SPARE-ICE: synergistic Ice Water Path from passive operational sensors

Article

Published Version

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

RESEARCH ARTICLE

Key Points:
- Combine IR and MW to retrieve IWP from passive operational sensors
- We use collocations to train with CloudSat CPR 2C-ICE as a reference data set
- Compares well to MODIS and seems to do well even in difficult conditions

Correspondence to:
G. Holl,
gerrit.holl@gmail.com

SPARE-ICE: Synergistic ice water path from passive operational sensors

G. Holl[1], S. Eliasson[1,2], J. Mendrok[1], and S. A. Buehler[1,3]

[1]Division of Space Technology, Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, Kiruna, Sweden, [2]Atmospheric Remote Sensing Unit, Research and Development Department, Swedish Meteorological and Hydrological Institute, Norrköping, Sweden, [3]Now at Meteorological Institute, University of Hamburg, Hamburg, Germany

Abstract This article presents SPARE-ICE, the Synergistic Passive Atmospheric Retrieval Experiment-ICE. SPARE-ICE is the first Ice Water Path (IWP) product combining infrared and microwave radiances. By using only passive operational sensors, the SPARE-ICE retrieval can be used to process data from at least the NOAA 15 to 19 and MetOp satellites, obtaining time series from 1998 onward. The retrieval is developed using collocations between passive operational sensors (solar, terrestrial infrared, microwave), the CloudSat radar, and the CALIPSO lidar. The collocations form a retrieval database matching measurements from passive sensors against the existing active combined radar-lidar product 2C-ICE. With this retrieval database, we train a pair of artificial neural networks to detect clouds and retrieve IWP. When considering solar, terrestrial infrared, and microwave-based measurements, we show that any combination of two techniques performs better than either single-technique retrieval. We choose not to include solar reflectances in SPARE-ICE, because the improvement is small, and so that SPARE-ICE can be retrieved both daytime and nighttime. The median fractional error between SPARE-ICE and 2C-ICE is around a factor 2, a figure similar to the random error between 2C-ICE ice water content (IWC) and in situ measurements. A comparison of SPARE-ICE with Moderate Resolution Imaging Spectroradiometer (MODIS), Pathfinder Atmospheric Extended (PATMOS-X), and Microwave Surface and Precipitation Products System (MSPPS) indicates that SPARE-ICE appears to perform well even in difficult conditions. SPARE-ICE is available for public use.

1. Introduction

The systematic and global observation of cloud properties is essential for the understanding of the climate system [*WMO*, 2010]. A fundamental parameter in estimates of atmospheric ice is Ice Water Path (IWP) [g/m^2], defined as the vertical integral of Ice Water Content (IWC) [g/m^3] or the atmospheric column density of ice. Climate model estimates of cloud IWP vary by more than an order of magnitude [*Waliser et al.*, 2009] and show deviating spatial distributions [*Eliasson et al.*, 2011]. One of the reasons for poor model performance is the lack of good constraints.

Only space-based remote sensing can provide global IWP measurements, but the remote sensing of atmospheric ice is severely underconstrained. Depending on the technique, the measured quantity (such as reflectance, radiance, or radar backscatter) depends in a complex manner on the vertical distribution of IWC, particle sizes, particle shapes, temperature, humidity, and other quantities. Individual measurements contain insufficient information to retrieve all those quantities directly and independently, and many assumptions need to be made [e.g., *Stephens and Kummerow*, 2007]. This leads to a considerable uncertainty in retrieved IWP. Moreover, the true uncertainty in IWP retrievals is very hard to quantify, not least because of the various difficulties of performing reliable in situ measurements.

Current IWP products are based on reflected solar radiation [e.g., *Rossow and Schiffer*, 1991; *King et al.*, 1997], passive microwave [e.g., *Ferraro et al.*, 2005; *Boukabara et al.*, 2011; *Gong and Wu*, 2013], or active sensors such as radar and lidar [e.g., *Austin et al.*, 2009; *Delanoë and Hogan*, 2010; *Deng et al.*, 2010]. Passive submillimeter measurements have been proposed for their ability to sample the size distribution and their sensitivity to relevant particle sizes, which allows for sensing the full ice column [e.g., *Buehler et al.*, 2012], but no down-looking instrument is currently in space. Active techniques have the capability to determine the vertical structure of atmospheric ice, and IWP retrieved from active sensors is likely more accurate than any IWP retrieval from passive measurements, although the error remains very difficult to characterize in

absolute terms. However, active sensors are costly and require a lot of energy, and the technology is less mature than for many passive technologies. They tend to have a small footprint and are (so far) exclusively carried on scientific platforms, which have a limited lifetime. On the other hand, passive sensors exist on both scientific and operational satellites. In particular, sensors on operational satellites provide a much better spatial and temporal coverage than active sensors. A nearly identical set of instruments is available on the National Oceanic and Atmospheric Administration (NOAA) 15–19 satellite series, as well as on MetOp A, MetOp B, and the future MetOp C satellite. Hence, any product based on those sensors can be readily processed from 1998 until the present (and beyond). Each individual satellite provides daily near-global coverage. Therefore, the potential data volume for such a product vastly outstrips the data volume from active sensors, allowing entirely different applications.

Eliasson et al. [2013] use strict collocations between radar, lidar, solar, and microwave sensors, in order to systematically compare spaceborne IWP retrievals on a footprint-level basis. Thus, they quantify where various techniques share IWP sensitivity with the active combined radar-lidar product raDAR/liDAR (DARDAR) [*Delanoë and Hogan*, 2010]. They find that IWP products based on solar bispectral methods share sensitivity with DARDAR from 40 g/m² upward, whereas the microwave-based Microwave Surface and Precipitation Products System (MSPPS) shares sensitivity with DARDAR only from 900 g/m² upward.

Existing products from passive operational sensors include Pathfinder Atmospheres Extended (PATMOS-X) [*Walther and Heidinger*, 2012] and MSPPS [*Ferraro et al.*, 2005]. PATMOS-X uses reflected solar radiation, which comes with the disadvantage of being available only in daylight. If a product performs equally well or better than PATMOS-X, but omits reflected solar radiation, then this has the distinct advantage of being available both daytime and nighttime, thus allowing climatological studies involving diurnal cycles.

A combination of different passive techniques may be capable of exploiting synergies, providing a more accurate IWP retrieval than by using any individual part of the spectrum. Synergies have been exploited in, e.g., retrievals of rainfall [e.g., *Kidd et al.*, 2003; *Marzano et al.*, 2004; *Rapp et al.*, 2009; *Kidd and Levizzani*, 2011] and liquid clouds [*Taylor and English*, 1995]. For retrievals of properties for atmospheric ice, synergies between radar and lidar have been used in at least two products [*Delanoë and Hogan*, 2010; *Deng et al.*, 2010], but synergies between different passive techniques appear underutilized.

This article presents SPARE-ICE, the first IWP product exploiting the synergy between terrestrial infrared and microwave. SPARE-ICE will be publicly available through the World Data Center for Remote Sensing of the Atmosphere (WDC-RSAT) under the Open Data Commons Attribution License.

The use of collocations between active, scientific instruments and passive, operational instruments to learn more about the latter has a limited history. *Holl et al.* [2010] introduce collocations between CloudSat Cloud Profiling Radar (CPR) and NOAA 18 Microwave Humidity Sounder (MHS) and briefly present three potential applications. *Liu and Seo* [2013] use radar reflectivities from CloudSat CPR as a proxy reference for the detection of snowfall over snow-covered surfaces. Other statistical retrievals do not use collocations but retrieval simulations [e.g., *Evans et al.*, 2012; *Jiménez et al.*, 2007] or output from general circulation models (GCMs) [e.g., *Surussavadee and Staelin*, 2008].

The retrieval presented in this study builds on a collocation toolkit developed by *Holl et al.* [2010]. Whereas *Eliasson et al.* [2013] used collocations between IWP retrievals to characterize different sensitivities, we use collocations between the same sets of satellites to develop a new retrieval algorithm. In the present study, we match reflectances from Advanced Very High Resolution Radiometer (AVHRR) solar reflected channels, radiances from AVHRR terrestrial infrared channels, and radiances from MHS humidity channels, with IWP as reported by the official CloudSat product 2C-ICE [*Deng et al.*, 2010] as distributed by the CloudSat Data Processing Center (CDPC). Then, we use an artificial neural network (ANN) to train an IWP retrieval from passive measurements (note that in our final product, SPARE-ICE, reflectances are not used). Thus, we reproduce actively retrieved IWP using only passive, operational sensors.

The primary aim of the study is to develop and analyze an improved IWP retrieval from passive operational sensors. A secondary aim of the study is to quantify synergies between different passive operational sensors.

The following sections describe the work in detail. In section 2, we describe the instruments used, the collocations, and the retrieval approach. Section 3 presents the results for the investigation into synergies and shows SPARE-ICE retrievals for three case studies and a 2007 gridded mean. In section 4, we discuss the results. Finally in section 5, we recommend tasks for further work and conclude the article.
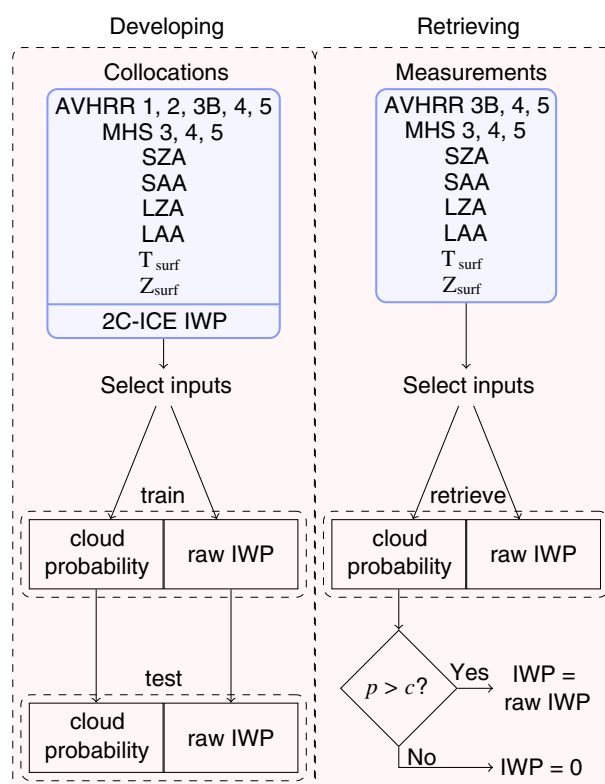
Figure 1. Diagram illustrating the approach for developing and retrieving SPARE-ICE and intermediate products. The horizontal line in the collocations box separates input measurements and auxiliary data from the target measurement. For a description of the input measurements and auxiliary data, including an explanation of the acronyms, refer to Table 1. All steps are described in detail in the text.

2. Methodology

In the following subsections, we will describe the SPARE-ICE retrieval algorithm in detail. The overall approach is summarized in Figure 1. Broadly, we consider three phases:

1. The development of the retrieval system. This consists of (i) obtaining collocations as described by *Holl et al.* [2010]; (ii) selecting various combinations of input channels and auxiliary information, considering solar, terrestrial infrared, and passive microwave, in order to investigate synergies; (iii) training of a pair of neural networks (one for cloud detection, one for IWP retrieval) for each of those selections; and (iv) finally, testing the performance of this pair of neural networks. This consists of estimating the error, as well as looking at the climatology.

These steps are repeated for many different combinations of input channels and auxiliary information.

2. Based on the results of part 1 of the study, we decide what information to include in the SPARE-ICE retrieval algorithm.
3. Processing measurements. This involves (i) collecting and combining measurements for the needed sensors and channels, (ii) selecting the input channels and auxiliary information, and (iii) applying the relevant set of neural networks to obtain IWP.

