2333955>.
- Kalluri, S., and Coauthors, 2018: From photons to pixels: Processing data from the Advanced Baseline Imager. *Remote Sens.*, **10**, 177, <https://doi.org/10.3390/rs10020177>.
- Kennedy, J., 1997: The particle swarm: Social adaptation of knowledge. *Proc. Int. Conf. on Evolutionary Computation*, Indianapolis, IN, Institute of Electrical and Electronics Engineers, 303–308, <https://doi.org/10.1109/ICEC.1997.592326>.
- Lawson, R. P., L. Angus, and A. Heymsfield, 1998: Cloud particle measurements in thunderstorm anvils and possible threat to aviation. *J. Aircr.*, **35**, 113–121, <https://doi.org/10.2514/2.2268>.
- Leroy, D., and Coauthors, 2017a: Ice crystal sizes in high ice water content clouds. Part II: Statistics of mass diameter percentiles in tropical convection observed during the HAIC/HIWC Project. *J. Atmos. Oceanic Technol.*, **34**, 117–136, <https://doi.org/10.1175/JTECH-D-15-0246.1>.
- , and Coauthors, 2017b: HAIC/HIWC field campaigns—Specific findings on ice crystal characteristics in high ice water content regions. *Eighth Atmosphere and Space Environment Conf.*, Washington, DC, American Institute of Aeronautics and Astronautics, 2016-4056, <https://doi.org/10.2514/6.2016-4056>.
- Mason, J. G., J. W. Strapp, and P. Chow, 2006: The ice particle threat to engines in flight. *44th Aerospace Sciences Meeting*, Reno, NV, American Institute of Aeronautics and Astronautics, 2006-206, <https://doi.org/10.2514/6.2006-206>.
- Minnis, P., and Coauthors, 2011: CERES edition-2 cloud property retrievals using TRMM VIRS and Terra and Aqua MODIS data—Part I: Algorithms. *IEEE Trans. Geosci. Remote Sens.*, **49**, 4374–4400, <https://doi.org/10.1109/TGRS.2011.2144601>.

- Politis, D. N., and J. P. Romano, 2012: The stationary bootstrap. *J. Amer. Stat. Assoc.*, **89**, 1303–1313, <https://doi.org/10.1080/01621459.1994.10476870>.
- Proctor, F., S. Harrah, G. F. Switzer, J. K. Strickland, and P. J. Hunt, 2017: High ice water concentrations in the 19 August 2015 coastal mesoconvective system. *Ninth Atmosphere and Space Environment Conf.*, Denver, CO, American Institute of Aeronautics and Astronautics, 2017-4370, <https://doi.org/10.2514/6.2017-4370>.
- Protat, A., and Coauthors, 2016: The measured relationship between ice water content and cloud radar reflectivity in tropical convective clouds. *J. Appl. Meteor. Climatol.*, **55**, 1707–1729, <https://doi.org/10.1175/JAMC-D-15-0248.1>.
- Puri, K., and Coauthors, 2013: Implementation of the initial ACCESS numerical weather prediction system. *Aust. Meteor. Oceanogr. J.*, **63**, 265–284, <https://doi.org/10.22499/2.6302.001>.
- Qu, Z., and Coauthors, 2018: Evaluation of a high-resolution numerical weather prediction model's simulated clouds using observations from *CloudSat*, *GOES-13* and in-situ aircraft. *Quart. J. Roy. Meteor. Soc.*, **144**, 1681–1694, <https://doi.org/10.1002/qj.3318>.
- Ratvasky, T., and Coauthors, 2019: Summary of the High Ice-Water Content (HIWC) radar flight campaigns. *Int. Conf. on Icing of Aircraft, Engines, and Structures*, Minneapolis, MN, SAE International, 2019-01-2027, <https://doi.org/10.4271/2019-01-2027>.
- Rogers, R., S. Aberson, J. Kaplan, and S. Goldenberg, 2002: A pronounced upper-tropospheric warm anomaly encountered by the NOAA G-IV aircraft in the vicinity of deep convection. *Mon. Wea. Rev.*, **130**, 180–187, [https://doi.org/10.1175/1520-0493\(2002\)130<0180:APUTWA>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<0180:APUTWA>2.0.CO;2).
- Skamarock, W. C., and Coauthors, 2008: A description of the Advanced Research WRF version 3. NCAR Tech. Note NCAR/TN-475+STR, 113 pp., <https://doi.org/10.5065/D68S4MVH>.
- Strapp, J. W., and Coauthors, 2016a: The High Ice Water Content (HIWC) study of deep convective clouds: Report on science and technical plan. FAA Rep. DOT/FAA/TC-14/31, 105 pp., www.tc.faa.gov/its/worldpac/techrpt/tc14-31.pdf.
- , L. E. Lilie, T. P. Ratvasky, C. R. Davison, and C. Dumont, 2016b: Isokinetic TWC evaporator probe: Development of the IKP2 and performance testing for the HAIC-HIWC Darwin 2014 and Cayenne field campaigns. *Eighth Atmospheric and Space Environments Conf.*, Washington, DC, American Institute of Aeronautics and Astronautics, 2016-4059, <http://arc.aiaa.org/doi/10.2514/6.2016-4059>.
- Yost, C. R., and Coauthors, 2018: A prototype method for diagnosing high ice water content probability using satellite imager data. *Atmos. Meas. Tech.*, **11**, 1615–1637, <https://doi.org/10.5194/amt-11-1615-2018>.
- Zipser, E., D. Cecil, C. Liu, S. Nesbitt, and D. Yorty, 2006: Where are the most intense thunderstorms on Earth? *Bull. Amer. Meteor. Soc.*, **87**, 1057–1072, <https://doi.org/10.1175/BAMS-87-8-1057>.