Diagonal Approximations to the Observation Error Covariance Matrix in Sea Ice Thickness Data Assimilation

by

Graham Stonebridge

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2017

Author's Declaration

I hereby declare that I am the sole author of this thesis.

I understand that my thesis may be made electronically available to the public.

Abstract

Data assimilation is a statistical technique for combining observations of a physical system with the state of a numerical model of that system. The procedure yields a new and ideally improved state estimate called the analysis. A critical component of data assimilation is the observation error covariance matrix, which describes the magnitude and the correlation of the errors in the observations. When the observation error correlation structure is unknown, an approximation can yield a poor analysis and an incorrect estimate of the quality of the analysis.

Little is known about the error correlation structure of remotely-sensed sea ice thickness observations. However, sea ice prediction centres are beginning to move forward with ice thickness assimilation under the simplifying assumption that the observation errors are uncorrelated. The assumption of uncorrelated observation errors is attractive because the errors can be represented by a diagonal observation error covariance matrix, which is inexpensive to invert. The purpose of this thesis was to develop an understanding of how the diagonal approximation might affect the quality of the sea ice state estimate.

This thesis describes a set of twin assimilation experiments that were conducted using a one-dimensional sea ice model. The twin experiment design enabled an investigation of the differences between the estimated and actual errors in the analysis state. The first part of this investigation explored how the diagonal approximation can impact the estimated mean analysis error standard deviation. The second component of the investigation explored the spatial scales of the errors present in the analysis.

The experimental results indicated that the diagonal approximation can be used without increasing the mean analysis error standard deviation so long as the observation error variances are multiplied by a sufficiently-large inflation factor. The results also indicated that the inflation factor can be conservatively overestimated without adversely impacting the analysis. For some of the experiments, the diagonal approximation resulted in an increase in the analysis error spectral variance at lower wavenumbers. The approximation had little effect at higher wavenumbers.

The main finding of this thesis is that diagonal approximations to the ice thickness observation error covariance matrix can likely be incorporated into ice prediction systems without adverse effects. One caveat of this statement is that an inflation factor should be used to increase the observation error variance estimates. A second caveat is that the analysis error covariance matrix may underestimate the correlation of analysis errors at the largest spatial scales. A final finding is that large improvements in analysis quality may be obtained if better approximations to the ice thickness observation error covariance matrix can be found and used in the analysis.

Acknowledgements

I would like to thank my supervisor, Dr. K. Andrea Scott, for providing helpful feedback and support throughout my time as her student. With her guidance, I was able to navigate through the uncertainty miring sea ice data assimilation toward the optimistic conclusion presented in this thesis.

I would also like to thank the Marine Environmental Observation Prediction and Response (MEOPAR) network for providing technical training and the opportunity to participate in an Arctic field program.

Table of Contents

List of Tables

List of Figures

List of Symbols

Symbol	Description
a	ice concentration
h	ice thickness
u	velocity
ρ	density
τ	quadratic drag term
ζ	bulk sea ice viscosity
P^*	sea ice strength parameter
S_h, S_a	thermodynamic source terms for ice thickness and concentration
$f(h)$	sea ice growth rate function
η	fluid free surface elevation
$\mathbf{x}^b$	background state vector
$\mathbf{x}^t$	true state vector
$\mathbf{y}$	observation vector
ε	error vector
σ	standard deviation
$\mathbf{x}^a$	analysis
$\mathbf{B}$	background error covariance matrix
$\mathbf{R}$	observation error covariance matrix
$\mathbf{A}$	analysis error covariance matrix

$H(\mathbf{x})$	observation operator
$\mathbf{H}$	observation operator (linear)
$\mathbf{K}$	Kalman gain matrix

List of Abbreviations

Acronym	Description
3DVar	Three-dimensional variational data assimilation
ACF	Autocorrelation Function
AEM	Airborne electromagnetic
BLUE	Best Linear Unbiased Estimate
CIS	Canadian Ice Service
CMC	Canadian Meteorological Centre
ECCC	Environment and Climate Change Canada
EnKF	Ensemble Kalman filter
EnOI	Ensemble Optimal Interpolation
GIOPS	Global Ice-Ocean Prediction System
MEOPAR	Marine Environmental Observation Prediction and Response Network
MIRAS	Microwave Imaging Radiometer with Aperture Synthesis
NWP	Numerical Weather Prediction
NWT	Northwest Territories
OI	Optimal Interpolation
RIPS	Regional Ice Prediction System
RMSD	Root-mean square deviation
RMSE	Root-mean square error
SMOS	Soil Moisture and Ocean Salinity (satellite)
ULS	Upward-looking sonar
VisIR	Visible Infrared

Chapter 1

Introduction

Accurate prediction of sea ice conditions is becoming increasingly important as global attention turns toward the Arctic Ocean [1]. Ongoing improvements to numerical sea ice models (e.g. the Los Alamos sea ice model [2]) have provided one means to improve forecasts. New satellite systems (e.g. CryoSat-2 [3]) have also helped improve forecast quality by providing valuable observational data. However, despite these efforts, there continue to be large errors in sea ice forecasts [4][5][6] and it remains difficult to plan safe operations in the Arctic Ocean [1].

This thesis explores how advances in sea ice data assimilation can help to improve sea ice prediction. Data assimilation is a set of statistical procedures for combining observations of a physical system with the state of a numerical model [7]. The product of data assimilation is a state estimate that may provide more accurate initial conditions for making forecasts. National forecasting centres such as the Canadian Meteorological Centre (CMC) at Environment and Climate Change Canada (ECCC) presently assimilate satellite observations of ice concentration [8], which is the fraction of the ocean surface that is frozen. However, there is now increasing motivation to begin assimilating observations of ice *thickness*, another critical state variable.

Ice thickness contributes to the ice strength and to the rate of heat transfer between the atmosphere and ocean [9]. Having a good initial estimate of ice thickness is therefore critical to producing a good forecast. Ice thickness can also vary significantly in space and time and the state of sea ice thickness typically poorly known [10]. By assimilating sea ice thickness observations, there is an opportunity to greatly improve forecasts [11][12].

The basis of any data assimilation system is prior knowledge of the model (background) and observation errors, which are described by the background and observation error covariance matrices, respectively [7]. However, the observation error covariance matrix is often either too difficult or too expensive to properly characterize [13]. Common simplifications, such as the

assumption that observation errors have no spatial correlation, facilitate implementation but can lead to overconfidence in what might be a poor state estimate [14][15]. Chapter 2.4 illustrates why it is especially difficult to parameterize sea ice thickness observation errors. Chapters 4 and 5 of the thesis explores the implications of the diagonal approximation to the observation error covariance matrix.

There are additional challenges compounding sea ice thickness data assimilation. First is the expense of conducting sea ice data assimilation experiments. Operational coupled sea ice-ocean-atmosphere models require large computational resources and expertise. Second, it is difficult to validate the results of a sea ice thickness assimilation experiment because there is very little *in situ* ice thickness data relative to the size of most model domains.

To help provide some insight for the design of operational sea ice prediction systems that ingest ice thickness observations, we developed a simplified one-dimensional dynamic sea ice model. The model contains the relevant terms necessary to simulate a realistic sea ice state and it is forced by synthetic but realistic forcing data. The primary application of the model is to conduct twin experiments, where observations are sampled from one model state using a synthetic sensor and assimilated into a second model state. This provides a means for assessing the impacts of the approximation of the observation error covariance matrix in a sea ice thickness data assimilation system.

1.1 Objectives

The princpal objective of this research project was to provide support for the design of an operational data assimilation framework that ingests sea ice thickness observations. More specifically, the objective was to quantify the effects of assuming an uncorrelated *diagonal* observation error covariance matrix when assimilating sea ice thickness observations having *correlated* errors. This is a pertinent and timely question due to the increasing availability of sea ice thickness observations and limitations of operational data assimilation systems. A secondary objective was to explore and communicate some of the challenges facing sea ice remote sensing, sea ice modeling, and data assimilation of observations with correlated errors.

1.2 Outline

This thesis is divided into five chapters. Chapter 2 describes in detail the unique challenges that sea ice thickness remote sensing presents, from a data assimilation perspective. We also

describe the limitations of recent experimental efforts to assimilate ice thickness data. Chapter 3 describes the toy sea ice model that was used as the basis of experiments in Chapters 4 and 5. Chapter 4 describes the twin experiment procedure, including the design of the background error covariance matrix and the synthetic sea ice thickness sensor. Chapter 5 explores the experimental results, followed by a pragmatic discussion of the future of sea ice thickness data assimilation.

Chapter 2

Background

Every data assimilation problem has unique challenges related to the sensor network and physical phenomenon being studied. This chapter explores the factors that complicate sea ice thickness data assimilation.

2.1 Sequential Data Assimilation

Geophysical systems including Earth's atmosphere and oceans are governed by dynamic, non-linear processes [16]. Using numerical models, it is possible to predict the future states of these systems. In this context, a *state vector*, or system state, is a vector containing the minimum amount of information at a given time to fully characterize that system [17]. Forecasts of future states can quickly become unrealistic due to model inaccuracies, errors in forcing data, and the propagation of errors in the models' initial states [16]. Ideally, observations of the geophysical system could be used as the initial conditions for a forecast. However, there are typically far fewer observations available than there are model variables [14]. Furthermore, these observations are prone to measurement and interpolation errors [14]. The statistical techniques of *data assimilation* provide a means to produce improved initial conditions, incorporating information from both the model state and the observations and taking into account the inaccuracies in each [14][7].

Sequential data assimilation describes a family of methods that follow a two-stage cyclical process [14], illustrated in Figure 2.1. In Figure 2.1, $\mathbf{x}^b$ represents the background state, our best estimate of the true system state, $\mathbf{x}^t$. The observation vector $\mathbf{y}$ contains a set of recent observations of the geophysical system. In the assimilation stage, the observations are combined

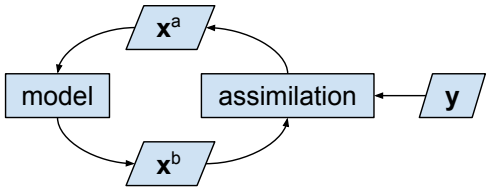

Figure 2.1: Schematic of an assimilation cycle. Observations, $\mathbf{y}$, are assimilated into a background state, $\mathbf{x}^b$, to produce an analysis, $\mathbf{x}^a$. The analysis is marched forward by a numerical model to produce the next background state.

with the background state to produce $\mathbf{x}^a$, the analysis. In the model stage, the analysis becomes the new model state and is marched forward in time using a numerical model, producing the background state for the next iteration.

In sequential data assimilation, the analysis is constructed as a linear combination of the background state and the observations [14]:

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{K}(\mathbf{y} - H(\mathbf{x}^b)) \tag{2.1}$$

where $\mathbf{K}$ represents the Kalman gain matrix and H represents the observation operator, which can be any linear or nonlinear user-defined function that maps the background state to the observations [14]. The linearized observation operator, $\mathbf{H}$, is the Jacobian of the observation operator:

$$\mathbf{H} = \left.\frac{\partial H}{\partial \mathbf{x}}\right|_{\mathbf{x}^b}. \tag{2.2}$$

The Kalman gain matrix, $\mathbf{K}$, is designed to minimize the trace of the analysis error covariance matrix, which describes the errors in the analysis [18]. The Kalman gain can be expressed as [18]

$$\mathbf{K} = \mathbf{B}\mathbf{H}^{\mathbf{T}}(\mathbf{R} + \mathbf{H}\mathbf{B}\mathbf{H}^{\mathbf{T}})^{-1} \tag{2.3}$$

where $\mathbf{B}$ and $\mathbf{R}$ are the background and observation error covariance matrices, respectively. The error covariance matrices are defined as

$$\mathbf{B} = < \varepsilon^b \varepsilon^{b^T} > \tag{2.4}$$

$$\mathbf{R} = < \varepsilon^o \varepsilon^{o^T} > \tag{2.5}$$

where ε^o and ε^b are error terms defined as

$$\boldsymbol{\varepsilon}^{\mathbf{b}} = \mathbf{x}^b - \mathbf{x}^t \tag{2.6}$$

$$\boldsymbol{\varepsilon}^{\mathbf{o}} = \mathbf{y} - \mathbf{y}^t \tag{2.7}$$

where $\mathbf{x}^t$ represents the true, but unknown state evaluated at the time of assimilation. Similarly, $\mathbf{y}^t$ represents the observations that might be obtained in the absence of sampling error [14]. This formulation yields the Best Linear Unbiased Estimate (BLUE) [7][14] when:

1. The errors are unbiased, i.e. $< \varepsilon^b >=< \varepsilon^o >= 0$; and when
2. The background errors and observation errors are uncorrelated, i.e. $< \varepsilon^o \varepsilon^{b^T} >= 0$.

There is also an implicit assumption here that the background and observation error covariance matrices are known exactly. Under these conditions, the analysis errors can be described by the analysis error covariance matrix, **A**. Equation 2.8 demonstrates how the background and observation error covariance matrices contribute to the analysis error covariance matrix.

$$\mathbf{A} = (\mathbf{I} - \mathbf{KH})\mathbf{B}(\mathbf{I} - \mathbf{KH})^T + \mathbf{KRK}^T \tag{2.8}$$

When the observation error covariance matrix and the observation operator are known exactly, Equation 2.8 simplifies to Equation 2.9 [7]. This more clearly demonstrates the impact of the Kalman gain matrix and the observation operator on the analysis, relative to the background error covariance matrix.

$$\mathbf{A} = (\mathbf{I} - \mathbf{KH})\mathbf{B} \tag{2.9}$$

In practical applications, it can be difficult to validate the assumptions listed above. Section 2.5 explains that large biases are known to exist between sea ice forecasts and observational datasets. Section 2.4 describes why for sea ice thickness it may not be reasonable to assume that the background and observation errors are uncorrelated. Stewart et al. [13] and Desroziers et al. [19] explain that many operational assimialtion systems do not correctly estimate the observation error covariance matrix. When any of these assumptions are not met, the analysis is effectively suboptimal and Equation 2.8 might not accurately describe the errors in the state estimate [7][13][14].

There exists a spectrum of sequential assimilation techniques that follow different approaches for estimating and updating the background error covariance matrix between cycles [7]. Optimal Interpolation (OI), for instance, makes the assumption that the background error covariance

matrix is static in time [20]. This is the simplest method and it is still used in some prediction systems today (e.g. [21]) but it is acknowledged that OI analyses may actually be suboptimal [21]. If the geophysical system is linear, the Kalman filter [22] can be used to propagate the background error covariance matrix through time [18]. However, most geophysical systems [16], including sea ice [23], are described by nonlinear equations and the Kalman filter is not well-suited to these systems [7][24]. More complicated techniques, such as the Ensemble Kalman filter (EnKF) [24], described in Section 2.2, can account for nonlinearities in the system dynamics at greater computational expense.

The estimation of the observation error covariance matrix has received comparatively less attention than the background error covariance matrix [13]. Sections 2.5 and 2.6 describe how many, if not all operational sea ice data assimilation systems blindly assume that the observation error covariance matrix is diagonal. This is largely due to challenges in estimating the observation error correlation structure [13][19].

An important challenge with the design of any new data assimilation system is to select an assimilation technique that yields acceptable analysis quality at an acceptable cost. Section 2.5 describes some of the design decisions that were made for sea ice prediction systems.

Note that only sequential methods are decribed above. There are also related variational methods, e.g. 3DVar [14][21], that approach data assimilation as an optimization problem. The variational approach is to find the state vector that minimizes a cost function, not presented herein. The variational problem can converge to the sequential analysis under certain circumstances [7]. A comparison of variational and sequential data assimilation is provided in [7]. We focused on sequential data assimilation because most of the sea ice thickness data assimilation experiements, summarized in Section 2.5, use sequential techniques.

2.2 The Ensemble Kalman Filter

The EnKF [24] and variants of the EnKF are applied in many of the most recent sea ice data assimilation experiments (e.g. [11]), which are described in detail in Section 2.6. An ensemble-derived background error covariance matrix is also used in the twin experiments described in Chapters 4 and 5. For these reasons, this section provides an introduction to the EnKF.

The EnKF uses an ensemble (a set) of model states to represent the distribution of background errors [24]. The advantage of this approach is that it can provide an accurate representation of the background errors, including any inter-variable correlations [7][24]. Equation 2.10 [24] demonstrates how the background error covariance matrix can be estimated using an ensemble of model states.

$$\mathbf{B} = \frac{1}{n_{ens}-1}\sum_{i=1}^{n_{ens}}(\mathbf{x}_i^b - \overline{\mathbf{x}^b})(\mathbf{x}_i^b - \overline{\mathbf{x}^b})^T \tag{2.10}$$

and where $\mathbf{x}_i^b$ represents one of n_{ens} model states in the ensemble and where the mean state is represented by

$$\overline{\mathbf{x}^b} = \frac{1}{n_{ens}}\sum_{i=i}^{n_{ens}}(\mathbf{x}_i^b) \tag{2.11}$$

During the analysis step, the observations are assimilated into each model state separately [7]:

$$\mathbf{x}_i^a = \mathbf{x}_i^b + \mathbf{K}(\mathbf{y}_i - H(\mathbf{x}_i^b)) \tag{2.12}$$

Note that the analysis also requires an ensemble of observations, $\mathbf{y}_i$. Typically there are not enough simultaneous observations to provide an ensemble-based estimate of the observation error covariance matrix [24]. Instead, the observations are perturbed for each ensemble member by sampling random errors from a Gaussian distribution with spatial correlations following a prescribed observation error covariance matrix [24].

Propagating forward each model state in time using the numerical model conveniently also marches forward the estimate of the background error covariances. Thereby the EnKF provides a reasonably inexpensive dynamic approximation of the background error covariance matrix.

The implementation of an EnKF presents several challenges. First is the choice of n_{ens}, the size of the ensemble. In most operational systems (e.g. [8]), fewer than fifty model states are used due to computational cost, with the drawback that the distribution of background errors may not be very well characterized. Error covariance *localization* is often performed (e.g. [12]) to filter out spurious long-range correlations that may be present. The second added challenge is that stochastic numerical models must be used in place of fully deterministic numerical models [7]. Without a stochastic component, the ensemble members can converge to a single state, implying overconfidence in the background state estimate and possibly leading to numerical errors [24]. The stochastic component requires careful tuning to ensure that there is sufficient spread between ensemble members [24].

background error covariance matrix provides a realistic approximations of the actual background errors.

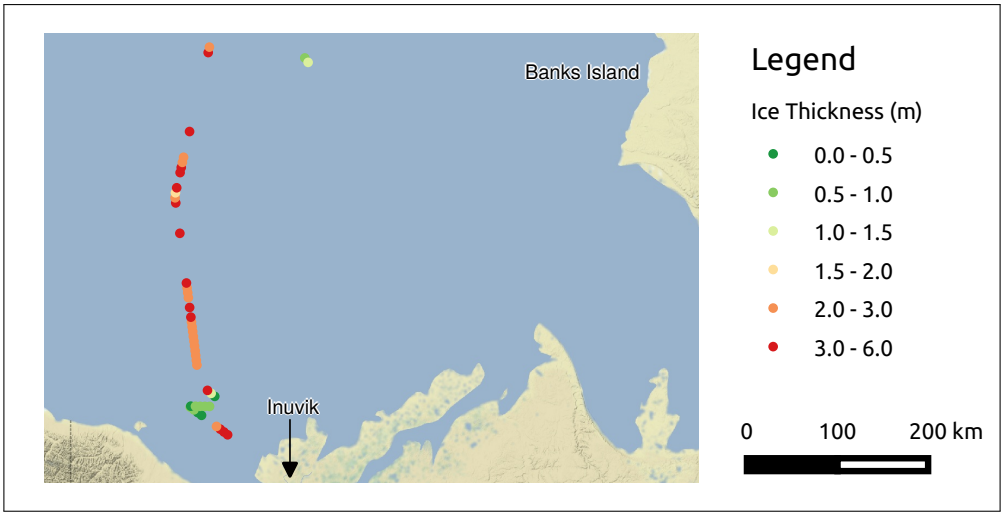

Figure 2.2: Route from a sea ice thickness AEM survey that occurred April 20, 2015. The survey began in Inuvik, NWT, and traversed approx. 400 km of the Beaufort Sea. The observed ice thicknesses were binned at a 5-km and are plotted by colour. This figure demonstrates the spatial distribution of thickness and the sparsity of observations that can be achieved from an AEM survey.

2.3 The Ice Thickness Distribution

This section presents observations from a recent sea ice field survey in an effort to describe the sea ice thickness distribution and how it is represented in numerical models and assimilation systems.

The field campaign was conducted on April 20th, 2015, and was funded by the Marine Environment Observation Prediction and Response (MEOPAR) network. The program was led by Dr. Christian Haas' research team at York University, with help from the University of Waterloo's Department of Systems Design Engineering. To measure sea ice thickness, an Airborne electromagnetic (AEM) sensor system mounted beneath a Basler BT-67 airplane [25] was flown North from Inuvik, Northwest Territories, over the Beaufort Sea. The resulting dataset was a linear track of ice plus snow thicknesses at a resolution of ~ 8 m, with a nominal thickness error standard deviation of 0.1 m [25]. Figure 2.2 illustrates the flight path with 5-km average thicknesses. Figure 2.3 depicts a typical scene from the airplane at an altitude of approximately thirty meters. Finally, Figure 2.4 illustrates the distribution of thicknesses observed during the survey.

In Figure 2.2, the 5-km mean sea ice thicknesses are fairly consistently 2-3 m. However, in Figure 2.3, there appears to be a variety of mixed thin and thicker ridged ice alongside a

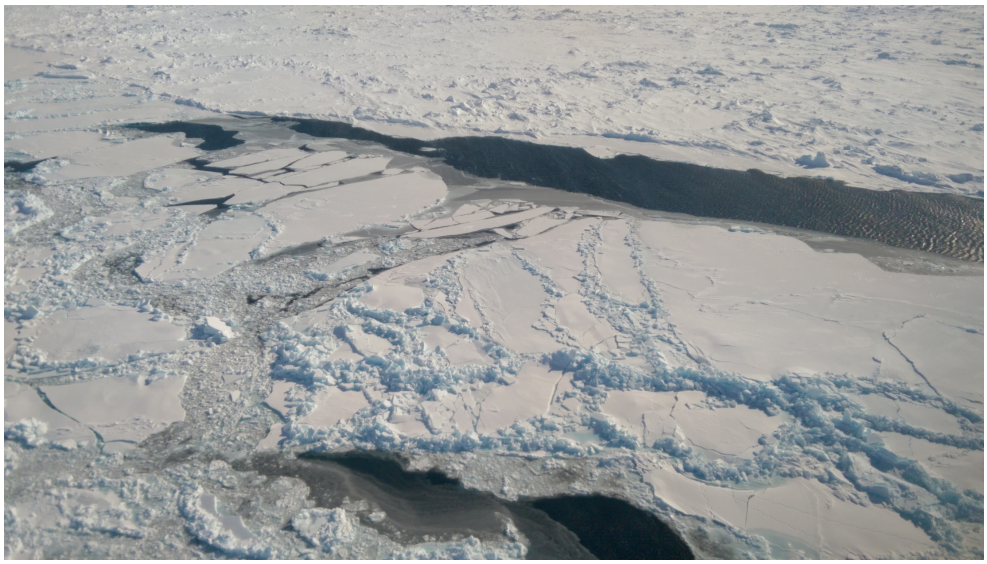

Figure 2.3: Photograph of sea ice in the Beaufort Sea, April 20, 2015. Depicted is a prominent lead in first-year ice. Within the lead is a mixture of open water, rubble, and new ice.

