Assimilation of ice and water observations from SAR imagery to improve estimates of sea ice concentration

By K. ANDREA SCOTT[1]*, ZAHRA ASHOURI[1], MARK BUEHNER[2], LYNN POGSON[3] and TOM CARRIERES[3], [1]*Department of Systems Design Engineering, University of Waterloo, Waterloo, Ontario, Canada*; [2]*Data Assimilation and Satellite Meteorology Research Section at Environment Canada, Dorval, Québec, Canada*; [3]*Marine and Ice Services, Environment Canada, Ottawa, Ontario, Canada*

ABSTRACT

In this paper, the assimilation of binary observations calculated from synthetic aperture radar (SAR) images of sea ice is investigated. Ice and water observations are obtained from a set of SAR images by thresholding ice and water probabilities calculated using a supervised maximum likelihood estimator (MLE). These ice and water observations are then assimilated in combination with ice concentration from passive microwave imagery for the purpose of estimating sea ice concentration. Due to the fact that the observations are binary, consisting of zeros and ones, while the state vector is a continuous variable (ice concentration), the forward model used to map the state vector to the observation space requires special consideration. Both linear and non-linear forward models were investigated. In both cases, the assimilation of SAR data was able to produce ice concentration analyses in closer agreement with image analysis charts than when assimilating passive microwave data only. When both passive microwave and SAR data are assimilated, the bias between the ice concentration analyses and the ice concentration from ice charts is 19.78%, as compared to 26.72% when only passive microwave data are assimilated. The method presented here for the assimilation of SAR data could be applied to other binary observations, such as ice/water information from visual/infrared sensors.

Keywords: sea ice, data assimilation, synthetic aperture radar, binary, Bayesian, probability, forward model

1. Introduction

Synthetic aperture radar (SAR) data are widely used for sea ice monitoring. This is due to the ability of SAR sensors to provide useful information over a range of weather and illumination conditions, and the relatively high resolution of the data (nominal pixel spacing ≈ 50 m for a wide swath RADARSAT 2 image). Visual interpretation of SAR imagery by trained analysts provides a key contribution to operational ice charts (Carrieres et al., 1996). These ice charts provide information regarding the ice concentration and stage of development (e.g. new ice, grey ice, grey-white ice, first-year ice, multi-year ice and fast ice) over a given region valid at a specific time. Ice concentration data from ice charts are currently assimilated in the operational ice forecasting system at Environment Canada (Buehner et al., 2013) in combination with data from passive microwave sensors and scatterometers (Buehner et al., 2014).

To improve the efficiency and consistency of information from SAR imagery entering the assimilation system, methods are being developed to assimilate information from SAR images without first visually interpreting the images. This is challenging because it is difficult to retrieve the parameters of interest (ice concentration, ice thickness, ice type) from SAR images automatically, and in a robust manner.

Several approaches have been proposed in the literature for classification of SAR sea ice imagery into ice and open water classes, and to estimate sea ice concentration from SAR imagery (Karvonen et al., 2005; Berg and Eriksson, 2012; Karvonen, 2012, 2014; Leigh et al., 2014). Berg and Eriksson (2012) used image backscatter autocorrelation as input to a neural network to separate ice and open water in SAR imagery. Spatial averaging was applied to the classification result to calculate sea ice concentration. Comparison of the ice concentration with that from operational ice charts found differences that could be either due to variability in surface characteristics unrelated to ice concentration, or details not captured by the ice charts. Karvonen et al. (2005) used a Gaussian mixture model

*Corresponding author.
email: ka3scott@uwaterloo.ca

1

(GMM) of image backscatter autocorrelation calculated for homogeneous segments in an algorithm to estimate ice concentration for each segment. The sea ice concentration was one minus the mixture fraction of the open water class in the GMM. The classified images were found to have an accuracy of approximately 80% in comparison with operational ice charts. More sophisticated methods from the computer vision community are also under development (Ochilov and Clausi, 2012; Leigh et al., 2014). An overall classification accuracy of 96% has been found in a recent study combining a pixel level support-vector machine with image contextual information to classify the image into ice and water regions (Leigh et al., 2014).

Assimilating image data in which the image has been segmented into homogeneous regions would require a different approach than that taken for continuous data, due to the segmented nature of the observations. An alternative to assimilating ice concentration or ice/water classes retrieved from the image would be to assimilate the backscatter values directly. Such an approach has been used to directly assimilate brightness temperatures from a passive microwave radiometer over sea ice (Scott et al., 2012) and SAR observations of snow over land (Phan et al., 2013). However, assimilating SAR backscatter directly over sea ice would require a physical model for the interaction of the SAR signal with the ice surface. Such a model would be a complex function of poorly known surface conditions, such as the sea ice salinity, snow cover and surface temperature. An alternative approach is used here to assimilate data from SAR imagery. The method begins with classifying groups of pixels over windows in the image as either ice or water. The resulting binary ice/water observations are then assimilated in a system to improve estimates of ice concentration, together with ice concentration observations from a passive microwave sensor.

To the best of the author's knowledge, there are few previous studies on assimilating binary data directly into continuous fields. Theoretical foundations regarding this problem have been developed in Storto and Tveter (2009), in which a Bayesian formulation is used to describe the probability of the (continuous) true state given a binary observation and the continuous background field. This theory is then applied to the problem of assimilating cloud cover observations. In application of the theory (Storto and Tveter, 2009), binary cloud/no cloud observations are combined to form a cloud fraction observation, similar to the concept of averaging ice/water points to develop an ice concentration. A method to map the cloud fraction observation to the state variable (humidity) was then developed for assimilation. Other studies on the use of binary snow cover observations (snow on/off) similarly often average points to develop a snow cover fraction (Pullen et al., 2011), or use the snow cover information in a rule-based

assimilation method (Rodell and Houser, 2004; Fletcher et al., 2012). The present study develops a novel approach, in which a mapping is developed between the continuous ice concentration state and the binary observation space to assimilate the binary observations directly.

The remainder of the paper is organised as follows; the data used in this study are described in Section 2, the method used to process observations from SAR imagery is described in Section 3, and the method used to assimilate these observations is described in Section 4. The experimental set-up and results are given in Sections 5 and 6, respectively, and conclusions are given in Section 7.

2. Data

2.1. SAR imagery

The sea ice imagery used in this study was acquired from RADARSAT 2, which obtains measurements at 5.3 GHz (C band). The data were acquired in ScanSAR wide mode, with a swath range of 500 km and a nominal pixel spacing of 50 m. For the results presented here, both HH-polarised images (horizontal transmit-horizontal receive) and HV-polarised images (horizontal transmit–vertical receive) were used. The dataset consists of 23 images, acquired from January 17th 2014 to February 10th 2014. All images were acquired over the Gulf of St. Lawrence, on the east coast of Canada. The chosen time period corresponds to freeze-up and ice growth in the Gulf of St. Lawrence.

2.2. Ice concentration from the ARTIST Sea Ice algorithm

Ice concentration calculated using the ARTIST Sea Ice (ASI) algorithm (Spreen et al., 2008) was assimilated either alone, or with the observations from SAR. The ice concentration data were downloaded from www.icdc.zmaw.de/seaiceconcentration_asi_amsre.html. The ice concentration in the ASI algorithm is calculated using data from the 89 GHz channels of the AMSR2 sensor and is available once per day on a polar stereographic grid with a spatial resolution of 3.125 km (Beitsch and Kern, 2014). It should be noted that actual spatial resolution of the data is limited by the instrument field of view for 89 GHz on the AMSR2 sensor, which is 3 km × 5 km. The ASI ice concentration was also used as a source of training data to obtain the ice and water observations. ASI ice concentration has been compared with ice concentration derived independently from ship-based observations (Spreen et al., 2008) and from Landsat images (Wiebe et al., 2009). The correlation with ship-based ice concentration was 0.80, with the lowest correlations being in areas of low ice concentration. In the comparison against Landsat data, it was found the errors

depended on the ice type, and were highest for new ice, with ASI underpredicting the ice concentration as compared to Landsat for this case. An extensive study of various algorithms for the calculation of ice concentration from passive microwave sensors found that an underestimate of ice concentration in regions of thin ice (thickness less than 30 cm) is common feature of these algorithms (Ivanova et al., 2013). The underestimation is, however, lower for algorithms using the high frequency channels, such as the ASI algorithm. The ASI algorithm was found to under-estimate the ice concentration only for ice less than 20 cm in thickness. Ice concentration calculated using ASI was chosen for this study because the region considered con-tains large quantities of thin ice; hence, any improvement from using SAR in addition to ASI, compared with assi-milating ASI data alone, will not be biased due to the choice of ice concentration data. In addition, the study region is enclosed by land and also contains several islands. This means that, due to proximity to coastlines, data need to be discarded due to land spillover. Since higher fre-quency channels on passive microwave sensors correspond to data at higher spatial resolution, this effect is minimised by choosing an ice concentration calculated using the 89 GHz channel.

2.3. *Image analysis charts from the Canadian Ice Service*

Image analysis charts prepared manually by expert analysts at the Canadian Ice Services are used to assess the quality of the assimilation, and are also used as an alternate source of training data. These charts represent a manual inter-pretation of the SAR imagery, and hence are not indepen-dent of the SAR imagery used in this study. In preparation of an image analysis chart, the analyst draws polygons indicating areas where the ice conditions appear to be either homogeneous, or to contain a mixture of up to three different ice types. The preparation of image analyses is subjective, and small-scale features, such as cracks in the ice cover, are typically not part of such an analysis. Errors in the ice concentration may also arise due to converting continuous image data to discrete categories.