These steps are described in more detail in the following subsections.

2.1. Instruments and Products

The instruments used in the study can be grouped in two categories: active and passive. In the following, we will first describe the passive observations that were considered for the retrievals, then auxiliary information, and then the active instruments and the 2C-ICE product that was used for training. The measurements considered for the retrieval development are listed in Table 1, and auxiliary information considered is listed in Table 2.

Table 1. Measurements Considered in the Development of SPARE-ICE[a]

Sensor	Channel	Spectral Range	Technique	Use in SPARE-ICE
AVHRR	1	0.58 to 0.68 μm	Solar	No
AVHRR	2	0.725 to 1 μm	Solar	No
AVHRR	3b	3.55 to 3.93 μm	Mixed	Both
AVHRR	4	10.3 to 11.3 μm	Terrestrial IR	Both
AVHRR	5	11.5 to 12.5 μm	Terrestrial IR	Both
MHS	3	183 ± 1 GHz	Microwave	Raw IWP only
MHS	4	183 ± 3 GHz	Microwave	Raw IWP only
MHS	5	190 GHz	Microwave	Raw IWP only

[a]In the column "use in SPARE-ICE," "both" means that the information is used both in the network to retrieve IWP and in the network to detect a cloud, as described in the text. See also Figure 1 and Table 2.

2.1.1. Passive Systems

Although active systems can retrieve many ice cloud properties more accurately than passive ones, there is a need for continued development of IWP products based on passive sensors. In particular, a product based on passive operational sensors can be processed far back in time and operational sensors are more or less guaranteed to remain working. In our approach, we use a combination of solar, terrestrial infrared, and microwave sensors, allowing for data to be processed back to 1998.

2.1.1.1. Reflected Solar Radiation and Terrestrial Infrared Radiation

AVHRR [*Cracknell*, 1997] was first launched on board NOAA 6 in 1979. Although the precise configuration has varied somewhat throughout the years, there may be no other meteorological satellite instrument with as many copies and as long a lifetime as AVHRR. Recent editions (from NOAA 15 onward) carry six channels, of which five are simultaneously operated. Three channels (1, 2, and 3A) measure reflected solar radiation. Two channels (4 and 5) measure terrestrial infrared radiation (sometimes referred to as "thermal"). Channel 3B, covering 3.55 to 3.93 μm, measures a signal containing both reflected solar and emitted terrestrial radiation. Channel 3 can be switched between 3A and 3B. See Table 1 for the spectral characteristics of AVHRR data considered for SPARE-ICE. Due to downlink limitations, global data for the NOAA satellites are only available at a limited resolution, through the global area coverage (GAC) product. AVHRR GAC data only contain every third AVHRR scan and average four adjacent measurements [*Robel et al.*, 2009], resulting in effective discontiguous footprints of $1 \cdot 4$ km^2, as illustrated in Figure 2. Both reflected solar and emitted terrestrial radiation contain information about clouds and other aspects of the Earth system. AVHRR has been used for a very wide variety of applications, including (but not limited to) the retrieval of cloud properties, such as in the PATMOS-X data set, [*Stowe et al.*, 2002; *Jacobowitz et al.*, 2003; *Heidinger and Pavolonis*, 2009] available from the early 1980s to present.

The High-Resolution Infrared Radiation Sounder (HIRS) also measures terrestrial radiation and does so in more spectral bands than AVHRR. *Holl et al.* [2010] briefly explored retrieving IWP from a combination of HIRS and MHS. However, as shown in Figure 2, HIRS footprints are highly noncontiguous, and therefore, HIRS is less suitable for usage in combination with other sensors and is not used in SPARE-ICE.

2.1.1.2. Microwave

Advanced microwave sounding unit (AMSU) and MHS are microwave sounders that have been carried on board NOAA satellites since the launch of NOAA 15 in 1998. AMSU-B [*Saunders et al.*, 1995] and MHS [*Kleespies and Watts*, 2007] measure at frequencies from 89 to 190 GHz. AMSU-B is carried on NOAA 15, 16, and 17, while MHS is carried on NOAA 18, NOAA 19, MetOp A, and MetOp B and will be carried on MetOp C. In this study, we use MHS on NOAA 18, but the retrieval algorithm can also be applied to AMSU-B, so we describe both.

Two channels, one at 89 GHz and one at 150 GHz (for AMSU-B) or 157 GHz (for MHS), are surface-sensitive channels under all conditions, and their usage requires characterizing surface

Table 2. Auxiliary Information Considered for SPARE-ICE[a]

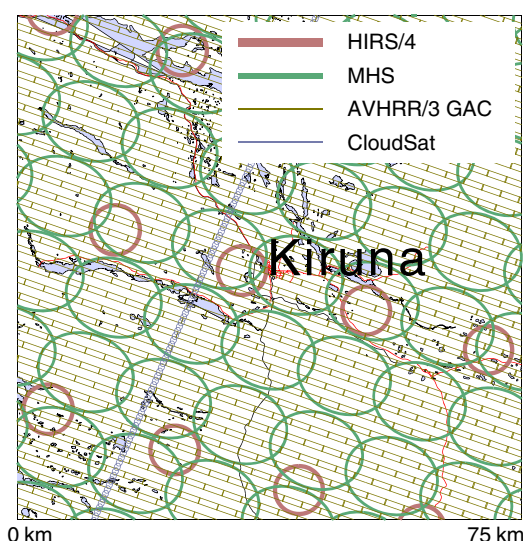| Information	Use in SPARE-ICE
Latitude	No
Solar zenith angle (SZA)	No
Solar azimuth angle (SAA)	No
Local zenith angle (LZA)	Both
Local azimuth angle (LAA)	Both
Surface temperature T_{surf}	No
Surface elevation z_{surf}	Both

[a]See also Table 1.

Figure 2. Illustration of various footprint sizes. Adapted from *Holl et al.* [2010]. Not shown is CALIPSO, which is smaller than CloudSat and falls mostly within the latter footprint.

emissivity. For this reason, we have chosen to focus on the remaining channels. The three remaining channels are near the 183 GHz water vapor absorption line, at 183 ± 1 GHz, 183±3 GHz, and finally 183±7 GHz (AMSU-B) or 190 GHz (MHS). If there is sufficient water vapor in the field of view, they are all sounding channels without any surface sensitivity. However, if the atmosphere is particularly dry (such as may happen at high latitudes), in particular the outermost channel becomes a surface-sensitive channel. At the frequencies around the 183 GHz water vapor absorption line, large ice particles can be detected from a spaceborne platform due to the scattering of radiation emitted by water vapor below the ice.

Both AMSU-B and MHS scan across-track with footprint diameters varying from a diameter of 16 km at nadir to approximately 52 km × 27 km at the edge of the scan [*Robel et al.*, 2009]. The path length through the atmosphere increases with an increasing nadir angle, which leads to Jacobians peaking higher up in the atmosphere, and therefore lower radiances (limb cooling). This needs to be taken into account by any retrieval utilizing these channels.

2.1.2. Auxiliary Information

For the SPARE-ICE development, we also consider auxiliary information, as summarized in Table 2. This includes instrument viewing angles, and solar angles when solar reflectances are considered. Additionally, we consider surface elevation from *Amante and Eakins* [2009] and surface temperatures from the National Centers for Environmental Prediction (NCEP) Community Forecast System Reanalysis (CFSR) [*Saha et al.*, 2010]. We also explore how adding latitude information affects the retrieval. See Table 2 for a complete overview of all auxiliary information considered.

2.1.3. Radar and Lidar

The CloudSat CPR is a 94 GHz radar launched in 2006 [*Stephens et al.*, 2002, 2008]. It measures vertical profiles of backscattering, in particular from clouds and precipitation. It observes only at (near) nadir with a footprint of 1.1 km. Retrieved profiles of IWC and liquid water content (LWC), and therefore retrieved IWP, have a considerable error, reported by *Austin et al.* [2009] to be around 40%, but that is likely an underestimate. The error may be large and is quite hard to quantify, due to poorly known ice particle sizes and shapes and due to a poor characterization of supercooled liquid and melting particles.

The Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP) is a lidar, carried on the Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation (CALIPSO), measuring backscattering at 532 nm and 1064 nm from aerosols and cloud tops [*Winker et al.*, 2003, 2009]. It measures only at nadir and has a footprint of 70 m.

CloudSat and CALIPSO fly close to each other in the A-Train satellite constellation.

A combination of radar and lidar, or specifically between CPR and CALIOP, provides a more complete measurement of the ice column than any other presently existing spaceborne measurement. Radar can penetrate thick systems of precipitating clouds but is mainly sensitive to large particles and does not detect small ones. The radar wavelength (3.2 mm in the case of CPR) is roughly an order of magnitude larger than many ice particles. This puts the scattering in the Rayleigh regime, where radar reflectivity is a function of the sixth power of particle size [*Austin et al.*, 2009]. Therefore, large particles dominate the returned signal. Lidar, on the other hand, operates at a wavelength smaller than most particles. Backscattering of laser energy is dominated by small particles (because they are most numerous), but it also gets attenuated more quickly. Therefore, radar and lidar are complementary.

Two products exploiting a radar-lidar synergy are publicly available at present. *Delanoë and Hogan* [2008] describe the radar-lidar product DARDAR, developed further in *Delanoë and Hogan* [2010]. A slightly newer product based on the same sensors is described by *Deng et al.* [2010] and published as an official CloudSat

product by the name 2C-ICE. Both DARDAR and 2C-ICE retrieve IWP on the CPR footprint. For the present study, we choose to use 2C-ICE, but the principle can also be carried out with DARDAR or any other reference data set, and there is no strong reason to choose one over the other. *Deng et al.* [2013] briefly review the differences between DARDAR and 2C-ICE.

Although DARDAR and 2C-ICE can, at least potentially, determine IWP better than any passive retrieval, their temporal and spatial coverage is inherently limited. Active systems require a lot more energy than passive ones, and both CPR and CALIOP observe only at an (almost) nadir geometry. The technology is less established, and the lifetime for an active system may be less than for a passive one. Although their lifetimes have already exceeded design and expectations considerably, both CPR and CALIOP will come to an end. EarthCare [*Bezy et al.*, 2005] is scheduled for launch in 2015 and will carry (among other instruments) a lidar and a cloud radar, and thus provide a continuation to CPR and CALIOP. However, active systems are not operational, and hence, their continued existence is not guaranteed.

Like any retrieval for ice clouds or frozen precipitation, the 2C-ICE retrieval is heavily dependent on forward model assumptions, both for the lidar and radar contributions. As 2C-ICE is used to train SPARE-ICE, the latter directly inherits all systematic errors and limitations in 2C-ICE. One limitation is that small particles are seen only near the top of the cloud. If small particles exist deep into the cloud (where lidar cannot reach), they go unnoticed by both radar and lidar [*Mace et al.*, 2009]. This may cause a dry bias, again hard to quantify due to the lack of a reference. *Deng et al.* [2010] report on the sensitivity of the 2C-ICE retrieval algorithm to multiple scattering (in the part of the column where lidar information is important), microphysical model assumptions, and simulated radar backscatter in the lidar-only region. They find that changing the assumed particle size distribution leads to a systematic 21% reduction in estimated IWC and that changing assumptions on particle shapes can affect retrieved IWC by 1 or 2 orders of magnitude. They also compare 2C-IWC retrieved from an airborne platform similar to the A-Train, to measurements by two independent in situ instruments, a comparison performed during the TC-4 measurement campaign [*Toon et al.*, 2010]. *Deng et al.* [2010] find that the ratio between 2C-IWC and in situ measurements is similar to the ratio between two independent in situ measurements, both being around a factor 2, and conclude that "the retrieval agrees well with in situ data." However, it should be kept in mind that this agreement is at best within a factor 2 (100% error). This puts a lower bound on the achievable accuracy in SPARE-ICE.