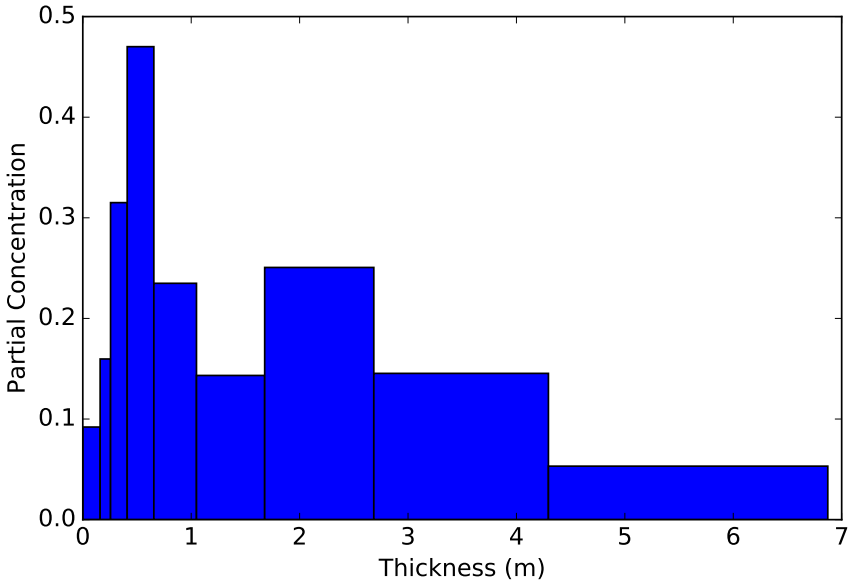

Figure 2.4: Histogram of sea ice plus snow thickness in the Beaufort Sea, April 20, 2015. The thicknesses are binned into categories that are typical of sea ice models, with more categories representing thinner ice.

prominent open water lead.

Considering the variability in Figure 2.3, what is the best way to represent ice thickness in a numerical model? Most numerical sea ice models (e.g [2], [26]) treat sea ice as a continuum, rather than modelling every segment of ice individually. In these models, ice is discretized not only in space and time but also in thickness space. Thickness is thereby represented as a histogram with anywhere from two to twenty categories, where each category has a mean thickness value and partial concentration. Figure 2.4 depicts the ice thickness histogram for the April 2015 ice thickness survey, with bins that represent typical ice thickness catagories.

In Figure 2.4, there are nine thickness bins. If these bins were used in a model, the model would be considered a ten-category model, since there is an additional category representing the concentration of zero-thickness ice, or open water. Notice that the mean and variance of the thickness poorly describe the distribution of thicknesses. Firstly there is a long tail that represents highly ridged ice. Secondly, there may be two separate peaks, representing the combination of thicker first/multi-year ice and newer ice that were observed during the survey.

In the context of data assimilation, special consideration needs to be given to the ice thickness distribution. A ten-category model increases the state size by one order of magnitude and the background error covariance by two orders of magnitude, as compared to a two-category model. Furthermore, there is an inherent discrepancy between thickness (units of meters) and thickness partial concentration (unitless). Section 2.6 describes how this problem is currently being approached in experimental sea ice data assimilation systems.

2.4 Remote Sensing of Sea Ice Thickness

Sea ice thickness can be measured reliably in the field by drilling through the ice with an auger and measuring tape [27]. Reasonably accurate estimates of sea ice thickness can also be obtained by airborne surveys [25] and submarine-based field programs [28]. Unfortunately, very few field surveys are conducted each year and airborne and ULS data are more frequently used for validation (e.g. [29]) than for assimilation. Until we can develop a field-based sensor network with sufficient spatial coverage and reliability, there is perhaps greater value in assimilating data collected by spaceborne sensors, which can provide better spatiotemporal coverage of the Arctic [7][9].

The trade-off with space-based observations of ice thickness is that there is greater uncertainty with the data quality [9]. To be clear, it is not actually possible to directly observe sea ice thickness remotely. Satellite sensors detect more basic electromagnetic signals that must be combined with ancillary data through a statistical or physically-based thickness model. Errors

Table 2.1: Satellite Sensors for Sea Ice Thickness Estimation

Sensor	Satellite	Type	Sensor Footprint (km)
SIRAL	CryoSat-2	Radar Altimeter	0.3*1.5
MODIS	Aqua	VisIR	1
VIIRS	Suomi NPP	VisIR	0.75
MIRAS	SMOS	Passive Microwave	35
AMSR-2	AMSR-2	Passive Microwave	14-62

in the ancillary data, model inaccuracies, and differences in spatial resolutions all contribute to errors in the thickness estimates.

There are three standard approaches for estimating sea ice thickness using observations from spaceborne sensors. The first approach is based on altimetry and Archimides principle [3]. The second approach is based on Visible-Infrared (VisIR) imagery and a thermodynamic model [30]. The third approach is based on passive mircowave imagery and an emissivity model [31][32]. Table 2.1 lists a selection of active satellites in the three sensor categories, along with the diameter of the sensor footprint. The sensor footprint is the area projected onto the Earth's surface that contributes to the observed signal. Passive microwave sensors typically produce imagery with a finer resolution than the sensor footprint. For instance, the SMOS sea ice thickness product has a spatial resolution of 12.5 km and a sensor footprint of 35 km [33].

Each of the satellites in Table 2.1 are in polar low-earth orbits. It is therefore possible to expect daily revisit times at most temperate latitudes, with increased frequency near the poles. The exception is CryoSat-2, which has a fine beam width and can take months to revisit the same location [3].

Table 2.2 lists the ancillary data required to produce a sea ice thickness estimate with each sensor type. In general, satellite sensor data comprise only a fraction of what is required to estimate ice thickness. Some of the required variables, like snow density, ice density, and ocean salinity, are often assumed to be constants [30]. Snow thickness is often specified as a piecewise function of the ice thickness [30]. Atmospheric variables can be interpolated from meteorological forecasts. Ice concentration estimates can be taken from remote sensing estimates of ice concentration or from the state of a sea ice model. In general, remote-sensing-based estimation of ice thickness is quite an involved process.

An important note is that each of these remote sensing-based thickness estimation techniques come with caveats on the condition of their use. Cryosat-2, for instance, cannot distinguish ice thicknesses less than 1 m because the freeboard (elevation of the ice above the ocean surface) is

Table 2.2: Required ancillary data for satellite sensor-based estimation of ice thickness

Sensor Type	Required Ancillary Data
Altimeter	Ice density, snow density, snow thickness, and ocean surface elevation [3]
VisIR	Albedo (for daytime retrievals), ice surface temperature, ice density, ice concentration (assumed to be 100%), cloud cover fraction, snow density, snow depth, air density, relative humidity, 2 m wind speed, 2 m air temperature, surface air pressure, and ocean salinity [30]
Passive Microwave	Snow thickness, ice temperature, ice salinity, ice concentration (assumed to be 100%), and the shape of the ice thickness distribution [32]

near the error standard deviation of the sensor [9]. MODIS-based ice thickness estimates become unreliable above 0.5-1 m [34]. Similarly, AMSR-2-based ice thickness estimates saturate at 0.2-0.3 m and SMOS-based ice thicknesses saturate above 0.5-1.0 m, depending on the sea ice salinity [32]. Another condition is that VisIR-based ice thickness estimates are only reliable for night-time imagery and all three methods are only reliable during freezing conditions, in the absense of melt ponds or wet snow. Finally, SMOS-based thickness observations are prone to contamination from electromagnetic interference [32]. These conditions are all theoretically manageable but there are clearly challenges facing the automated retrieval algorithms that would be required for an operational data assimilation syste.

Validation of satellite-based ice thickness products typically involves comparing the product to a set of field observations from the same date [9]. For SMOS-based ice thickness estimates, relative standard thickness errors are estimated to be 20% for ice < 30 cm and 100% for ice > 1 m thick [32]. This corresponds to error standard deviations of of 0.06 and 1 m for the respective thicknesses. Estimates of error standard deviations for MODIS-based ice thickness estimates range from 0.15 to 0.5 m for ice thicknesses ranging from 0.05 to 0.8 m, and only for air temperature $< -20C$ and windspeeds < 5 m/s [34]. These error estimates can be several times larger for warmer air temperatures and greater windspeeds [34]. CryoSat-2 errors are more consistent, with an error standard deviation of 0.1 m when the snow and ice density are known [32].

In a comparison of sea ice thickness products from six different sources in the Arctic Ocean, Wang et al. [6] demonstrated that there are often large biases between thickness products. For

instance, thicknesses from the satellite sensor AVHRR (Vis-IR) were 0.19 m less than the corresponding thicknesses from CryoSat-2, in regions where both sensors predicted ice thinner than 1 m [6]. Compared to a set of AEM survey data, AVHRR and CryoSat-2 both had positive biases of 0.18 and 0.29 m, respectively. The authors investigated the spatial correlation of errors but only qualitatively [6]. From a visual inspection of the thickness maps the thickness error correlations appeared to often span hundreds of kilometers.

We can conclude that it *is* possible to produce remote sensing-based sea ice thickness observations. It may also be also be possible to estimate the observation error variances, at least under certain conditions. However, in an operational assimilation system, we lack he ability to estimate the observation error covariance structure.

2.5 Operational Sea Ice Assimilation in Canada

ECCC has two operation sea ice forecasting systems: Regional Ice Prediction System (RIPS) [21] and Global Ice Ocean Prediction System (GIOPS) [5]. RIPS provides 48-hour forecasts at 5-km resolution in Canadian coastal waters [21]. The objective for this system is to support the Canadian Ice Service (CIS) in its mandate to support safe naval operations in Canadian waters. GIOPS provides 10-day forecasts on a coarser 10-km grid [5]. GIOPS forecasts are used as input for numerical weather prediction (NWP) systems.

GIOPS uses the Los Alamos sea ice model (CICE) coupled to the Nucleus for European Modelling of the Ocean (NEMO) ocean model [5]. RIPS also uses CICE, forced by an ocean state interpolated from GIOPS [21]. CICE is a multicategory ice model, meaning that ice thickness is represented by up to ten categories, where each category is represented by a thickness value and a partial concentration [2]. Atmospheric forcing data are interpolated from ECCC's Regional and Global Deterministic Prediction Systems (RDPS and GDPS).

GIOPS has two separate data assimilation systems. The first system assimilates remotely-sensed sea surface temperature (SST) and sea surface anomaly observations, along with field-based observations of the same variables, into the ocean model state using a Singular Evolutive Extended Kalman (SEEK) filter [5]. The second system uses 3DVar to assimilate ice concentration observations from passive microwave satellite sensors and from manually-classified ice charts from the CIS into the sea ice model. RIPS also uses 3DVar to assimilate sea ice concentration observations from passive microwave sensors and CIS ice charts [21]. Both systems are now moving toward the ensemble-based EnVar [8].

Both sea ice concentration assimilation systems share a few specific details [5][21]:

- Passive microwave concentration observations are rejected when the air temperature is greater than zero and when the SST is above $-4^{\circ}C$;
- Since the passive microwave observations have a much larger footprint size than the model resolution, a many-to-one linear observation operator (a footprint operator) is used to interpolate the background concentrations to the observations;
- The analysis occasionally produces concentrations greater than one or less than zero, which is non-physical. These concentrations are truncated to the range [0,1];
- The observation error covariance matrix is diagonal but the error variances are inflated by a factor greater than one to account for the lack of off-diagonals; and
- Sea ice concentration is the only variable in the background state.

The reason that GIOPS and RIPS both use a diagonal observation error covariance matrix is two-fold. Firstly, ECCC's variational data assimilation system requires the inverse of the observation error covariance matrix, which becomes too poorly conditioned to invert directly when non-zero spatial covariances are included [21]. The second reason is that it is simply too challenging to estimate the actual error covariance structure in an operational setting. Inflating the observation error variances ensures that the assimilation system does not underestimate the analysis error covariances.

In GIOPS, the ice state is represented by ten partial concentrations of ice thickness, yet only the total concentration is used in the analysis [5]. After updating the total concentration with 3DVar, the partial concentrations must be adjusted. However, linearly scaling each partial concentration by the change in total concentration is not reflective of the physical processes that led to the errors. Instead, a more sophisticated redistribution scheme is used. The details of the redistribution scheme are beyond the scope of this thesis but they are described in detail in [5].

In RIPS and GIOPS, the root-mean-square errors (RMSEs) for ice concentration are estimated to be approximately 0.2-0.3 (20-30%), based on comparisons with manually-classified Synthetic Aperture Radar image analysis datasets [4][5]. Typical biases were ± 0.1 (10%), although bias is not handled explicitly by the system [4]. Thickness error estimates were not reported for either system. These error estimates provide a benchmark upon which we might be able to improve through ice thickness assimilation.

2.6 Experimental Sea Ice Thickness Assimilation Systems

Sea ice thickness is not currently being assimilated operationally at any NWP centre. However, there is active experimental research being conducted to support the development of such systems. This section describes three distinct areas of research in this field.

2.6.1 Seasonal Forecast Initialization Using Field Observations

The quality of numerical sea ice models degrades severely during the summer when higher temperatures prevent passive microwave sensors from producing reliable sea ice concentration observations. Late summer is also the period when it is possible for ships to travel through the Northwest Passage [1]. There is therefore great value in ensuring that seasonal forecasts are initialized with the best possible model state.

In 2012, Lindsay et al. [35] used the thickness products from twelve AEM surveys to help set the initial conditions for a seasonal forecast spanning the summer months. The AEM surveys covered the Beaufort and Chuchki Seas, collecting ice thickness observations at fourty-meter resolution. With this high-resolution field data, it was possible to estimate the partial concentrations for each thickness category [35]. Lindsay et al. used OI to combine the partial concentration estimates with the June 1, 2012 state of PIOMAS. The result was an improvement in the predicted September total ice extent, as compared to a control forecast.

While field surveys provide accurate observations that are well-suited to assimilation into ice models, these are not cost-effective for operational use. The AEM surveys used in [35] did not span the entire arctic ocean, nor did they cover important regions such as the Canadian Arctic Archipelago. A pan-arctic sea ice survey has never been done and such an effort cannot be expected every year. That said, several regional AEM surveys are conducted each year and these might be used effectively for tasks like validation of satellite-based ice thickness products and for assisting with regional forecasts.

In a similar study using an ensemble of climate models, Day et al. [36] showed that exactly knowing the initial thickness state could have impacts extending up to eight months into the future. While it is not possible to exactly determine the thickness state, this investigation set the upper bound on potential long term impacts from ice thickness assimilation.

2.6.2 Twin Experiments with Synthetic Observations

Lysaeter et al. [11] conducted perhaps the first sea ice thickness data assimilation experiment. In this experiment, the authors ran two sets of two-category arctic sea ice models in parallel.

Observations were sampled from one model state using a synthetic sensor that simulated the observations that CryoSat would produce. The sampling errors were prescribed to be spatially uncorrelated, but the error variances were a function of the ice concentration and thickness in the model state. These observations were then periodically assimilated into an ensemble of other model states using an EnKF. In their assimilation system, the background state vector included not only mean ice thickness but also concentration, ocean salinity and SST.

The authors demonstrated that biweekly assimilation of CryoSat data reduced the RMSE from 0.8 m to 0.4 m [11]. Improvements to SST and ice concentration were also observed, although these improvements only occurred at the ice edge (the marginal ice zone). While this experiment was overall somewhat idealistic, it demonstrated and quantified the potential benefits of routine sequential sea ice thickness assimilation.

2.6.3 Quasi-Operational Systems

In the most recent sea ice thickness data assimilation experiments, researchers are sequentially assimilating sets of satellite-based sea ice thickness observations into operational-quality coupled ice-ocean models forced with atmospheric reanalysis data. The analyses and forecasts aren't being used directly by industry, but insights from these experiments are applicable to the design of next-generation NWP systems.

Yang et al. [12] investigated the impacts of assimilating SMOS-based ice thickess observations along with passive microwave sea ice concentration observations. The authors used the daily SMOS-based thickness product that is produced by the University of Hamburg [33]. These data are subject to rigorous quality control and are provided with an estimate of the observation uncertainty. This provides a convenient starting point for operational data assimilation because the observation uncertainties can be used to estimate the observation error variances.

In their study, Yang et al. used a two-category sea ice model coupled to a multi-level ocean model [12]. The ice model is two-category, meaning that the only thickness variable is the mean ice thickness. From a thickness assimilation perspective, this is convenient, although the numerical model may be less accurate. The authors used a background state that included the thickness, concentration, and ocean prognostic variables.

Yang et al. used a sophisticated Local Singular Evolutive Interpolated Kalman (LSEIK) filter, which is an extension of the EnKF, for their analyses [12]. The authors described in detail the construction of the background error covariance matrix, which is generated using the ensemble of model states and then localized to a decorrelation length of 125 km. The observation operator implemented linear interpolation to map the model thicknesses to the observation locations. The

observation uncertainties provided with the SMOS thickness dataset were used to create observation error covariance matrix, which was specified to be diagonal. Thickness observations greater than 1 m were given an error variance of 5 m to prevent these observations from impacting the analysis.

The authors reported favourable results, with root-mean-squared differences (RMSDs) between the SMOS thickness and model thickness decreasing from 0.8-0.85 to 0.35-0.54 m over the course of a season in the first-year ice zone [12]. The RMSD in the early winter was less than the RMSD at the end of the growth season, in part because there was more thin ice early in the season. An important note is that 80% of the grid points experienced negative increments during the assimilation, suggesting that the analysis was in fact largely just a bias correction.

SMOS-based observations will likely be the first thickness observations assimilated sequentially. Despite the relatively coarse 35 km sensor footprint size and 12.5 km spatial resolution, SMOS is able to produce pan-Arctic daily coverage, unlike CryoSat-2. Furthermore, SMOS data is less sensitive to weather and light conditions than VisIR-based observations. There is also increased confidence in the data now that a daily SMOS-based thickness product is being produced. The efforts from these investigations confirm that the assimilation of SMOS-based observations has the potential to improve sea ice model state quality during the ice growth season and at the ice-water boundary.

2.7 Synthesis

The NWP community is nearing the point where it is possible to assimilate satellite-based sea ice thickness observations into coupled ice-ocean models. 3DVar and/or EnKF systems ingesting SMOS-based thickness data will likely form the first operational systems.

Despite this progress, there remain several barriers to the immediate implementation of a sea ice thickness data assimilation system. Ice thickness biases in both the model states and in the observations are known to exist but neither operational nor experimental NWP systems can adequately correct this. Remote sensing-based thickness estimates are also notoriously poor and little has been done to quantify ice thickness observation error covariances. It seems likely that the first operational prediction systems that will assimilate ice thickness observations will do so using diagonal approximations to the observation error covariance matrix. If short-term sea ice forecasts are to play a role in ship navigation and other safety-critical operational activities in Canada's North, there is a demonstrated need to understand how the state of a numerical model will change when a diagonal approximation to the observation error covariance matrix is used in ice thickness data assimilation.

Chapter 3

A One-Dimensional Dynamic-Thermodynamic Sea Ice Model

Experimental data assimilation methods are often tested on lower-dimensional models before moving to full-scale prediction systems. For instance, the development of the Ensemble Kalman Filter (EnKF) was supported by experiments with the Lorenz equations and the shallow water equations [24]. Using these simpler tools it is possible to rapidly and inexpensively screen hypotheses and quantify the impact of alternative data assimilation methodologies. While the simplified models do not exactly reproduce the dynamics of the geophysical system, they can be used effectively to reveal insights that may be applicable to larger systems.

The basis of this thesis is a simplified sea ice model, coupled to simplified models of the atmosphere and ocean. The sea ice model implements a two-category representation of sea ice in one-dimension. The model is not bound to any particular coastlines or latitude and can be configured to mimic the evolution of a sea ice state in any season. Starting with a set of random initial conditions, the model provides a random simulation of the ice pack. The model is referred to herein as the *toy model.*

This chapter describes the design and validation of the toy model. The toy model was designed to meet two requirements: first, that it be inexpensive and suitable for use in twin data assimilation experiments on a desktop computer; and second, that it produces a dynamic sea ice state with thickness and velocity distributions that are similar to observational datasets. This second feature of the toy model, the realistic simulated forcing, is a novel contribution that provides a different perspective than other one-dimensional ice models, which are more commonly used for testing and validating new model processes and algorithms [2][26]. The use of the toy model in a twin experiment is described in Chapter 4.

3.1 State Variables

Hibler [23] developed one of the first dynamic-thermodynamic sea ice models. Hibler's state variables were concentration (a), thickness (h), and the meridional and zonal velocities (u and v). This represents a *two-category* model, where a given model cell has a fraction a of *thick* ice with thickness h and a fraction $(1-a)$ of either open water or very thin ice that that has minimal effect on the model processes. While multi-category models are more realistic, the two-category model was considered better suited to the scope of this study.

There can be many additional state variables in addition to concentration, thickness and velocity. Examples include albedo, snow density, snow thickness, ice density, ice temperature, ice salinity, and melt pond area. Full scale multi-category models such as the Los Alamos Sea Ice Model (CICE) [2] and the Louvain-la-Neuve Ice Model (LIM) [26] propagate each of these variables at increased expense and complexity. For the purposes of the toy model, we chose not to include any more variables than Hibler modeled. Instead, variables such as the ice density were set to be constant parameters. We also removed the second velocity component, v, effectively making the model one-dimensional.

3.2 Simplified Sea Ice Equations

In the continuum representation, sea ice behaves like a viscous plastic that is forced mechanically by the wind and ocean currents. The balance between internal and external stresses is represented in the momentum equation [23]:

$$\rho_{ice} ah \frac{Du}{Dt} = -\rho_{ice} ahfk \times u + \tau_a + \tau_w - \rho_{ice} ahg\nabla H + F \tag{3.1}$$

where ρ_{ice} represents the density of the ice, k is a unit vector, f is the Coriolis parameter, τ_a represents the external stress caused by the wind, τ_w represents the external stress caused by shallow ocean currents, ∇H represents the surface gradient of the ocean free surface, and F represents the internal stresses within the ice sheet. Note that this equation is two-dimensional, and only a one-dimensional implementation is considered herein.

A scaling exercise based on field observations in the Baltic Sea [37][38] revealed that the atmospheric, oceanic, and internal stresses are typically orders of magnitude larger than the other terms for ice conditions with a mean thickness of approximately 1 m. Retaining these terms, constraining dynamics to one dimension, and expanding the source terms we obtain a

simplified momentum equation:

$$\rho_{ice} ah \frac{\partial u}{\partial t} = \rho_a a C_a u_a |u_a| + \rho_o a C_o (u_o - u)|u_o - u| + \frac{\partial}{\partial x}\left(\zeta \frac{\partial u}{\partial x}\right) \tag{3.2}$$

where C_a and C_o represent the quadratic drag coefficients for the wind and ocean, respectively; u_a and u_o represent the wind and ocean velocities, respectively; ρ_a and ρ_o represent the wind and ocean densities, respectively; and ζ represents the bulk viscosity of the sea ice. The drag coefficients are typically set to constant values based on field observations, but they can be parameterized based on the ice roughness. The air and ocean velocities are typically interpolated from ocean and atmospheric models that are being integrated in parallel (e.g. [5]).

The viscous-plastic ice rheology produces a viscosity term that is a nonlinear function the ice concentration, thickness, and velocity. The equations describing the bulk viscosity in one dimension are [23]

$$\zeta = \frac{P}{2\Delta} \tag{3.3}$$

$$\Delta = \left|\frac{\partial u}{\partial x}\right| (1 + e^{-2}) \tag{3.4}$$

$$P = P^* h \exp(-C(1-a)) \tag{3.5}$$

where P represents the ice strength. P^*, C, and e are empirically derived/tuned constants. Note that the ice strength is linearly proportional to thickness and exponentially proportional to concentration. To avoid singularities, the viscosity is constrained within the range of $\zeta_{min} = 4 * 10^8$ kgs^{-1} and $\zeta_{max} = 2.8 * 10^8 P$ [23]. At the minimum viscosity, ice forms flows that behave like a solid. At higher viscosities, ridging or fracturing occurs.