3. Calculation of observations from SAR imagery

3.1. *Demonstration with a toy model*

To motivate the method used to generate the ice and water observations from the SAR imagery, a simple example is presented in this section. A true sea ice state is generated that consists of binary values, zero and one. This can be

thought of as a very high resolution representation of sea ice concentration, in which the pixels are small enough to contain pure samples of ice and water. The true sea ice state is designed such that the locations of ones and zeros in a specified interval (20 points) are chosen randomly, but the concentration of ones in each interval increases linearly. In this way, an ice concentration can be generated by averaging the true state over the interval, which will be referred to as a spatial average. The true state, IC_{true}, is shown in Fig. 1a.

The true state is then used to simulate high-resolution training data. For this purpose, observations are sampled from the true state according to

$$y = \begin{cases} 10 + \epsilon_{obs}, & \text{if } IC_{true} = 0 \\ 20 + \epsilon_{obs}, & \text{if } IC_{true} = 1 \end{cases} \qquad (1)$$

The observation error, ϵ_{obs}, is chosen to follow a Gaussian distribution $N(0, 2.5^2)$. A random set of these observations is shown in Fig. 1b. The chosen Gaussian parameterisation leads to the probability density functions (PDFs) shown in Fig. 1c, which represent the likelihood of observing a given value, knowing that the point is either ice or water. The likelihoods are then used to calculate the probability of ice or water given an observed value using Bayes theorem. Assuming two classes, A_1, A_2 (ice and water), and a uni-form prior for each class, the probability of a class given the observed value, B, using Bayes theorem is

$$P(A_n|B) = \frac{P(A_n)P(B|A_n)}{\sum_j P(A_j)P(B|A_j)}. \qquad (2)$$

Note that because there are only two classes, and these two classes are mutually exclusive (it is assumed a point cannot be covered by both ice and water), the probabilities of ice and water must sum to 1.

Probabilities of ice and water are then calculated for two sets of observations. The first set consists of perturba-tions of the true state, following eq. (1), and they were left at the resolution of the true state. These observations will be referred to as the high-resolution observations. For this case, even when using only a single set of observations, the probability of ice, $P(ice)$, tends to be high where IC_{true} indicates ice, and low where IC_{true} indicates water (see Fig. 1d). Applying a threshold to these probabilities, such that points with $P(ice) > 0.9 = 1$ and points where $P(ice) < 0.1 = 0$ generates a binary ice/water state, very similar to true state. If the desired observation for assimilation was ice concen-tration, a spatial averaging operator could be applied to this estimated ice/water state, with the fraction of points corresponding to ice within the support of the operator, divided by the total, corresponding to the ice concentra-tion. Carrying this out for a single set of observations yields the ice concentration shown as the solid line in Fig. 1e.
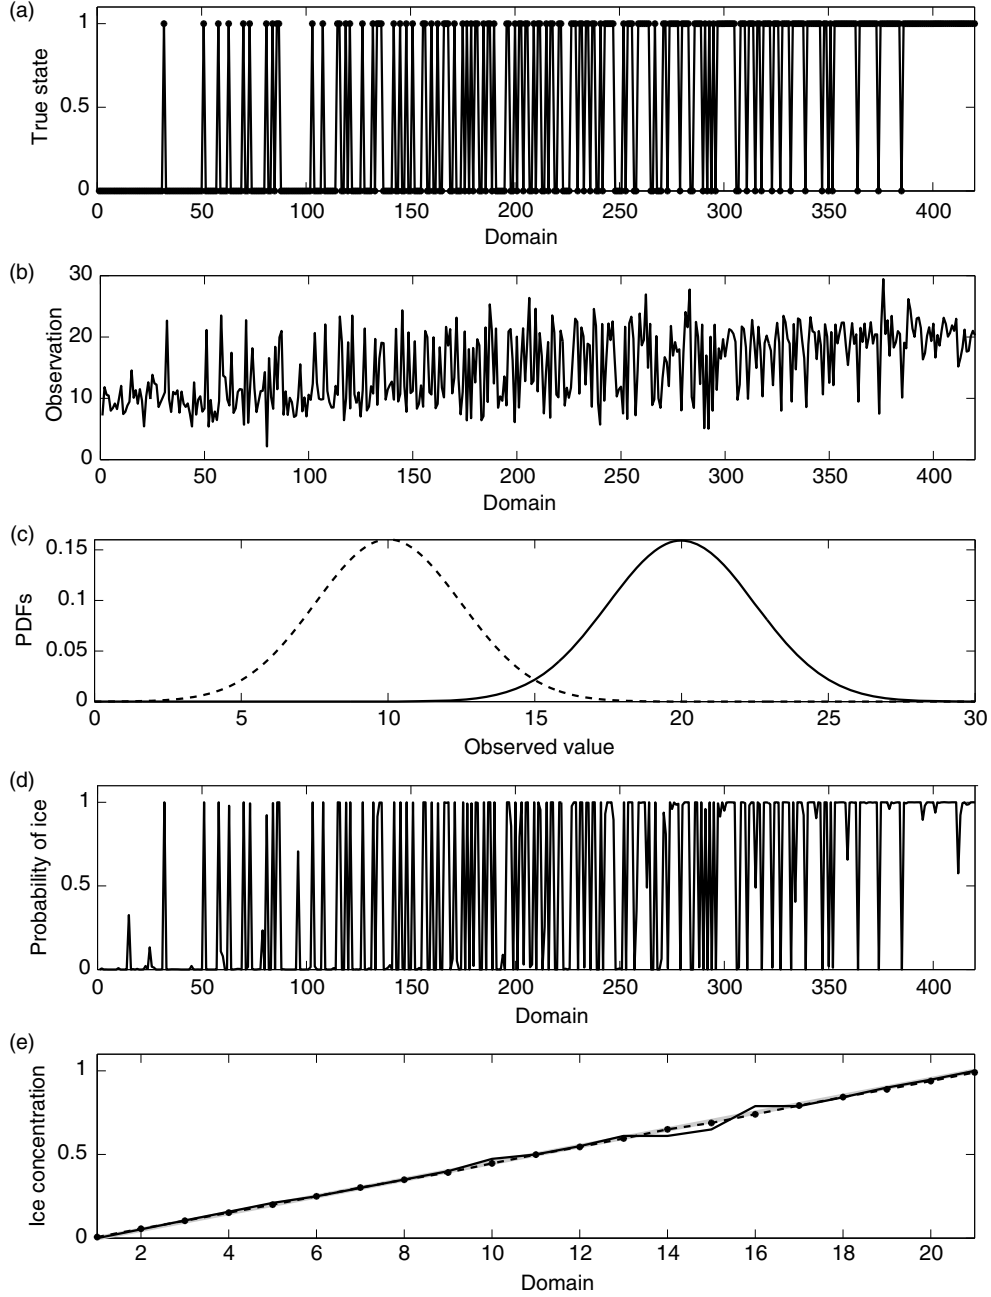
Fig. 1.   True sea ice state (a), a random set of high-resolution observations (b), likelihoods of ice (solid) and water (dashed) given an observed value (c), probability of ice corresponding to the high-resolution observations (d), ice concentration derived from the probability of ice, solid line is from a single set of observations, dashed line is averaged over many sets, thick grey line is the true ice concentration (e).

Note that to generate this ice concentration, points for which the probability of ice is between 0.1 and 0.9 are not included in the spatial averaging operation, with the support of the averaging operation adjusted accordingly. Averaging results over many sets of observations yields a smoother ice concentration, shown as the dashed line in Fig. 1e, which corresponds very closely to the true ice concentration, shown as the thick grey line. These results

are expected because the training data and the high-resolution observations were generated using the same sampling procedure and are at the same resolution, and the spatial averaging to obtain ice concentration was carried out over the same interval of 20 points that was used to design the true state. The purpose is to show that given high-resolution observations and training data, a binary ice/water state can be recovered that can subsequently be

spatially averaged to obtain an ice concentration representative of the true ice concentration.

The second set of observations was generated to more closely represent the present situation, in which the observations from the SAR image are not of sufficiently high resolution to represent pure pixels of ice and water. These observations are generated by first sampling the true state according to eq. (1) and then spatially averaging these observations over intervals of 20 points that follow those used to design the true state. These observations will be referred to as the low-resolution observations. A single set of these observations is shown in Fig. 2a, with the corresponding probability of ice for this set shown as the solid line in Fig. 2c. The likelihood functions used to calculate the probability of ice are the same as those used for the high-resolution observations (i.e. the same training data are used), and are shown in Fig. 2b to guide the interpretation of Fig. 2c.

The probability of ice in Fig. 2c is high over a range of points in the true state that correspond to a high ice concentration. If the true state were to be spatially averaged to generate an ice concentration, it would be found that the probability of ice is 1 when this true ice concentration is greater than 0.8. Applying a threshold to the probability of ice, such that the points where $P(ice) > 0.9$ correspond to 1, and points where $P(ice) < 0.1$ correspond to 0, would generate a binary ice/water state, with missing values reflecting the locations where the probability of ice has intermediate values between 0.1 and 0.9. However, a value of 1 in this state can only indicate that the given point has a high probability of having a high ice concentration. Calculating an ice concentration from these ice/water points would not make sense. This situation represents the present case, in which the observations generated from the SAR imagery have a resolution of 2.1 km and an observation with high probability of ice cannot be interpreted as having an ice concentration of 100%. For this reason, the observations assimilated consist of binary values (0 and 1) and are not averaged to generate an ice concentration before assimilation. Note that reducing the observation error standard deviation from 2.5 to 1 does not alter this conclusion. With a reduced observation error, there is less overlap between the ice and water PDFs, shown in Fig. 2d, resulting in a sharper transition in the probability of ice, shown in Fig. 2e.