2.2. Collocations

In the context of this study, a collocation is an event where two spaceborne sensors observe the same place at almost the same time. *Holl et al.* [2010] have shown that due to the proximity of NOAA 18 to the A-Train (the constellation including CloudSat, CALIPSO, and other satellites), collocations between sensors on NOAA 18 and CloudSat occur globally. The collocation algorithm is described in *Holl et al.* [2010] and updated in *John et al.* [2012].

In this study, we use collocations to obtain a database for training artificial neural networks. Figure 1 illustrates the role of collocations in the SPARE-ICE retrieval development, and Table 1 gives an overview of all measurements used.

For obtaining the collocations, we use the flexible and sophisticated toolkit initially developed by *Holl et al.* [2010] but since developed considerably further. The collocation algorithm is fully based on time and geolocation information contained in the data files. Others, such as *Nagle and Holz* [2009], use orbital parameters to calculate when collocations occur, before looking at data files.

We collocate NOAA 18 MHS (the primary) with CloudSat CPR (the secondary). To this, we add AVHRR measurements and auxiliary information. From CPR, we use the 2C-ICE product. Although this uses both CPR and CALIOP, it is retrieved on the CPR footprint, so we do not need to explicitly collocate with CALIOP.

As shown in Figure 2, the footprints for CPR and AVHRR are much smaller than the ones for MHS. Therefore, we "collapse" CPR and AVHRR onto the MHS, for all secondary footprints within 7.5 km and 10 min of the primary, as defined by the center of the measurement. For each MHS footprint, we store the number of collocated CPR and AVHRR footprints, as well as the mean and standard deviation for each AVHRR channel measurement (reflection or radiance), as well as for IWP from both 2C-ICE and DARDAR. We also store the fraction of 2C-ICE or DARDAR footprints with IWP > 10 g/m^2. Even when the CloudSat ground track passes exactly through the center of an MHS footprint, CPR covers less than 10% of MHS, and an error from imprecise collocations is unavoidable. As long as this error is random, the effect on the SPARE-ICE retrieval will be

Figure 3. If collocations were limited to maximum 1 min and 1 km, they would only occur at very specific angles and latitudes, as shown in this figure obtained from *Holl et al.* [2010]. The numbers in this figure relate to collocations before collapsing (see text), for January 2007 only. The collocations used in this study have much larger thresholds for distance and time difference and therefore fill a much larger part of the space.

limited. However, if this introduces a systematic difference, then this will affect the final IWP product. This systematic error is partly unavoidable, because a larger footprint acts similar to a low-pass filter on measured IWP.

The collocations form a training set for a neural network (see next section). A training set for an artificial neural network must have "a sufficient number of properly distributed data records that adequately resolve the functional complexity of the target mapping" [*Krasnopolsky*, 2007]. In our case, that means all latitudes, angles, cloud types, etc. must be represented, although the distribution of the training data does not need to match the distribution of the real world [*Swingler*, 1996, section 5.2.2].

Although collocations occur globally, they occur only at particular ranges of MHS viewing angles, where the exact range is a function of latitude. This effect is strongest for collocations with very strict requirements. Figure 3 illustrates an extreme case, with statistics for collocations occurring with at most 1 min and 1 km in-between. In this example, collocations for a particular latitude occur only at two particular angles (one corresponding to the ascending, one to the descending node). With longer requirements (15 min, 15 km), a much larger part of the latitude/viewing-angle space is filled, but even then, collocations do not fill the entire space of latitude and MHS viewing angle, which would be desirable for our neural network training approach described below. Additionally, collocations are not equally distributed over the globe but occur more frequently near the poles than around the equator (like any measurements from polar orbiting satellites). Due to the nonuniform distribution, collocations also contain correlations that do not occur in noncollocated data. Of course, even in noncollocated data, quantities like latitude and surface temperature have correlations (of either sign) with IWP. To mitigate the nonuniform distribution of collocations as a function of latitude, we make a subselection where we use an increasingly small fraction of collocations poleward from the equator, so that the density of collocations (per square kilometer) is approximately constant.

In the present study, we have roughly $1.3 \cdot 10^6$ collapsed collocations with valid radiances occurring through 2007. After thinning to get a constant geographical distribution, $7.9 \cdot 10^5$ collocations are left. Of those, approximately $3.4 \cdot 10^5$ have an MHS-averaged 2C-ICE IWP larger than zero and can be used to train the IWP retrieval (which occurs in log space—see below). For details on the collocations, refer to *Holl et al.* [2010].

2.3. Retrieval Development
The aim of the SPARE-ICE retrieval is to make an IWP product that is globally available and at the full combined scan range of AVHRR and MHS. We consider the full atmospheric column, including cloud ice, graupel, and snowfall—we do not attempt any separation. In developing SPARE-ICE, we also develop intermediate products using only a subset of channels and auxiliary information, as indicated in Figure 1. The process of training and testing the retrieval algorithm described below is repeated for each of these subsets. In the following pages, we describe our retrieval approach.

2.3.1. Neural Networks
We use the collocations as a retrieval database for training ANNs, as illustrated in Figure 1. An ANN is a network of interconnected processing units called nodes. We use multilayer perceptrons (MLPs), feed-forward neural networks where the nodes are divided in multiple layers. MLPs are commonly used to statistically characterize the complex relation between a set of inputs and one or more quantities of interest, called targets. This makes MLPs popular for geophysical retrievals [e.g., *Krasnopolsky*, 2007]. We use an MLP with three layers. For the first layer, called the input layer, we use one node for each input quantity. See Table 1 for an overview of all input quantities used. The second layer is the so-called hidden layer, in which we use 10 nodes. The third and final layer is the output layer and contains a single node corresponding to our target, the logarithm of 2C-ICE IWP. All nodes in layer n are connected to all nodes in layers $n + 1$ via activation functions. The training algorithm iteratively assigns weights and biases to each activation function. This is

done by minimizing the cost function, which is defined as the mean square error difference between the neural network output and the targets. The trained neural network forms a nonlinear regression, mapping the set of inputs to one or more target quantities. Such networks are prone to "overfitting," where the network learns the statistical relation not only between the inputs and the target but also between input noise and the target. To prevent overfitting, part of the data is set apart as validation data. If, for N subsequent iterations, the fitting improves (i.e., the cost function becomes smaller) with the training data, but gets worse with the validation data, then the training is considered finished (this approach is called the "early stopping" criterion). There is no perfect choice for N. If N is too small, the cost function may prematurely reach a local minimum, and if N is too large, training might take unnecessarily long. In our implementation, we choose $N = 5$. Note that there are stochastic elements in the training (for example, the initial values of weights and biases) and that there is no guarantee that the fitted network is the best possible one. For the implementation, we use the Mathworks MATLAB® Neural Network toolbox V8.0.1 (R2013b).

For the training, we use all collocations where all inputs and the target have valid measurements, occurring during 2007 (see Table 1 for an overview of input measurements). We divide the collocations in three subsets: training, validation, and testing. The training data are used to minimize the cost function, and the validation data are used to prevent overfitting as described above. The testing data are not used in the training and serve as independent data to characterize the performance of the regression. From the collocation database, we draw 200,000 random samples according to a uniform distribution. From those samples, we randomly assign 2/3 to be used for training and 1/3 to be used for validation. The remaining collocations are used for testing the neural network.

There is no guarantee that a trained network is optimal, and the minimization of the error function may very well get stuck in a local minimum. Therefore, we perform the training 5 times every time we need to obtain a trained neural network and select the network that performs best for further use.

We split up the retrieval in two neural networks, as illustrated in Figure 1. The two networks operate independently, and the final IWP value is a result of combining the output from both nets.

1. One network classifies a scene as either cloudy or noncloudy. During the training, we define a scene as cloudy if 2C-ICE IWP averaged over the covered part of the MHS footprint is larger than a threshold value t. From the sensitivity study by *Eliasson et al.* [2013] and confirmed later in this study (see Figure 8), $t = 10$ g/m^2 is the lower sensitivity limit we can expect for SPARE-ICE and is therefore a good choice as a threshold. During the retrieval, the classification network retrieves a value $p \in (0, 1)$ that we interpret as a cloud probability. For each value $p > c$, where c is the cutoff value, we consider the scene as cloudy and proceed to retrieve IWP. For $p < c$, we consider a scene free of ice. For the classification network, we leave out microwave measurements. *Holl et al.* [2010] show that for MHS channel 5 (at 190 GHz), there is no significant radiance signal at IWP < 100 g/m^2. Any cloud thick enough to be detectable by microwave is also easily detectable by solar and terrestrial infrared measurements.

2. The other network retrieves "raw" IWP. Because of the high dynamic range of IWP, and because we independently assess whether or not a scene is cloudy, we retrieve IWP in log space (^{10} log IWP). Training and validation data are limited to cases with IWP > 0 but are otherwise a subset of the data used in the cloud classification network. The distribution of IWP is much closer to lognormal than to normal [*Eliasson et al.*, 2013], which is beneficial to the neural network fitting algorithm.

When performing the actual retrieval, both the classification network and the network retrieving "raw" IWP are applied for all measurements. Then, we determine *IWP* $= 0$ where the classification network retrieves $p < c$, and *IWP* $=$ "raw" IWP otherwise.

2.3.2. Error Analysis

Neural network retrievals do not provide a direct uncertainty estimate for an individual retrieval. Other methods do, but even for optimal estimation-based retrievals, the reliability of the uncertainty estimate is only as good as the forward model and the error characterization therein, and may very well be an underestimate of the true error, i.e., the error compared to the unknown truth. A reliable error estimate would require a ground truth, that is difficult to obtain for IWP. In our approach, we use the aforementioned testing data, which were not used in the training, to get an estimate of the error as a function of IWP. For the error estimate, we consider the two steps—classification and quantification—separately.

2.3.2.1. Cloud Classification Error

For the cloud classification network, we explore the rate of false positives and the rate of false negatives. A false positive is an occasion where the neural net obtains a cloud ($p > c$) while the independent reference states that there is not (according to the cloud definition described above). A false negative is an occasion where the neural net obtains no cloud ($p < c$) while the reference states there is. A higher cutoff will decrease the rate of false positives but increase the rate of false negatives. The final choice for c is somewhat subjective, depending on the balance desired for false positives and false negatives.

2.3.2.2. IWP Retrieval Error

We compare the IWP retrieved by the neural network against the reference IWP. For the IWP retrieval, we define the *fractional error*,

$$FE = \exp \left| \ln \frac{IWP_{retrieved}}{IWP_{reference}} \right| - 1 . \tag{1}$$

For example, if the reference IWP is 200 g/m^2, then a retrieved value of 50 g/m^2 would have a fractional error of 3 (300%), just like a retrieved value of 800 g/m^2 would. This should be kept in mind when comparing errors to other sources (such as *Deng et al.* [2013]) that are not always explicit in their definitions of the error but that may consider a 50 g/m^2 retrieval for a 200 g/m^2 reference to be a 75% error rather than 300% one. For small errors (less than approximately 10%), an error according to our fractional error definition is close to the more classical definition of $|b - a|/a$ (where a is the reference and b the retrieved value). However, errors in IWP and IWC are often much larger than that and may easily exceed a factor 2 (100%). With the classical definition, a retrieval algorithm that systematically retrieves a factor 3 too high would have a larger error (200%) than one that unconditionally retrieves 0 g/m^2 (error 100%), although the former is clearly more useful. With our fractional error definition, these retrievals would have errors of 200% and infinity, respectively. Moreover, IWC and IWP are often retrieved in logarithmic space, and a fractional error in linear space corresponds to a "classical" error in logarithmic space.