In addition to the momentum equation, there are transport equations for both thickness and concentration [23]:

$$\frac{\partial a}{\partial t} = -\frac{\partial ua}{\partial x} + S_a \tag{3.6}$$

$$\frac{\partial h}{\partial t} = -\frac{\partial uh}{\partial x} + S_h \tag{3.7}$$

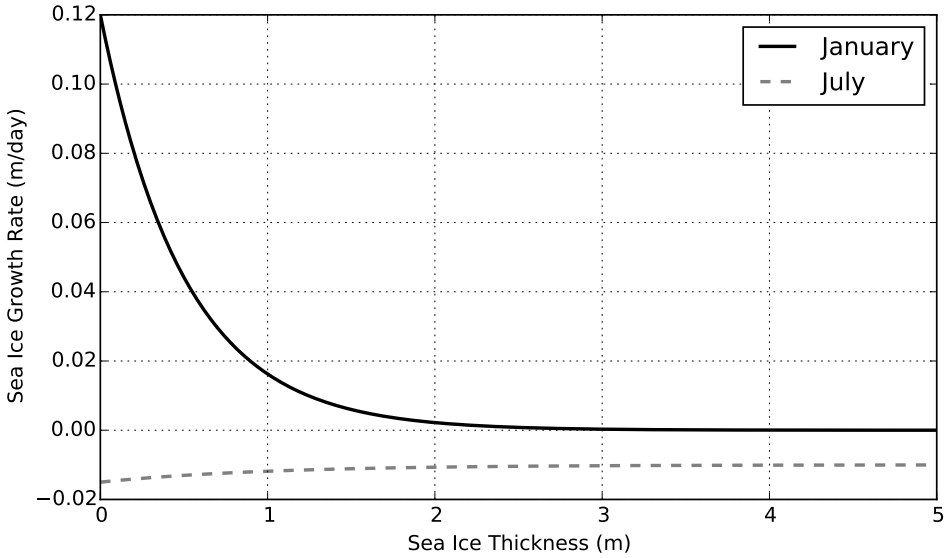

Figure 3.1: Empirical functions for the seasonal sea ice growth rate in the central Arctic [23][39]. The growth rate function $f(h)$, found in equations 3.8 and 3.9, is greatest for thin ice, decreasing with increasing ice thickness. In the summer, the growth rate is negative.

where S_a and S_h are thermodynamic terms given by

$$S_h = f(h/a)a + (1-a)f(0) \tag{3.8}$$

$$\begin{aligned} S_a = &\begin{cases} f(0)(1-a)/h_0, & \text{if} \quad f(0) > 0, \\ 0, & \text{if} \quad f(0) < 0 \end{cases} \\ &+ \begin{cases} 0, & \text{if} \quad S_h > 0, \\ 0.5ah^{-1}S_h, & \text{if} \quad S_h < 0 \end{cases} \end{aligned} \tag{3.9}$$

where $h^0 = 0.1$ m is a fixed demarcation between thin and thick ice, and $f(h)$ is the growth rate. Hibler [23] used an exponential function of thickness for growth rate based on field experiments conducted by Thorndike et al. [39]. Figure 3.1 illustrates this growth function. From Figure 3.1 and equations 3.8 and 3.9, we can see that winter growth rates are considerably higher when the ice is thin and has low concentration. When the ice is thicker or the concentration is 100%, there is little growth, reflecting the insulative properties of ice thickness.

The January and July growth rate curves depicted in Figure 3.1 may represent more extreme growth rates than might be observed at lower latitudes or at other times during the year. Furthermore, there may be daily fluctuations in growth rate from the monthly average rate. To allow

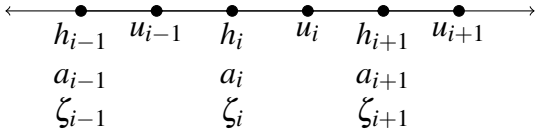

Figure 3.2: The staggered grid for the toy sea ice model. The velocities are offset by one half-grid cell from the other variables.

for intermediate growth rates, the toy model was configured so that the January growth function $f(h)$ can be multiplied by a growth rate scaling factor.

Note also that CICE and LIM use a prognostic thermodynamic model for simulating heat transfer through the ice and snow. While this adds a level of realism, it requires additional state variables and input from sophisticated ocean and atmospheric models. By using Hibler's empirical growth rate function, we preserved the simplicity of the toy model while not ignoring thermodynamics entirely.

3.3 Numerical Solution

The numerical solution to the sea ice momentum equations is an active area of research (e.g. [40], [41]). Since ice velocity is a function of the ice viscosity, and viscosity is a function of the ice velocity, the momentum equation is highly nonlinear. Viscosity can take on values spanning two orders of magnitude, effectively representing the sea ice as a fluid at low viscosities and a solid at high viscosities. The solution of the toy model follows the mixed implicit-explicit approach described by Lemieux et al. [40]. In this approach, the momentum equation is solved implicitly at each time step, while concentration and thickness equations are solved explicitly.

Algorithm 1 [40] describes the procedure for advancing the sea ice state at time t. The outer k iterations are required for the ice state to reach a stable viscosity ζ and velocity u. The use of u_l^{k-1} instead of simply u^{k-1} damps the increment, leading to better convergence. Following this approach, a time step of six hours can be used accurately on a 10-km grid [40]. Lemieux et al. [40] found that a reasonable level of convergence can be achieved in up to 5 iterations of the outer loop. Our selection of ten outer loop iterations was conservative and we did not observe visible differences in ice states generated with either fewer or more iterations.

The model domain is a one-dimensional periodic grid spanning 1000 km with a grid spacing of 1 km. While this represents a higher resolution than RIPS (5 km), the grid may be reflective of next-generation ice prediction systems. The grid is staggered in such a way that the velocity u_i is stored one-half grid cell away from the concentration, thickness and viscosity. Figure 3.2

Data: u^{t-1}, a^{t-1}, h^{t-1}, u_o^t, u_a^t, $f(h)^t$
$u^0 = u^{t-1}$;
$u^1 = u^{t-1}$
for $k = 2 \rightarrow 10$ **do**
    Compute bulk viscosity ζ;
    $u_l^k = \frac{u^{k-2}+u^{k-1}}{2}$;
    Solve the momentum equation (3.2), substituting u_l^{k-1} for u^{t-1};
    Result: u^k
end
$u^t = u^k$;
Solve the transport equations (3.6 and 3.7) to update a and h;
Result: u^t, a^t, h^t

Algorithm 1: Numerical Solution to Sea Ice Equations [40]

illustrates the grid. A forward-in-time, centered-in-space finite difference discretization of the momentum equation (3.2) yields:

$$\alpha_1 u_{i-1}^t + \alpha_2 u_i^t + \alpha_3 u_{i+1}^t = b \tag{3.10}$$

where

$$\alpha_1 = \frac{1}{(\Delta x)^2}\zeta_i$$

$$\alpha_2 = \frac{1}{(\Delta x)^2}(\zeta_i + \zeta_{i+1}) - \left(\frac{\rho_i}{2\Delta t}\right)(a_{i+1}h_{i+1} + a_i h_i) - \frac{\rho_w C_w}{2}(a_i + a_{i+1})\left|u_{w_i} - u_i^{t-\Delta t}\right|$$

$$\alpha_3 = \frac{1}{(\Delta x)^2}\zeta_{i+1}$$

$$b = -\frac{\rho_{ice}}{2\Delta t}(a_{i+1}h_{i+1} + a_i h_i)u^{t-\Delta t} - \frac{\rho_w C_w}{2}(a_i + a_{i+1})u_{w_i}\left|u_{w_i} - u_i^{t-\Delta t}\right| - \frac{\rho_a C_a}{2}(a_i + a_{i+1})u_{a_i}|u_{a_i}|$$

Evaluating these coefficients at every grid point produces a series of equations:

$$\mathbf{Au} = \mathbf{b} \tag{3.11}$$

Solving the system of momentum equations (3.11) can be done using any number of numerical solvers. Lemieux et al. [40] used the *Generalized Minimum RESidual* (GMRES) solver. With

our coefficient matrix **A** being only 1000x1000, it was faster to solve this system of equations directly.

Once the ice velocities are computed, the concentration and thickness can be advected. This was accomplished using the Adams-Bashforth-3 time-stepping scheme, which is second-order accurate in time [42]. For ice thickness:

$$h^t = h^{t-1} + \Delta t \left(\frac{23}{12} f_h^t - \frac{4}{3} f_h^{t-1} + \frac{5}{12} f_h^{t-2} \right) \tag{3.12}$$

where f_h^t represents the thickness flux at time t, given by

$$f_h^t = -\frac{1}{2\Delta x} \left(h_i u_i + h_{i+1} u_i - h_{i-1} u_{i-1} - h_i u_{i-1} \right) + S_h \tag{3.13}$$

Substituting concentration a for thickness h in Equations 3.12 and 3.13 yields equivalent equations for advecting concentration.

The default time step for the toy model was set to 0.5 hours. Based on experimentation with the model, this was found to be an approximate lower bound on stability. Note also that under compressive stresses, ice concentration can exceed 100%, which is non-physical. To limit the concentration at 100% and conserve ice mass, a *redistribution* stage is necessary. With redistribution, the concentration is reduced to 100% and the thickness is increased by a factor that conserves the ice mass. Under diverging conditions, if either the concentration or thickness fall below zero, both variables are set to zero.

3.4 Forcing

The novel component of the toy model is that it is driven by simulated forcing data to produce ice dynamics with a realistic velocity field. Operational ice models (e.g. [21]) are driven by atmospheric datasets taken from full-scale numerical weather prediction models. As there was no equivalent dataset for the one-dimensional toy model, we constructed our own simplified models of the ocean and atmosphere to produce this forcing data.

Single-layer, one-dimensional shallow water models were selected provide the wind and ocean currents that drive the ice model. The shallow water equations are simplifications of the Navier-Stokes equations where vertical pressure gradients are assumed to be hydrostatic and horizontal pressure gradients are caused by displacement of the free surface [16]. These equations were selected to force the ice model because they are relatively simple and because preliminary

experiments revealed that they can produce velocity fields with a wide range of spatial scales and with magnitudes that match observational data.

The governing equations for the shallow water model can be written as [16]

$$\frac{\partial u}{\partial t} = -\frac{\partial}{\partial x}\left(g\eta + \frac{1}{2}u^2\right) \tag{3.14}$$

$$\frac{\partial \eta}{\partial t} + \frac{\partial \eta u}{\partial x} = 0 \tag{3.15}$$

where η represents the fluid depth, u represents the fluid velocity, and g represents acceleration due to gravity.

These equations are discretized and solved on the same grid as the ice model, with velocity and fluid elevation staggered. The same Adams-Bashforth-3 timestepping scheme used for concentration and thickness is used to march forward the velocity and fluid depth. The timestep for the shallow water model was set to one second to ensure numerical stability.

The initial conditions for the shallow water models include a velocity field set to zero. The fluid surface is displaced from its mean elevation by adding a simulated Gaussian Markov Random Field (GMRF) [43]. By displacing the fluid surface, we effectively inject energy across the model domain. The length scale, variance, and spatial covariance model of the random field can be modified to produce different dynamics. All experiments conducted in this thesis used a decorrelation length of 10 km and a Gaussian correlation function. We considered the decorrelation length to represent the distance at which the correlation is 0.05. We selected the decorrelation length of 10 km because it produced velocity features with the widest range of scales but without producing sharp-fronts that led to numerical errors in the velocity field. Similarly, the Gaussian correlation structure was chosen because it prevented numerical errors in the model that we observed with other correlation structures.

The model uses a Fast Fourier Transform (FFT)-based approach for simulating the random surface displacement field [43]:

$$\eta = FFT^{-1}\left(\sqrt{FFT(\mathbf{c})} \odot FFT(\mathbf{w})\right) \tag{3.16}$$

where $\mathbf{c}$ denotes the covariance vector as a function of distance separating the points, $\odot$ represents the Schur product, and $\mathbf{w}$ represents a vector of independantly and identically distributed noise. The Gaussian covariance function was computed as [44]:

$$c(d) = exp\left(-3\left(\frac{d}{r}\right)^2\right) \tag{3.17}$$

Table 3.1: Model Files

File	Summary
`world.py`	Defines classes for Ice, Ocean, and Atmosphere models.
`assimilation.py`	Defines a class for performing Optimal Interpolation.
`main.py`	Main code for running the model.
`main_oi.py`	Conducts a twin experiment using the assimilation class.
`utils.py`	Contains various utility functions (e.g. random field simulation).

Two additional sets of parameters are included to adjust the properties of the shallow water models. First, the velocities were multiplied by a constant to produce more realistic distributions of wind and ocean current velocities. This constant was tuned to produce wind velocities between approximately -15 and 15 m/s and ocean velocities between -0.3 and 0.3 m/s. The timesteps were also multiplied by constants to produce slower wave speeds, in other words decreasing the temporal scale of velocity. These parameters were tuned to produce wind and ocean currents fields with a temporal decorrelation length of approximately ten days. Overall, these shallow water models produced nonlinear wave-like mechanical forcing with regions of mixed convergence and divergence.

3.5 Implementation

The model was written in Python (version 3.5) and can be found at `https://github.com/gastoneb/sea_ice_model`. The model requires the *numpy*, *scipy*, and *matplotlib* libraries. Table 3.1 summarizes the five included Python files and Table 3.2 lists the parameter values that define a standard model run.

Table 3.2: Default model parameter values

Parameter	Units	Default Value
C_a	-	0.0015
C_w	-	0.0015
C	-	20
e	-	2
h_0	m	0.05
P^*	$N\,m^{-1}$	5000
ρ_i	$kg\,m^{-3}$	910
ρ_w	$kg\,m^{-3}$	1035
ρ_a	$kg\,m^{-3}$	1.3
Δx	m	1000
Δt		0.5

3.6 Validation

The generality and significance of the experimental results described later in this thesis were contingent on the sea ice model being reasonably realistic. To validate the model, statistics generated from a long model run were compared to identical statistics gathered from historical field datasets and other model runs (e.g. long-term climate studies). Six measures were generated and compared herein. These measures were selected based on availability of data available for comparison [45].

1. Histograms of wind and ocean velocities
2. Histogram of sea ice velocity
3. Histogram of sea ice divergence, ∂_u/∂_x
4. Temporal autocorrelation function (ACF) of sea ice velocity
5. Spatial ACF of sea ice thickness
6. Histogram of sea ice thickness distribution

To generate these plots, the model was initialized with a thickness of 10 cm, a concentration of 50 %, and growth rate factors sampled from the range $[0.5, 0.9]$. The ice was allowed to

grow for a period of 30 days. Following this, the model was run for another thirty days with a growth factor randomly selected each day from a range of $[-0.05, 0.05]$, representing milder thermodynamic conditions. For this second period, the model state and forcing data were saved at each time step. The saved states were used to create the histograms and ACFs listed above. Figure 3.3 illustrates the model state at an intermediate time step.

Following the sixty-day integration, the mean sea ice thickness and concentration were approximately 0.72 m and 0.95, respectively. Note that greater variability in ice velocity was observed in regions with relatively low ice thickness and concentration. This can be attributed to the ice strength being a function of ice thickness and concentration (Equation 3.5).

Figures 3.4a and 3.4b illustrate the distributions of wind and ocean velocities simulated during the model run. The distribution of wind velocities was approximately Gaussian and comparable to the distribution of wind velocities collected from a 10 m meterological tower during a field program over the Beaufort Sea [45]. A similar study in the Beaufort Gyre observed typical maximum ocean surface currents of 5-10 cm/s [46]. Overall we felt that the shallow water models used to force the sea ice model would provide a reasonable distribution of surface stresses on the ice.

Following a verification of the forcing data, we examined the sea ice drift velocity fluctuations. A velocity fluctuation is the recorded velocity minus the long-term average velocity [45]. The long-term average sea ice velocity was approximately zero. Figure 3.4c illustrates the distrbution of sea ice velocity fluctuations observed during the model run. The sea ice drift velocity fluctuations roughly followed a Gaussian distribution with a mean of zero and with maximum observed speeds of 0.2 m/s. This distribution was compared to a distribution of velocity fluctuations gathered from buoys deployed during an Arctic Field campagin [47]. Both distributions were comparable. However, the buoy data contained very infrequent velocity fluctutations up to 1 m/s that were not observed in the toy model. These large velocities were be attributed to turbulent-like features [47]. We did not expect our simplified model to reproduce these large velocities.

The divergence rate is the horizontal derivative of velocity, $\partial u/\partial x$, representing the deformation being experienced by the ice [48]. Figure 3.4d illustrates the divergence rates observed by the model. The divergence rate distribution was compared to the distribution of divergence rates copmuted from a set of Radarsat Geophysical Processor System (RGPS) observations [48]. Both distributions approximately followed a power-law-like distribution, with a large amount of ice that exhibits very little deformation and smaller amounts of ice exhibiting fracturing or ridging events. In the RGPS data, there were more divergence rates near zero and fewer greater than 0.2 day^{-1}, as compared to the toy model. A possible explanation for this discrepancy is that the RGPS data was obtained from first-year and multi-year Arctic sea ice that is much thicker and

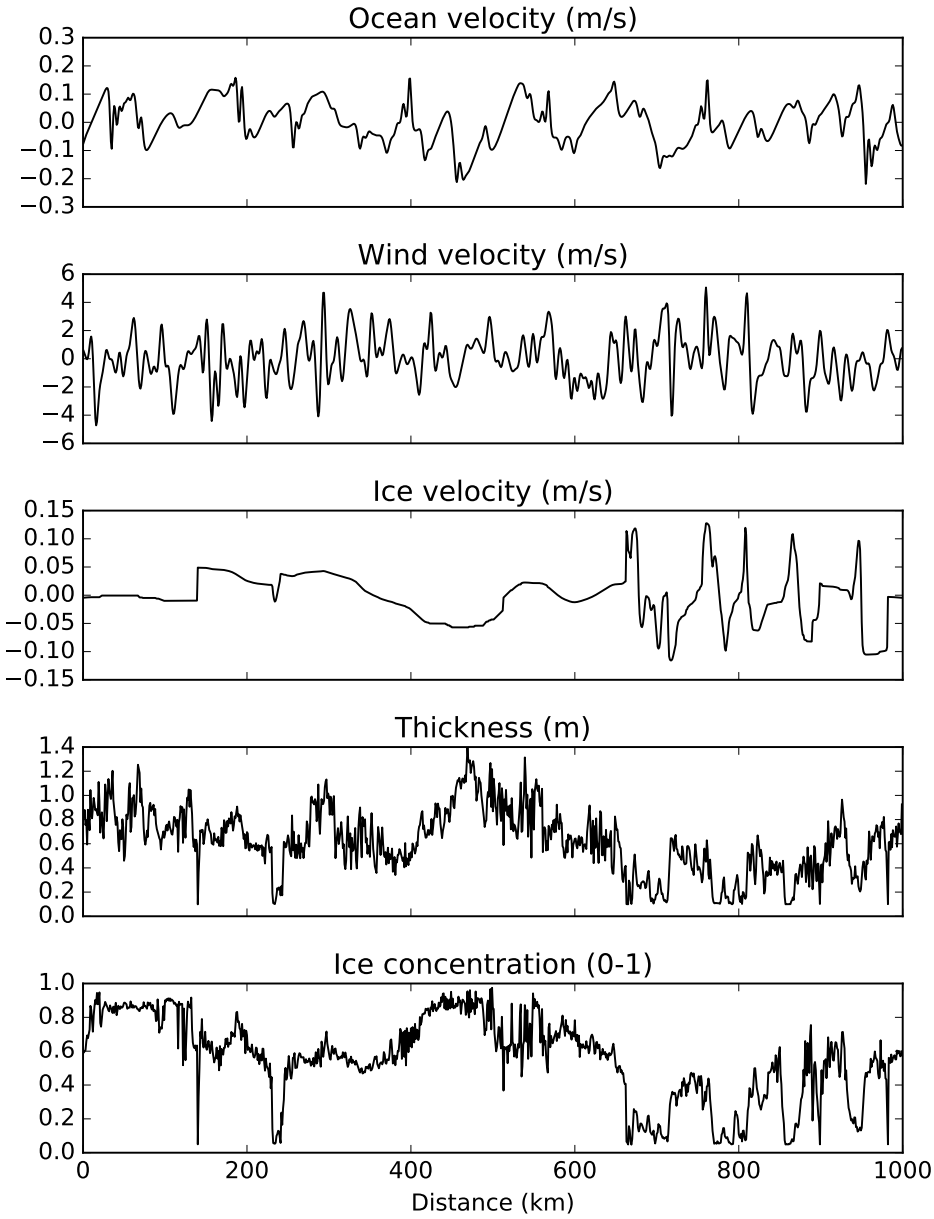

Figure 3.3: Toy sea ice model state after a sixty-day integration from nearly ice-free conditions. The mean thickness was 0.72 m and the mean ice concentration was 0.95. Note that in regions where there is relatively low thickness and concentration, the variability of ice velocity is higher. This demonstrates the sea ice processes that the model is able to simulate.

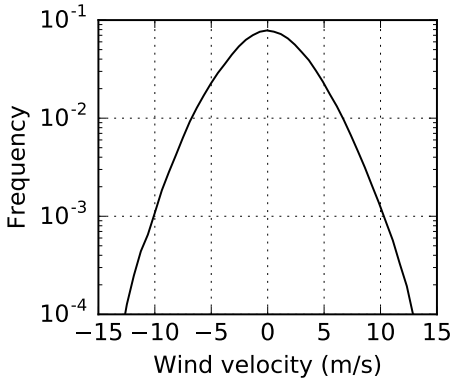

(a) Distribution of wind velocities

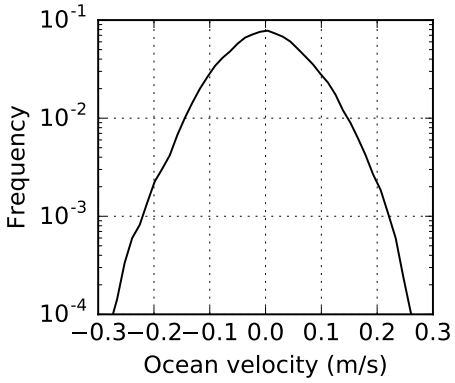

(b) Distribution of ocean velocities

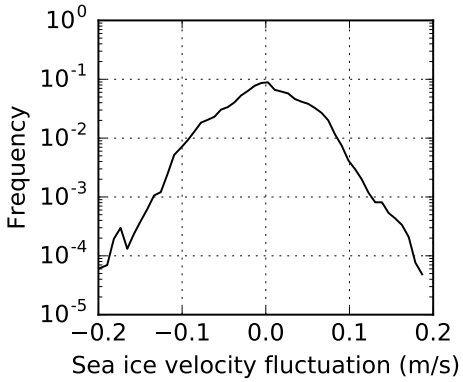

(c) Distribution of sea ice fluctuating velocities

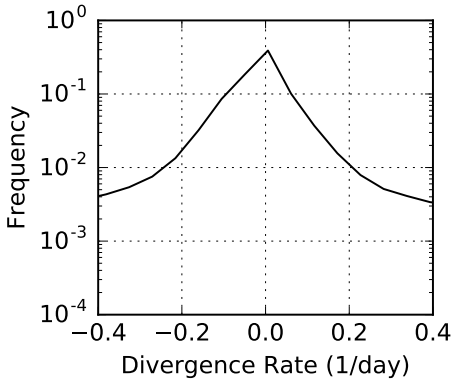

(d) Distribution of ice divergence rates

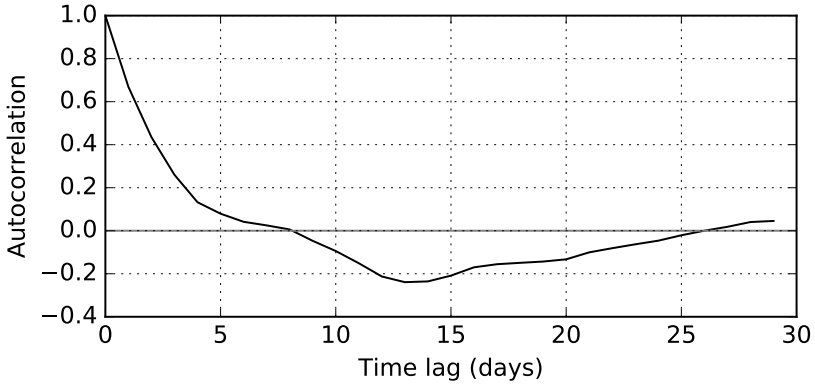

(e) Temporal ACF of ice velocities

Figure 3.4: Ice, ocean and wind velocities observed during the sixt-day validation run. The velocity distributions were all approximately Gaussian. These produced a PDF of ice divergence rates that approximately followed a Power-Law distribution. The temporal decorrelation length of ice velocities was approximately ten days. These distributions are comparable to the distributions of field observations presented in [45]

stronger than the ice simulated by the toy model.