The toy model uses a true state consisting of zero and ones that are averaged over intervals containing 20 points. To relate this model to the real situation, the typical floe size in the region is compared with the spatial scale used to calculate the SAR image features. Inspection of the polygon information in the CIS daily ice charts for the period January 17th–February 10th indicates that a typical floe size is 100 m. Averaging over 20 points corresponds to

an average over a spatial distance of 2 km, which is approximately the same as the distance over which the SAR image backscatter is taken into account in the calculation of the image features (mean HV backscatter and HH correlation), as discussed in Section 3.2.

3.2. Selection and calculation of image features

Visual inspection of a SAR image reveals changes in image tone and texture that are related to changes in the surface conditions and imaging geometry. For example, the tone of the HH backscatter over open water changes with incidence angle, and texture features are sometimes visible over open water due to wind roughening that can be similar to what is seen over ice. Consequently, the tone from an HH SAR image cannot be used to discriminate between ice types or ice and water without additional information. HV imagery has been found to be more useful toward discriminating between ice and water (Leigh et al., 2014). However, the HV signal is close to the noise floor, and should be combined with additional information for a more robust result. In this study, HV backscatter is combined with a texture feature from the HH image. Texture features were calculated using the grey level co-occurrence matrix (GLCM) (Soh and Tsatsoulis, 1999). Before calculating the texture features, image backscatter values are averaged in 2×2 blocks to reduce speckle noise. In the calculation of the texture features, a window size of 21×21 pixels was used. In combination with the 2×2 block averaging, this means that each pixel of the resulting image corresponds to 42×42 pixels in the original image. This results in a nominal spatial resolution of 2.1 km. The GLCM is calculated using 32 quantisation levels, a displacement of 1 and results using different orientations are averaged together to produce an isotropic representation of the texture in the GLCM window. Discussion of the roles of quantisation, displacement and orientation in the calculation of GLCM texture features can be found in Clausi (2002); Shokr (1991); Soh and Tsatsoulis (1999). While various texture features were investigated, such as contrast, entropy and correlation, it was decided to use only one HH texture feature in combination with HV mean backscatter, averaged over the 2.1-km window, to maintain simplicity. The chosen HH texture was the GLCM correlation, because it visually led to reasonable ice/water separation. Results were not significantly altered when a different texture feature was chosen.

3.3. Generation of probabilities

Ice and water observations are calculated from the SAR imagery using a simple Bayesian maximum likelihood estimator (MLE). The likelihood of a class (e.g. ice or water)
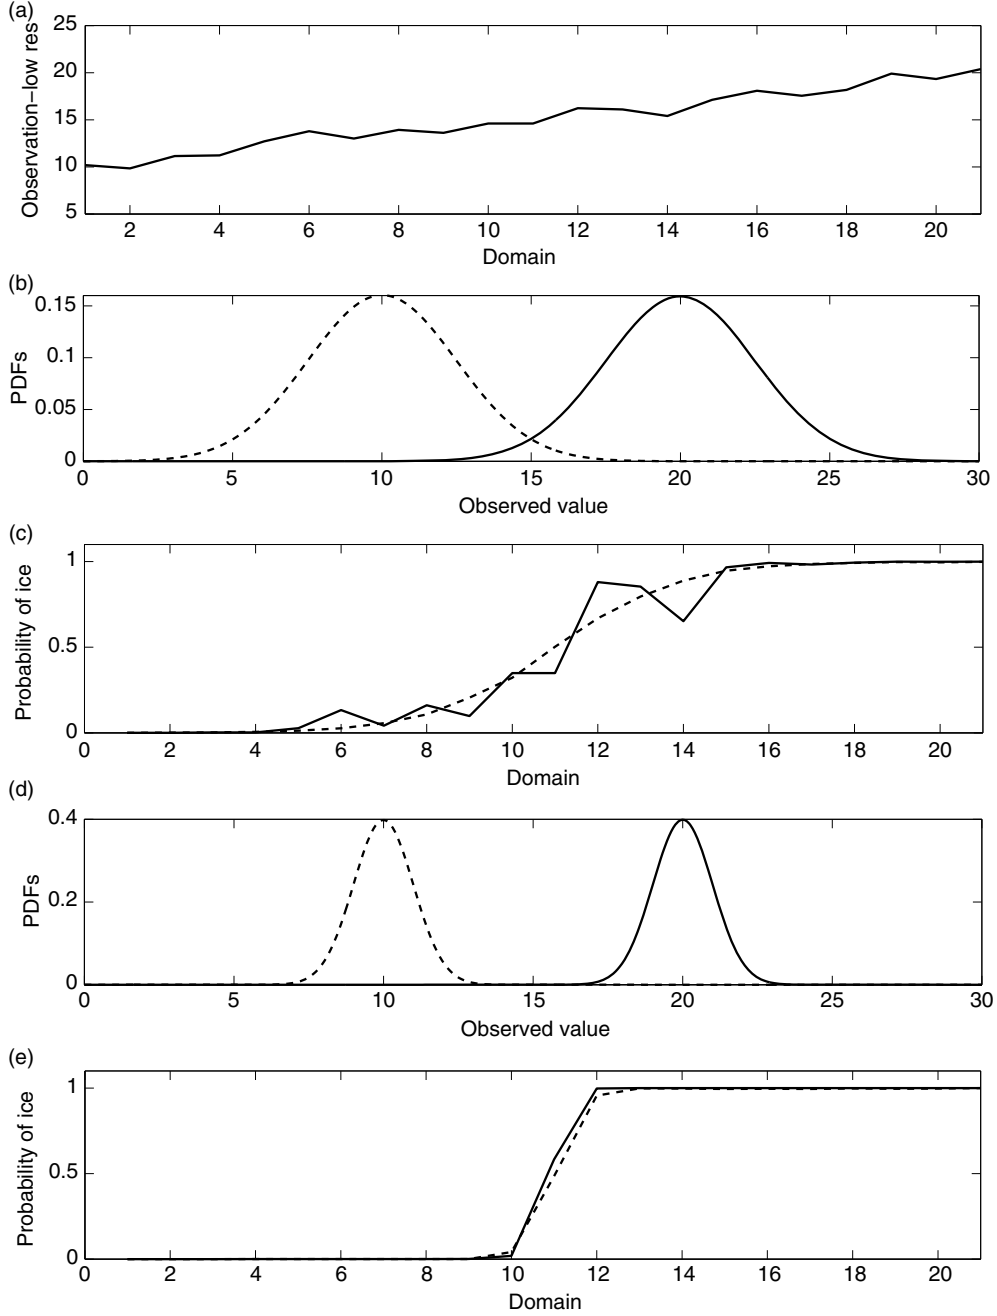
Fig. 2. A random set of low-resolution observations (a), likelihoods of ice (solid) and water (dashed) given an observed value, panel (b) $\sigma_o = 0.25$, panel (d) $\sigma_o = 0.1$, probability of ice corresponding to the low-resolution observations, solid line is from a single set of observations, dashed line is averaged over many sets, panel (c) $\sigma_o = 0.25$, panel (e) $\sigma_o = 0.1$.

given an image feature (e.g. backscatter, texture) is found using a supervised method. The pixels in each image are binned according to the ice concentration in the training data to form two sets of image features, one corresponding to ice and one corresponding to water.

From this binned data the joint pdf that represents the likelihood function for the image features given the class is

calculated. The probability of the class given the image features is then given by

$$P(A_n|B_1 \cap B_2) = \frac{P(A_n)P(B_1 \cap B_2|A_n)}{\sum_j P(A_j)P(B_1 \cap B_2|A_j)}, \qquad (3)$$

where A_n refers to the class (ice, water) and B_i refers to different features calculated from the SAR imagery (here

this is the GLCM correlation from the HH image and mean backscatter from the HV image). The prior probability of each class is assumed to be equal, that is, $P(A_1) = P(A_2)$, which leads to a MLE. For the results shown in this paper, a non-parametric Parzen window method was used to represent the likelihood functions (Botev et al., 2010).

Two different sets of data were tested to bin the image pixels into ice or water, for computing $P(B_1 \cap B_2|A_n)$. The first set used was the ASI ice concentration, while the second set was the ice concentration from the image analysis. In both cases, all 23 available SAR images are used in training. The purpose of training is to identify image features representative of ice and water, respectively. To this end, points in the training data that correspond to ice and water need to be identified. For the ASI data, if the ice concentration at the pixel location is greater than 0.98, the pixel is considered to be ice, and if the ice concentration is less than 0.02, the pixel is considered to be open water. For the image analysis charts, pixels contained in polygons with ice concentration 9^+ are considered to be ice, while those in polygons labelled as water are considered to be open water. In both cases, points in the training data that do not meet these criteria are not used.

Due to the fact that ASI tends to underestimate the ice concentration in areas of thin ice (Wiebe et al., 2009; Ivanova et al., 2013), thin ice regions have a low ice concentration and do not meet the criteria to be included in the ice training data. For example, the ASI ice concentration and the daily ice chart indicating the stages of development for January 24th are shown in Fig. 3a and b, respectively. The new ice (defined as ice with thickness less than 0.1 m) corresponds to the light pink regions in the ice chart that have concentrations generally near 0.9 (as indicated by the top number in the corresponding 'egg' code given in units of tenths). These regions typically have very low ice concentration in ASI, and at times may be included instead in the open water training data. This is reflected in the PDF for the water class, shown in Fig. 4b. The tail of the PDF, which is associated with low values of HV backscatter and moderate values of HH GLCM correlation (around 0.4), represents the thin ice that is being included in the open water training data. The probability of ice for the SAR image acquired at 10:23 UTC on January 24, 2014, is shown in Fig. 5a. It can be seen that the areas of thin ice that are not captured in the ASI data, are associated with a low probability of ice, as expected. In comparing Fig. 5b with 9f, it can be seen that when the ice charts are used for training, there is a strong correspondence between locations with a high probability of ice, and the occurrence of ice in the ice charts, as expected.