We divide the testing data (i.e., collocations not used in the training) in bins according to the reference IWP and calculate the median fractional error for each bin. The bins are logarithmically spaced between 10^{-1} and 10^4 g/m^2. Note that when shown as such, a low median fractional error for a certain IWP does not necessarily mean the retrieval is useful. In the hypothetical situation where a retrieval always results in a constant 10 g/m^2, then the median fractional error (or, indeed any other sensibly defined error) as a function of the reference IWP would be very small close to 10 g/m^2 (even a stationary clock states the correct time once or twice per day). Therefore, we also generate scatter plots for each retrieval, in order to identify the range where the retrieval has sensitivity to IWP.

2.3.3. Analyzing Multispectral Synergies

We systematically explore synergies between three techniques: solar reflected, terrestrial infrared, and terrestrial microwave radiation. For each of the techniques, as well as for any combination of two or three techniques, we train a retrieval (consisting of two networks, as described above). We also explore how adding other data affects the retrieval, as described in section 2.1.2 above. See Tables 1 and 2 for a complete overview of all data considered. Based on the synergy-exploration and other considerations, we make an initial decision of what information to use for the SPARE-ICE retrieval.

2.3.4. Characterizing SPARE-ICE

After choosing the input data based on the error analysis and other considerations, we study the resulting SPARE-ICE IWP in more detail. For a selection of atmospheric scenes, we compare SPARE-ICE IWP against IWP from two other products. This comparison is the first step away from the collocation world. Therefore, this comparison can indicate whether the retrieval can be extrapolated to cases where no collocations exist, such as toward the edge of the scan, or in geographic regions where retrievals may be difficult (such as mountain areas) that are too small to show up when doing global statistics. Then, we compare global maps of the 2007 climatological mean for a selection of products.

First, we compare scenes against the MHS-based MSPPS product [*Ferraro et al.*, 2005]. MSPPS measurements are based on radiation scattered by large (precipitation-sized) ice particles at 89 GHz and 150 GHz. The retrieval requires an accurate characterization of surface temperature and emissivity (SPARE-ICE uses channels centered around 183 GHz, where this problem is much less present). MSPPS retrievals are at the MHS footprint size, with a diameter of approximately 16 km at nadir.

Figure 4. Performance for a global retrieval based only on solar reflected measurements from AVHRR channels 1 and 2 (see Table 1), with or without additional angle or latitude information used in the retrieval. The fractional error is defined by equation (1).

Second, we compare scenes against the Moderate Resolution Imaging Spectroradiometer (MODIS) Aqua MYD08_M3 monthly gridded IWP product [*King et al.*, 1997]. The retrieval is based on four channels measuring reflected solar radiation, including one visible channel (0.645 μm) and three near-infrared channels (0.858 μm, 1.24 μm, and 2.13 μm). The MODIS MYD08$_{M3}$ product is derived from measurements with a footprint size of approximately 1 km.

Then, we compare maps of a 2007 gridded mean IWP for a selection of products. First, we compare with MODIS, described above. Second, we compare with 2C-ICE. The comparison with 2C-ICE is not an independent comparison, because 2C-ICE was used to train SPARE-ICE. However, this comparison can still 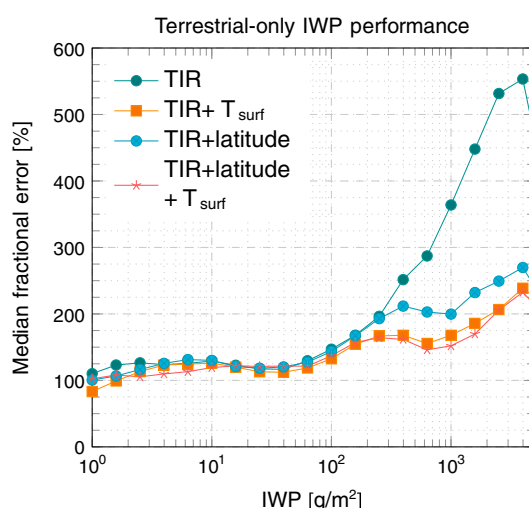give a useful insight into SPARE-ICE's ability to obtain a climatology similar to 2C-ICE. Finally, we compare SPARE-ICE with PATMOS-X. PATMOS-X is an operational retrieval based on AVHRR measurements, providing daytime estimates of IWP and other quantities [*Heidinger and Pavolonis*, 2009].

The comparisons are meant to indicate differences between SPARE-ICE and other products but are not meant as a detailed intercomparison. A full intercomparison of different IWP products (without SPARE-ICE) is presented by *Eliasson et al.* [2013], and a comparison focussing on spatial distributions of climatological means is shown by *Eliasson et al.* [2011]. Note that this comparison is not a validation of SPARE-ICE.

3. Results

The results are presented in two parts. First, we explore synergies and choose the inputs to SPARE-ICE. Then, we analyze the resulting product for individual scenes and gridded means.

3.1. Exploring Synergies and Choosing the Inputs to SPARE-ICE

We will first show the performance for single-technique retrievals and then for retrievals using a combination of two or three techniques. All results are for global retrievals for all angles for which we have collocations. For these results, we look only at the network retrieving IWP, i.e., the part of the retrieval assuming the presence of (cloud) ice. We will explore the performance of the cloud filter separately.

3.1.1. Single-Technique Retrievals

Figure 4 shows the performance of a retrieval using only measurements from AVHRR channels 1 (0.58 to 0.68 μm) and 2 (0.725 to 1 μm), using reflected solar radiation. As indicated before, for all retrievals, we observe a local minimum in the fractional error somewhere between 1 and 10 g/m^2. This minimum does not indicate that the retrieval is truly performing well for these IWP values. Rather, this local minimum indicates that the neural network retrieves these values when it lacks information. If the neural network lacks information, it tends to the mean state of similar measurements for which it lacks information. In this case, that is the geometric mean (because retrievals are in log space) for clouds too thin to be detected, hence the values between 1 and 10 g/m^2. Therefore, one should be cautious in interpreting errors for small values of IWP based on figures like this one and also look at scatter plots such as shown later in Figure 8.

For a global retrieval using only AVHRR channels 1 and 2 and no additional information, the performance for all but the very thickest clouds is poor. Clouds are essentially nonabsorbing at the short wavelengths of AVHRR channels 1 and 2. Lacking the absorption at higher IR frequencies, IWP must be determined purely from reflection. For moderately thick clouds, the reflectivity is affected by both the cloud visible optical depth τ and effective radius r_e. However, for thick clouds, reflectivity is mostly affected by τ and not much by r_e. Therefore, retrieving IWP is easier for thicker clouds [*McFarquhar and Heymsfield*, 1998], explaining the decreasing IWP error for larger values of IWP. Adding latitude information improves the solar-only retrieval performance a little bit, and adding angular information improves this more, although errors remain high

Figure 5. Performance for a global retrieval based only on terrestrial emission measurements from AVHRR channels 4 and 5 (see Table 1), with or without additional surface temperature or latitude information used in the retrieval. TIR is short for terrestrial infrared radiation.

at over 600% compared to 2C-ICE. The improvement from angular information is not surprising, because solar angles are essential to interpret information from solar reflectances. The small improvement from adding latitude information is likely due to correlations in the collocated data set. The finding that latitude information does not achieve much if angular information is already present suggests that latitude information is useful only due to its correlation with viewing angles (see also Figure 3). Despite the improvements by adding auxiliary information, variations in surface brightness, as well as other factors, mean that solar reflected information from AVHRR channels 1 and 2 alone is not enough to obtain a quantitative estimate of IWP.

Figure 5 shows the performance of a retrieval using only the two AVHRR terrestrial infrared channels. Without surface temperature information, a global retrieval based on only the two terrestrial radiation AVHRR channels performs very poorly (errors reach almost 600%). Adding latitude information improves results, but adding surface temperature information (from CFSR) improves them more. Adding latitude and surface temperature information does not help beyond adding only surface temperature information. The addition of latitude information might help mainly due to its correlation with surface temperature. In the atmospheric window region, terrestrial infrared measurements essentially give information on the target temperature. A large difference between target brightness temperature and surface temperature indicates the presence of a cloud. The error remains around 200%, because this temperature difference alone is still not sufficient to quantify the amount of atmospheric ice.

Figure 6 shows the performance of a retrieval using only the three MHS water vapor channels around 183 GHz. This retrieval is similar to the one presented in *Holl et al.* [2010], although *Holl et al.* [2010] did not attempt to do global retrievals. Without additional information, a global microwave-only retrieval performs even poorer than a global solar-only or a global terrestrial infrared-only retrieval. We see a significant improvement by adding latitude or angular information and even more improvement when we add both, with the median fractional error going down to 100% for very thick clouds. A retrieval at these frequencies is based on scattered radiation, so ice is detected through a brightness temperature depression for a down-looking sensor. At high latitudes, where the atmosphere is very dry compared to the tropics, these water vapor sounding channels become surface channels, in particular when the surface elevation is high (such as on Antarctica). The low surface emissivity (and low surface temperature) near 183 GHz then results in a low brightness temperature that, without additional information, cannot be distinguished from a brightness temperature depression due to a thick ice column, at least not from a single channel. At off-nadir scan angles, the sounding altitude is increased due


Figure 6. Performance for a global retrieval based only on MHS channels 3–5 near the 183 GHz water vapor absorption line (see Table 1), with or without additional latitude and viewing geometry information used in the retrieval. The viewing geometry information ("angles" in the legend) includes the zenith and azimuth angles for the satellite as seen from the surface. Note that the *y* axis has different scales for the upper and lower parts of the axes.

Figure 7. Comparison between the performance of the IWP retrieval networks using different combinations of input data. In this figure, solar refers to AVHRR channels 1 and 2, TIR refers to AVHRR channels 4 and 5, and MW refers to MHS channels 3–5 (see Table 1 for channel details). Here, all retrievals except SPARE-ICE include latitude information, retrievals using solar or MW include information on solar and satellite angles, and retrievals including IR use surface temperature information. For details on information included in SPARE-ICE, refer to Table 1. Note that the y axis has different scales for the upper and lower part of the axes.

to the longer path length through the atmosphere, resulting in a lower radiance. This too is observed as a low brightness temperature, indistinguishable from scattering due to atmospheric ice. Therefore, and as Figure 6 shows, information on latitude and viewing geometry is valuable for any global microwave-based retrieval. Even more valuable should be a combination of surface temperature and emissivity, but emissivities were not explored in the present study.

3.1.2. Multitechnique Retrievals

Figure 7 compares the performance of single-technique and multiple-technique retrievals, again according to the median fractional error as defined in equation (1). For this comparison, for each single-technique retrieval, we chose the best-performing (lowest error) combination of measurements and auxiliary information. The multitechnique retrievals are then simply combinations of the different single-technique ones. For IWP values between approximately 10 and 1000 g/m², out of the single-technique retrievals, solar performs worst, followed by microwave, and then terrestrial infrared. The combination of solar and terrestrial infrared performs better than either of those alone for any IWP between 10 and 4000 g/m². For the combination of solar and microwave, we

also see a strong synergy between 30 and 2000 g/m². Out of the dual-technique retrievals, terrestrial infrared and microwave performs the best and considerably better than either technique alone throughout the entire range of IWP values. The triple-technique retrieval, combining all three techniques, performs roughly as well as the infrared-microwave retrieval. However, a day-night retrieval using only terrestrial infrared and microwave performs almost as well as a daytime-only retrieval that uses solar.

SPARE-ICE uses terrestrial infrared and microwave but is slightly different from the TIR+MW retrieval shown in the figure. SPARE-ICE includes AVHRR channel 3B (3.55 to 3.93 μm) that is neither exclusively solar nor exclusively terrestrial infrared and that is not included in any other retrievals shown in Figures 4–7. Table 1 summarizes the data used in SPARE-ICE. Figure 7 shows that for most values of IWP, the median fractional error in SPARE-ICE is slightly below 100%. It is below 110% throughout the entire range of IWP and briefly drops off to 75% for IWP between 1000 and 2000 g/m².