Figure 3.4e illustrates the temporal autocorrelation of sea ice velocity experienced at a randomly-selected grid point. The decorrelation time was between five and ten days. For comparison, the Arctic buoy data [47] had a decorrelation time of ten days. Figure (3.5a) presents the histogram of ice thickness produced by the toy model. This histogram can be compared to the AEM-based thickness histogram presented in Figure 2.4. In both histograms there is a thickness mode less than 1 m and a long tail indicating thicker ice. In the AEM dataset, there is a longer tail with thicknesses extending up to 7 m. This is in part related the higher resolution of the AEM dataset, since sea ice can exhibit greater variability at smaller scales [9]. It is also related to the fact that there was first- and multi-year ice covered in the AEM survey, which may have experienced many more ridging events. Overall, the ice thickness distribution produced by the toy model was reasonable for young ice. Figure 3.5b illustrates the spatial autocorrelation of thicknesses produced by the toy model. The ice thickness had a decorrelation length of approximately 30 km. This ACF was comparable to the ACF of ice draft observations collected using a ULS beneath the Beaufort Sea [49].

One important concluding note is that the results depended largely on both the initial sea ice state and the forcing. When the inital ice concentration was set to a constant 100%, the range of velocities was much smaller due to the greatly increased ice strength. Girard et al. [48] noted that when the ice concentration was 100%, the ice divergence rates no longer followed a power law-like distribution. The toy model was more representative of new, thinner ice than the thick first- and multi-year ice studied in [47]. While this leads to confidence in the ability of the toy model to adequately simulate sea ice, it also means that results from experiments with the model should be considered unique to the model state and forcing. This caveat was carried forward in later chapters.

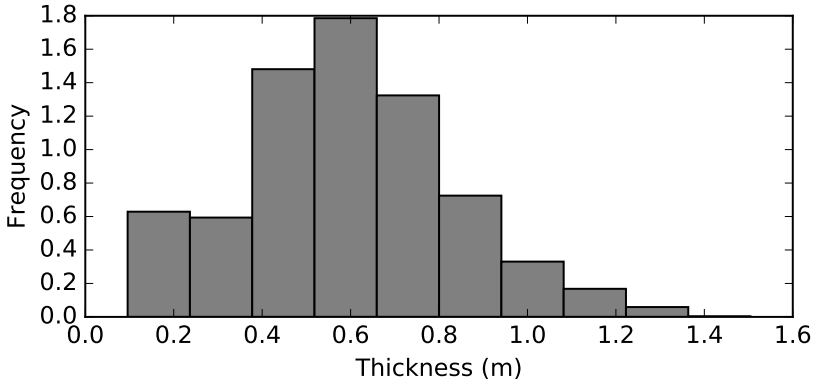

(a) Distribution of sea ice thicknesses generated during the sixty-day validation run of the toy model

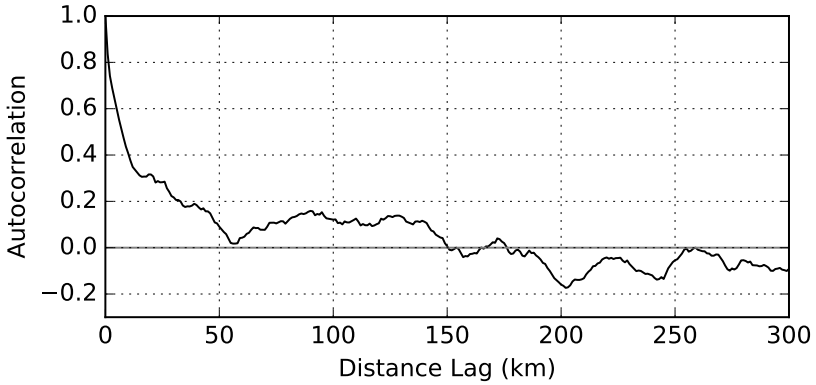

(b) Spatial ACF of sea ice thickness

Figure 3.5: Distribution and spatial autocorrelation of sea ice thickness generated during the sixty-day validation run. The ice thickness distribution had a slight skew and the ACF had a decorrelation length of approximately 30 km. Both plots are comparable to field-based observations [25][45].

3.7 Summary

This Chapter presented a toy sea ice model that was developed to simulate sea ice dynamics and thermodynamics on a one-dimensional grid. The toy model is different from other simplified models in that it is forced by simulated but realistic forcing data. The intended use of the toy model was to support the twin sea ice thickness data assimilation experiments in Chapter 4.

Overall, the toy model provides a reasonable representation of sea ice dynamics and thermodynamics. There are a few noted deficiencies, such as the lack of shear and the two-category representation of thickness. Another deficiency is that the ice model and shallow water models are only coupled in one direction, whereas the ice exerts equiavlent drag forces on the wind and ocean. Despite these shortcomings, the toy model produces a sea ice simulation that has characteristics comparable to observational data sets.

The model is well-suited to thickness assimilation because thickness is stored as a single variable, avoiding the challenges facing multi-category models. The model is also well-suited to assimilation experiments because the model domain with 1000 grid points and three state variables produces a background error covariance matrix that easily fits in computer memory.

A final note is that the model state evolves differently under different initial conditions and thermodynamic forcing. Any insights gathered from assimilation experiments should be considered unique to the given conditions. Because the model was validated primarily on a relatively young ice pack, similar conditions should be used in the assimilation experiments.

Chapter 4

Twin Experiment Formulation

In the context of data assimilation, a twin experiment involves assimilating observations sampled from one model state designated the *true* model state, into a second model state that represents the background [50]. Often, multiple twin experiments are conducted, each with a different assimilation technique or observation type, as a means to assess the value of that technique or sensor (e.g. [11]). Compared to experiments with actual observational data, a twin experiment permits greater flexibility and control over unknowns. It also allows for comparison of the analysis to the true system state, which is not possible in operational systems. Lisaeter et al. [11] and Yang et al. [12] used twin experiments to demonstrate the value of assimilating sea ice thickness observations.

Twin experiments with the toy model comprised the main analytical component of this thesis. The objective of these experiments was to demonstrate and quantify the impacts to the sea ice state that might occur when sea ice thickness is assimilated with incorrectly-characterized observation error statistics. This chapter describes the methodology followed in the twin experiments. Chapter 5 describes the results of the investigation.

4.1 Base Experiment Design

The twin experiments described in this chapter each followed the same basic formulation. For the base experiment, we first sampled a background state, $\mathbf{x}^b$, from a true state, $\mathbf{x}^t$, using a background error covariance matrix, $\mathbf{B}$. The state vectors contained three variables: thickness, concentration, and velocity. The background error covariance matrix, which included cross-covariance terms, was obtained from an ensemble of model runs. Section 4.2 describes the

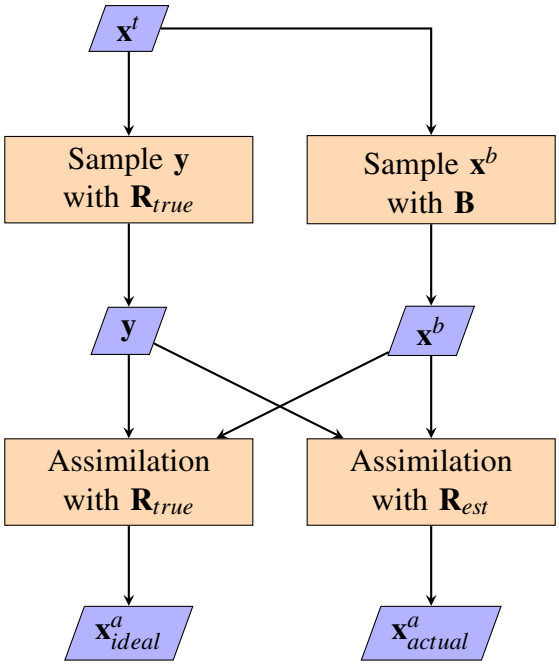

Figure 4.1: Diagram of the twin experiment procedure. From the true state, we sample a background state and a set of observations. The observations are then assimilated into the background state an approximation to the observation error covariance matrix.

generation of the true state, the background error covariance matrix, and the sampling of the background state.

We sampled a set of thickness observations from the true model state using a *true* observation error covariance matrix, denoted $\mathbf{R}_{true}$. We then assimilated these observations into a background state, $\mathbf{x}^b$, using an approximation to the observation error covariance matrix, $\mathbf{R}_{est}$, generating the analysis, $\mathbf{x}^a_{actual}$. We then compared this analysis to the *ideal* analysis, $\mathbf{x}^a_{ideal}$, obtained when the true observation error covariance matrix was used in the assimilation. Figure 4.1 illustrates the experimental procedure. The assimilation steps were conducted using the sequential assimilation equations presented in Section 2. Effectively, the twin experiments comprised a single cycle of an EnKF [24].

We devised five scenarios corresponding to five true observation error covariance matrices. The true observation error covariance matrices, presented in Section 4.3 each had different correlation structures, ranging from diagonal to a dense Gaussian structure. Each of these matrices contained the same error variances.

For each scenario, we conducted the twin experiment three times. For the first experiment, the

correct observation error covariance matrix was used in the analysis. For the second experiment, a diagonal approximation was used. The third experiment was identical to the second except that the error variances were multiplied by an inflation factor of 2.4. The choice of inflation factor is described in Chapter 5. The inflation of observation error variances is applied in RIPS and GIOPS at ECCC [21][5]. However, the actual inflation factor was not reported for these systems.

The decision to focus on diagonal approximations was based on the reality that in the near future, it is not likely that sea ice prediction systems will begin to incorporate correlated observation errors. By only considering diagonal approximations to the true observation error covariance matrices, we can provide more practical recommendations to ECCC and other forecasting centres.

Section 4.4 describes in detail the procedure that was followed in the comparison of the analyses. The objective was to quantify the reduction (or increase) in error for each state variable that accompanied the diagonal approximation to the true observation error covariance matrix. It was possible to accomplish this task without actually sampling a background state or observations.

4.2 Background State

The background error covariance matrix in the twin experiments was constructed from an ensemble of ice model states. This section describes the procedure followed to create the background error covariance matrix.

4.2.1 Freeze-up Procedure

To generate the ensemble of model states for estimating the background error covariance matrix, we initialized five hundred identical model states, plus a *true* model state, and ran toy the sea ice model for a period of thirty days. The procedure is described in detail below:

Initial Conditions

We first created an ensemble, $\mathbf{X}^b$, of $n_{ens} = 500$ identical model states, $\mathbf{x}_i^b$, plus an additional model state, $\mathbf{x}^t$, deemed the true state. Each model state had thickness of 0.1 m, concentration of 10 %, and a zero velocity field across the entire domain. Each state was paired with an instance of the toy model. For each model state in the background ensemble, we perturbed the corresponding numerical model parameters, following the approach described in Section 4.2.2. This procedure

was similar to the way the sea ice ensemble was generated in Shlyaeva et al. [8]. Note that the model parameters were not perturbed for the true state. This follows the procedure presented by Houtekamer and Mitchell [51], who performed one of the first ensemble-based twin experiments.

Freeze-Up Period

Each model state was propagated forward in time for thirty days. Every twenty-four hours, the growth scaling factor was resampled from a uniform distribution spanning [-0.2,0.5]. We found that when a single positive growth scaling factor was used, the ice concentrations were nearly all 100 % and the spread of ice concentration was much less than reported for RIPS and GIOPS [4][5]. By alternating between melting and freezing conditions, we achieved a more realistic spread. After the first fifteen days, we began sampling from a uniform distribution spanning [-0.2,0.2], representing milder thermodynamic conditions. Each model was supplied with identical forcing data. However, each model applied an independent perturbation to the forcing, as described in Section 4.2.2.

In Schlyaeva et al. [8], the ice models were spun-up between twenty and forty days before assimilating ice concentration observations. The thirty-day freeze-up period in our study was selected so that the background error variances were reasonably large while not fluctuating greatly in time. Section 4.2.3 illustrates the growth of the background error variances.

4.2.2 Stochastic Perturbations in the Toy Model

A critical requirement in ensemble data assimilation is that the numerical models must have a stochastic component [7]. The use of carefully-tuned stochastic models instead of purely deterministic models ensures that the distribution of model states about the mean is approximately equal to the distribution of errors about the true state [7]. To carry out ensemble data assimilation, the deterministic toy sea ice model therefore requried a stochastic component.

A common practice in sea ice ensemble forecasting is to perturb the model parameters and forcing data rather than adding errors directly to the ice thickness, concentration, and velocity [8][52]. Sea ice parameters are typically poorly known and numerical models of the ocean and atmosphere, which force the sea ice, contain their own independent background errors [8][52]. The perturbation of model parameters and the use of an ensemble of forcing data has been shown to yield a satisfactory ensemble spread [8][52]. In this study, we perturbed

1. The ice strength parameter, P^*;

2. The drag coefficients, C_a and C_o;

3. The ice density, ρ_i;

4. The thermodynamic growth scaling factor;

5. The ocean velocities, u_w; and

6. The wind velocities, u_a.

The strength, drag and density parameters were perturbed by multiplying the default parameter value by a factor sampled from a Gaussian distribution with a mean of 1 and a standard deviation of 0.1. This produced some model states that were stiffer and other states that experienced more deformation.

Factors that influence thermodynamic growth include the albedo, snow depth, air temperature, atmospheric moisture content, and cloud cover, all of which may vary in space and time [2] and none of which are modeled explicitly by the toy model. These variables are subject to spatially and temporally correlated errors just like ice thickness, and we wanted to account for that despite the limitations of our simplified thermodynamic growth term. Onto the growth scaling factor we added a random vector sampled from a Gaussian distribution with a mean 0 and standard deviation of 0.2, with a Gaussian correlation structure that had a decorrelation length of 50 km. This accounted for potential spatially correlated errors in variables that influence ice growth. The growth scaling factor was resampled every twenty-four hours for each model state. This accounted for potential temporal variability in the errors in variables such as humidity and snow depth.

The ocean and wind velocities were also multiplied each time step by a factor sampled from a Gaussian distribution with mean 1 and standard deviation 0.2. This approach was implemented successfully in Shlyeva et al. [8]. In addition, the velocities were perturbed spatially (i.e. shifted) by a Gaussian random variable with mean 0 and standard deviation of 5 km, rounded to the nearest kilometer. This introduced spatially correlated errors into the ensemble. We considered forcing each ice model with a separate shallow water model but this was found to be too computationally expensive.

Excluding the thermodynamic term, these perturbation factors were sampled only once at the start of each model run. Furthermore, they were sampled independently for each ensemble member and for the true state.

Note that the true model state was integrated purely deterministically. No perturbations were applied to this state.

4.2.3 Evolution of Ensemble Spread during Freeze-Up

At the end of the thrity-day growth period, the mean concentration in the true state had increased to a value of 96%. The thickness grew from 0.1 m to a mean of 0.50 m. The minimum and maximum observed ice thicknesses were 0.07 m and 1.2 m in the true state. This relatively thin ice state was well-suited to experiments with a SMOS-based thickness sensor. Figure 4.2 illustrates the mean and 10th and 90th percentiles of the ensemble after the thirty-day growth period. The true state is also plotted for reference.

At each timestep during the freeze-up, the ensemble spread was calculated for each state variable. Figure 4.3 illustrates the evolution of this statistic over the thirty-day period. In this study, spread was defined as [53]

$$spread = \sqrt{\frac{1}{(n_x - 1)(n_{ens} - 1)} \sum_{i=1}^{n_x} \sum_{j=1}^{n_{ens}} (\mathbf{x}_{ij} - \overline{\mathbf{x}_i})^2} \tag{4.1}$$

Several observations can be made regarding the evolution of the sea ice state. Firstly, the thickness spread increased steadily with time, reaching a value of 0.22 m after twenty-five to thirty days.

The spread of concentration did not experience the same monotonic increase. The spread after thirty days was 0.073. In regions where the concentration was near 100%, the error was much lower. Conversely, in areas where the concentration was lower, there was greater error. This is consistent with operational systems, where the best estimates of concentration occur when the concentration is either 100% or 0% (e.g. [5]).

The spread of ice velocities experienced an immediate jump, but over the first ten days the spread gradually decreased. At the beginning of the freeze-up period, the ice was relatively weak and the range of ice velocities was larger. As the ice thickened, the strength increased and the velocities accordingly decreased.

4.2.4 Background Error Covariance Matrix

Following the freeze-up, we found large biases in the ensemble. Any given ensemble member, compared to the ensemble mean (Equation 2.11), had a different mean thickness, concentration and velocity. The standard deviations of these biases were 0.074 m, 4.1 %, and 0.034 m/s, respectively. ECCC also found large ice concentration biases in RIPS and GIOPS, relative to

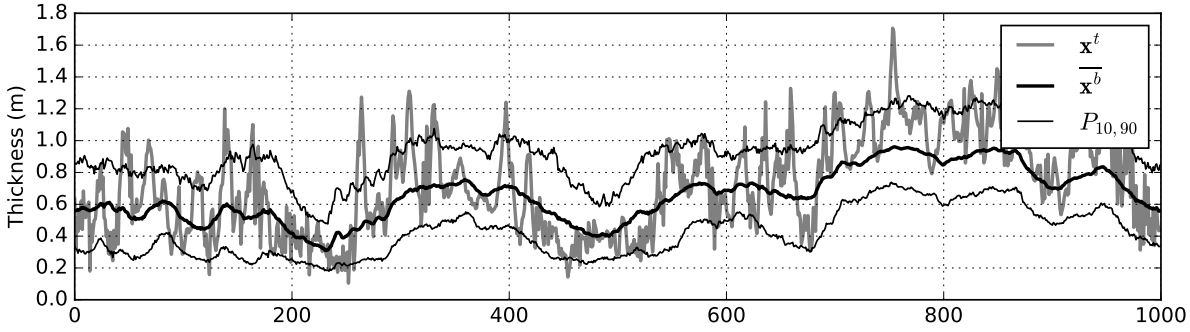

(a) Ensemble of ice thicknesses

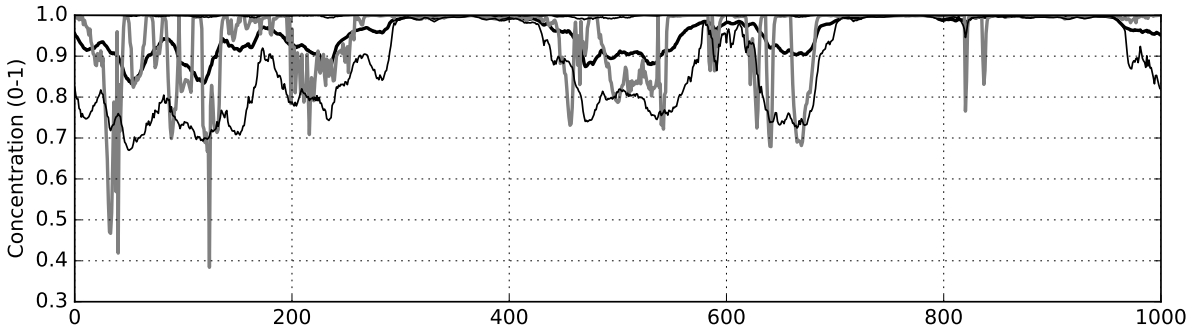

(b) Ensemble of ice concentrations

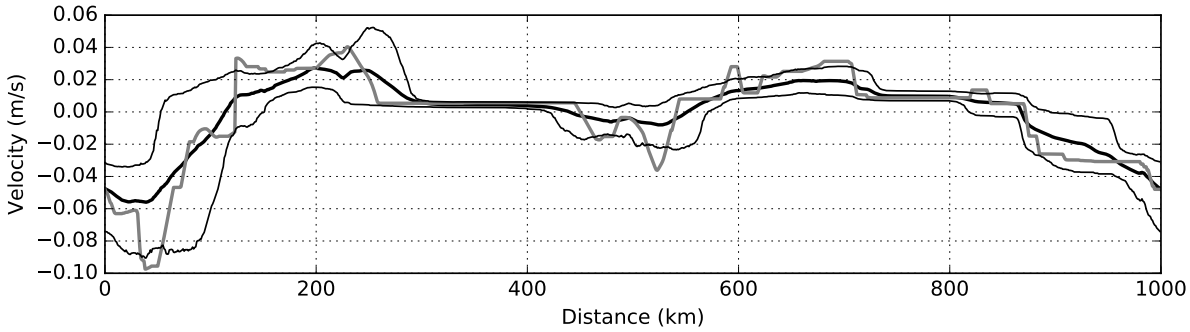

(c) Ensemble of ice velocities

Figure 4.2: Ensemble of sea ice model states with the true state following the thirty-day freeze-up period. The mean and the 10th/90th percentiles of the ensemble are plotted to illustrate the spread. Note that where the true ice concentration was near 100 %, the spread of the ice concentration and the velocity errors were the least but there was little correlation between the state and the thickness spread.

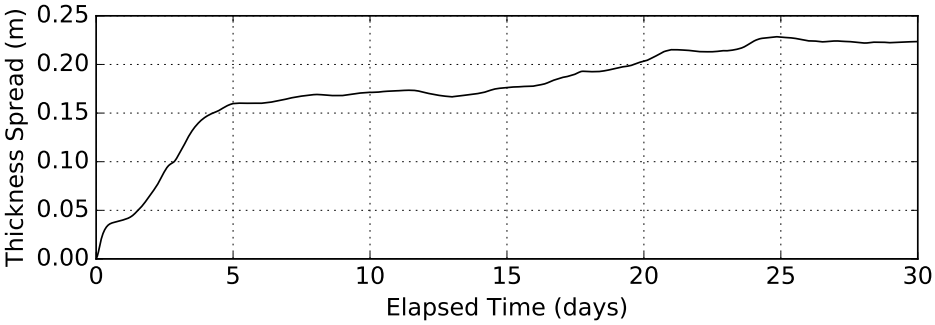

(a) Evolution of sea ice thickness spread

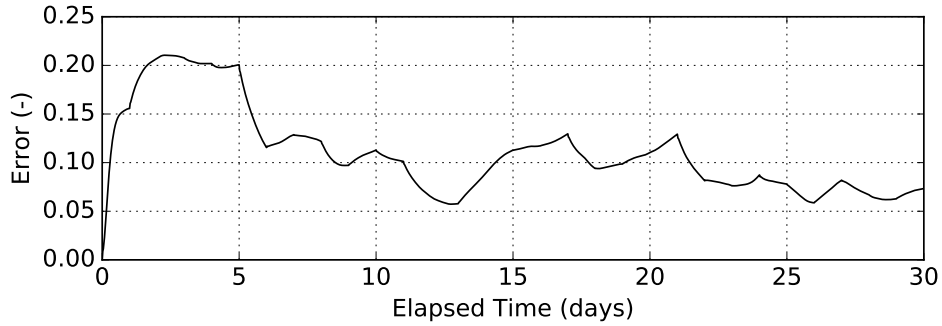

(b) Evolution of sea ice concentration spread

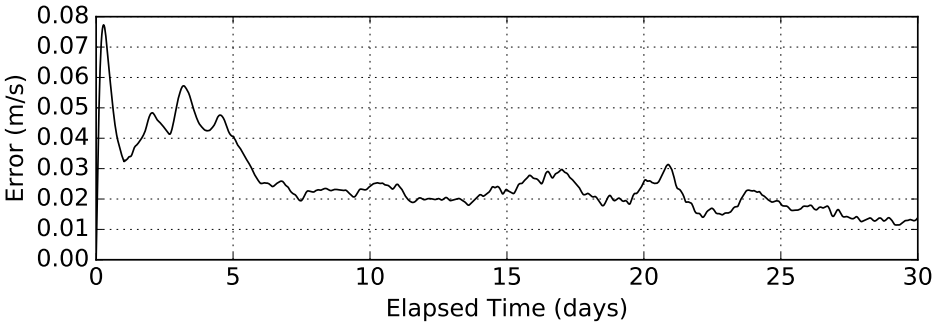

(c) Evolution of sea ice velocity spread

Figure 4.3: Temporal evolution of spread in the five-hundred-member ensemble during the thirty-day freeze-up period. The thickness spread increased to a value of approximately 0.2 m, which is reasonably high considering the mean thickness of 0.5 m. The concentration spread experienced an initial increase but slowly decreased as the mean concentration approached 100 %. Similarly, at the beginning of the growth period, the ice was weak and there was a greater velocity spread than at the end, when the sea ice was more developed and stronger.

observational data [5][4]. Massonnet et al. [54] demonstrated that modelled sea ice velocity biases, relative to RGPS data, can be partly attributed to the model parameters, including the ice strength parameter and the drag coefficients. The perturbation of these parameters was perhaps the source of the observed biases in our experiment. Bias is important to consider in operational sea ice prediction systems and merits future work beyond the scope of this thesis.