In contrast, the PDFs for the ice and water classes are shown in Fig. 4c and d when image analyses are used as the training data. It can be seen that the water PDF no longer has the tail corresponding to thin ice. The ice PDF is now shifted to lower values of HV backscatter, and higher values of correlation. This shift reflects the large portion of thin ice in the training data. The probability of ice for the image on January 24th is shown in Fig. 5b for the case when the charts are used as the training data. It can be seen that the probability of ice is now relatively high over both the thin ice and the thicker ice, because both types of ice are included in the training data.

3.4. Generation of observations from probability maps

The water observations were chosen to be those pixels for which the probability of water is greater than 0.98. For the ice observations, two different thresholds were tested, corresponding to pixels where the probability of ice is greater than 0.90 and 0.98, respectively. Observations corresponding to these conditions for the image acquired at 10:23 UTC on January 24th 2014 are shown in Fig. 5c and e for the case when ASI data are used for training and in Fig. 5d and f for the case when the image analysis charts are used for training. It can be seen that there are more observations when the image analysis charts are used for training, in particular in the areas where there is thin ice. It should be noted that this day is not typical, in that quantity of SAR observations available is larger than usual. The number of SAR observations available on each day is shown in Fig. 6. It can be seen that the number of SAR observations increases substantially when the probability threshold is lowered from 0.98 to 0.90.

4. Assimilation method

The data assimilation method is based on finding the analysis that corresponds to the minimum of a cost function that measures the difference between the analysis and the background, weighted by the inverse of the background error covariance matrix, B, and the difference between the analysis and the observations, weighted by the inverse of the observation error covariance matrix, R. This cost function is written as

$$J(x) = \frac{1}{2}(x - x_b)^T B^{-1}(x - x_b) + \frac{1}{2}(y - H(x))^T R^{-1}(y - H(x)),$$

where x is the analysis state that is being determined, x_b is the background state, y is the vector of observations and H is the observation operator (also referred to as the forward model) that maps the state vector to the observation space. The cost function is minimised iteratively, with the final analysis state denoted as x_a, and the analysis
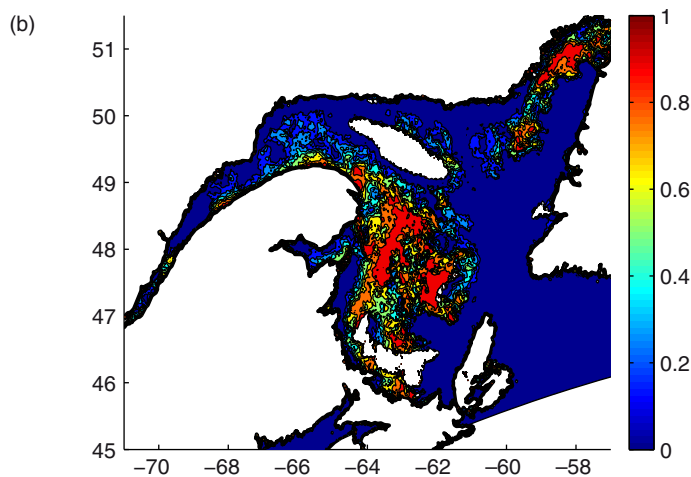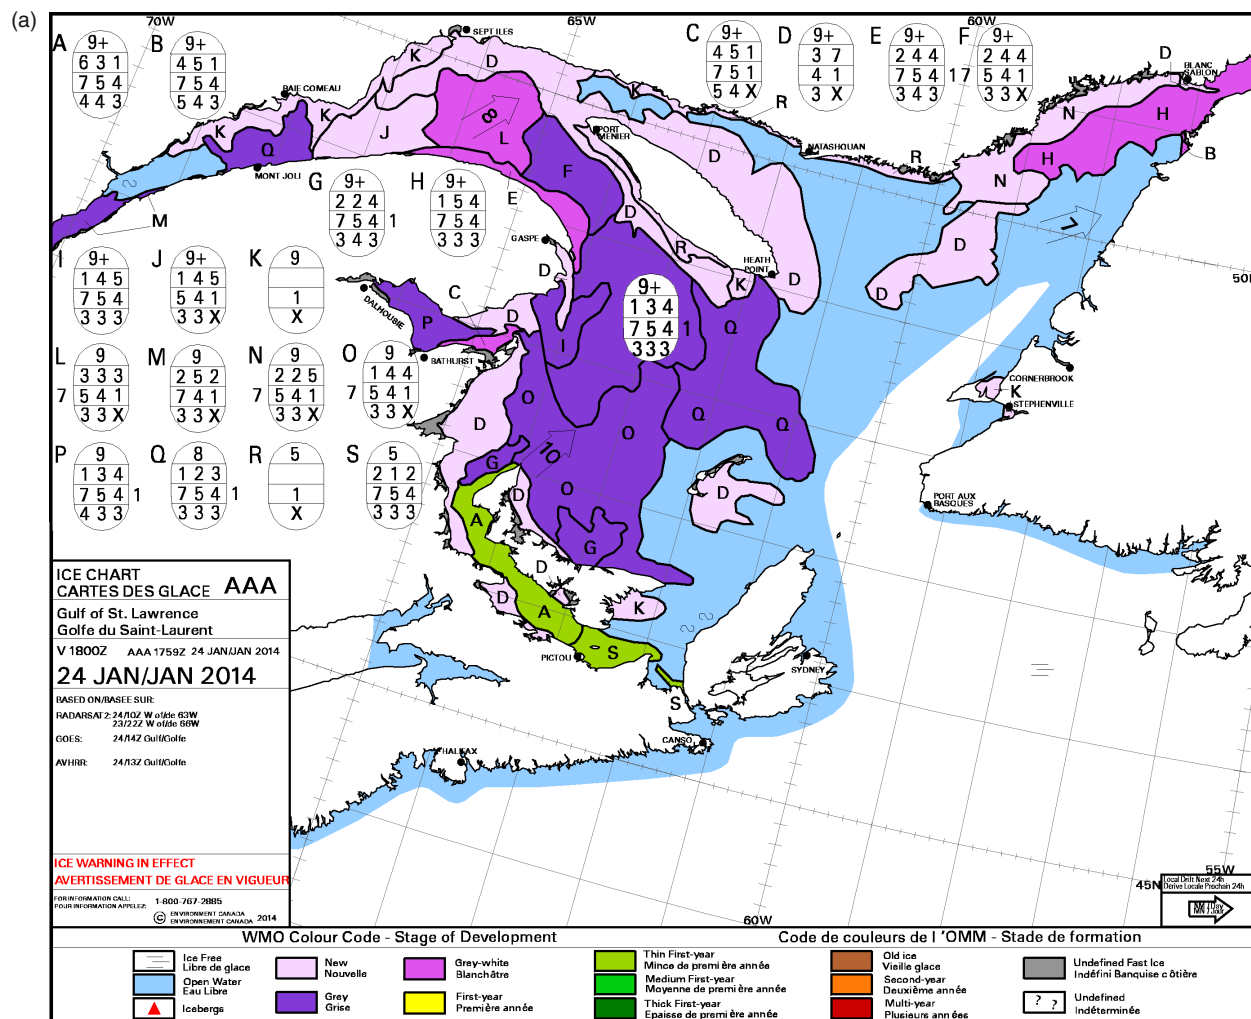
Fig. 3.   Ice concentration on January 24th 2014. Panel (a) is the ice chart from the Canadian Ice Service, with colour coding corresponding to stage of development of the ice. The light pink colour corresponds to new ice (ice of thickness less than 10 cm). Panel (b) is the ASI ice concentration for January 24th 2014.
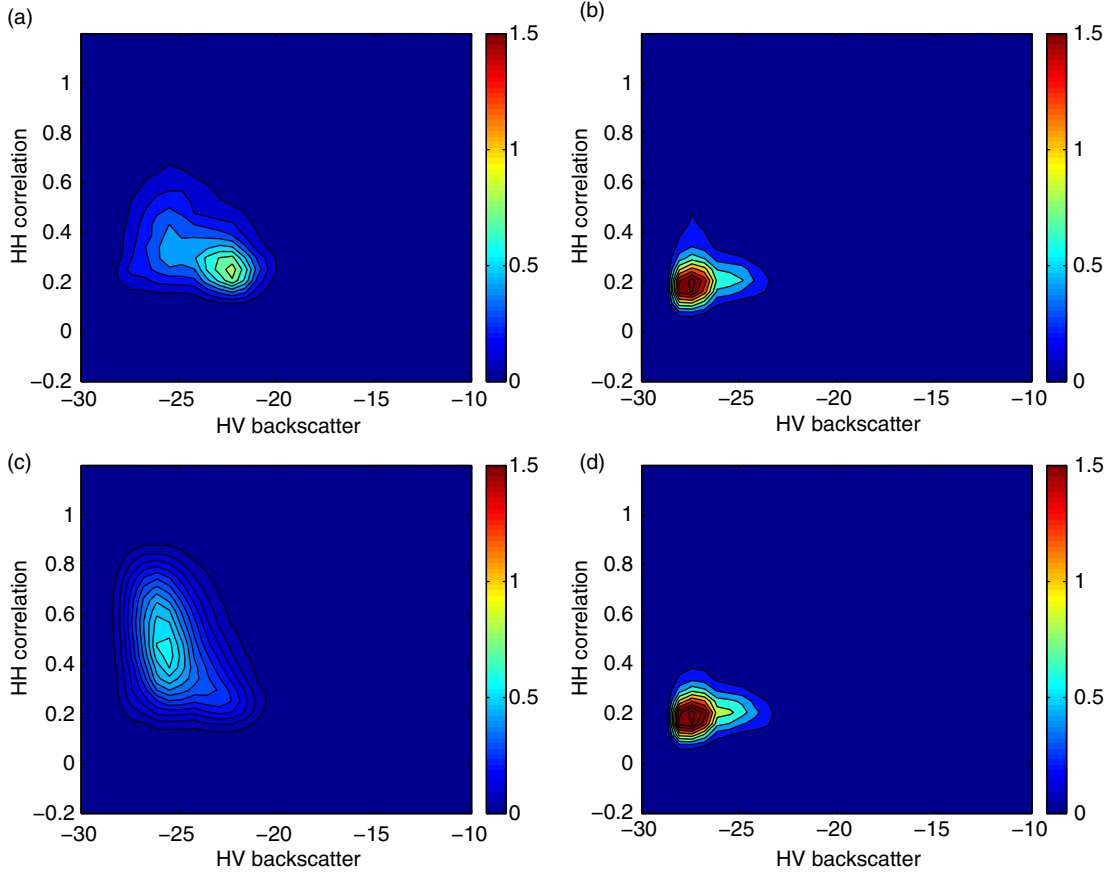
Fig. 4. Joint PDF of the probability of (a, c) ice and (b, d) water given HV mean backscatter and HH GLCM correlation over a 2.1-km window, calculated using ASI training data (a, b) and training data from image analysis charts (c, d).