Figure 8 shows a scatter density plot for SPARE-ICE versus 2C-ICE. Note that just like for the previous figures, this comparison does not yet take into account the cloud filter but retrieves nonzero IWP everywhere. As expected, the correlation between 2C-ICE and SPARE-ICE is very good for IWP > 10 g/m². Our retrieval was trained against 2C-ICE IWP. Although the testing data were not used in the training, they come from the same data set and therefore exhibit the same statistics as the training and validation data. Therefore, a good correlation between the retrieved and reference IWP is expected. For smaller values of IWP, we can see that SPARE-ICE IWP tends to be larger than 2C-ICE IWP and levels off at a median of around 1–2 g/m² for 2C-ICE IWP < 1 g/m². SPARE-ICE is not sensitive to such small values of IWP. This justifies the choice of 10 g/m² as a cutoff in the cloud filter. This also illustrates the phenomenon visible in Figures 4–7, where the median fractional error drops off as IWP approaches the sensitivity limit, which does not indicate a good retrieval but rather a lack of information, as explained before. Figure 8 also shows that the random error remains

Figure 8. Comparison between 2C-ICE and SPARE-ICE IWP. All testing data were binned first according to the reference 2C-ICE IWP and then according to the retrieved SPARE-ICE IWP. The color in each bin indicates the absolute number of retrievals with a particular combination of 2C-ICE / SPARE-ICE IWP. The black line represents the diagonal 1:1-line. The grey lines show the median for one product when the other product is kept fixed; for example, when 2C-ICE is 1 g/m², the median SPARE-ICE is 2 g/m². Any vertical or horizontal intersection is a histogram; for example, when 2C-ICE is 10 g/m², SPARE-ICE varies roughly between 1 and 100 g/m². Note that this figure considers "raw" IWP, meaning IWP before any cloud filter is applied.

quite large; this is quantified by the aforementioned fractional error, shown in Figure 7, which has a median around 100% for SPARE-ICE, but with outliers up to 2 orders of magnitude away from the reference value.

3.2. Cloud Filter Performance

Because IWP is retrieved in log space, the retrieval is always nonzero. This is not physically correct. Therefore, as described earlier, a separate network is used for the cloud filter.

Figure 9 shows the performance of the cloud filter. The cloud filter uses the same inputs as SPARE-ICE, minus the microwave channels—microwave adds no information when the purpose is to detect clouds. A scene is considered cloudy if the cloud probability p is larger than the defined cutoff. The sum of false positives and false negatives has a minimum for a cutoff of 0.40, and the rate of false positives and false negatives is equal at a cutoff of approximately 0.31. The figure also shows reference and retrieved IWP values for the false positives and false negatives. For a cutoff of 0.5, if the cloud filter falsely detects a cloud where

there really is not, SPARE-ICE retrieves a median IWP of around 25 g/m², while the reference median IWP is close to 0 (recall that the definition of "cloudy" means the reference IWP > 10 g/m²). Similarly, when the cloud filter misses a cloud, SPARE-ICE retrieves a "raw" value of less than 10 g/m², while the true IWP has a median of around 25 g/m². Either way, the values of IWP for clouds where the cloud filter is wrong are quite small. SPARE-ICE uses a cutoff of 0.5.

3.3. Characterizing SPARE-ICE

We have processed measurements for NOAA 18 for 2007, 2008, and 2009, resulting in 3 years of retrieved SPARE-ICE IWP. We show a selection of three case studies, all from NOAA 18.


Figure 9. Performance of the SPARE-ICE cloud filter depending on the selected cutoff value. A false positive occurs when the cloud filter considers a scene to be cloudy (IWP > 10 g/m²) when it is not, and a false negative occurs when a scene is cloudy, but the filter concludes it is not. The solid lines show the rate of false positives, false negatives, and the total error rate, according to the left axis line. The dashed lines show the retrieved and reference IWP values for false positives (f.p.) and false negatives (f.n.), respectively, according to the right axis line.

We show composites of three AVHRR channels and three MHS channels. For AVHRR, we follow *Dybbroe et al.* [2005] and combine channel 1, channel 2, and the inverse of channel 4 (so that clouds are white in all channels) and put those in the red, green, and blue planes of a composite image. Similarly, we make a composite image of MHS channels 3, 4, and 5 (where clouds are black in all panels). Although AVHRR channels 1 and 2 have not been used in SPARE-ICE, the 1-2-4' combination allows for a convenient visual inspection of the scene.

Figure 10 shows a fragment of a swath for a midlatitude winter scenario in the North Atlantic Ocean, off the coast of Newfoundland, Canada. The retrievals from MODIS and SPARE-ICE show similar spatial structures, although there are some differences. The exact values of retrieved IWP differ considerably between MODIS and SPARE-ICE. For example,
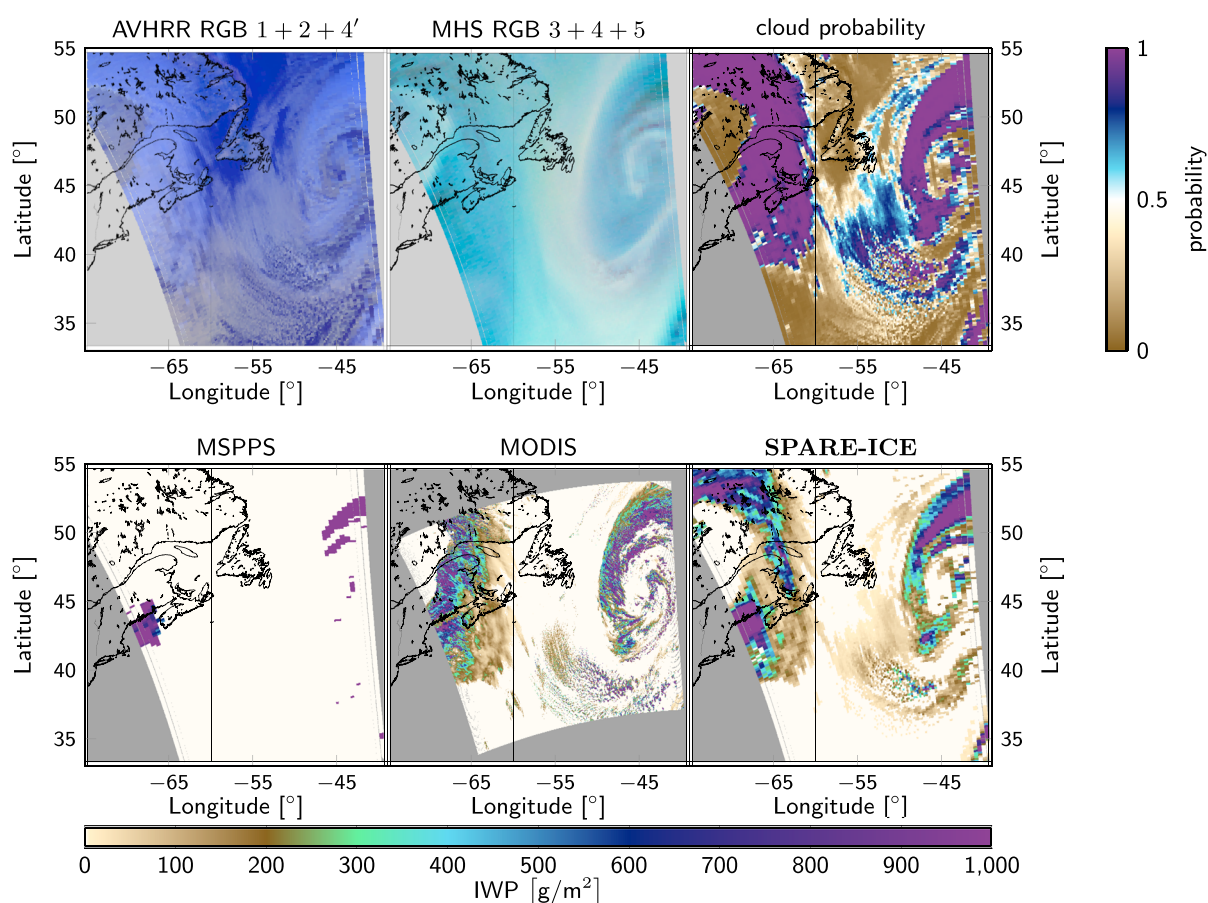
Figure 10. (top, right) Snapshot of a NOAA 18 swath at 1 January 2007, 16:40 UTC, off the coast of Canada. (top) Color composites for AVHRR and MHS, respectively, as described in the text. The cloud probability, retrieved from the AVHRR channels, is shown in Figure 10. (bottom) Retrieved IWP according to three different products: MSPPS, MODIS (Aqua) and, finally, SPARE-ICE. Inputs to the product not shown in the figure are satellite and solar angles, surface elevation, and the individual AVHRR and MHS channels.

for a particular region (48–52°W, 38–42°N), the all-sky mean for MODIS IWP is 228 g/m², whereas the all-sky mean SPARE-ICE IWP is 116 g/m². This may be partly due to the different measurement resolutions, where a SPARE-ICE footprint is more than 100 times larger than a MODIS footprint. Lacking an independent reference, there is no way of determining which one is more accurate. The retrieval from MSPPS is clearly lacking in comparison to either MODIS or SPARE-ICE, detecting only a small fragment of the system (but note that MSPPS retrieves using 89 and 150 GHz and is known to retrieve mostly precipitating ice).

All features visible in SPARE-ICE are also visible to some degree in the AVHRR and MHS composites, or both, and in the individual channels (not shown). The cloud probability network mostly retrieves either probabilities close to 1 or close to 0. In the region where it is close to the cutoff, the corresponding SPARE-ICE retrieval is small, so the cutoff has no major impact on the overall IWP (see also Figure 9).

Another scene is shown in Figure 11. Here, the observation is of Tropical Cyclone Indlala in the south-west Indian Ocean, near Madagascar, on 14 March 2007. For this system, the microwave signature is significant enough for MSPPS to also obtain most of the spatial structure. This structure is also clearly visible in all input channels, although not nearly as cold in the 3.55 to 3.93 μm channel as in the longer wavelength AVHRR channels. This channel contains a mixed signal of solar reflected and terrestrial emitted, and for a very deep ice cloud, the solar reflected signal increases, while the terrestrial emitted signal decreases. Also another system on this map, at the coast of Africa, is seen in all input channels and recognized by MODIS, MSPPS, and SPARE-ICE.