Preliminary computations of the background error covariance matrix, using Equation 2.10 and the ensemble of model states, revealed large non-zero error correlations that spanned the entire domain. We found that these long-range correlations were a product of the model biases. By subtracting the bias from each model state, relative to the ensemble mean state, these long-range correlations disappeared. This procedure resulted in background error decorrelation lengths of 20-100 km that were more consistent with literature [21][7].

The background error covariance matrix was 3x3 block diagonal. The first 1000x1000 block along the diagonal represents the thickness error covariances. The second block represents the concentration error covariances and the third block represents the velocity error covariances. The off-diagonal blocks correspond to the cross-covariances between the three state variables. We noted that while most error covariances decayed to zero after approximately 20-100 km, there were fluctuations in the error covariances about a mean of zero at much greater distances. These covariances were likely spurious, given that the dimension of each state vector was three-thousand but the ensemble contained only five-hundred model states. Note that in operational systems, for much larger domains, ensembles are typically run with only fifteen [12] to one-hundred [11] ensemble members.

To produce a more realistic estimate of background errors, we localized the background error covariance matrix. Localization is commonly performed in ensemble data assimilation to remove the undesired effect of spurious long-distance correlations [7]. To perform localization, we took the element-wise product of each of the nine 1000x1000 submatrices of the background error covariance matrix with a 1000x1000 circulant correlation matrix, $\mathbf{L}$. The correlation structure in $\mathbf{L}$ was Gaussian with a decorrelation length of 150 km. This effectively damped the long-range correlations while minimally affecting the correlations up to a distance of 100 km. The localized background error covariance matrix was computed as

$$\mathbf{B}_{ij} = \mathbf{B}^{*}_{ij} \odot \mathbf{L} \tag{4.2}$$

where $\mathbf{B}^{*}$ represents the background error covariance matrix prior to localization and where $\odot$ represents the element-wise product.

To provide a better illustration of spatial correlations of background errors, the background

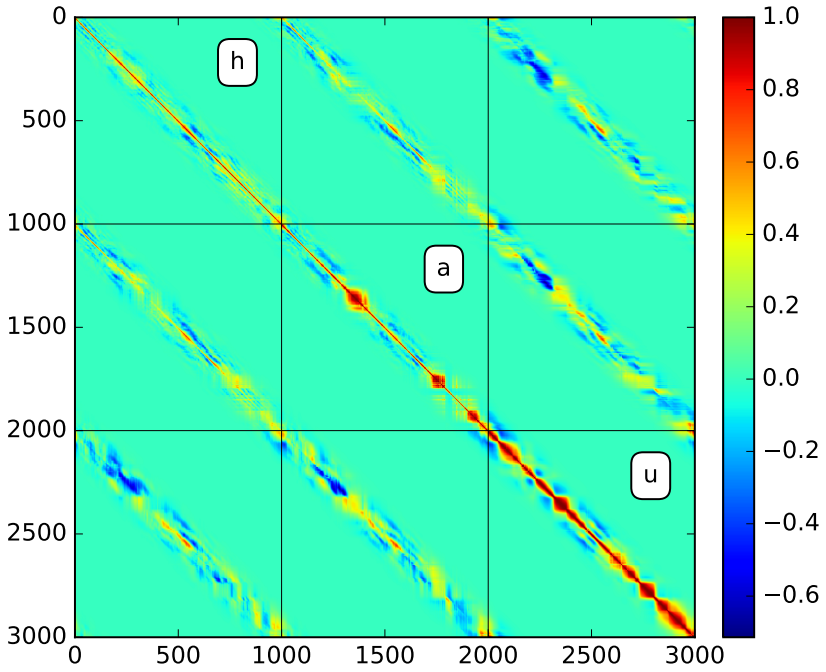

Figure 4.4: Background error correlation matrix corresponding to the estimated background error covariance matrix. Observe that the spatial correlations are variable across the domain and across variables. This may be a result of the nonlinear response of the model to the forcing data. The significance of this correlation matrix is that some regions may a experience greater analysis increment than the others.

error covariance matrix was decomposed and its underlying correlation matrix extracted. The matrix decomposition, presented in [50], is

$$\mathbf{B} = \mathbf{D}^{1/2}\mathbf{C}\mathbf{D}^{1/2} \tag{4.3}$$

$$\mathbf{C} = \mathbf{D}^{-1/2}\mathbf{B}\mathbf{D}^{-1/2} \tag{4.4}$$

where $\mathbf{D}$ is a diagonal matrix of the error variances, taken from the diagonal of the background error covariance matrix, and where $\mathbf{C}$ represents the corresponding background error correlation matrix. Figure 4.4 depicts the resulting background error correlation matrix.

The matrix was symmetric, with entirely positive eigenvalues and positive diagonal elements. This met the requirements of a covariance matrix [43]. The matrix had a large condition number

of $1.1 * 10^{16}$. Fortunately, the inversion in Equation 2.3 is performed only on $\mathbf{HBH}^T + \mathbf{R}$. The matrix $\mathbf{HBH}^T$ was much smaller and had a condition number of only 483.

4.2.5 Skewness of State Variable Error Distributions

The derivation of the Kalman filter [7][22] does not make any assumption that the distribution of background errors must follow a Gaussian distribution. However, it is assumed that the the filter that minimizes the error variance is the *best* filter. When there is skewness, kurtosis, or multiple modes in the errors, the variance alone may no longer be the best descriptor of the background error. Sequential assimilation under these circumstances may be beneficial but an alternative filter may be better-suited to the problem [55].

Because thickness and concentration are both bounded variables, we considered it important to investigate the distributions of the errors for each state variable in the ensemble. We calculated the skewness across the ensemble for each state variable at each grid cell. Figure 4.5 presents the empirical probability distributions of skewnesses, which were calculated following [56].

$$s = \frac{\frac{1}{n_{ens}} \sum_{i=1}^{n_{ens}} (\psi_i - \overline{\psi})^3}{\left[\frac{1}{n_{ens}-1} \sum_{i=1}^{n_{ens}} (\psi_i - \overline{\psi})^2\right]^{3/2}} \tag{4.5}$$

where ψ_i represents the particular state variable corresponding to a particular ensemble member i at a given grid point. $\overline{\psi}$ represents the ensemble mean at that grid point.

From Figure 4.5, only the velocity errors had an approximately zero skew. The ice thickness errors were consistently skewed positively, indicating a long tail to the right at each grid point. The concentration error distributions were entirely negatively skewed. The positive thickness error skewnesses can be attributed to the lower bound of 0 m that ice thickness can take. Similarly, the negative concentration error skewnesses can be attributed to the upper concentration bound of 100%. The implication is that while sequential data assimilation may improve the sea ice state estimate, there may be alternative filters that provide improved analyses.

These skewnesses, along with the biases described in Section 4.2.3, can be attributed to the bounding of the variables and to the combined effects of nonlinear model processes and perturbed model paramters. The assumption of nonlinearity presented in Chapter 2 is challenging because the BLUE can be derived without making this assumption. The role of nonlinearity is perhaps indirect. Firstly, nonlinear model processes can lead to model bias, which does in fact violate the critical assumption of unbiased errors. Secondly, nonlinear model processes appear to lead

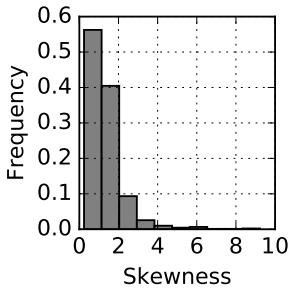

(a) Skewness of ice thickness errors

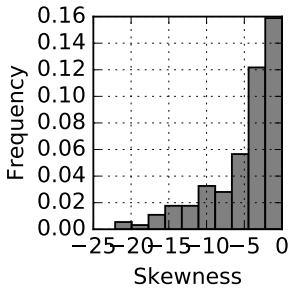

(b) Skewness of ice concentration errors

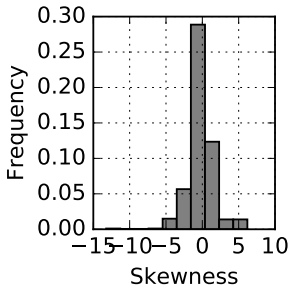

(c) Skewness of ice velocity errors

Figure 4.5: Empirical probability distribution functions of sample skewness of errors for each state variable in the sea ice model ensemble. The velocity errors have approximately zero skewness but there are significant positive and negative skewnesses for the thickness and concentration errors. With large skewnesses, the performance of sequential data assimilation may not be ideal.

to non-Gaussian error distributions, in which case the Kalman Filter may no longer be the *best* linear unbiased estimator (BLUE) [55].

In an operational setting, analyses might occur every six hours rather than every thirty days as was the case with this experiment. With assimilation cycling at a shorter interval, it is possible that skewnesses and biases would not have the opportunity to develop. The actual distribution of errors in sea ice state estimates is a topic that has not received much attention in literature, meriting additional consideration in follow-up work.

4.3 Sensor Design and Sampling of Synthetic Sea Ice Thickness Observations

We designed a synthetic SMOS-like sea ice thickness sensor to generate observations for the twin experiments. The observations were defined in terms of a true observation state vector $\mathbf{y}^t$ and an observation error vector sampled using the true observation error covariance matrix.

$$\mathbf{y} = \mathbf{y^t} + \varepsilon^o \tag{4.6}$$

This section describes how we determined of the true observation state, $\mathbf{y}^t$, along with the observation error covariance matrices and the sampling of the error vector, ε^o.

4.3.1 Observation Operator

The true observation state was obtained by mapping the true model state to the observation space. The mapping was achieved through a linear observation operator.

$$\mathbf{y^t} = \mathbf{H}\mathbf{x}^t \tag{4.7}$$

The sensor had an observation footprint diameter of 35 km and an observation spacing of 12.5 km. This resulted in an observation operator with dimensions of 80 by 3000. The observation operator was largely sparse, with non-zero entries, H_{jk}, in locations where thickness entries k fell inside the footprint of observation j. The value of the non-zero coefficients in the observation operator was a constant 1/35, meaning that each grid cell within the footprint contributed equally to the observation. This footprint operator was similar to the operator used in RIPS for the assimilation of coarse-resolution passive microwave sea ice concentration observations [21].

The observation operator that was used to generate the observations was also used to assimilate the observations. In practice this may not be possible because the impulse response of the sensor may not be known exactly. The actual impulse response of SMOS is not actually a uniform step function as described herein. The use of an incorrect observation operator introduces *representativity errors*, which are commonly grouped with the sensor observation error [7].

4.3.2 Specification of Observation Error Variances

The observation error standard deviations were specified as a function of the true thickness, $\mathbf{y}^t$. SMOS-based observation standard errors are reportedly 0.06 m for ice less than 0.3 m thick and 1.0 m for ice 1.0 m thick [32]. We fit a piecewise function to these values. The error standard deviation was constant for $h \leq 0.3$ and set to a quadratic function of thickness for $h > 0.6$, with a vertex (0.3,0.06) and passing through (1.0,1.0).

$$\sigma_o(h) = \begin{cases} 0.06 & h \leq 0.3, \\ 1.92(h-0.3)^2 + 0.06 & h > 0.3, \end{cases} \tag{4.8}$$

Figure 4.6 illustrates the observation error standard deviation function evaluated for thicknesses ranging from 0 to 1.5 m. For thicknesses less than 0.5 m, the observation error is relatively

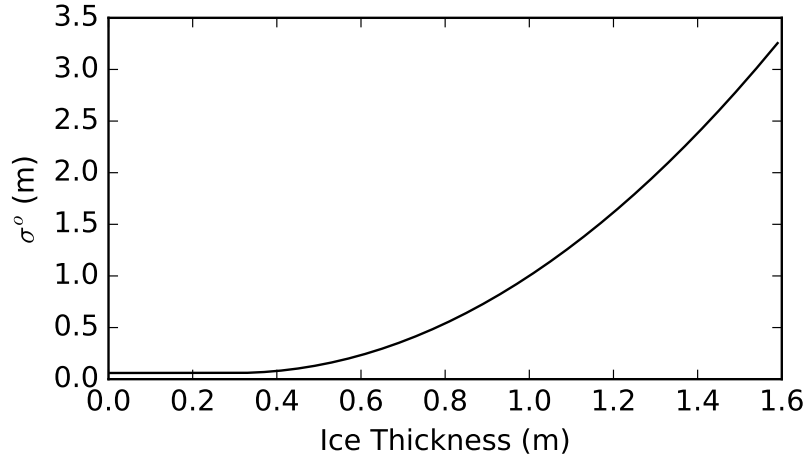

Figure 4.6: Prescribed true observation error standard deviation as a function of ice thickness. As the true ice thickness increases past 0.3, the observation error standard deviation increases quadratically, meaning that there is low observation accuracy for regions of thick ice.

low and these observations may have an impact on the analysis. As the thickness increases, the observations have a much smaller impact on the analysis.

The result of the quadratic function was that any observed thicknesses greater than 1 m had error variances high enough that they would not contribute significantly to the analysis. Yang et al. [12] outright did not assimilate thickness observations greater than 1 m. It was not immediately clear what might be the impact of clipping the data versus assimilating the complete observation set but with large error variances where appropriate. A preliminary hypothesis is that clipping the data above a given value may affect the bias in the analysis state.

Figure 4.7 illustrates the background and observation error standard deviations plotted alongside $\mathbf{H}\mathbf{x}^t$, which represents the true state mapped to the observation locations. For most of the domain, the background error standard deviation is less than the observation error variance. However, where the true ice thickness is thinnest, the observation error standard deviations are comparable or less than those of the background state. If the true background and observation error covariance matrices were diagonal, we would see the majority of the analysis increment in these areas.

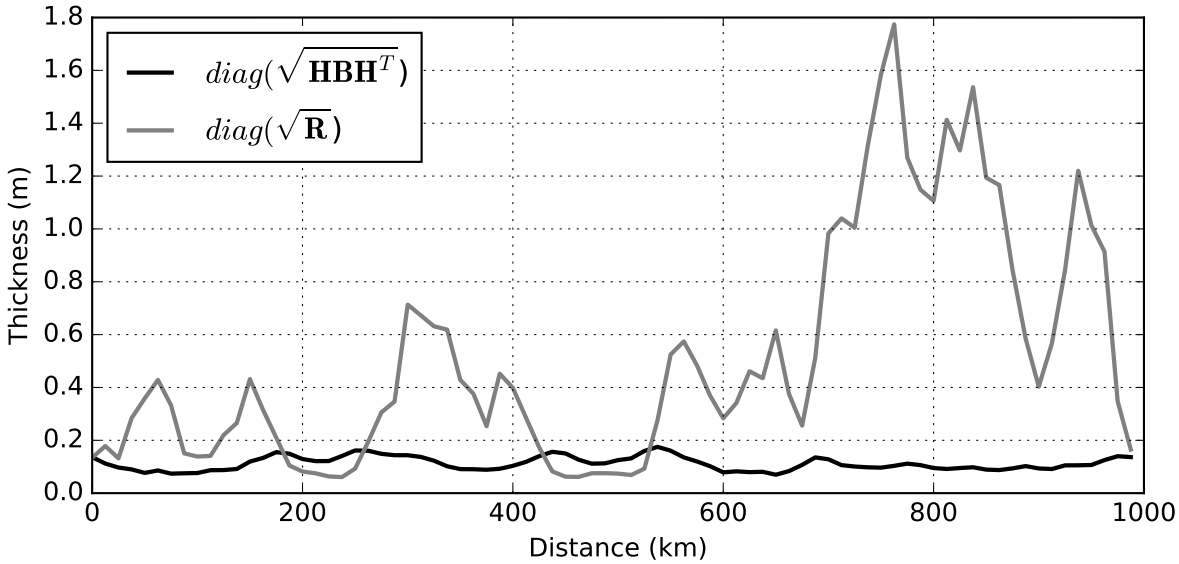

Figure 4.7: Evaluation of the observation error standard deviation function at the true state mapped to the observation locations. The background error standard deviations (solid black) can be compared to the observation error standard deviations (grey), revealing that the background error standard deviations are relatively constant across the domain while the observation error standard deviations are more variable. We can expect greater impact from the assimilation in two isolated regions where the observation error standard deviations are less than the background error standard deviations.

Table 4.1: Correlation structure of each observation error covariance matrix

Matrix	Correlation Structure	Decorrelation Length (km)
$\mathbf{R}_{true_1}$	Gaussian	150
$\mathbf{R}_{true_2}$	Gaussian	50
$\mathbf{R}_{true_3}$	Exponential	150
$\mathbf{R}_{true_4}$	Exponential	50
$\mathbf{R}_{true_5}$	Diagonal	0

4.3.3 Observation Error Covariance Matrices

Stewart et al. [50] presented a matrix decomposition approach for constructing prescribed observation error covariances. We constructed $\mathbf{R}_{true}$ following their approach:

$$\mathbf{R}_{true} = \mathbf{D}^{1/2}\mathbf{C}\mathbf{D}^{1/2} \tag{4.9}$$

where the diagonal entries in $\mathbf{D}^{1/2}$ were defined by $\sigma_o(\mathbf{y}^t)$.

The error correlation matrix, $\mathbf{C}$, was constructed using a deterministic function of the distance between observation points. For these experiments, we used five error correlation matrices to produce five $\mathbf{R}_{true}$ matrices. We decided to use five error correlation matrices to rerpesent a range of possible observation error scenarios. Table 4.1 describes the five true obsevation error covariance matrices.

The Gaussian correlation functions were defined as [44]

$$c(d) = exp\left(-3\left(\frac{d}{r}\right)^2\right) \tag{4.10}$$

and the exponential correlation function was defined as

$$c(d) = exp\left(-3\left(\frac{d}{r}\right)\right) \tag{4.11}$$

where d represents the distance between points and r represents the length at which the correlation decreases to a value of 0.05.

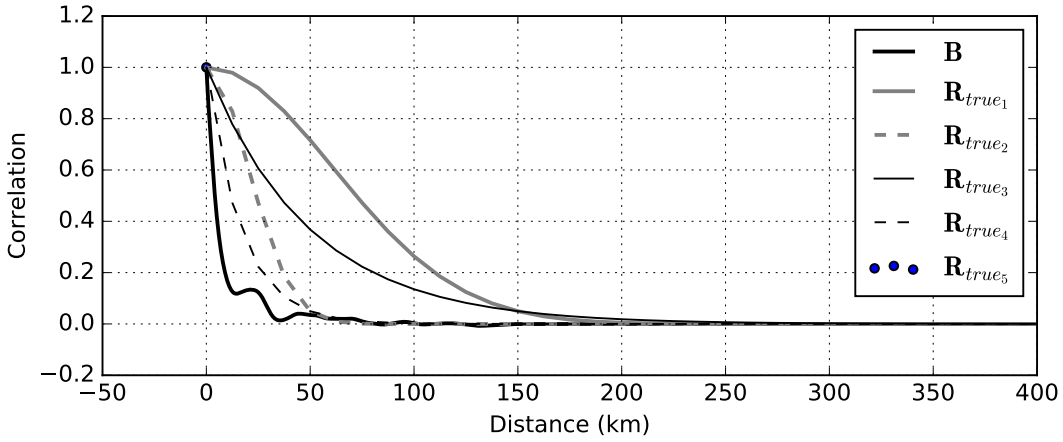

Figure 4.8: Correlation structures of the true (prescribed) observation error covariance matrices and of the background error covariance matrix. Each of the observation error correlation functions except for $\mathbf{R}_{true_5}$ (no spatial correlation) have a longer correlation length scale than the background.

The observation error covariance matrices were symmetric and positive definite, with entirely positive eigenvalues and with condition numbers less than $1.7 * 10^6$. The matrix $\mathbf{HBHT} + \mathbf{R}$ had a condition number less than $4.0 * 10^4$.

Figure 4.8 illustrates the correlation structures of each observation error covariance matrix alongside the correlation structure extracted from the background error covariance matrix. The background error correlation structure appears similar to the exponential (50 km) structure of $\mathbf{R}_{true_4}$. We considered also creating additional $\mathbf{R}_{true}$'s with decorrelation lengths that were less than the background error decorrelation length. However, with an observation spacing of 12.5 km, this would have produced a nearly-diagonal observation error covariance matrix.

Figure 4.9 illustrates $\mathbf{R}_{true_1}$ for example. Unfortunately, due to the extreme error variances observed in some regions of the domain, the structure of the observation error covariance matrix was difficult to visualize.

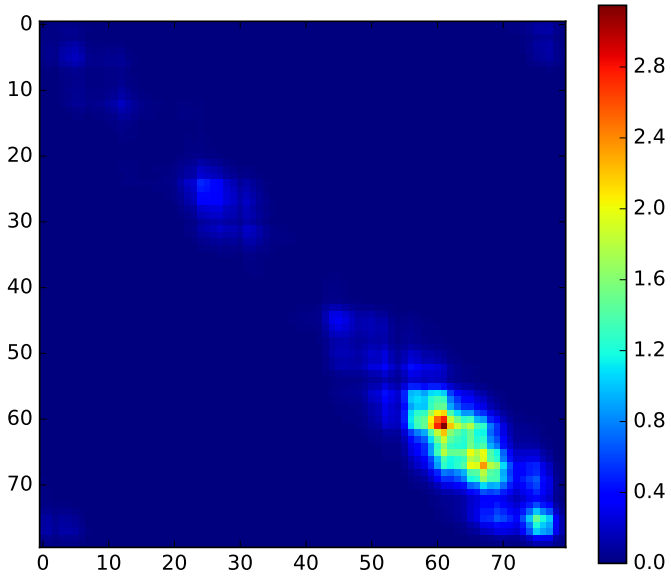

Figure 4.9: True (prescribed) observation error covariance matrix, $\mathbf{R}_{true_1}$. Observe first that the error variance is not constant along the matrix diagonal. Also not that the error covariance structure is not smooth, a result of correlations in the true ice thickness state. The implication is that some observations will have a greater impact in the analysis than others.

4.4 Comparing Experimental Results

In order to compare the analyses, it was not necessary to actually sample observations and assimilate them into a sampled background state. Instead, we compare the analysis error variances, corresponding to diagonals of the analysis error covariance matrices. The analysis error covariance matrix is described in Chapter 2 by Equation 2.8. By comparing the analyses based solely on the analysis error covariance matrices, we avoided errors associated with the sampling of the background state and observations.

For each experiment scenario, it was possible to compute four variants of the analysis error covariance matrix, described in Table 4.2. For the first analysis error covariance matrix, $\mathbf{A}_{ideal}$, we computed both the Kalman gain, $\mathbf{K}$, and the analysis error covariance matrix using the true observation error covariance matrix. This represents the analysis error covariances for the ideal case where the exact observation error covariance matrix was used in the assimilation.

For the second analysis error covariance matrix variant, $\mathbf{A}_{\mathbf{est}}$, we computed both the Kalman gain and the analysis error covariance matrix using a diagonal approximation to the observation error covariance matrix. This matrix represents the naïve estimate of the analysis error covariance matrix that would be obtained in an operational system where the true observation error correlation structure is unknown. Based on Stewart et al. [13], for a diagonal approximation to the true observation error covariance matrix, we would expect $\mathbf{A}_{\mathbf{est}}$ to underestimate the analysis error variances for the cases where true observation error covariance matrix is not diagonal.

The third matrix, $\mathbf{A}_{actual}$, represents the actual analysis error covariances for the case where an incorrect observation error covariance matrix was used in the assimilation.

For the third matrix, $\mathbf{A}_{actual}$, an incorrect observation error covariance matrix was used in the computation of the Kalman gain. The analysis error covariance matrix was then constructed using the correct observation error covariance matrix and and the incorrect Kalman gain matrix that was compuetd with the incorrect observation error covariance matrix. Effectively, this *actual* analysis error covariance matrix is what we could obtain if we used an incorrect observation error covariance matrix in the assimilation while also knowing the exact true model state. It would not be possible to calculate this analysis error covariance matrix, $\mathbf{A}_{actual}$, in an operational system where the observation error covariance matrix was not known exactly. Healy and White [15] introduced the concept of $\mathbf{A}_{actual}$, using a different but equivalent formulation. Rainwater et al. [57] referred to this matrix as the *performance* analysis error covariance matrix. By comparing $\mathbf{A}_{est}$ and $\mathbf{A}_{actual}$, we were able to explore how the diagonal approximation to the observation error covariance matrix affects the perceived analysis quality.