increment denoted as x_a–x_b. To allow a clear examination of the impact of the observations, the background and observation error covariance matrices were assumed diagonal, with equal error variances, $\sigma_o^2 = \sigma_b^2$. In this way, each observation can only impact the location where it is assimilated, i.e. the usual spreading of observation information to neighbouring locations by the spatial correlations within the background error covariance matrix is not included.

Due to the fact that the spatial resolution of the SAR observations, 2.1 km, is higher than that of the background state, 3.125 km, in some cases there is more than one SAR observation at a given point in state space. When this occurs a majority vote rule is applied to determine if the SAR observation assimilated is a zero or a one. In the case that the number of ones at the given point in state space is equal to the number of zeros, no SAR observation is assimilated due to the fact that the analysis increment from assimilating a one would cancel with that from assimilating a zero.

4.1. Assimilation of ice and water observations

In order to assimilate the ice and water observations, a forward model, denoted as $H(x)$, that maps the state vector, x, consisting of ice concentration, to the observation space, y, consisting of zeros and ones, is needed. Due to the binary nature of the observations, a different forward model is used depending on whether the observation is ice (one) or water (zero).

For the case when the observation is ice, a forward model is desired such that when the ice concentration from the background state is high, the observation will have little or no impact on the final analysis state, whereas when the ice concentration from the background state is low, the observation can lead to an increase in the ice concentration relative to the background state. This behaviour is desired because an ice observation does not provide information on the exact details of how high the ice concentration is, only that the pixel is most likely ice covered (see Section 3.1). When the ice concentration in the background state is already relatively high, and the observation is ice, the two
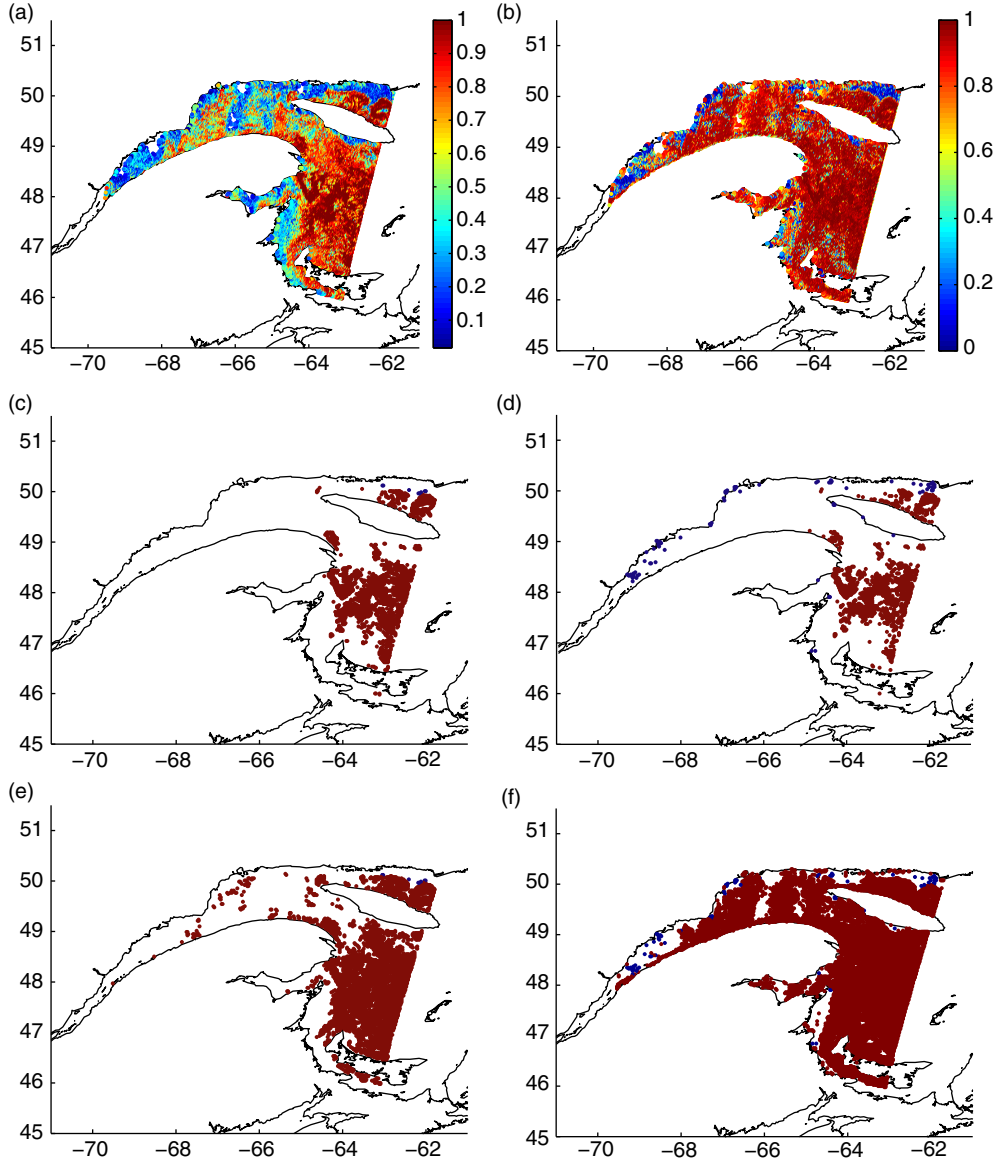
Fig. 5. Probabilities of ice for January 24th 2014 when (a) ASI data are used for training; and (b) when image analysis charts are used for training. Observations generated from the probabilities when (c, e) ASI data are used as the training data, and (d, f) when image analysis charts are used as the training data. Panels (c, d) correspond to a threshold of 0.98 and panels (e, f) correspond to a threshold of 0.90.

are considered to be in agreement, and the assimilation of the SAR ice observation should not change the background state. Similarly, when the observation is water, the forward model should be chosen such that it cannot remove ice when the ice concentration in the background state is already low.

4.1.1. Non-linear forward models. The first forward model considered that has the desired properties is shown in Fig. 7a. It should be noted that the slope of $H(x)$ is one

for low ice concentration, which will allow a non-zero cost function gradient to generate an ice concentration increment. At higher ice concentration the slope of $H(x)$ approaches zero, which means the cost function has no gradient, and hence no ice increment can be generated. The forward model for the open water observations is designed following a similar logic, and is shown in Fig. 7b. The functions chosen for the forward models shown in Fig. 7 are given in Appendix.

The behaviour of the forward model can be seen more clearly by examining the analysis state as a function of the
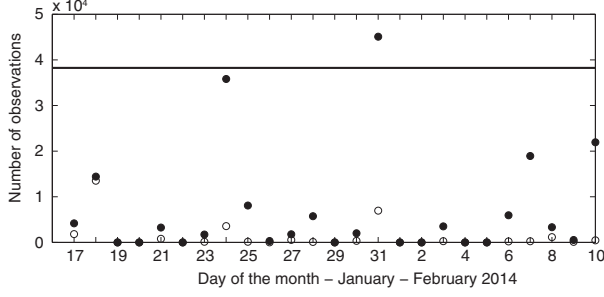
Fig. 6. Number of observations available on each day of the study. The horizontal line indicates the number of ASI observations, which is the same each day, while the circles indicate the SAR observations. Solid circles indicate the number of observations when a probability threshold of 0.90 is used, while unfilled circles indicate the number of observations when a probability threshold of 0.98 is used.

background state for the case when an ice observation is assimilated, and for the case when a water observation is assimilated at each point in the state space. This is shown in Fig. 8a for the forward models in Fig. 7. When the observation is ice, and the background state is water ($x_b = 0$), the ice concentration analysis is 0.5. As the background state ice concentration increases, the analysis increases smoothly to approach the background state, reflecting the fact that an ice observation has diminishing impact. Similarly, when the observation is water, and the background state is ice ($x_b = 1$), assimilation of the observation reduces the ice concentration to 0.5. Again, the impact of the observation decreases smoothly to allow the analysis to approach the background state, this time with decreasing ice concentration. Note that there is no impact from assimilating ice observations for background ice concentrations above approximately 0.8, and similarly no impact from assimilating water observations for background ice concentrations below approximately 0.2.