For the final example, in Figure 12, we look at a system above Antarctica. Antarctica is a region where retrieving atmospheric ice is difficult with any passive technology. Using microwave humidity channels, it is difficult because the high surface elevation and the low temperature both result in a very low specific
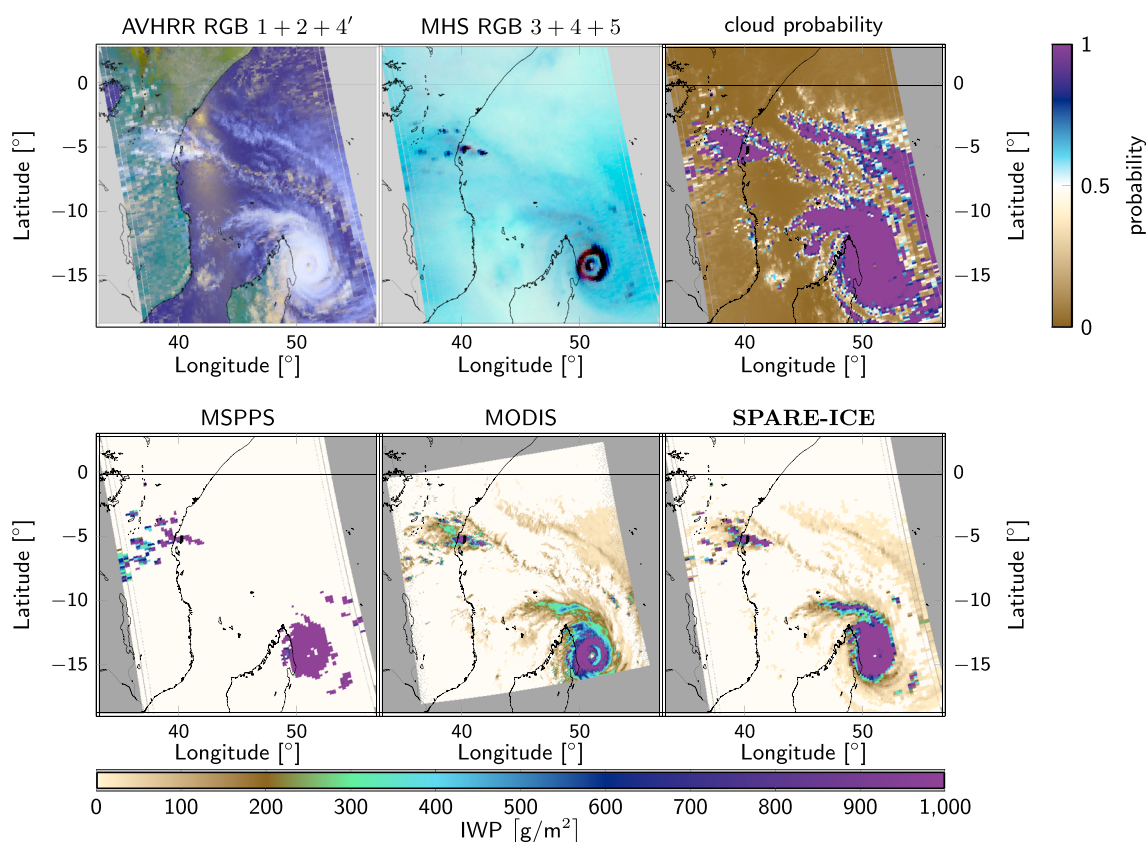
Figure 11. Observation of 14 March 2007, 10:40 UTC, of Tropical Cyclone Indlala, making landfall at Madagascar in the south-west Indian Ocean. See Figure 10 for an explanation of the individual panels.

humidity throughout the atmospheric column. Combined with a low surface emissivity, this leads to a very low radiance. We lack the humidity background against which ice scattering is visible around 183 GHz. The color gradient of the MHS composite shows how the lack of humidity first affects the outermost channel (channel 5) first and the innermost channel (channel 3) last. The MHS composite has regions that are black (all channels cold) and regions that are red (i.e., channel 3 still warm, but channels 4 and 5 cold). Brightness temperatures for channel 5 drop to 133 K. Solar channels are also hard to use, because the reflectance by ice is similar to the reflectance by ice clouds. Finally, the ice cloud signature from terrestrial infrared usually derives from the temperature difference compared to the surface; when the surface is very cold, this difference is no longer a clear indicator of clouds. Even so, SPARE-ICE seems to still perform quite well in this difficult region—at least the retrieved IWP is not clearly unphysical. From AVHRR channels 4 and 5, there is a region near the South Pole with a radiance of around 240 to 250 K and a region near the coast with a similar radiance. At the coast, this region exists against a warmer background, and indeed, this is the area where a high cloud probability is obtained (larger than 0.7) and where SPARE-ICE retrieves nonzero IWP. MODIS and SPARE-ICE agree that there is nonzero IWP in this region, although SPARE-ICE sees a larger area with atmospheric ice than MODIS does. Even closer to the South Pole, SPARE-ICE observes clouds with low IWP. In the same region, MODIS retrieves some very high IWP values, up to $5.5 \cdot 10^3$ g/m^2 as far south as 88°. This is not realistic, and likely due to measurement noise, since the signal-to-noise ratio is low at such low temperatures. For comparison, the highest retrieved IWP by SPARE-ICE south of 85° in this swath is 261 g/m^2. Although we do not know whether or not this is correct, at least it is not obviously wrong. Above Antarctica, MSPPS sees virtually nothing (if these very high MODIS IWP retrievals were true, MSPPS should also detect IWP).

Figure 13 shows a gridded map for SPARE-ICE for 2007 on a projection according to *Mollweide* [1805] (due to the equal-area, pseudocylindrical nature of this projection, it is more suitable for geophysical visualizations than equirectangular projections). For this map, all daytime measurements between 1 January and 31 December were sorted into 1° × 1° bins, and the mean IWP was calculated for each bin. Although
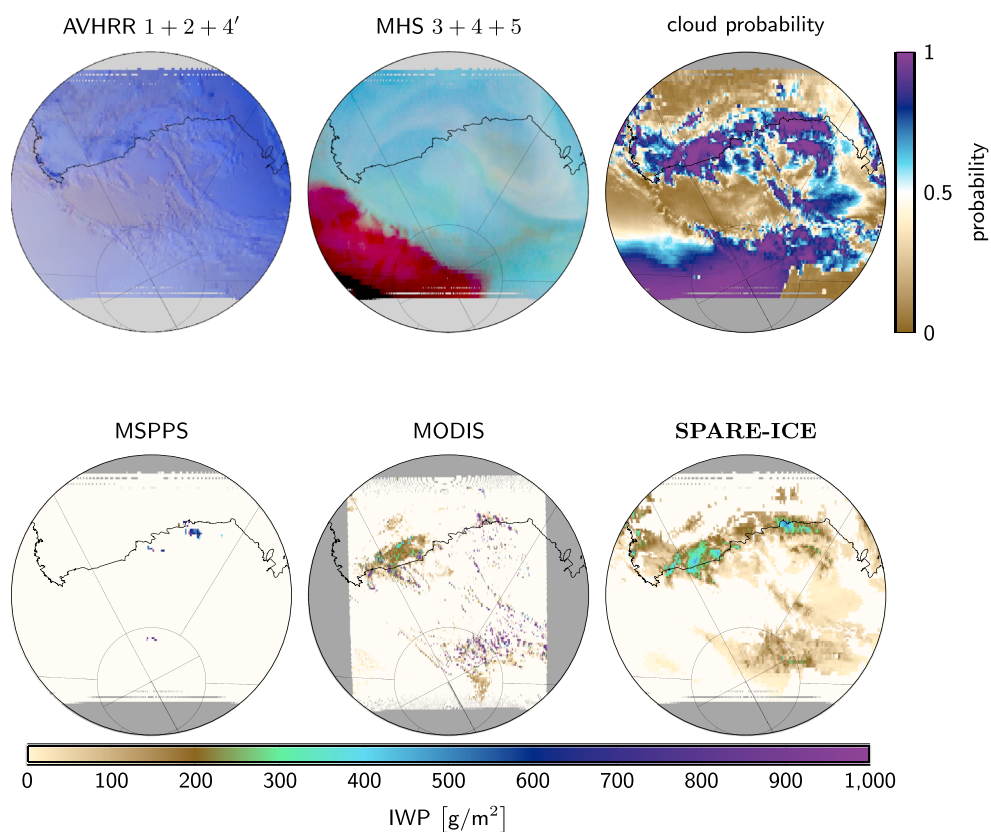
Figure 12. Example swath passing over Antarctica, 21 January 2007, 5:45 UTC. The South Pole is visible as a cross near the bottom of the maps. Quantities shown are as in Figure 10.

SPARE-ICE can be retrieved day and night, only day measurements were considered for a consistent comparison with MODIS and PATMOS-X. The IWP distribution appears in line with what one would expect, and there are no regions where SPARE-ICE is clearly unphysical. However, we can not verify that our IWP measurement is correct, either.
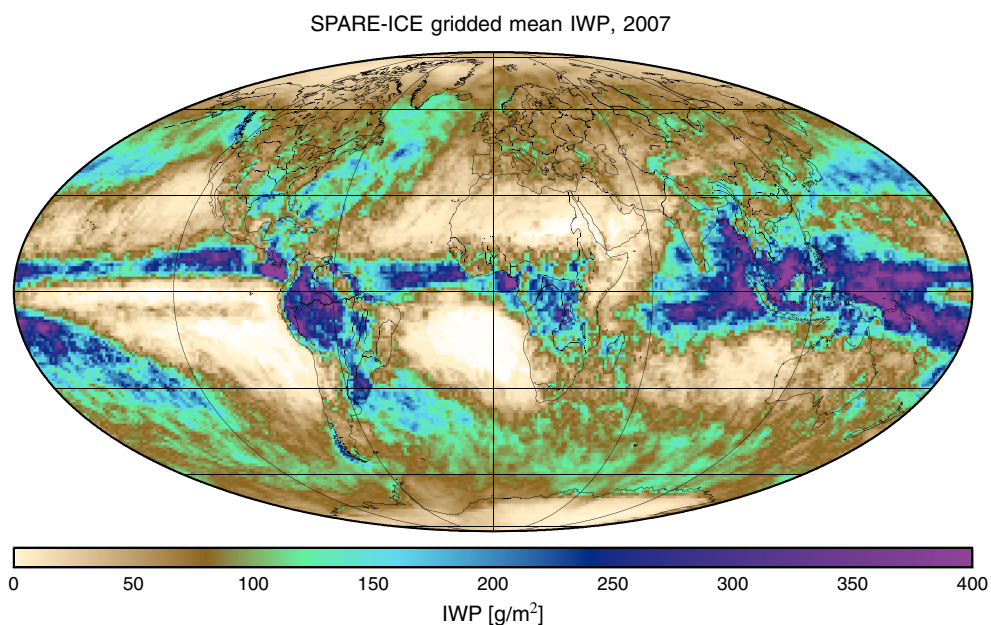