The background error covariance matrix was also included in Table 4.2 to represent the case where no data was assimilated. We considered it especially important to identify instances where

Table 4.2: Four variants of the analysis error covariance matrix that were computed for each twin experiment

A	Calculation
$\mathbf{A}_{ideal}$	$\mathbf{A}_{ideal} = (\mathbf{I}-\mathbf{KH})\mathbf{B}(\mathbf{I}-\mathbf{KH})^T + \mathbf{KR}_{true}\mathbf{K}^T$ $\mathbf{K} = \mathbf{BH}^T(\mathbf{HBH}^T + \mathbf{R}_{true})^{-1}$
$\mathbf{A}_{est}$	$\mathbf{A}_{est} = (\mathbf{I}-\mathbf{KH})\mathbf{B}(\mathbf{I}-\mathbf{KH})^T + \mathbf{KR}_{est}\mathbf{K}^T$ $\mathbf{K} = \mathbf{BH}^T(\mathbf{HBH}^T + \mathbf{R}_{est})^{-1}$
$\mathbf{A}_{actual}$	$\mathbf{A}_{actual} = (\mathbf{I}-\mathbf{KH})\mathbf{B}(\mathbf{I}-\mathbf{KH})^T + \mathbf{KR}_{true}\mathbf{K}^T$ $\mathbf{K} = \mathbf{BH}^T(\mathbf{HBH}^T + \mathbf{R}_{est})^{-1}$
B	No assimilation

the estimated analysis error covariance matrix contained larger error variances than the background error covariance matrix, since this would indicate that the assimilation had a negative impact on the state estimate.

With these analysis error covariance matrices, it was possible to compute the mean analysis error standard deviation, σ^a, for each of the four analysis error covariance matrices and for each state variable. The error term was defined as:

$$\sigma^a(\mathbf{A}) = \sqrt{\frac{1}{n_x} trace(\mathbf{A})} \tag{4.12}$$

where $\mathbf{A}$ is the block submatrix corresponding to either thickness, concentration or velocity for any one of $\mathbf{A}_{ideal}$, $\mathbf{A}_{est}$, or $\mathbf{A}_{actual}$. Refer back to Figure 4.4 for an interpretation of these submatrices.

With these measurements, we were able to quantify the total decrease or increase in background error when thickness observations were assimilated using an approximation to the observation error covariance matrix. The error standard deviations for the background state (σ^b) were 0.200 m, 7.28 %, and 0.0138 m/s, respectively for thickness, concentration and velocity.

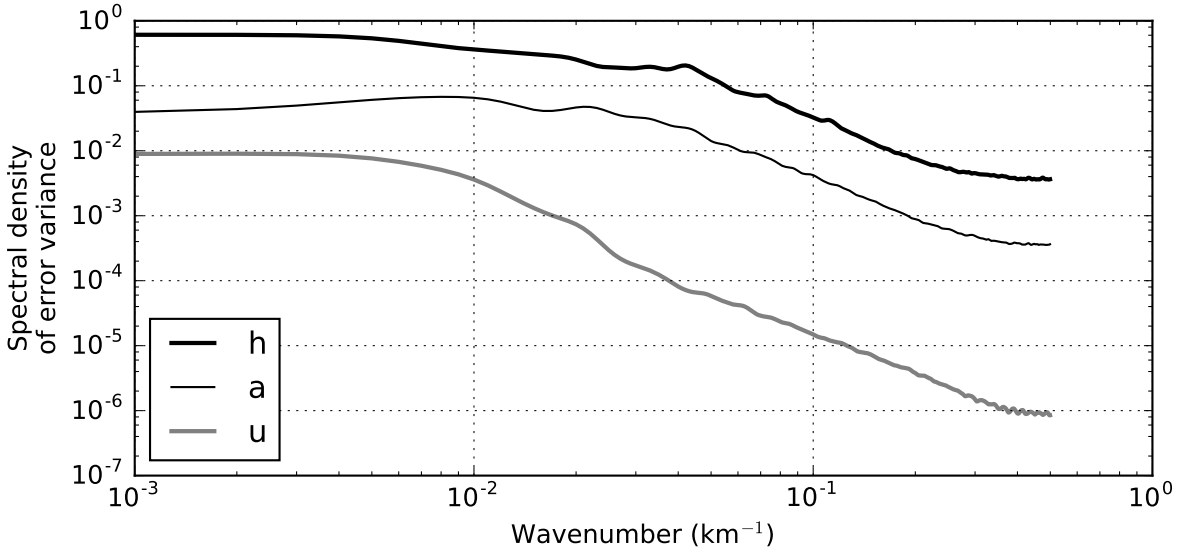

Figure 4.10: Spectral density of errors in the background error covariance matrix. The three curves have overall negative slopes, indicating that the majority of the spectral density exists at low wavenumbers.

To compare σ^a and σ^b, we computed the relative change in mean background error standard deviation as $(\sigma^b - \sigma^a)/\sigma^b$.

We also investigated the spatial scales in the background and analysis errors. To do this, we computed the real component of the discrete Fourier transform of the mean row of $\mathbf{A}_{ii}$, where each row j was first rotated (rolled) to the left by j units. This allowed us to plot the spectral density of the analysis error covariance structure. This approach is adapted from Shlyaeva et al. [58]. Figure 4.10 illustrates the spectral densities of the background error covariances for the three state variables.

As a final note, it is important to consider that the assimilation of thickness observations may produce an analysis with negative thicknesses or with concentrations outside of [0,1]. At ECCC, these non-physical values are accounted for by clipping the state at the variable boundaries [5]. The implication for this thesis is that the analysis error covariance matrices may not exactly reflect the actual errors that would exist after this clipping.

4.5 Summary

This chapter presented an experimental methodology that can be followed to investigate the impacts of using mischaracterized observation error statistics in sea ice thickness data assimilation. The base twin experiment procedure consisted of a single EnKF cycle with the toy sea ice model where synthetic observations sampled from a true model state were assimilated into an ensemble of background model states. The base twin experiment was repeated for five scenarios corresponding to different realizations of a prescribed true observation covariance matrix. For each scenario, the assimilation was performed for three approximations to the corresponding true observation error covariance matrix.

The synthetic ice thickness sensor was designed to imitate SMOS-based thickness retrievals, with error variances parameterized as a function of the true ice thickness. The true decorrelation lengths of the observation error were set to either 150 km, 50 km, or 0 km. The true observation errors followed both Gaussian and exponential correlation structures. While the

The spread of the ensemble background error was approximately 0.20 m after the freeze-up period. This was comparable to the observation error variances in regions where the true ice thickness state was relatively thin. For most of the model domain, the background error variances were less than the observation error variances. We also observed large skewnesses and biases in the background sea ice thickness and concentration errors. These were attributed to the propagation of Gaussian noise through the nonlinear model. It would be sensible to further investigate the distribution and bias of sea ice model errors using an ensemble of full-scale models.

To compare the analyses from the twin experiments we required only the analysis error covariance matrices, which could be obtained without sampling a background state or observations. With just the analysis error covariance matrices we were able to compare the analysis error standard deviations and the spatial scales of the anlaysis errors for each experiment scenario.

Chapter 5

Experimental Results

This chapter presents the results from the twin experiment methodology described in Chapter 4. Following the presentation and discussion of the results, we propose several recommendations for the sea ice prediction community.

5.1 Analysis Error Standard Deviations

Tables 5.1 through 5.5 list the mean analysis error standard deviations for each of the experiment scenarios.

We first investigated the effect of spatial correlation in the observation errors by comparing the analysis error standard deviations for the cases where true observation error covariance matrix was used in the analysis. These analysis error standard deviations (σ^a, Equation 4.12) correspond $\mathbf{A}_{ideal}$ described in Chapter 4. In the first set of experiments (Table 5.1), the relative differences between the mean background and analysis error standard deviations were 29 %, 17 %, and again 17 % for thickness, concentration, and velocity, respectively. This first experiement experienced the largest decrease in background errors. This first experiment also corresponds to the Gaussian observation error correlation structure with the 150 km decorrelation length.

The observation error covariance matrix that yielded the smallest relative difference between background and analysis errors was $\mathbf{R}_{true_5}$, having a diagonal structure. From Table 5.5, the mean background error standard deviations fell by only 10 %, 6 %, and 9 % for this case. This assimilation effectively had a three-fold decrease in effectiveness compared to first experiment, which had correlated observation errors. The other observation error covariance matrices produced analysis increments of intermediate magnitude. Generally speaking, the observation error

Table 5.1: Estimated and actual analysis errors for the twin experiments with $\mathbf{R}_{true_1}$. $\mathbf{R}_{true_1}$ had a Gaussian correlation structure and a decorrelation length of 150 km

$\mathbf{R}_{est}$		h (10^{-1} m)	a (10^{-2})	u (10^{-2} m/s)
-	σ^b	2.00	7.28	1.38
$\mathbf{R}_{diag}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.89	6.93	1.34
$\mathbf{R}_{diag} * 2.4$	σ^a_{est}	1.86	7.00	1.30
	σ^a_{actual}	1.87	6.96	1.31
$\mathbf{R}_{true_1}$	σ^a_{est}	1.43	6.01	1.14
	σ^a_{actual}	1.43	6.01	1.14

Table 5.2: Estimated and actual analysis errors for the twin experiments with $\mathbf{R}_{true_2}$. $\mathbf{R}_{true_2}$ had a Gaussian correlation structure and a decorrelation length of 50 km

$\mathbf{R}_{est}$		h (10^{-1} m)	a (10^{-2})	u (10^{-2} m/s)
-	σ^b	2.00	7.28	1.38
$\mathbf{R}_{diag}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.87	6.97	1.31
$\mathbf{R}_{diag} * 2.4$	σ^a_{est}	1.86	7.00	1.31
	σ^a_{actual}	1.86	6.97	1.30
$\mathbf{R}_{true_2}$	σ^a_{est}	1.67	6.49	1.24
	σ^a_{actual}	1.67	6.49	1.24

Table 5.3: Estimated and actual analysis errors for the twin experiments with $\mathbf{R}_{true_3}$. $\mathbf{R}_{true_3}$ had an exponential correlation structure and a decorrelation length of 150 km

$\mathbf{R}_{est}$		h (10^{-1} m)	a (10^{-2})	u (10^{-2} m/s)
-	σ^b	2.00	7.28	1.38
$\mathbf{R}_{diag}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.87	6.94	1.32
$\mathbf{R}_{diag} * 2.4$	σ^a_{est}	1.86	7.00	1.30
	σ^a_{actual}	1.86	6.96	1.30
$\mathbf{R}_{true_3}$	σ^a_{est}	1.78	6.76	1.26
	σ^a_{actual}	1.78	6.76	1.26

Table 5.4: Estimated and actual analysis errors for the twin experiments with $\mathbf{R}_{true_4}$. $\mathbf{R}_{true_4}$ had an exponential correlation structure and a decorrelation length of 50 km

$\mathbf{R}_{est}$		h (10^{-1} m)	a (10^{-2})	u (10^{-2} m/s)
-	σ^b	2.00	7.28	1.38
$\mathbf{R}_{diag}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.84	6.93	1.29
$\mathbf{R}_{diag} * 2.4$	σ^a_{est}	1.86	7.00	1.30
	σ^a_{actual}	1.85	6.95	1.29
$\mathbf{R}_{true_4}$	σ^a_{est}	1.82	6.88	1.28
	σ^a_{actual}	1.82	6.88	1.28

Table 5.5: Estimated and actual analysis errors for the twin experiments with $\mathbf{R}_{true_5}$. $\mathbf{R}_{true_5}$ is diagonal.

$\mathbf{R}_{est}$		h (10^{-1} m)	a (10^{-2})	u (10^{-2} m/s)
-	σ^b	2.00	7.28	1.38
$\mathbf{R}_{diag}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.81	6.87	1.26
$\mathbf{R}_{diag} * 2.4$	σ^a_{est}	1.86	7.00	1.30
	σ^a_{actual}	1.83	6.92	1.27
$\mathbf{R}_{true_5}$	σ^a_{est}	1.81	6.87	1.26
	σ^a_{actual}	1.81	6.87	1.26

covariance matrices with Gaussian spatial correlations resulted in smaller analysis error standard deviations than the observation error covariance matrices with exponential structures. The implication is that observation errors with greater spatial correlations can yield larger analysis increments, compared to observations with less- or uncorrelated errors.

Stewart et al. [13][50] made similar observations in a twin experiment with a shallow water model. The best analyses occurred when the observation error covariance matrix had the largest off-diagonals, which can be represented by the Frobenius norm [50]. These authors explained this phenomenon from an information theory perspective. Effectively, the use of spatially correlated observation incorporates more information into the analysis than the use of observations with uncorrelated errors [13][50].

After investigating the impact of correlated observation errors in these ideal scenarios, we looked at the impact of the diagonal approximation. The impact of diagonal approximations to the true observation error covariance matrices varied across the five experiments. For the fifth experiment (Table 5.5), $\mathbf{R}_{true_5}$ was diagonal and accordingly σ^a_{est} was equal to σ^a_{actual}. Here, the diagonal approximation had no effect. For ice thickness, $\sigma^a_{est} = \sigma^a_{actual} = 0.181$ m, which represents a 10% improvement over the background state. The experiment that produced the largest σ^a_{actual} under the diagonal approximation was the first experiment (Table 5.1). Again, $\mathbf{R}_{true_1}$ had the greatest Frobenius norm. For the ice thickness estimate in this experiment, $\sigma^a_{actual} = 0.189$ m, which represents only a 6% improvement in background error versus the estimated 10% improvement. This reflects overconfidence in the analysis. If an EnKF were used in an operational setting, extra attention would need to be paid to avoid filter divergence. Again we saw intermedi-

ate impacts in the other experiments, with more sensitivity when the observation error correlation structure was Gaussian and for lonegr decorrelation lengths. In none of the experiments did we see mean actual analysis error standard deviations greater than the background error standard deviation.

Inflating the diagonal observation error covariance matrices had a generally positive effect. When the inflated diagonal approximation to the observation error covariance matrix was used in the assimilation, the estimated mean analysis error standard deviations were approximately 0.186 m, 7.00%, and $1.30{*}10^{-2}$ m/s, for thickness, concentration, and velocity, respectively. This reflected relative improvements over σ^b of 7%, 4%, and 6%, respectively. While these were not large analysis increments, the main benefit is that σ^a_{actual} was approximately equal to σ^a_{est}. For each of the experiments, σ^a_{est} was consistently slightly larger than σ^a_{actual}, which is perhaps desireable since it represents a conservative underestimate of the analysis quality.

For inflation experiments, an interesting note is that σ^a_{actual} was nearly identical across experiments, even though only a single inflation factor of 2.4 was used. This suggests that perhaps analysis error is not sensitive to the choice of inflation factor. The choice of inflation factor is discussed in greater detail in Section 5.3.

The cross-variable impact of assimilating thickness observations was surprisingly strong. Compared to the decreases in σ^b observed for ice thickness, the decrease in σ^b for the other state variables had relative magnitudes between 0.5 and 0.8. One benefit of this cross-variable effect is that thickness assimilation can clearly be used to improve the secondary state variables. A second benefit is that the assimilation of ice concentration or velocity observations might have a similar effect on the ice thickness state. The corresponding drawback of the effect of cross-correlation is that if a poor estimate of the observation error covariance matrix is used in the assimilation, the analysis error standard deviations for the secondary state variables will also be incorrect. This was demonstrated in Tables 5.1 through 5.4. For instance, in Table 5.1, for ice velocity, σ^a_{est} had a value of $1.26 * 10^{-2}$ m/s when σ^a_{actual} was in fact $1.34 * 10^{-2}$ m/s. This corresponds to estimated and actual relative decreases in mean background error standard deviation of 9% and 3%, which are considerably different.

5.2 Spectral Densities of Analysis Errors

For an initial demonstration, we first considered the case presented in Table 5.5, where the estimated and true observation error covariance matrices were diagonal. Figure 5.1 depicts the analysis error correlation matrix for this scenario, computed using Equation 4.4. The analysis error correlation matrix can be compared to the background error correlation structure presented

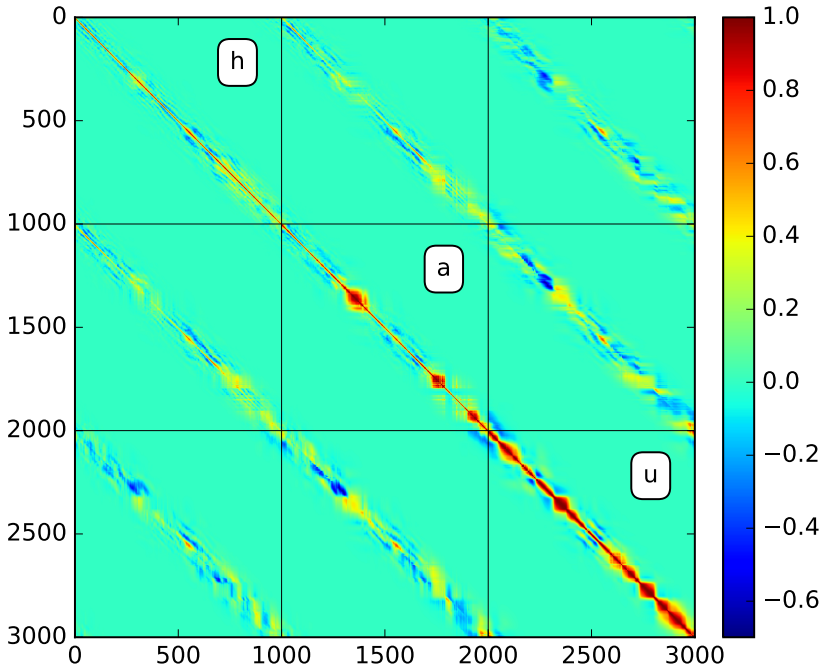

Figure 5.1: Analysis error correlation matrix corresponding to the analysis error covariance matrix for the case where the observation error covariance matrix was diagonal. This figure can be compared with the background error correlation matrix presented in Figure 4.4. The differences bewteen these two matrices are very subtle, indicating that the analysis has not largely changed the error correlation structure

in Figure 4.4. The differences between these two matrices are subtle, which is reflective of the fact that the assimilation of observations here had only a small effect on the background errors.

A clearer perspective can be obtained by comparing the spectra of the background and analysis error covariance matrices. Figures 5.2a and 5.2b illustrate the differences between the background and analysis error variance spectral densities. Here we plot the spectra of the mean and median rows of the background, observation, and analysis error covariance matrices.

The spectral density function of the mean row of the diagonal observation error covariance matrix here was flat, which was expected considering that a diagonal covariance matrix represents a white-noise process. The spectral densities for the observation error covariance matrix did not exist for any wavenumbers greater than $4 * 10^{-2}$ km^{-1}. This wavenumber corresponds to a wavelength of 25 km, which is twice the observation spacing. The implication here is that the

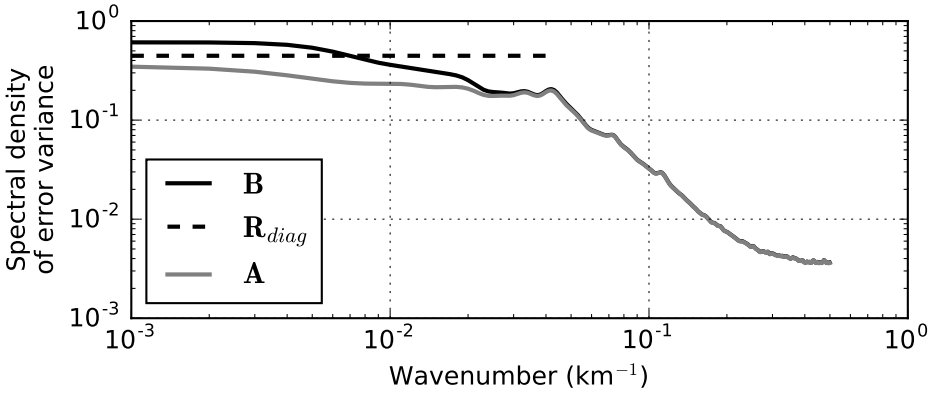

(a) Spectral densities of the mean rows of three error covariance matrices

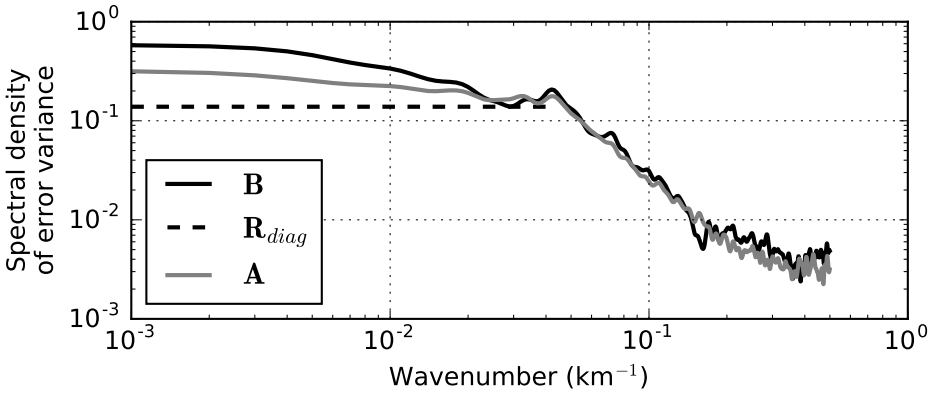

(b) Spectral densities of the median rows of three error covariance matrices

Figure 5.2: Spectral densities of the mean and median rows of the background, observation and analysis error covariance matrices for ice thickness and for $\mathbf{R}_{true_5}$. The largest improvements occured at the lowest wavenumbers and the analysis provided no improvement at wavenumbers higher than 0.04 km^{-1} (wavelength of 25 km).

assimilation of SMOS data cannot improve the high-frequency errors in the background state.

Daley [14] and Shlyaeva et al. [58] demonstrated that when the spectral density of the background errors was less than the spectral density of the observation errors, the spectral density of the analysis was closer to that of the background. Similarly, when the spectral density of the observation errors was less than that of the background errors, the spectral density of the analysis errors was less than the background. We did not observe this phenomenon but perhaps this can be attributed to the non-constant error variances in our study. Shlyaeva et al. [58] used constant error variances for both the background and the observation error covariance matrices. We plotted the spectral densities of the median rows of these matrices to provide alternative perspectives. In Figure 5.2b, the spectral density of the median row of the observation error covariance matrix was less than that of the background error covariance matrix over the entire domain. Hence the spectral density function of the analysis error covariance matrix did not intersect that of the backgroudn error covariance matrix except in one small region.

Appendix A contains a set of figures representing the spectra of the mean rows of the background and analysis covariance matrices for each of the twin experiments and for each state variable. This is equivalent to taking the column-wise mean of the real component of the Fourier transform of each row of these covariance matrices. From these plots we can see where the estimated and actual analysis error covariance matrices depict an improvement (or degradation) over the background errors. Each figure contains three plots. The first plot represents the case where we assumed a diagonal observation error covariance matrix. The second plot represents the case were we assumed a diagonal observation error covariance matrix with inflated variances. The final plot depicts the scenario where the true observation error covariance matrix was used in the assimilation.

Several general observations can be interpreted from the plots in Appendix A. The first is that assimilating the observations with the correct observation error covariance matrix resulted in the smallest analysis error spectral density at all scales. This was not unexpected considering the findings of Stewart et al. [50].

The second general observation is that for most of the analyses, the greatest decrease in spectral density occurred at the lowest wavenumbers. As the wavenumber approached $0.04\ \text{km}^{-1}$ (25 km), the assimilation had a decreased impact. There were a few exceptions. When the true observation error covariance matrix had a Gaussian structure, there were several unexplained decreases in the spectral density at intermediate frequencies. These could have been caused by the particular construction of the observation error covariance matrix, recalling Equation 4.9. Effectively, some of the observations had much lower error variances than others. Furthermore, the spatial correlation of the true ice thickness state contributed to the spatial concentration of the observation errors.