4.1.2. Linear forward models. A benefit to using a linear forward model is that it leads to a cost function that is quadratic. The linear forward models used in this study correspond to $H(x) = x$. The first linear forward model that was used simply treats the ice and water observations as though they corresponded to ice concentration values of 0 and 1. The analysis for this forward model is shown in Fig. 8a for the case when an ice (water) observation is assimilated, shown as the solid red(blue) lines. It can be seen that the analysis increments (difference between coloured lines and dotted black line) are larger than those from using the non-linear model.

The problem with the previous approach is that it naively treats the observations as though they corresponded to ice concentration values, which means that even when the ice concentration in the background state is already high, the ice observation can still add ice to the state, in spite of the fact that it is not known if the ice observation corresponds to an ice concentration value greater than that in the background state. To assimilate ice and water observations using a simple linear forward model while still maintaining the desired behaviour of having no analysis increment above a certain saturation value, the approach tested here is to assimilate an observed value equal to the saturation value. For example, if it is desired that the non-linear ice observation operator saturates (has zero slope) at a value of 0.7, the ice observation is assigned a value of 0.7. This observation is assimilated using the linear forward model $H(x) = x$ if the background state is less than 0.7, and is not assimilated if the background state is greater than 0.7. An analogous situation occurs for a water observation, with the difference being that the observation assimilated is 0.3. The analyses for these cases are shown as the lines with open circles in Fig. 8b. It can be seen that for the case when an ice observation is assimilated, and the background state indicates water ($x_b = 0$) the ice concentration analysis is
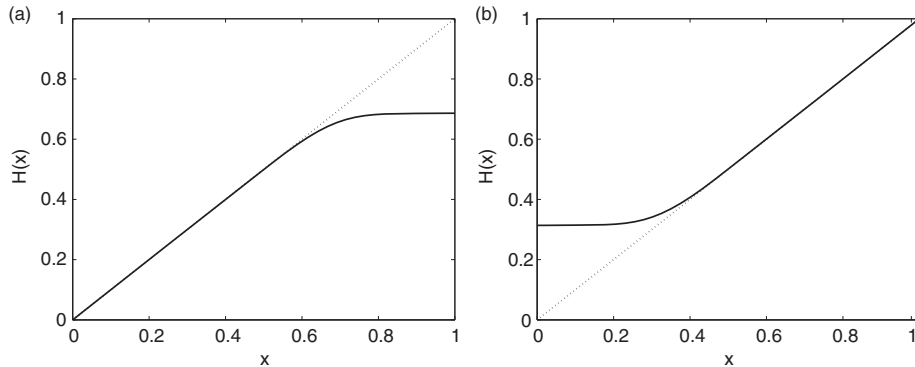

Fig. 7. Observation operator for the case when the observation is ice (a) and for the case when the observation is water (b). The linear observation operator is shown as the dotted line.
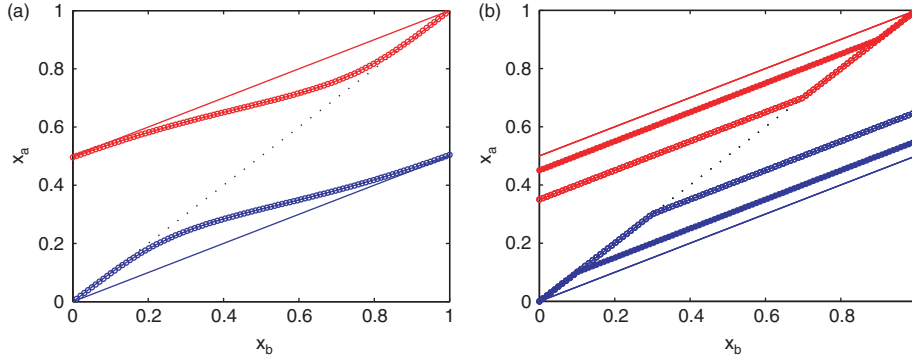
Fig. 8.   Background state and analysis state for the case where water observations are assimilated (blue) and for the case where ice observations are assimilated (red). In panel (a), the solid lines correspond the linear forward model assimilating observed values of 0 and 1, and the lines with circles correspond to the non-linear forward model. In panel (b), all models used are linear, solid lines correspond to assimilating observed values of 0 and 1, closed circles correspond to assimilating observed values of 0.9 and 0.1, and open circles correspond to assimilating observed values of 0.7 and 0.3. For both figures, $\sigma_o = 0.1$. and $\sigma_b = 0.1$.

0.35, which is less than that from when an observed value of 1 is assimilated. To also investigate a case that will have higher increments over regions where the background state indicates water and ice observations are assimilated, observed values of 0.9 and 0.1 were also tested for ice and water, respectively. Analysis increments for this case are also shown in Fig. 8b.

5. Experimental set-up

Each data assimilation experiment was initialised with the ASI ice concentration retrieval on January 16th, 2014. A sea ice model was not incorporated in the data assimilation cycle. Instead, persistence was used, which means that the analysis from the previous cycle provided the background state for the next assimilation cycle. Assimilation cycles with a prognostic sea ice model will be carried out in a future study. Observations were assimilated once per day, using a 24-hour assimilation window centred at 12 UTC. On each day, there is one set of ASI observations assimilated in combination with the observations from the SAR imagery available for the given day. The number of SAR images available on a daily basis ranges from 0 to 2, and the number of SAR observations available for each day is given in Fig. 6.

The observation vector is formed each day from the available ice/water observations from SAR, and ice concentration observations from ASI. Thus, for a day when both observation types are available, the observation vector contains entries consisting of 1's and 0's for ice/water observations from SAR, and ice concentration observations from ASI (continuous values between 0 and 1). The background state, consisting of ice concentration, is interpolated to each observation location using a nearest neighbour method. This approach, in combination with the

assumption of diagonal background and observation error covariance matrices, means the ice concentration estimated by the 3D-Var analysis at each location is the weighted average of the ASI ice concentration, the SAR observation (0 or 1) and the background ice concentration. The relative weight for each depends on the forward model used for each, and relative accuracy of each, where the relative accuracy is defined by the inverse of the respective error variance (Bouttier and Courtier, 1999). On a day for which only ASI observations are available, the analysis at each location is the weighted average of the ASI ice concentration and the background ice concentration, where again the relative weight is given by the inverse of their respective error variances.

Several experiments were conducted in which the source of training data and the forward model $H(x)$ were varied. The experiments for which results are presented are described in Table 1. The forward models linear07 and linear09 correspond to assimilating observations with a value of 0.7/0.3 for ice/water and 0.9/0.1 for ice/water, respectively. Experiments were also carried out using a threshold of 0.98 on the probability of ice to select ice observations, but the impact of assimilating SAR data was not as significant for these cases as compared to using a threshold of 0.90.

6. Results

6.1. Ice concentration analyses

Results from assimilating ice and water observations are shown in Fig. 9 for January 24th, 2014. The analysis increments for the case when ASI data are used as training, and when the image analysis charts are used as training with the non-linear forward model (EXP2 and EXP3) are shown in Fig. 9a and b, respectively. The analysis

Table 1.   Description of experiments carried out

Experiment	Observations assimilated	Training data	Threshold for ice observations	H(x)
EXP 1	ASI	None	N/A	Linear
EXP 2	ASI+SAR	ASI	0.90	Non-linear
EXP 3	ASI+SAR	Charts	0.90	Non-linear
EXP 4	ASI+SAR	Charts	0.90	Linear1
EXP 5	ASI+SAR	Charts	0.90	Linear07
EXP 6	ASI+SAR	Charts	0.90	Linear09

increment is larger in both magnitude and area covered for the case when the charts are used for training, due to the fact that there are more ice observations, and that many of these ice observations are in areas where the background state and the ASI data indicate a low ice concentration. Hence, both the slope of the observation operator, H, and the difference between the observation and $H(x_b)$, are relatively large, leading to a large analysis increment. The analyses for each of these cases is shown in Fig. 9c and d. It can be seen that the analysis for the case when charts are used for training is closer to the verification data (Fig. 9f) than the case when ASI data are used as training.
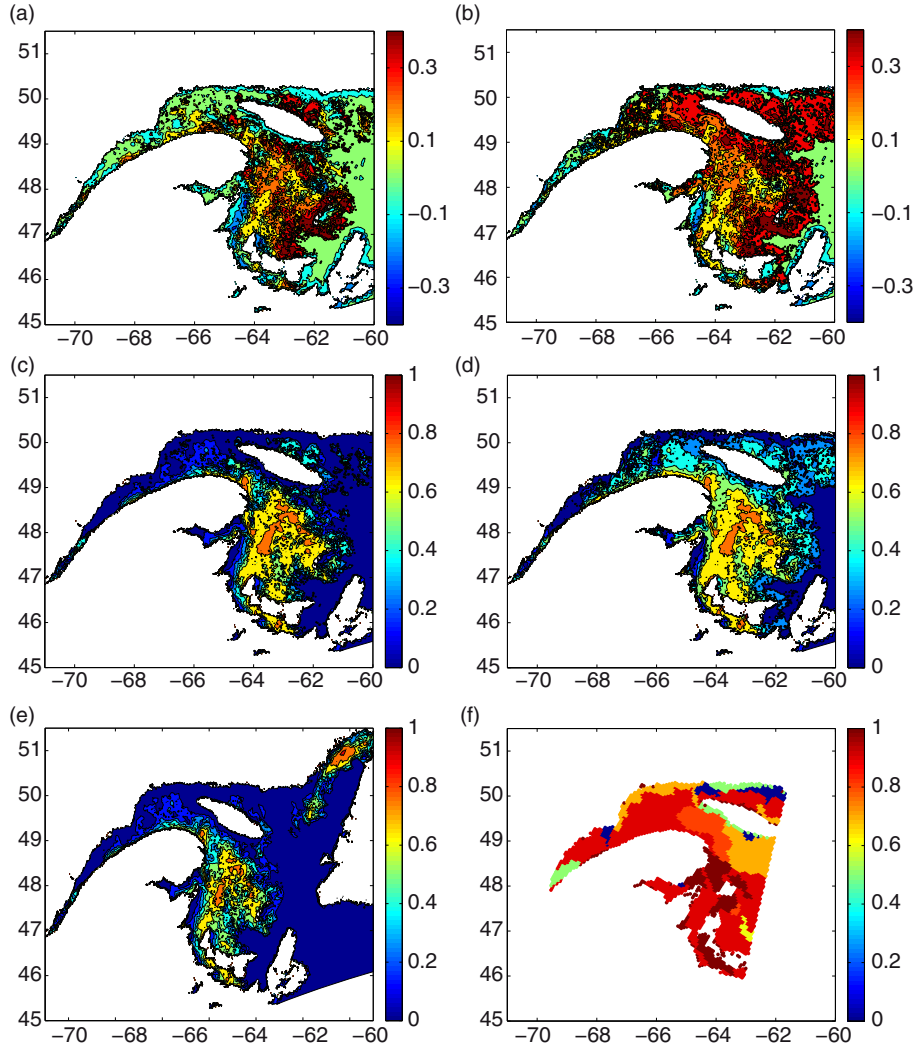