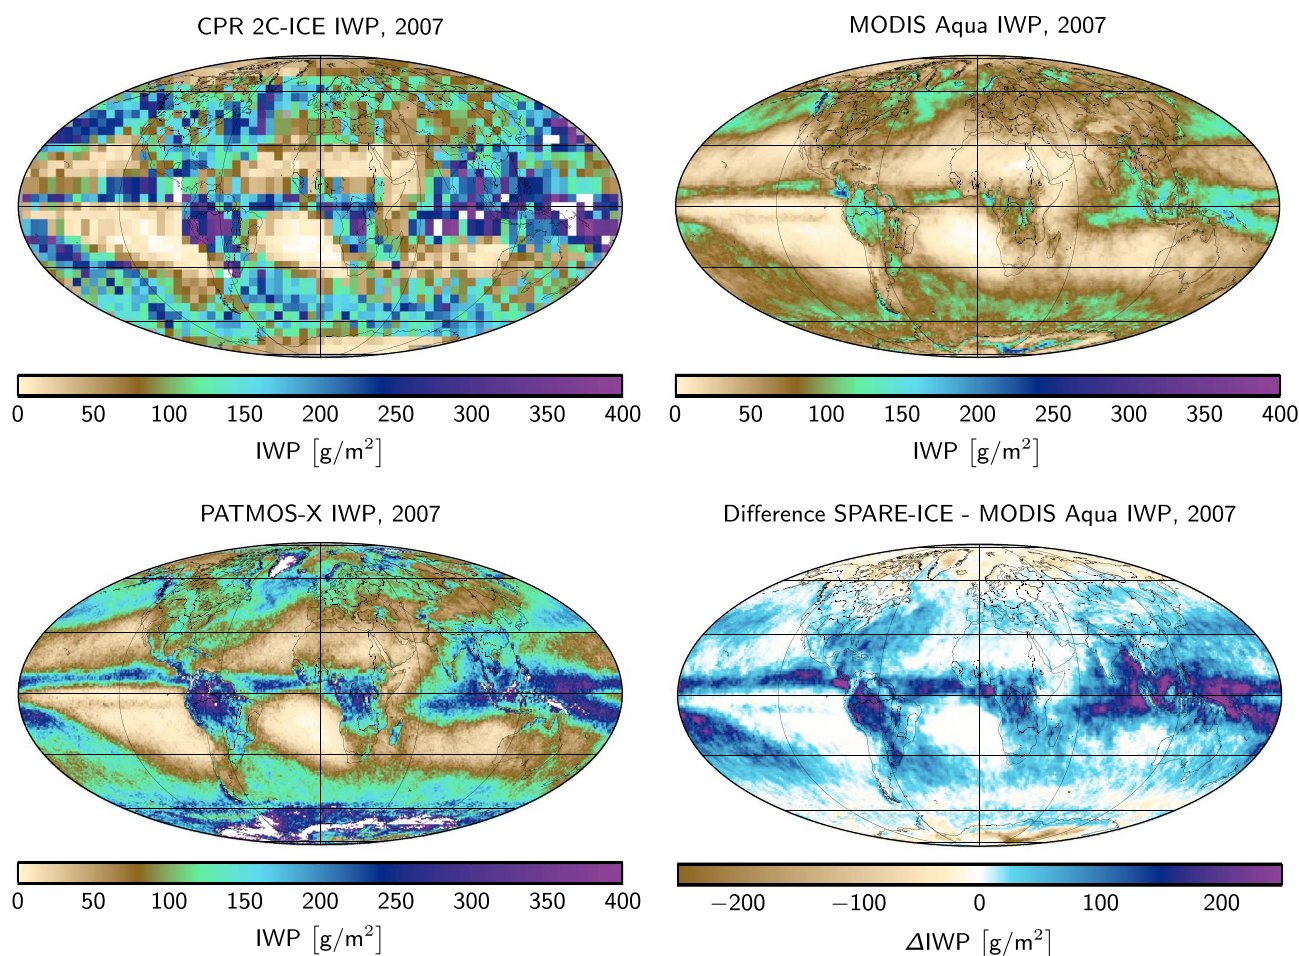**Figure 13.** SPARE-ICE IWP gridded mean for 2007, using NOAA 18.

Figure 14. The 2007 gridded mean IWP for MODIS Aqua, CPR 2C-ICE, PATMOS-X, and the difference between MODIS Aqua and SPARE-ICE. 2C-ICE is gridded on a 5° grid, while the other products are all gridded on a 1° grid.

We compare the 2007 SPARE-ICE mean IWP with three other products. Figure 14 shows level-3 IWP for MODIS Aqua (product MYD08), 2C-ICE, PATMOS-X, and the difference between SPARE-ICE and MODIS. For all maps, only daytime measurements are considered. In the discussion here, we will focus on SPARE-ICE in relation to other products. A thorough comparison of different level-3 IWP products (predating SPARE-ICE) is given by *Eliasson et al.* [2011], while a systematic comparison on a footprint-level basis is presented by *Eliasson et al.* [2013] and for MODIS in particular, by S. Eliasson et al. (The uncertainty in MODIS ice water path retrievals depending on cloud scenario, submitted to *Journal of Geophysical Research*, 2013).

There are considerable differences between yearly mean SPARE-ICE and IWP from the other products. Over most of the planet, SPARE-ICE retrieves higher values of IWP than MODIS does. The difference is largest in tropical convective regions. MODIS IWP is obtained from a retrieval of optical depth and effective radius, which, by the nature of MODIS, are obtained for the upper layers of the cloud [*King et al.*, 1997]. The effective radius in the upper layers of the cloud is typically smaller than the effective radius for the entire column, because particles near the top are typically smaller than particles further down. Therefore, we consider that MODIS likely underestimates IWP for these cases (see also *Eliasson et al.* [2013]). On the other hand, SPARE-ICE includes microwave information, which helps to quantify cases with very thick clouds.

In polar areas, as well as in some parts of the midlatitudes, MODIS IWP is higher than SPARE-ICE IWP, something that can also be seen in the Antarctica observation in Figure 12. MODIS has difficulties over bright surfaces, such as those covered by snow or ice, and likely retrieves spuriously high IWP due to measurement noise (for example, it retrieves several thousand g/m² above the South Pole, despite a high surface and very low temperatures). To the contrary, SPARE-ICE considers the polar areas to have very low IWP.

PATMOS-X has higher values than SPARE-ICE in most areas, but not everywhere. PATMOS-X has serious problems above cold, bright surfaces, such as over the poles—retrieved IWP is extremely high, a problem that has been reported before [e.g., *Eliasson et al.*, 2011]. The explanation given for MODIS above applies. Throughout the midlatitudes, PATMOS-X also retrieves higher IWP than SPARE-ICE but not unrealistically so. In the tropics, PATMOS-X has higher IWP above land, whereas SPARE-ICE has higher IWP above the oceans. It is not easy to point to the exact cause of these differences, nor to determine which product is closer to the truth. An oversimplified model of mixed phase clouds could mean that 2C-ICE, and therefore SPARE-ICE, overestimates IWP over ocean and underestimates it over land, but this remains speculation.

The gridded map for CPR 2C-ICE exhibits significant noise due to the poor sampling resolution, even if the grid resolution is 25 times larger (5×5 instead of 1×1 degree). As illustrated in Figure 2, measurements are only obtained at (near) nadir. Even for a full year of data, the number of measurements per grid box is quite small. Products based on radar and/or lidar, therefore, are of limited use for studying geographic patterns in a particular year.

As mentioned before, lacking an independent measurement, we can not state with certainty which one is more correct (or less wrong). This would at least require in situ measurements of IWP. Such measurements would be a major and novel undertaking in itself and would need to have a sufficiently small error to be useful for validation.

4. Discussion

Compared to 2C-ICE, SPARE-ICE has a median fractional error of around 100% or a factor 2. This can be compared to the scatter found by *Deng et al.* [2013], who compared IWC retrievals between 2C-ICE and in situ measurements, among other things. *Deng et al.* [2013] found that "[IWC is] strongly correlated with the in situ data (…) although the scatter is around a factor of 2." Although *Deng et al.* [2013] do not quantify the random error more rigorously and use IWC rather than IWP, it does indicate that the error between SPARE-ICE and active retrievals may be of similar magnitude as the error between active retrievals and the truth. If both are random, uncorrelated, and 100%, that would put the true error for SPARE-ICE at $\sqrt{1+1} \cdot 100\% = 141\%$ according to the law of error propagation.

From the case studies, it appears that SPARE-ICE does well in a wide variety of cases. Drawing information from both terrestrial infrared and microwave not only leads to synergies as seen directly by the median fractional error (Figure 7) but also appears to extend the IWP retrieval to climate regimes that normally pose large difficulties, such as over Antarctica (Figure 12), although the latter is currently impossible to verify. It is difficult to establish whether the system that SPARE-ICE concludes to be atmospheric ice, really is such and not humidity or something else. However, unlike MODIS, SPARE-ICE appears not to obtain any IWP values that are clearly unphysical.

SPARE-ICE can be used for reasonably long time series, because the combination of sensors that it uses has been around since 1999. At the time of writing, AVHRR and MHS are carried on NOAA 15 through NOAA 19, MetOp A, and MetOp B, and similar instruments are carried on other polar orbiting satellites. Therefore, SPARE-ICE has many potential applications for climatological studies.

A number of variations and improvements to SPARE-ICE can be made. In the training, SPARE-ICE does not take into account the error in the reference data set 2C-ICE. To take an error into account in neural network training is not completely trivial. One way that could be investigated is to adapt calculation of the cost function and use a weighted error sum rather than a simple root mean square sum, while giving higher weight to observations with a small error in 2C-ICE. This might improve the retrieval somewhat and could be investigated in future work. In the present study, we did not test whether the retrieval is improved by further constraining it with integrated water vapor from reanalysis model output. This would be interesting to study.

The chosen target, 2C-ICE, could very well be replaced by DARDAR. This is not likely to lead to large changes for SPARE-ICE but could be worth investigating nevertheless.

The development of a synergistic retrieval from a database need not be from collocations but can also be done using simulations. By combining a sophisticated method of generating synthetic atmospheric profiles, such as presented by *Evans et al.* [2012], and a suitable radiative transfer model, one can also build a retrieval database. This could use either the same input channels, which would give a product using the same

information as SPARE-ICE but potentially independent from either SPARE-ICE or 2C-ICE (depending on how the retrieval database is created), or additional inputs using channels that do not yet exist, such as those at submillimeter frequencies. In general, the addition of submillimeter frequencies should improve the SPARE-ICE retrieval. After 2022, European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) plans to fly the Ice Cloud Imager (ICI) on the MetOp-SG-B1 satellite. ICI is very similar to instruments that have been proposed before, such as CloudIce [*Buehler et al.*, 2012].

Another aspect that is worthwhile to study is to more formally investigate the information content for each of the channels used in SPARE-ICE, along the lines of *Cooper et al.* [2006]. This might lead to a more specific and stronger recommendation on what information to include or not for an IWP retrieval algorithm.

The principle by which the SPARE-ICE retrieval works can be applied to other geophysical quantities. In general, it is very beneficial if scientific satellites fly in orbits near operational ones, because operational retrievals can learn from scientific ones, as SPARE-ICE illustrates.

5. Summary and Conclusion

In this article, we have presented SPARE-ICE, a new IWP product based entirely on passive, operational sensors available day and night. By collocating NOAA 18 with the CloudSat 2C-ICE IWP product, we obtained a training database of AVHRR and MHS measurements on the one hand and joint radar-lidar IWP on the other hand. With this database, we have trained a set of neural networks for the detection of atmospheric ice and the retrieval of IWP, using 2C-ICE IWP as a reference. By using three AVHRR channels, three MHS channels, and auxiliary information (satellite angles, surface temperature, and surface elevation), SPARE-ICE can retrieve IWP, while the median fractional error compared to 2C-ICE IWP is around 100% for IWP > 10 g/m². Because we use only terrestrial infrared and microwave measurements, SPARE-ICE is available day and night, unlike products such as MODIS and PATMOS-X, which are available daytime only.

Overall, SPARE-ICE should be a relevant new member to the family of IWP products. SPARE-ICE will be publicly available through the WDC-RSAT under the Open Data Commons Attribution License.

References

Amante, C., and B. Eakins, (2009), ETOPO1 1 arc-minute global relief model: Procedures, data sources and analysis, *Tech. Rep.*, NOAA NESDIS NGDC.

Austin, R. T., A. J. Heymsfield, and G. L. Stephens (2009), Retrieval of ice cloud microphysical parameters using the CloudSat millimeter-wave radar and temperature, *J. Geophys. Res.*, *114*, D00A23, doi:10.1029/2008JD010049.

Bezy, J.-L., W. Leibrandt, A. Heliere, P. Silvestrin, C.-C. Lin, P. Ingmann, T. Kimura, and H. Kumagai (2005), The ESA Earth Explorer EarthCARE mission, *Proc. of SPIE*, *5882*, 143–154, doi:10.1117/12.619438.

Boukabara, S.-A., et al. (2011), MiRS: An all-weather 1DVAR satellite data assimilation and retrieval system, *IEEE Trans. Geosci. Remote*, *49*(9), 3249–3272, doi:10.1109/TGRS.2011.2158438.

Buehler, S. A., et al. (2012), Observing ice clouds in the submillimeter spectral range: The CloudIce mission proposal for ESA's Earth Explorer 8, *Atmos. Meas. Tech.*, *5*, 1529–1549, doi:10.5194/amt-5-1529-2012.

Cooper, S. J., T. S. L'Ecuyer, P. Gabriel, A. J. Baran, and G. L. Stephens (2006), Objective assessment of the information content of visible and infrared radiance measurements for cloud microphysical property retrievals over the global oceans. Part II: Ice clouds, *J. Appl. Meteorol. Clim.*, *45*, 42–62, doi:10.1175/JAM2327.1.