When we used a diagonal approximation to the true observation error covariance matrix by neglecting the off-diagonals, the naïve estimate of the analysis error covariance matrix, $\mathbf{A}_{est}$, predicted a positive benefit from the assimilation at all frequencies less than 0.04 km^{-1}. However, when we consider the actual analysis error covariance matrix, $\mathbf{A}_{actual}$, the spectral density of the analysis was in fact larger than predicted at low frequencies. For the cases where the decorrelation length was 150 km, the actual spectral densities at the lowest frequencies were actually greater than those of the background state. The implication is that the analysis state may contain a greater amount of low-frequency errors when a diagonal observation error covariance matrix is blindly specified.

The use of inflation also had a generally positive effect. The estimated and actual analysis error spectral densities were much closer than they were when there was no inflation. However, there were still differences between the estimated and actual analysis error spectral densities. For the case where the true observation error covariance matrix was diagonal, we observed a slight increase in the actual analysis spectral densities at all wavenumbers less than 0.04 km^{-1}.

The results for concentration and velocity were again similar to those for thickness but with decreased magnitude. For instance, comparing Figures A.1c and A.6c, the decrease in analysis error spectral density is not as large for concentration as it is for ice thickness. The effects on the spectral densities of the velocity errors seem to be even less pronounced even those for ice concentration. This includes both the positive effects of the assimilation and the negative effects of the diagonal approximation.

Overall, there did not appear to be any results from this spectral density study that might prohibit the use of diagonal approximations to the observation error covariance matrix. A negative result may have been if the diagonal approximation effectively increased the density of actual analysis errors at the wavelengths corresponding to length scales of 25-50 km. This could have led to the prediction of navigable channels through the ice that might not actually exist. Another negative result might have been if the diagonal approximation resulted in high-frequency errors in the concentration or ice velocity fields, which could lead to predictions of a more dynamic ice field.

5.3 Optimal Inflation Factors

A single inflation factor of 2.4 was used in the experiments described above. However, for every assimilation problem there is a possibly unique inflation factor that minimizes the analysis error standard deviation [59]. Iterative methods can be applied to identify the optimal inflation factor (e.g. [59]). However, these methods require repeatedly solving the sequential assimilation

equations, which may be too costly for an operational system. It would be more convenient to use a single inflation factor in practice. In this section we investigate the impact of the choice of inflation factor.

To investigate the impact of the inflation factor, and to identify the inflation factor that minimizes the actual analysis error variance, we repeated the twin experiment for each scenario with a range of inflation factors spanning [0.1,6]. Figures 5.3a to 5.3b illustrate how the estimated and actual analysis error standard deviations vary as a function of the inflation factor.

For the inflation factor experiment with $\mathbf{R}_{true_5}$, which is diagonal, the minimum value of actual analysis error standard deviation occured when the inflation factor had a value of 1. At this value, the actual and estimated analysis error standard deviation curves intersected. This makes sense for this scenario because when the true observation error covariance matrix is diagonal, inflation of the error variances should only increase the analysis error.

For the experiments with the other true observation error covariance matrices, at an inflation factor of 1, the estimated analysis error standard deviation was less than the actual analysis error standard deviation. The difference between the two error terms varied depending on the structure of the true observation error covariance matrix. When the observation error covariance matrix had the highest Frobenius norm (Gaussian, 150 km), the difference was the greatest. For $\mathbf{R}_{true_4}$, which had an exponential shape with a 50 km decorrelation length, the impact was the smallest.

For each experiment there was a different optimal inflation factor that minimized the actual analysis error standard deviation. For $\mathbf{R}_{true_1}$, the greatest optimal inflation factor of 2.4 was observed. $\mathbf{R}_{true_5}$ had an optimal inflation factor of 1. Intermediate optimal inflation factors were observed for the other experiments. These optimal inflation factors are consistent with Liang et al. [59], who found typical optimal inflation factors between one and five, and with Stewart et al. [13], who used inflation factors between two and four.

When the inflation factor decreased below a value of 0.5, the actual analysis error standard deviation increased asymptotically past the background error standard deviation while the estimated analysis error standard deviation decreased asymptotically. In this region, we are effectively grossly overconfident about a very poor analysis. This was observed in each of the experiments. The implication is that it is very important not to underestimate the observation error variance when assuming a diagonal observation error covariance matrix.

When the inflation factor was greater than its optimal value, both the estimated and actual error standard deviations increased at very similar low rates. It appears that the system was very insensitive to overinflating the observation error covariance matrix. This is a promising observation because it means there is some flexibility in setting the inflation parameter. This also indicates value in specifying a conservatively high inflation factor.

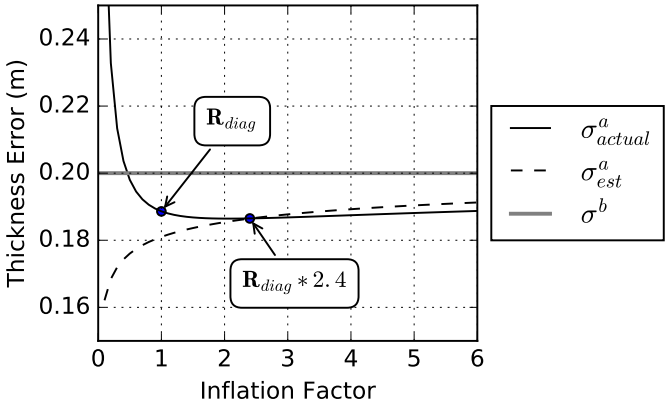

(a) σ^a for $\mathbf{R}_{true_1}$

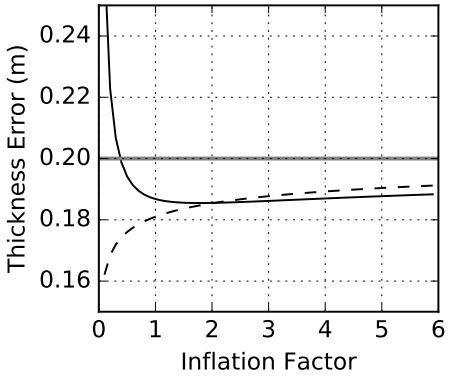

(b) σ^a for $\mathbf{R}_{true_2}$

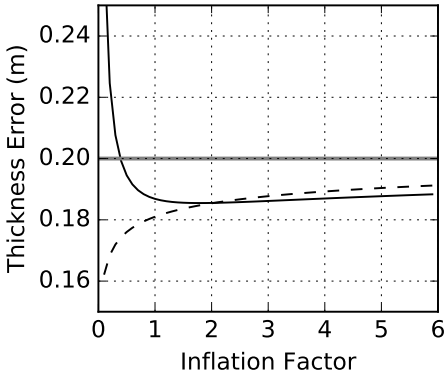

(c) σ^a for $\mathbf{R}_{true_3}$

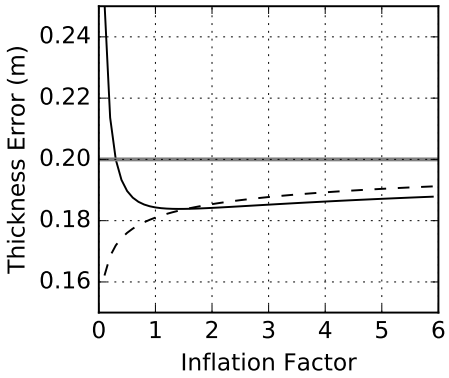

(d) σ^a for $\mathbf{R}_{true_4}$

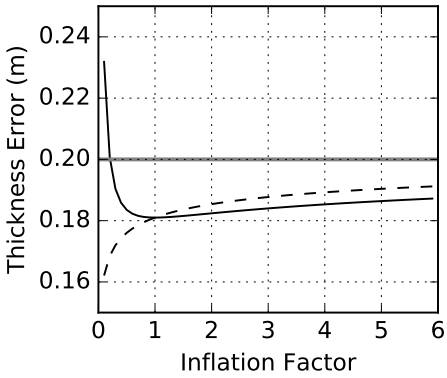

(e) σ^a for $\mathbf{R}_{true_5}$

Figure 5.3: Estimated and actual thickness analysis error standard deviation as a function of the inflation factor. Each experiment had an optimal inflation factor between 1 and 2.4. Underestimation of the inflation factor resulted in a divergence between the estimated and actual analysis error standard deviation. Overestimation of the inflation factor had a lesser effect.

5.4 Discussion

The results from these experiments suggest that it may be safe, although not optimal, to assimilate SMOS-based ice thickness observations using a diagonal observation error covariance matrix so long as the observation error variances are inflated by a sufficiently large factor. This is a promising overall result but it should be interpreted with cautious optimism. This section presents a discussion of the limitations of the twin experiments. We later propose a series of recommendations that might lead to more optimal use of the observational data.

5.4.1 Limitations

Chapters 1 and 2 described how the size and cost of quasi-operational prediction systems have limited the depth and thoroughness of assimilation experiments in the literature. We used this to justify the creation of the toy model and the twin experiment framework. Using this tool, we were able to conduct a range of experiments and to investigate the results in greater depth than we may have been able to accomplish using a different experimental framework. However, the results were still not completely general. Here we discuss how the twin experiments could have been improved.

The first limiation was related to the true state. The observation error variances and covariances were defined as a function of the thicknesses in the true state. However, sea ice can exhibit a wide range of topographies depending on many factors including the coastal geometry, the ice age, and the weather. Additionally, the ice sensor is much more effective for for thinner ice states. It is possible that entirely different results may have been obtained with a thinner or thicker true ice state or with an ice state that had a different spatial correlation of thicknesses. There may be value in repeating the experiments for different initial ice conditions. With a mean thickness of approximately 0.7 m and with many concentrations less than 100%, the results from the experiments presented herein may be reflective of assimilation in the marginal ice zone or in regions where there is no established first- or multi-year ice pack.

A similar limitation is related to the background error covariance matrix. The background error covariance matrix generated in Chapter 4 had an approximately exponential correlation structure and what seemed to be strong correlations across state variables. However, as with the true state, different results may have been obtained with a different background error covariance matrix. The background error covariance matrices for RIPS and GIOPS at ECCC were created by solving the diffusion equation on the model grid, producing a smoother error covariance shape [21][5]. However, there was limited justification provided for using that structure. In the sea ice EnKF literature, the authors did not present their mean error correlation structures [8][12][11],

and thus there was little to compare against. The authors did present their mean background error standard deviations, for sea ice thickness, which ranged from 0.4 to 0.8 m [8][12][11]. These errors were much greater than our background error standard deviation of 0.2 m. However, our experiment represented solely new/marginal ice whereas the experiments in the literature spanned the entire Arctic, which in general has much thicker ice. Again, the implication is that the results presented herein represent what *may* happen in operational systems.

Beyond the background variances and covariance structure, the results from Chapter 4 identified significant bias and skewness in the freeze-up ensemble. Large biases are acknowledged to exist in ice prediction systems (e.g. [21]) and they are not being actively accounted for operationally. We explicitly removed bias in both the background and observation errors and it is not clear how the results may have varied if biases were actually present. Regarding the observed skewnesses, we did not find any material in the literature to indicate what the expected distribution of errors should be. The distribution of errors is perhaps worth attention in future experiments. As explained in Chapter 4, a non-Gaussian error distribution does not prohibit use of the Kalman Filter but rather there may exist alternative filters (e.g. the particle filter) that could provide a better analysis.

In Sections 5.1 and 5.2, we assessed the analysis quality in terms of the analysis error standard deviation and in terms of the spectral density of error variances in the analysis error covariance matrix. Literature tends to focus more on just the RMSE or the RMSD and the bias (e.g. [21], [12]). To that end we have provided a novel perspective but it remains unclear what metrics are actually most valuable to know about a sea ice state. For instance, what information would lead a ship captain to successfully navigate a ship through an ice field? What types of model state errors might lead a captain to make decisions that would result in his or her ship becoming immobilized in ice? Section 5.2 demonstrated that diagonal approximations can increase the low-frequency errors in the ice state. But what are the impacts of this? From Chapter 3, the viscosity of sea ice is a function of the spatial derivatives of concentration and thickness. Following this definition, a smoother ice state may be physically much stiffer and less dynamic. In Chapter 3, we also presented five metrics (e.g. absolute divergence) that can be used to gauge the quality of an ice state. It may be worthwhile to further evaluate which qualities of an ice model state are the most valuable to preserve in a sub-optimal assimilation.

This leads to the fact that we did not investigate the temporal effects of assimilating ice thickness observations with diagonal observation error covariance matrices. Using the toy model it would be feasible to either perform forecasts or to perform daily analyses over some period to see the longer-term effects of the diagonal approximation. However, we felt that for the purposes of this thesis it was more worthwhile to more rigorously assess the quality of the analyses. An EnKF experiment might be the most worthwhile future experiment. However, we note that running an EnKF, which requires careful tuning to avoid filter divergence, would comprise a

substantial effort.

We also did not consider the variety of thickness retrieval algorithms, even when considering only the SMOS sensor (e.g. [29] and [60]). Each of these algorithms might produce different estimates of the observation error variances. We did not explicitly consider the effect of mischaracterizing the error variances. However, the optimal inflation factor results presented in Section 5.3 provide some indication of what might happen if the observation error variances are under- or over-estimated uniformly across the entire domain.

An additional limitation relates to the assumption in the derivation of the Kalman Filter that the background and observation errors are uncorrelated [7]. Chapter 2 explains that it is possible that there are correlations between SMOS-based ice thickness observations and the ice state. The correlations are related to the retrieval algorithms requiring ancillary data including snow depth, the ice thickness distribution, the ice salinity, and most importantly the ice concentration. Our synthetic sensor did not rely on any variables apart from the true ice thickness, meaning that we effectively avoided the need to question this assumption.

The final and the most obvious limitation of these results are that they were obtained using a simplified 1-D ice model, rather than using a quasi-operational model and actual ice thickness data. This limitation was expected and desired. What we have presented herein may be used to spark conversation and to guide future experiments on full-scale prediction systems. In fact, a secondary contribution of this thesis is that by looking at this problem in detail, we have identified a number of unrelated topics (e.g. potential skewness in the background error distribution) that may lead to better numerical prediction in the future.

Despite the many limitations identified above, nothing seems to indicate that the main findings of this thesis should be outright rejected. Working with diagonal approximations to the ice thickness observation error covariance matrix is likely safe so long as an appropriate inflation factor is used. Using the correct observation error covariance matrix might provide a substantial improvement in the analysis quality but the increment from assimilation with an inflated, diagonal observation error covariance matrix does not appear to be negligible no matter the true observation error covariance structure.

5.4.2 Recommendations

The previous section presented a number of recommendations that can be followed to improve the twin experiments with the toy sea ice model. The current section combines the findings from the twin experiments with the material presented in previous Chapters to present some additional practical recommendations that might be worth exploring in future work.

Estimation of Observation Error Covariance Structure

Until we can properly estimate the true observation error covariance matrix, predictions of the sea ice state will continue to be suboptimal. This applies not only to sea ice thickness but also to ice concentration and ice velocity observations.

We propose two methods for estimating these error covariance terms. The first method is the *Desroziers diagonstic* approach, which involves estimating the observation error covariance matrix as [19]

$$\mathbf{R}_{est} = \frac{1}{N_{ens} - 1} \sum_{i=1}^{N_{ens}} \mathbf{d}_i^a \mathbf{d}_i^{bT} \tag{5.1}$$

where $\mathbf{d}^a = \mathbf{y} - H(\mathbf{x}^a)$, $\mathbf{d}^b = \mathbf{y} - H(\mathbf{x}^b)$, and $\mathbf{y}$ represents a vector of real observations (i.e. not simulated). Here, the observations must first be assimilated using an approximation to the observation error covariance matrix, perhaps diagonal. The resulting estimate is a better approximation to the true observation error covariance amtrix. This procedure would require actual data and a full scale sea ice model. While this still only provides an approximation to the observation error covariance structure, it would at least provide some indication of what scales are present in the observation errors.

An alternative approach might involve passing the error covariance matrices of the ancillary data through the retrieval algorithm to produce an estimate of the ice thickness error covariances that are attributable to the ancillary data. This could be accomplished in a Monte Carlo simulation. We could accomplish this given ensembles of model states of the ice, ocean and atmosphere. This would not account for errors in the actual MIRAS sensor aboard SMOS. However, again, this would provide an initial idea of what the observation error covariances might look like.

Approximate Solutions to the Sequential Assimilation Equations

One of the main findings of this thesis is that it is perhaps extremely valuable to identify the true observation error covariance structure and to incorporate this into an operational assimilation system. However, for numerical reasons it is perhaps naïve to assume that a complete observation error covariance matrix with long-scale correlations could be used in an operational setting. ECCC uses diagonal observation error covariance matrices in RIPS and GIOPS in part because the variational assimilation problem requires the inverse of the observation error covariance matrix. This inversion is trivial for a diagonal matrix but for a covariance structure similar to that of $\mathbf{R}_{true_1}$, the matrix becomes too-poorly conditioned to invert. Rather than inverting these matrices exactly, it is possible to find approximate solutions to the system of linear equations [43]. The

discussion of preconditioning and efficient numerical solution to the sequential and variational data assimilation equations is beyond the scope of this thesis but may require some consideration in the future.

Alternatives

We have described in this section and in previous chapters, from a data assimilation perspective, a breadth of issues associated with the ice thickness retrieval algorithms and ancillary data. However, we note here that a critical presumption of this thesis has been that the only way to use SMOS data (or VisIR data, etc.) is to feed the sensor data through a retrieval algorithm and to assimilate the resulting thickness estimates into a model state. In fact this is not correct and there are alternatives to using these retrieval algorithms.

The first alternative is to directly assimilate the satellite sensor data (brightness temperature, etc.) into the model state through the use of the nonlinear observation operator, $H(\mathbf{x})$. This approach was applied by Scott et al. [61] to directly assimilate AMSR-E (a passive microwave satellite sensor) data into a combined sea ice-ocean state. The approach is to specify $H(\mathbf{x})$ as a radiative transfer model that converts the ice state into a satellite brightness temperature [61]. Conveniently, these transfer functions already exist, created during the development of the SMOS-based ice thickness retrieval algorithms [29][31]. The transfer function is effectively the inverse of the retrieval algorithm.

The key benefit of this alternative approach to data assimilation is that we no longer need to know the error covariances of the ancillary data such as snow depth or ice density. We would only need to understand the brightness temperature error covariances. The secondary sea ice state variables would become part of the background state vector and could be generated using an EnKF.

This approach would solve a number of other issues plaguing ice thickness assimilation. Firstly, it would now be reasonable to assume that there is no correlation between the observation errors and the background errors. Secondly, this would solve the discrepancy between with the ice thickness distribution (a series of partial concentrations) and mean ice thickness. The ice thickness distribution partial concentrations could be incorporated into the background state vector without any sort of transformation.

A second alternative that is not mutually exclusive to the first is to make use of the cross-correlation terms in the ensemble-derived background error covariance matrix and to focus on assimilation of improved ice concentration datasets. It is now possible to derive high-resolution ice concentration products from synthetic aperture radar imagery using Convolutional Neural Networks (e.g. [6]). Supposing that we can correctly estimate the observation error covariance

matrix for these products, it is possible that the cross-variable (i.e. concentration to thickness) increment might be greater than what could be achieved with SMOS, especially considering the difference in spatial resolution of the data.

For this second alternative, we note that little is known about the structure of the ice concentration observation error covariances. Operational ice prediction systems (e.g. RIPS [21]) continue to assimilate ice concentration data under the diagonal approximation. The final recommendation of this thesis is therefore an appeal to place a larger emphasis on the estimation of observation error covariances.

5.5 Summary

This chapter presented the results from five sets of twin experiments conducted with the toy sea ice model. We investigated the actual and perceived analysis error quality that is observed when the corresponding observations are assimilated using diagonal approximations to the observation error covariance matrices. To compare the analyses we computed the expected and actual analysis error standard deviations and the spectral densities of the error variances in the analysis error covariance matrices. We also investigated the sensitivity of the analyses to the chosen diagonal observation error covariance matrix inflation factor.

Several observations were interpreted from the results. The first was that using the correct observation error covariance matrix, especially when there is a long-distance error correlation structure, yields by far the best analyses. The second conclusion was that when a diagonal approximation is used, the estimated analysis error variance is less than the actual analysis error standard deviation. In this scenario, we saw an increase in the low-frequency errors in the state estimate. When an inflation factor was used, the actual and estimated analysis errors were approximately equal, even for the case where the true observation error covariance matrix was diagonal. We found that the system was very sensitive to underestimating the inflation factor but not nearly as sensitive to overestimation. We considered this promising for the design of operation assimilation systems that are restricted to the use of diagonal observation error covariance matrices.

Despite the fact that sea ice forecasting systems are quickly moving toward the assimilation of thickness observations, the discussion in this chapter revealed that there is perhaps a great deal of basic research to be conducted before ice thickness observations can be ingested optimally and confidently. That said, it appears possible to safely assimilate SMOS-based ice thickness observations using diagonal approximations to the observation covariance matrix, provided that a sufficiently-large inflation factor is used.

Chapter 6

Conclusions

The state of sea ice in the Arctic Ocean is not known with great accuracy and it is difficult to make reliable predictions of future conditions [1]. There is motivation now to begin assimilating ice thickness observations into ice model states to provide improved nowcasts and forecasts.

The retrieval of ice thickness observations from spaceborne sensors is not straightforward and there is lingering uncertainty about the usefuless of the data. When datasets are assimilated without exactly knowing the observation error covariances, it is possible that the assimilation can actually degrade the model state estimate [7]. It seems likely that when ice thickness observations are first assimilated in operational sea ice prediction systems, the observation error covariance matrix will be approximated by a diagonal matrix [21][5][12][62]. The objective of this thesis was to examine how this approximation might affect the quality of the analyses.

6.1 Summary

Chapter 1 introduced the concepts of both sea ice thickness and data assimilation. We then described at a high level some of the challenges facing sea ice thickness data assimilation. This led to the presentation of the motivations and objectives of this thesis.

Chapter 2 provided a more comprehensive summary of the challenges facing sea ice thickness remote sensing, data assimilation, and forecasting. We described the assumptions underlying sequential data assimilation and then demonstrated how these assumptions may be violated once sea ice thickness observations are ingested. We described why it is diffcult to estimate the observation error covariance matrix for satellite-based ice thickness estimates. We then described some of the state-of-the-art operational ice prediction systems in Canada and the experiments that

have recently been conducted to support ice thickness data assimilation. Finally, we described why it is difficult to conduct fully-rigorous data assimilation experiments with quasi-operational ice prediction systems.

Chapter 3 presented a simplified *toy* sea ice model that was developed to provide the basis for twin ice thickness data assimilation experiments. The toy sea ice model could be used in future experiments.

Chapter 4 described the design of the twin experiments, which comprised the analytical component of this thesis. The procedure was to assimilate thickness observations into a background ice state with and without the correct observation error covariance matrix, and then to compare the analysis error estimates. We described how an ensemble-based background error covariance matrix was obtained. We then described the synthetic ice thickness sensor, which was based on the SMOS satellite. Finally, we described the difference between the true analysis error covariance matrix and the estimated error covariance matrix, and how to compare the two.

Chapter 5 presented results from the twin experiments. There were several key findings. The first was that it is entirely reasonable to use a diagonal approximation to the observation error covariance matrix so long as the observation error variances are inflated. An inflation factor of at least 2.4 was found to be suitable in this study. The analysis quality was insensitive to the overestimation of this parameter. The diagonal approximation resulted in an increase in the low-wavenumber errors in the analysis state, relative to what could be obtained if the actual error covariance matrix was known. There was little impact at higher wavenumbers. A final finding was that the ice thickness analysis error can be greatly reduced when the actual observation error covariance matrix is known.

6.2 Future Work

The findings from this thesis point to two distinct regions for future work. The first is to move toward operational use of SMOS-based ice thickness observations under the diagonal approximation to the observation error covariance matrix. The second area of future work is to make more-optimal use of the satellite data.

As a first step toward operational assimilation of SMOS data, we recommend first repeating a selection of the twin experiments on a quasi-operational prediction system. This would help to reveal whether any of the results presented in this thesis were unique to the background state and background error covariance matrix. Following a positive result, it may be worthwhile to begin ingesting ice thickness observations within operational prediction systems.