Fig. 9.   Results for January 24th. Analysis increments (a, b) and analyses (c, d). Panels (a, c), ASI data used for training (EXP2) (b, d); image analysis charts used for training (EXP3) (e) analysis when only ASI data are assimilated (EXP1) (f) image analysis chart.
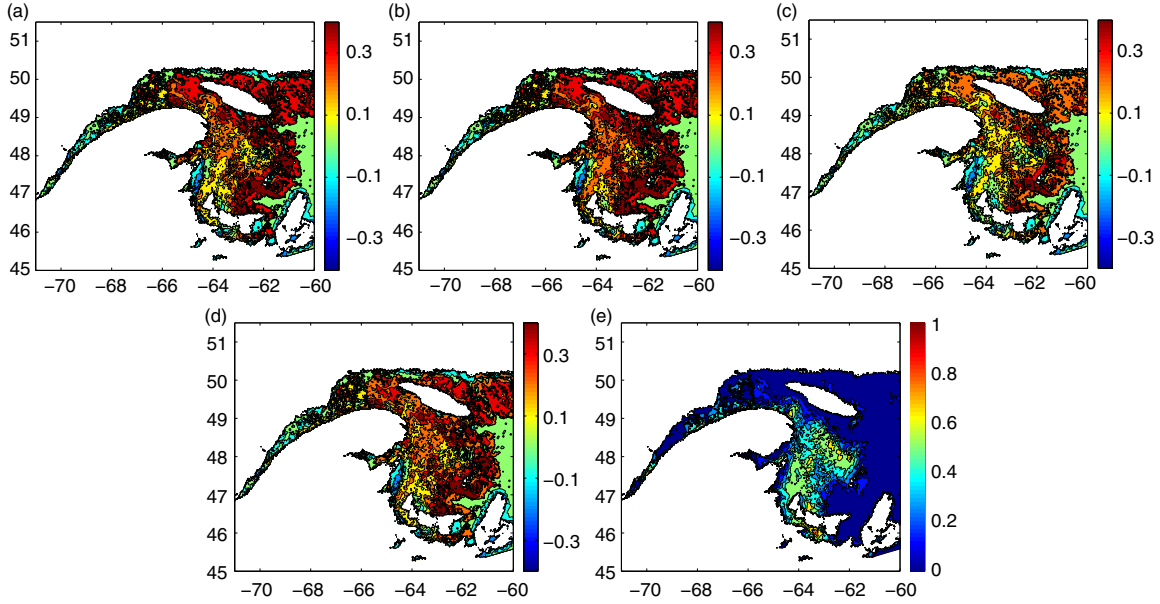
Fig. 10. Analysis increments for January 24th, impact of observation operator; (a) EXP3, (b) EXP4, (c) EXP5, (d) EXP6 and (e) background state for EXP3.

Both cases that assimilate SAR data are in closer agreement with the verification data than the case that assimilates only ASI data (EXP1) shown in Fig. 9e. The ice concentration analysis shown in Fig. 9e has less ice than the ASI observations (Fig. 3b) for this day due to the fact that the background state does not contain enough ice. This is because the experiment takes place during freeze-up, and large-scale areas of ice cover only begin to appear in the ASI ice concentration on January 22nd, 2014. This underestimation in ice concentration is still seen in the background state for EXP1 on January 24th.

6.2. Analysis increments

Analysis increments are shown in Fig. 10 for January 24th, 2014, for the different observation operators. For this figure, the training data used was image analysis charts. January 24th was chosen because prior to this day the assimilation of SAR data had little impact on the background state, hence the background states for each case are almost identical. A representative background state is shown in Fig. 10e. It can be seen in Fig. 10a that when the non-linear forward model is used there are significant positive analysis increments over a large portion of the domain. When this is compared with a linear model (Fig. 10b), the linear model produces slightly higher increments over those areas where the background state has a high ice concentration. Comparing these analysis increments with those found when observations of 0.7 and 0.3 are assimilated with a linear forward model, shown in Fig. 10c, it can

be seen that for the latter case the analysis increments are smaller, in particular over the portions of the background state where the ice concentration is zero but a SAR ice observation is present. This is due to the fact that an ice observation of 0.7 has a smaller value for $y-H(x_b)$ than an ice observation of 1, and is consistent with what is seen in Fig. 7. Modifying the value of the observations assimilated to 0.9 and 0.1 for ice and water, respectively, leads to the analysis increments shown in Fig. 10d, which are similar to those in Fig. 10b.

6.3. Verification against image analysis charts

The ice concentration analyses and background states from each data assimilation experiment are compared quantitatively against the ice concentration from the image analysis charts. The bias and standard deviation of each of the experiments are given in Table 2. When ASI data are used for training, it can be seen that there is a small positive impact from assimilating the SAR data. This is reflected in that fact that the ice concentration in both the background state and the analysis are closer to the ice charts than for the case when only the ASI data are assimilated. For the case when image analysis charts are used for training, it can be seen that the positive impact of assimilating the SAR data is larger, and the differences with image analysis charts are lower than for the case when ASI data are used as training, as expected. In comparing the linear model with the non-linear model, the results are very similar. The main potential drawback of using the linear model is that it

Table 2. Summary of bias (bias-b) and standard deviation (std-b) between background ice concentration and ice concentration from image analyses (values given are in%) over the study period, and bias (bias-a) and standard deviation (std-a) between ice concentration analysis and ice concentration from image analyses over the study period

| Experiment	bias-b	bias-a	std-b	std-a	Total number of SAR observations
EXP 1	−27.69	−26.72	27.75	26.69	∅
EXP 2	−27.20	−25.07	27.75	26.06	24 194
EXP 3	−25.73	−20.64	26.58	22.82	82 344
EXP 4	−25.59	−19.78	26.61	23.01	82 344
EXP 5	−26.43	−23.07	27.01	23.95	82 344
EXP 6	−25.88	−20.98	26.73	23.22	82 344

could yield an analysis increment that is too high when an ice observation is assimilated, or too low when a water observation is assimilated. However, for this study there were very few water observations, meaning that the latter case would not likely be seen in the analysis. For the former case, it may be difficult to detect if the ice concentration analysis is too high due to the nature of the verification data, which does not contain the small-scale details of the ice cover (e.g. openings in the ice) over a given polygon.

In Fig. 11, the difference between ice concentration from the image analysis charts and that from the analyses is shown for each day of the experimental period. It can be seen that both the bias and the standard deviation of the difference between the analysis from the assimilation and the image analysis are reduced for the experiments assimilating the SAR data (EXP2, EXP3 and EXP4) relative to the one that does not assimilate the SAR data (EXP1), and that this is the case for each day in the experimental period.

The differences may be larger toward the latter part of the experimental period as the Gulf of St. Lawrence became increasingly ice covered.