Cracknell, A. P. (1997), *The Advanced Very High Resolution Radiometer*, CRC Press, Taylor and Francis, Ltd., London, U. K.

Delanoë, J., and R. J. Hogan (2008), A variational scheme for retrieving ice cloud properties from combined radar, lidar, and infrared radiometer, *J. Geophys. Res.*, *113*, D07204, doi:10.1029/2007JD009000.

Delanoë, J., and R. J. Hogan (2010), Combined CloudSat-CALIPSO-MODIS retrievals of the properties of ice clouds, *J. Geophys. Res.*, *115*, D00H29, doi:10.1029/2009JD012346.

Deng, M., G. G. Mace, Z. Wang, and H. Okamoto (2010), Tropical composition, cloud and climate coupling experiment validation for cirrus cloud profiling retrieval using CloudSat radar and CALIPSO lidar, *J. Geophys. Res.*, *115*, D00J15, doi:10.1029/2009JD013104.

Deng, M., G. G. Mace, Z. Wang, and R. P. Lawson (2013), Evaluation of several A-Train ice cloud retrieval products with in situ measurements collected during the SPARTICUS campaign, *J. Appl. Meteorol. Clim.*, *52*(4), 1014–1030, doi:10.1175/JAMC-D-12-054.1.

Dybbroe, A., K.-G. Karlsson, and A. Thoss (2005), NWCSAF AVHRR cloud detection and analysis using dynamic thresholds and radiative modelling—Part I: Algorithm description, *J. Appl. Meteorol.*, *44*, 39–54.

Eliasson, S., S. A. Buehler, M. Milz, P. Eriksson, and V. O. John (2011), Assessing observed and modelled spatial distributions of ice water path using satellite data, *Atmos. Chem. Phys.*, *11*, 375–391, doi:10.5194/acp-11-375-2011.

Eliasson, S., G. Holl, S. A. Buehler, T. Kuhn, M. Stengel, F. Iturbe-Sanchez, and M. Johnston (2013), Systematic and random errors between collocated satellite ice water path observations, *J. Geophys. Res.*, *118*, 1–14, doi:10.1029/2012JD018381.

Evans, K. F., J. R. Wang, D. O. Starr, G. Heymsfield, L. Li, L. Tian, R. P. Lawson, A. J. Heymsfield, and A. Bansemer (2012), Ice hydrometeor profile retrieval algorithm for high-frequency microwave radiometers: Application to the CoSSIR instrument during TC4, *Atmos. Meas. Tech.*, *5*(9), 2277–2306, doi:10.5194/amt-5-2277-2012.

Ferraro, R. R., F. Weng, N. C. Grody, L. Zhao, H. Meng, C. Kongoli, P. Pellegrino, S. Qiu, and C. Dean (2005), NOAA operational hydrological products derived from the advanced microwave sounding unit, *IEEE Trans. Geosci. Remote*, *43*(5), 1036–1049.

Gong, J., and D. L. Wu (2013), CloudSat-constrained cloud ice water path and cloud top height retrievals from MHS 157 and 183.3 GHz radiances, *Atmos. Meas. Tech. Discuss.*, *6*, 8187–8233, doi:10.5194/amtd-6-8187-2013.

Heidinger, A. K., and M. J. Pavolonis (2009), Gazing at cirrus clouds for 25 years through a split window, part 1: Methodology, *J. Appl. Meteorol. Clim.*, *48*(6), 1100–1116.

Holl, G., S. A. Buehler, B. Rydberg, and C. Jiménez (2010), Collocating satellite-based radar and radiometer measurements—Methodology and usage examples, *Atmos. Meas. Tech.*, *3*, 693–708, doi:10.5194/amt-3-693-2010.

Jacobowitz, H., L. L. Stowe, G. Ohring, A. Heidinger, K. Knapp, and N. R. Nalli (2003), The advanced very high resolution radiometer pathfinder atmosphere (PATMOS) climate dataset: A resource for climate research, *Bull. Am. Meteorol. Soc.*, *84*(6), doi:10.1175/BAMS-84-6-785.

Jiménez, C., S. A. Buehler, B. Rydberg, P. Eriksson, and K. F. Evans (2007), Performance simulations for a submillimetre wave cloud ice satellite instrument, *Q. J. R. Meteorol. Soc.*, *133*(S2), 129–149, doi:10.1002/qj.134.

John, V. O., G. Holl, S. A. Buehler, B. Candy, R. W. Saunders, and D. E. Parker (2012), Understanding inter-satellite biases of microwave humidity sounders using global simultaneous nadir overpasses, *J. Geophys. Res.*, *117*, D02305, doi:10.1029/2011JD016349.

Kidd, C., and V. Levizzani (2011), Status of satellite precipitation retrievals, *Hydrol. Earth Syst. Sci.*, *15*(4), 1109–1116, doi:10.5194/hess-15-1109-2011.

Kidd, C., D. R. Kniveton, M. C. Todd, and T. J. Bellerby (2003), Satellite rainfall estimation using combined passive microwave and infrared algorithms, *J. Hydrometeorol.*, *4*, 1088–1104.

King, M. D., S.-C. Tsay, S. E. Platnick, M. Wang, and K.-N. Liou, (1997), Cloud retrieval algorithms for modis: Optical thickness, effective particle radius, and thermodynamic phase, *Tech. Rep.*, MODIS Science Team, mODIS Algorithm Theoretical Basis Document No. ATBD-MOD-05.

Kleespies, T. J., and P. Watts (2007), Comparison of simulated radiances, jacobians and linear error analysis for the Microwave Humidity Sounder and the Advanced Microwave Sounding Unit-B, *Q. J. R. Meteorol. Soc.*, *132*, 3001–3010.

Krasnopolsky, V. M. (2007), Neural network emulations for complex multidimensional geophysical mappings: Applications of neural network techniques to atmospheric and oceanic satellite retrievals and numerical modeling, *Rev. Geophys.*, *45*, RG3009, doi:10.1029/2006RG000200.

Liu, G., and E.-K. Seo (2013), Detecting snowfall over land by satellite high-frequency microwave observations: The lack of scattering signature and a statistical approach, *J. Geophys. Res. Atmos.*, *118*(3), 1376–1387, doi:10.1002/jgrd.50172.

Mace, G. G., Q. Zhang, M. Vaughan, R. Marchand, G. Stephens, C. Trepte, and D. Winker (2009), A description of hydrometeor layer occurance statistics derived from the first years of merged Cloudsat and CALIPSO data, *J. Geophys. Res.*, *114*, D00A26, doi:10.1029/2007JD009755.

Marzano, F. S., M. Palmacci, D. Cimini, G. Giuliani, and F. J. Turk (2004), Multivariate statistical integration of satellite infrared and microwave radiometric measurements for rainfall retrieval at the geostationary scale, *IEEE Trans. Geosci. Remote*, *42*(5), 1018–1032, doi:10.1109/TGRS.2003.820312.

McFarquhar, G. M., and A. J. Heymsfield (1998), The definition and significance of an effective radius for ice clouds, *J. Atmos. Sci.*, *55*, 2039–2052.

Mollweide, K. (1805), Ueber die vom Prof. Schmidt in Gießen in der zwerten Abtheilung seines Handbuch der Naturlehre s. 595 angegebene Projection der Halbkugelfläche, *Mon. Corr.*, *13*, 151–163.

Nagle, F. W., and R. E. Holz (2009), Computationally efficient methods of collocating satellite, aircraft, and ground observations, *J. Atmos. Oceanic Technol.*, *26*, 1585–1595.

Rapp, A. D., G. Elsaesser, and C. Kummerow (2009), A combined multisensor optimal estimation retrieval algorithm for oceanic warm rain clouds, *J. Appl. Meteorol. Clim.*, *48*, 2242–2256.

Robel, J., et al. (2009), NOAA KLM User's Guide with NOAA-N, -N' supplement, *User Guide*, NESDIS, NOAA.

Rossow, W. B., and R. A. Schiffer (1991), ISCCP cloud data products, *Bull. Am. Meteorol. Soc.*, *72*, 2–20.

Saha, S. et al. (2010), The NCEP climate forecast system reanalysis, *Bull. Am. Meteorol. Soc.*, *91*(8), 1015–1057, doi:10.1175/2010BAMS3001.1.

Saunders, R. W., T. J. Hewison, S. J. Stringer, and N. C. Atkinson (1995), The radiometric characterization of AMSU-B, *IEEE Trans. Microw. Theory*, *43*(4), 760–771.

Stephens, G. L., and C. D. Kummerow (2007), The remote sensing of clouds and precipitation from space: A review, *J. Atmos. Sci.*, *64*(11), 3742–3765, doi:10.1175/2006JAS2375.1.

Stephens, G. L., et al. (2002), The CloudSat mission and the A-train, *Bull. Am. Meteorol. Soc.*, *83*, 1771–1790.

Stephens, G. L., et al. (2008), CloudSat mission: Performance and early science after the first year of operation, *J. Geophys. Res.*, *113*, D00A18, doi:10.1029/2008JD009982.

Stowe, L. L., H. Jacobowitz, G. Ohring, K. R. Knapp, and N. R. Nalli (2002), The advanced very high resolution radiometer (avhrr) pathfinder atmosphere (patmos) climate datatset: Initial analyses and evaluations, *J. Clim.*, *15*(11), 1243–1260, doi:10.1175/1520-0442(2002)015<1243:TAVHRR>2.0.CO;2.

Surussavadee, C., and D. H. Staelin (2008), Global millimeter-wave precipitation retrievals trained with a cloud-resolving numerical weather prediction model, part I: Retrieval design, *IEEE T. Geosci. Remote*, *46*, 99–108, doi:10.1109/TGRS.2007.908302.

Swingler, K. (1996), *Applying Neural Networks: A Practical Guide*, Morgan Kaufmann, San Francisco, USA.

Taylor, J. P., and S. J. English (1995), The retrieval of cloud radiative and microphysical properties using combined near-infrared and microwave radiometry, *Q. J. R. Meteorol. Soc.*, *121*, 1083–1112.

Toon, O. B., et al. (2010), Planning, implementation, and first results of the tropical composition, cloud and climate coupling experiment (TC4), *J. Geophys. Res.*, *115*, D00J04, doi:10.1029/2009JD013073.

Waliser, D. E., et al. (2009), Cloud ice: A climate model challenge with signs and expectations of progress, *J. Geophys. Res.*, *114*, D00A21, doi:10.1029/2008JD010015.

Walther, A., and A. K. Heidinger (2012), Implementation of the daytime cloud optical and microphysical properties algorithm (DCOMP) in PATMOS-x, *J. Appl. Meteorol. Clim.*, *51*, 1371–1390, doi:10.1175/JAMC-D-11-0108.1.

Winker, D. M., J. R. Pelon, and M. P. McCormick (2003), The CALIPSO mission: Spaceborne lidar for observation of aerosols and clouds, in *Lidar Remote Sensing for Industry and Environment Monitoring III, Proceedings of SPIE*, vol. 4893, edited by Z. L. U. N. Singh and T. Itabe, pp. 1–11, SPIE, Hangzhou, China, doi:10.1117/12.466539.

Winker, D. M., M. A. Vaughan, A. Omar, Y. Hu, K. A. Powell, Z. Liu, W. H. Hunt, and S. A. Young (2009), Overview of the CALIPSO mission and CALIOP data processing algorithms, *J. Atmos. Oceanic Technol.*, *26*, 2310–2323, doi:10.1175/2009JTECHA1281.1.

WMO (2010), *Implementation Plan for the Global Observing System for Climate in Support of the UNFCCC*, vol. GCOS-138 (WMO/TD No. 1532), WMO, Geneva, Switzerland.