7. Concluding remarks

Results presented in this study demonstrate a positive impact associated with assimilating ice and water observations from SAR imagery. The bias between the ice concentration analyses and the ice concentration from ice charts is 19.78% when assimilating SAR observations with passive microwave ice concentration, as compared to 26.72% when assimilating only passive microwave ice concentration. The method used to obtain and assimilate the SAR observations is simple and scalable and does not require significant computing resources. The impact of the assimilation was greater when image analysis charts were used as the training data due to the fact that these
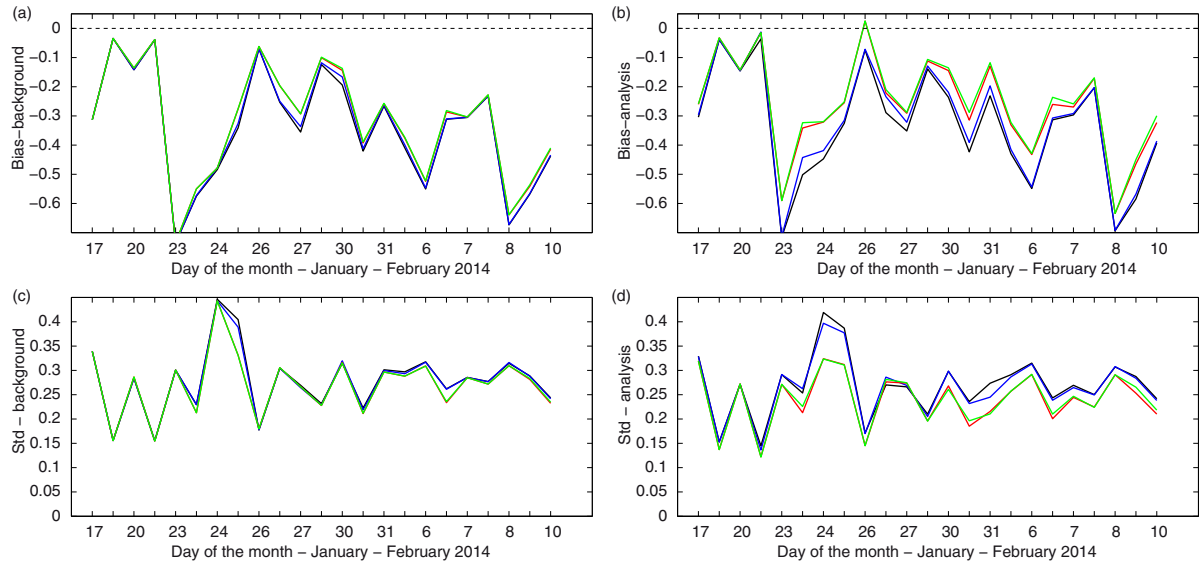
Fig. 11. Differences between ice concentration from the background state and that from image analyses, (a) bias – background state; (b) bias – analysis; (c) standard deviation – background state; (d) standard deviation – analysis state. In each panel, the black lines are results from EXP1, the red lines are from EXP2, the green lines are from EXP3 and the blue lines are from EXP4.

charts contain large areas of high concentration thin ice that appear as ice of very low concentration in ASI data. However, even in the case when the ASI data are used for training, there is a positive impact of the assimilation of SAR data when evaluated with data from image analysis charts. It should be kept in mind that both the image analysis charts and the assimilated results are based on SAR imagery, although the assimilated results also contain information from ASI ice concentration.

The assimilation of SAR data was tested using both linear and non-linear forward models. It was found that it is possible to approximate the saturation behaviour of the non-linear forward model by changing the value of the observation assimilated to the saturation value, and applying a linear model when the background state is less than the saturation value. While this method reduces the analysis increments at locations where there is a large difference between the state and the ice or water observation, it can be viewed as a more conservative approach than representing the ice and water observations as corresponding to ice concentrations of 100 and 0%, respectively, and applying a linear model over the entire range of state space.

The SAR observations are available at a spatial resolution of 50 m. In this study, the observations assimilated from SAR use GLCM texture and hence consider information over a spatial window, which is 2.1 km. While it may be advantageous to assimilate SAR observations at a higher resolution, the calculation of texture allows the data from the HH channel to be used without an incidence angle correction. Using the data at the full resolution of 50 m would not allow texture features to be calculated and the data from the HH channel would need to be dealt with much more carefully to account for the large across-swath variation in HH backscatter over open water, and in some cases, over ice.

The impact of assimilating SAR data in an assimilation system with a sea-ice or ice-ocean model will be assessed in a future study. In addition, the assimilation of SAR data in different geographic regions and for different seasons should be investigated. Due to the fact that SAR backscatter from sea ice is very sensitive to the ice conditions, the training would likely need to be carried out separately for different geographic regions. In addition, the performance of passive microwave retrieval algorithms varies with ice conditions. In the Gulf of St. Lawrence, the main disadvantage associated with ice concentration from passive microwave imagery may be an under prediction of ice concentration in areas covered with thin ice. For a future study, a region in which the quantity of open water is consistently under predicted by either passive microwave data or models may be used to more clearly evaluate the impact of the open water observations from SAR.

8. Acknowledgements

9. Appendix

To map the state space, x, to the binary observation space a pair of functions were chosen for $H(x)$ as modified rectilinear functions. For the case when the observation is ice, the chosen function has the form

$$H(\boldsymbol{x}) = 0.5 - (1/A)log(B + exp(-A(\boldsymbol{x} - 0.5))). \quad (A1)$$

For the case when the observation is water, the chosen function has the form

$$H(\boldsymbol{x}) = 0.5 + (1/A)log(B + exp(A(\boldsymbol{x} - 0.5))). \quad (A2)$$

For both cases, $A = 21$ and $B = 0.02$ were used.

References

Beitsch, A., Kaleschke, L. and Kern, S. 2014. Investigating high-resolution AMSR2 sea ice concentrations during the February 2012 fracture event in the Beaufort Sea. *Remote Sens.* **6**(5), 3841–3856.

Berg, A. and Eriksson, L. E. B. 2012. SAR algorithm for sea ice concentration – evaluation for the Baltic Sea. *IEEE Geosci. Remote Sens. Lett.* **9**(5), 938–942.

Botev, Z., Grotowski, J. and Kroese, D. 2010. Kernel density estimation via diffusion. *Ann. Stat.* **38**(5), 2916–2957.

Bouttier, F. and Courtier, P. 1999. *Data Assimilation Concepts and Methods, ECMWF Meteorological Training Course Lecture Series*. Technical Report 14, European Center for Medium-Range Weather Forecasting, Reading, England.

Buehner, M., Caya, A., Carrieres, T. and Pogson, L. 2014. Assimilation of SSMIS and ASCAT data and the replacement of highly uncertain estimates in the Environment Canada Regional Ice Prediction System. *Q. J. Roy. Meteorol. Soc.* DOI: http://dx.doi.org/10.1002/qj.2408

Buehner, M., Caya, A., Pogson, L., Carrieres, T. and Pestieau, P. 2013. A new environment Canada regional ice analysis system. *Atmos. Ocean.* **51**(1), 18–34.

Carrieres, T., Greenan, B., Prinsenberg, S. and Peterson, I. K. 1996. Comparison of Canadian ice charts with surface observations off Newfoundland, winter 1992. *Atmo. Ocean.* **34**, 207–236.

Clausi, D. A. 2002. An analysis of co-occurrence texture statistics as a function of grey level quantization. *Can. J. Remote Sens.* **28**(1), 45–62.

Fletcher, S., Liston, G., Hiemstra, C. and Miller, S. 2012. Assimilating MODIS and AMSR-E snow observations in a snow evolution model. *J. Hydrometeorol.* **13**, 1475–1492.

Ivanova, N., Pedersen, L. and Tonboe, R. 2013. *D2.5 Product Validation and Algorithm Selection Report (PVSAR). Sea Ice Concentration*. Technical Report SICCI-PVSAR Version 1.1, European Space Agency, Paris, France.

Karvonen, J. 2012. Baltic sea ice concentration estimation based on C-Band HH-polarized SAR data. *IEEE Trans. Geosci. Remote Sens.* **5**(6), 1874–1884.

Karvonen, J. 2014. Baltic sea ice concentration estimation based on C-Band dual-polarized SAR data. *IEEE Trans. Geosci. Remote Sens.* **52**(9), 5558–5566.

Karvonen, J., Simila, M. and Makynen, M. 2005. Open water detection from Baltic sea ice Radarsat-1 SAR imagery. *IEEE Geosci. Remote Sens. Lett.* **2**(3), 275–279.

Leigh, S., Wang, Z. and Clausi, D. A. 2014. Automated ice-water classification using dual polarization SAR satellite imagery. *IEEE Trans. Geosci. Remote Sens.* **52**(9), 5529–5539.

Ochilov, S. and Clausi, D. A. 2012. Operational SAR sea-ice image classification. *IEEE Trans. Geosci. Remote Sens.* **50**, 4397–4408.

Phan, X. V., Ferro-Famil, L., Gay, M., Durand, Y., Dumont, M. and co-authors. 2013. 3D-VAR multilayer assimilation of X-band SAR data into a detailed snowpack model. *Cryosphere Discuss.* **7**, 4881–4912.

Pullen, S., Jones, C. and Rooney, G. 2011. Using satellite-derived snow cover to implement a snow analysis in the Met Office global NWP model. *J. Appl. Meteorol. Climatol.* **50**, 958–973.

Rodell, M. and Houser, P. 2004. Updating a land surface model with MODIS-derived snow cover. *J. Hydrometeorol.* **5**, 1064–1075.

Scott, K. A., Buehner, M., Caya, A. and Carrieres, T. 2012. Direct assimilation of AMSR-E brightness temperatures for estimating sea ice concentration. *Mon. Weather Rev.* **140**, 997–1013.

Shokr, M. E. 1991. Evaluation of second-order texture parameters for sea-ice classification from radar images. *J. Geophys. Res.* **96**, 10625–10640.

Soh, L.-K. and Tsatsoulis, C. 1999. Texture analysis of SAR imagery using grey level co-occurrence matrices. *IEEE Trans. Geosci. Remote Sens.* **37**(2), 780–795.

Spreen, G., Kaleschke, L. and Heybster, G. 2008. Sea ice remote sensing using AMSR-E 89 GHz channels. *J. Geophys. Res.* **113**(C2), C02303.

Storto, A. and Tveter, F. 2009. Assimilating humidity pseudo-observations derived from the cloud profiling radar aboard CloudSat in ALLADIN 3D-Var. *Q. J. Roy. Meteorol. Soc.* **16**, 461–479.

Wiebe, H., Heygster, G. and Markus, T. 2009. Comparison of the ASI ice concentration algorithm with Landsat-7 ETM+ and SAR imagery. *IEEE Trans. Geosci. Remote Sens.* **47**(5), 3008–3015.