There are several ways that we can make more effective use of satellite data in ice thickness data assimilation. The first is to attempt to estimate the error covariance structure for ice thickness observations. We demonstrated that much larger analysis increments can be achieved when the true observation error covariance structure is known. An alternative, should this prove infeasible, would be to investigate the direct assimilation of satellite brightness temperatures using transfer functions that incorporate the ice thickness. Other unrelated areas for future work include investigating the presence of skewness in the background error distribution.

To conclude this thesis, we restate that we are now in a position where we can begin making improved sea ice predictions through the use of remotely-sensed ice thickness observations. The diagonal approximation to the observation error covariance matrix provides a suitable starting point. However, we also stress that there is much opportunity to make better use of this satellite data.

References

[1] H. Eicken, "Artic sea ice needs better forecasts," *Nature*, vol. 497, May 2013.

[2] E. Hunke, W. Lipscomb, A. K. Turner, N. Jeffery, and S. Elliott, *CICE: the Los Alamos sea ice model documentation and software user's manual, version 5.1*. Los Alamos National Laboratory, 2015.

[3] D. Wingham, C. Francis, S. Baker, C. Bouzinac, D. Brockley, R. Cullen, P. de Chateau-Thierry, S. Laxon, U. Mallow, C. Mavrocordatos, L. Phalippou, G. Ratier, L. Rey, F. Rostan, P. Viau, and D. Wallis, "CryoSat: A mission to determine the fluctuations in earth's land and marine ice fields," *Advances in Space Research*, vol. 37, no. 4, pp. 841 – 871, 2006. Natural Hazards and Oceanographic Processes from Satellite Data.

[4] J.-F. Lemieux, C. Beaudoin, F. Dupont, F. Roy, G. C. Smith, A. Shlyaeva, M. Buehner, A. Caya, J. Chen, T. Carrieres, L. Pogson, P. DeRepentigny, A. Plante, P. Pestieau, P. Pellerin, H. Ritchie, G. Garric, and N. Ferry, "The regional ice prediction system (RIPS): verification of forecast sea ice concentration," *Quarterly Journal of the Royal Meteorological Society*, vol. 142, no. 695, pp. 632–643, 2016.

[5] G. C. Smith, F. Roy, M. Reszka, D. Surcel Colan, Z. He, D. Deacu, J.-M. Belanger, S. Skachko, Y. Liu, F. Dupont, J.-F. Lemieux, C. Beaudoin, B. Tranchant, M. Drévillon, G. Garric, C.-E. Testut, J.-M. Lellouche, P. Pellerin, H. Ritchie, Y. Lu, F. Davidson, M. Buehner, A. Caya, and M. Lajoie, "Sea ice forecast verification in the Canadian global ice ocean prediction system," *Quarterly Journal of the Royal Meteorological Society*, vol. 142, no. 695, pp. 659–671, 2016.

[6] X. Wang, J. Key, R. Kwok, and J. Zhang, "Comparison of Arctic sea ice thickness from satellites, aircraft, and PIOMAS data," *Remote Sensing*, vol. 8, no. 9, p. 713, 2016.

[7] W. Lahoz, B. Khattatov, and R. Menard, "Data assimilation and information," in *Data Assimilation: Making Sense of Observations* (W. Lahoz, B. Khattatov, and R. Menard, eds.), pp. 3–12, Springer Berlin Heidelberg, 2010.

[8] A. Shlyaeva, M. Buehner, A. Caya, J.-F. Lemieux, G. C. Smith, F. Roy, F. Dupont, and T. Carrieres, "Towards ensemble data assimilation for the Environment Canada regional ice prediction system," *Quarterly Journal of the Royal Meteorological Society*, pp. 1090–1099, 2015. QJ-15-0078.R1.

[9] D. Lubin and R. Massom, *Polar Remote Sensing: Volume I Atmsophere and Oceans*. Springer, 2006.

[10] D. A. Rothrock, J. Zhang, and Y. Yu, "The arctic ice thickness anomaly of the 1990s: A consistent view from observations and models," *Journal of Geophysical Research: Oceans*, vol. 108, no. C3, 2003. 3083.

[11] K. A. Lisaeter, G. Evensen, and S. Laxon, "Assimilating synthetic CryoSat sea ice thickness in a coupled ice-ocean model," *Journal of Geophysical Research: Oceans*, vol. 112, no. C7, 2007. C07023.

[12] Q. Yang, S. N. Losa, M. Losch, X. Tian-Kunze, L. Nerger, J. Liu, L. Kaleschke, and Z. Zhang, "Assimilating SMOS sea ice thickness into a coupled ice-ocean model using a local SEIK filter," *Journal of Geophysical Research: Oceans*, vol. 119, no. 10, pp. 6680–6692, 2014.

[13] L. M. Stewart, S. L. Dance, and N. K. Nichols, "Correlated observation errors in data assimilation," *International Journal for Numerical Methods in Fluids*, vol. 56, no. 8, pp. 1521–1527, 2008.

[14] R. Daley, *Atmospheric Data Analysis*. Cambridge University Press, 1991.

[15] S. Healy and A. White, "Use of discrete Fourier transforms in the 1d-var retrieval problem," *Quarterly Journal of the Royal Meteorological Society*, vol. 131, no. 605, pp. 63–72, 2005.

[16] G. K. Vallis, *Atmospheric and Oceanic Fluid Dynamics*. Cambridge, U.K.: Cambridge University Press, 2006.

[17] P. J. Antsaklis and A. N. Michel, *A Linear Systems Primer*. Birkhäuser Basel, 1st ed., 2007.

[18] R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Trans. ASME, Ser. D, J. Basic Eng*, p. 109, 1961.

[19] G. Desroziers, L. Berre, B. Chapnik, and P. Poli, "Diagnosis of observation, background and analysis-error statistics in observation space," *Quarterly Journal of the Royal Meteorological Society*, vol. 131, no. 613, pp. 3385–3396, 2005.

[20] A. C. Lorenc, "A global three-dimensional multivariate statistical interpolation scheme," *Monthly Weather Review*, vol. 109, no. 4, pp. 701–721, 1981.

[21] M. Buehner, A. Caya, L. Pogson, T. Carrieres, and P. Pestieau, "A new Environment Canada regional ice analysis system," *Atmosphere-Ocean*, vol. 51, no. 1, pp. 18–34, 2013.

[22]

[23] W. D. Hibler, "A dynamic thermodynamic sea ice model," *Journal of Physical Oceanography*, 1979.

[24] G. Evensen, "The Ensemble Kalman Filter: theoretical formulation and practical implementation," *Ocean Dynamics*, vol. 53, no. 4, pp. 343–367, 2003.

[25] C. Haas, S. Hendricks, H. Eicken, and A. Herber, "Synoptic airborne thickness surveys reveal state of Arctic sea ice cover," *Geophysical Research Letters*, vol. 37, no. 9, 2010. L09501.

[26] M. Vancoppenolle, S. Bouillon, T. Fichefet, H. Goosse, O. Lecomte, M. Maqueda, and G. Madec, *LIM: The Louvain-la-Neuve sea ice model*. Nucleus for European Modelling of the Ocean, 2012.

[27] C. Haas and M. Druckenmiller, "Ice thickness and roughness measurements," in *Field Techniques for Sea-Ice Research* (H. Eicken and M. Salganek, eds.), University of Chicago Press Books, 2010.

[28] R. Moritz and M. Wensnahan, *Sea-ice Thickness and Draft Statistics from Submarine ULS, Moored ULS, and a Coupled Model*. National Snow and Ice Center, Boulder, Colorado, 2006. `http://nsidc.org/data/g01360`.

[29] L. Kaleschke, N. MaaB, S. Hendricks, G. Heygster, and R. T. Tonboe, "A sea-ice thickness retrieval model for 1.4 GHz radiometry and application to airborne measurements over low salinity sea-ice," *The Cryosphere*, 2010.

[30] X. Wang, J. Key, and Y. Liu, "A thermodynamic model for estimating sea and lake ice thickness with optical satellite data," *Journal of Geophysical Research*, vol. 115, December 2010.

[31] L. Kaleschke, X. Tian-Kunze, N. MaaB, M. Makynen, and M. Drusch, "Sea ice thickness retrieval from SMOS brightness temperatures during the arctic freeze-up period," *Geophysical Research Letters*, vol. 39, no. L05501, 2012.

[32] X. Tian-Kunze, L. Kaleschke, N. Maaß, M. Mäkynen, N. Serra, M. Drusch, and T. Krumpen, "SMOS-derived thin sea ice thickness: algorithm baseline, product specifications and initial verification," *The Cryosphere*, vol. 8, no. 3, pp. 997–1018, 2014.

[33] X. Tian-Kunze, L. Kaleschke, and N. Maass, "SMOS daily sea ice thickness, updated 2016." `http://icdc.cen.uni-hamburg.de/1/daten/cryosphere/l3c-smos-sit.html`, 2013. University of Hamburg.

[34] M. Makynen, B. Cheng, and M. Simila, "On the accuracy of thin-ice thickness retrieval using MODIS thermal imagery over Arctic first-year ice," *Annals of Glaciology*, vol. 54, no. 63, pp. 87–96, 2013.

[35] R. Lindsay, C. Haas, S. Hendricks, P. Hunkeler, N. Kurtz, J. Paden, B. Panzer, J. Sonntag, J. Yungel, and J. Zhang, "Seasonal forecasts of arctic sea ice initialized with observations of ice thickness," *Geophysical Research Letters*, vol. 39, no. 21, 2012. L21502.

[36] J. J. Day, E. Hawkins, and S. Tietsche, "Will Arctic sea ice thickness initialization improve seasonal forecast skill?," *Geophysical Research Letters*, vol. 41, no. 21, pp. 7566–7575, 2014.

[37] J. Dempsey and H. Shen, eds., *Scaling Laws in Ice Mechanics and Ice Dynamics*, IUTAM, Springer, 2000.

[38] M. Leppäranta, *The drift of sea ice*. Springer Berlin Heidelberg, 2 ed., 2011.

[39] A. S. Thorndike, D. A. Rothrock, G. A. Maykut, and R. Colony, "The thickness distribution of sea ice," *Journal of Geophysical Research*, vol. 80, no. 33, pp. 4501–4513, 1975.

[40] J. Lemieux, B. Tremblay, S. Thomas, J. Sedlacek, and L. Mysak, "Using the preconditioned generalized minimum residual (GMRES) method to solve the sea-ice momentum equation," *Journal of Physical Oceanography*, 2008.

[41] J.-F. Lemieux, D. A. Knoll, M. Losch, and C. Girard, "A second-order accurate in time implicit–explicit (IMEX) integration scheme for sea ice dynamics," *Journal of Computational Physics*, vol. 263, pp. 375 – 392, 2014.

[42] W. Gautschi, *Numerical Analysis*. Springer, 2 ed., 2012.

[43] P. Fieguth, *Statistical image processing and multidimensional modeling*. Information science and statistics, New York: Springer, 2011.

[44] J. Chilès and P. Delfiner, *Modelling Spatial Uncertainty*. John Wiley and Sons, Inc., 2008.

[45] J. Weiss, *Drift, Deformation, and Fracture of Sea Ice*. Springer, 2013.

[46] A. J. Plueddemann, R. Krishfield, T. Takizawa, K. Hatakeyama, and S. Honjo, "Upper ocean velocities in the Beaufort Gyre," *Geophysical Research Letters*, vol. 25, no. 2, pp. 183–186, 1998.

[47] P. Rampal, J. Weiss, D. Marsan, and M. Bourgoin, "Arctic sea ice velocity field: General circulation and turbulent-like fluctuations," *Journal of Geophysical Research: Oceans*, vol. 114, no. C10, 2009. C10014.

[48] L. Girard, J. Weiss, J. M. Molines, B. Barnier, and S. Bouillon, "Evaluation of high-resolution sea ice models on the basis of statistical and scaling properties of Arctic sea ice drift and deformation," *Journal of Geophysical Research: Oceans*, vol. 114, no. C8, 2009. C08015.

[49] D. B. Percival, D. A. Rothrock, A. S. Thorndike, and T. Gneiting, "The variance of mean sea-ice thickness: Effect of long-range dependence," *Journal of Geophysical Research: Oceans*, vol. 113, no. C1, 2008. C01004.

[50] L. Stewart, S. Dance, and N. Nichols, "Data assimilation with correlated observation errors: experiments with a 1-D shallow water model," *Tellus A*, vol. 65, no. 0, 2013.

[51] P. Houtekamer and H. Mitchell, "Data assimilation using an ensemble Kalman filter technique," *Monthly Weather Review*, vol. 126, no. 3, pp. 796–811, 1998.

[52] Q. Yang, S. N. Losa, M. Losch, T. Jung, and L. Nerger, "The role of atmospheric uncertainty in arctic summer sea ice data assimilation and prediction," *Quarterly Journal of the Royal Meteorological Society*, vol. 141, no. 691, pp. 2314–2323, 2015.

[53] E. Simon and L. Bertino, "Application of Gaussian anamorphosis to assimilation in a 3-D coupled physical-ecosystem model of the North Atlantic with the EnKF: a twin experiment," *Ocean Science*, vol. 5, no. 4, pp. 495–510, 2009.

[54] F. Massonnet, H. Goosse, T. Fichefet, and F. Counillon, "Calibration of sea ice dynamic parameters in an ocean-sea ice model using an ensemble Kalman filter," *Journal of Geophysical Research: Oceans*, vol. 119, no. 7, pp. 4168–4184, 2014.

[55] J. Amezcua and P. van Leeuwen, "Gaussian anamorphosis in the analysis step of the EnKF: a joint state-variable/observation approach," *Tellus A*, vol. 66, no. 0, 2014.

[56] W. Navidi, *Statistics for Engineers and Scientists*. McGraw-Hill, 3 ed., 2010.

[57] S. Rainwater, C. H. Bishop, and W. F. Campbell, "The benefits of correlated observation errors for small scales," *Quarterly Journal of the Royal Meteorological Society*, vol. 141, no. 693, pp. 3439–3445, 2015.

[58] A. Shlyaeva, M. Buehner, and C. Snyder, "Correlated observation errors in the perturbed observation ensemble data assimilation," in *International symposium on data assimilation 2016*, (Reading, UK), 2016.

[59] X. Liang, X. Zheng, S. Zhang, G. Wu, Y. Dai, and Y. Li, "Maximum likelihood estimation of inflation factors on error covariance matrices for ensemble Kalman filter assimilation," *Quarterly Journal of the Royal Meteorological Society*, vol. 138, no. 662, pp. 263–273, 2012.

[60] L. Kaleschke, X. Tian-Kunze, N. Maas, R. Ricker, S. Hendricks, and M. Drusch, "Improved retrieval of sea ice thickness from SMOS and CryoSat-2," in *2015 IEEE International Geoscience and Remote Sensing Symposium, IGARSS 2015, Milan, Italy, July 26-31, 2015*, pp. 5232–5235, 2015.

[61] K. A. Scott, M. Buehner, A. Caya, and T. Carrieres, "Direct assimilation of AMSR-E brightness temperatures for estimating sea ice concentration," *Monthly Weather Review*, vol. 140, no. 3, pp. 997–1013, 2012.

[62] G. Stonebridge, K. A. Scott, M. Buehner, and A. Caya, "Twin sea ice thickness data assimilation experiments with a toy model," in *Sea Ice Modeling Working Group (SIMWG) Meeting 2016*, (Montreal), 2016.

Appendix A

Analysis Error Covariance Spectral Density Plots

Spectral Densities of Analysis Error Variances – Ice Thickness

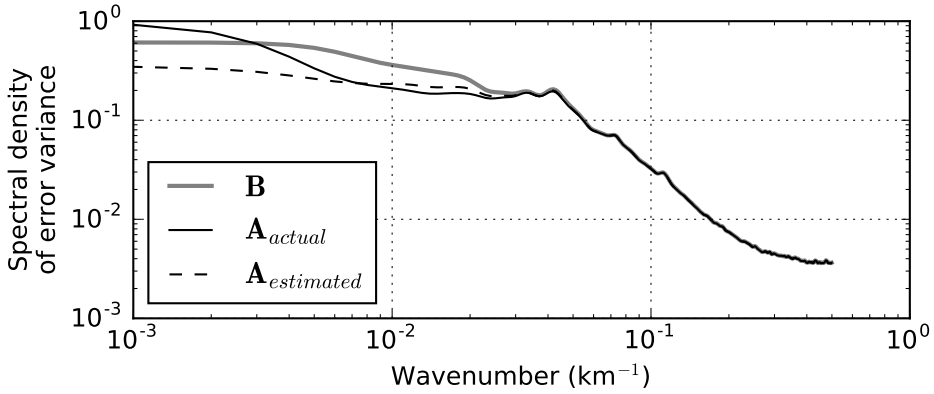

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

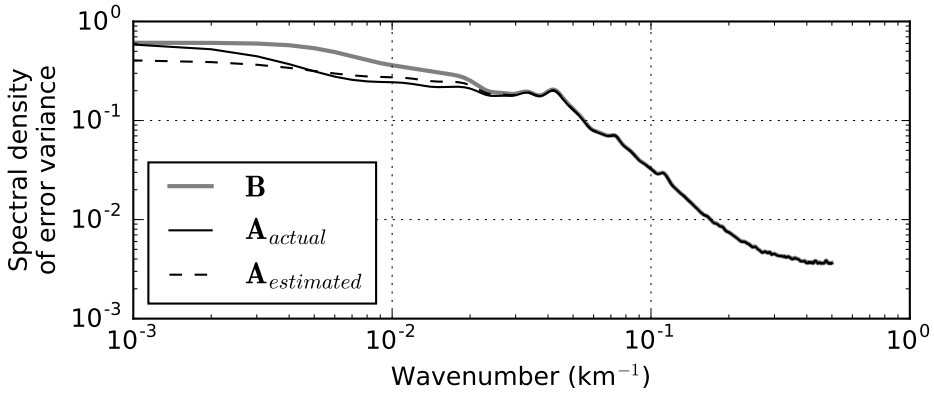

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

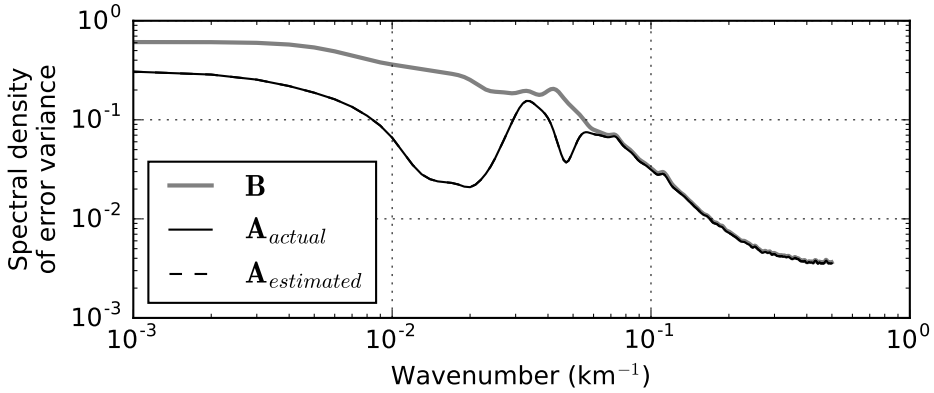

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.1: Spectral densities of thickness analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 150 km

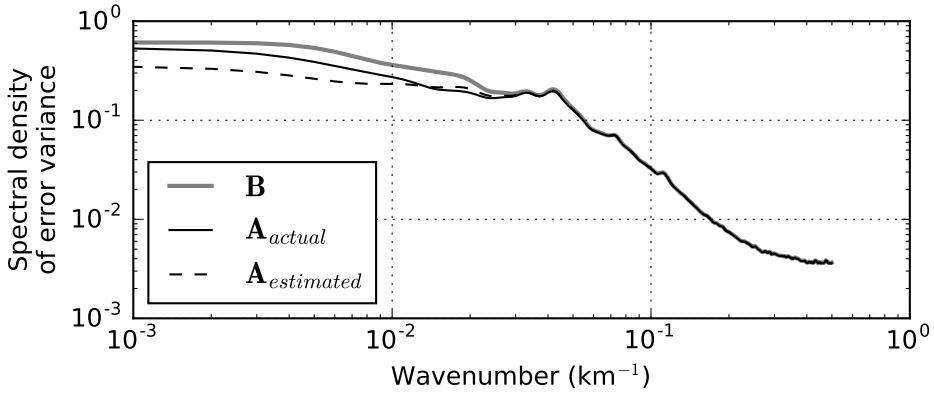

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

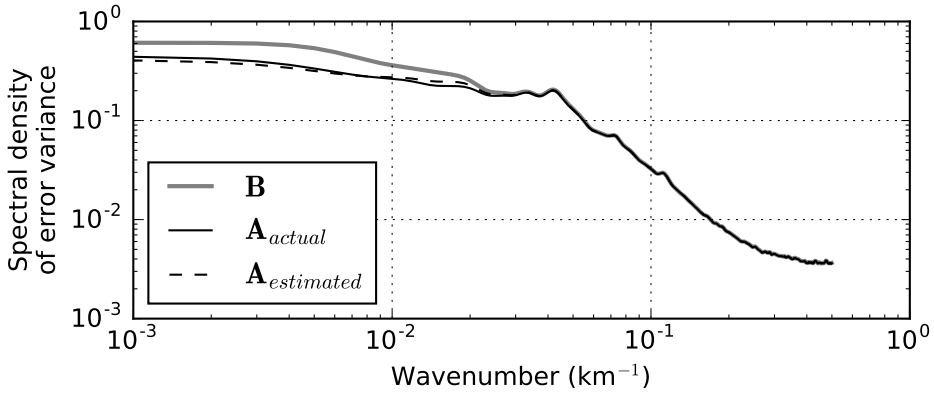

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

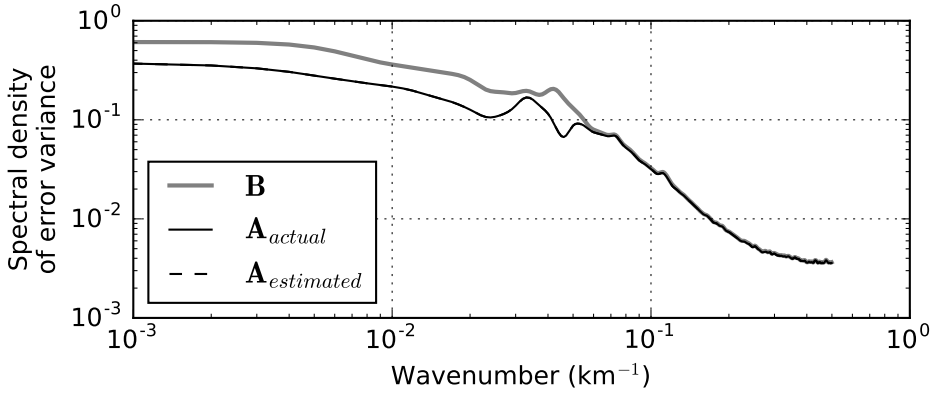

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.2: Spectral densities of thickness analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 50 km

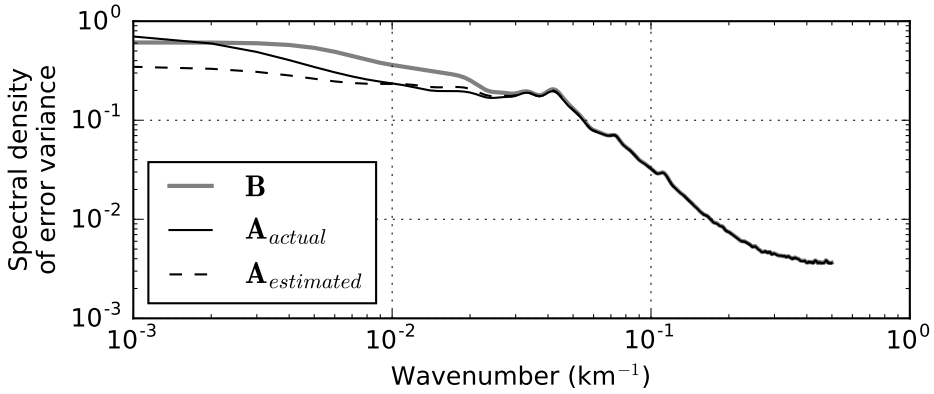

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

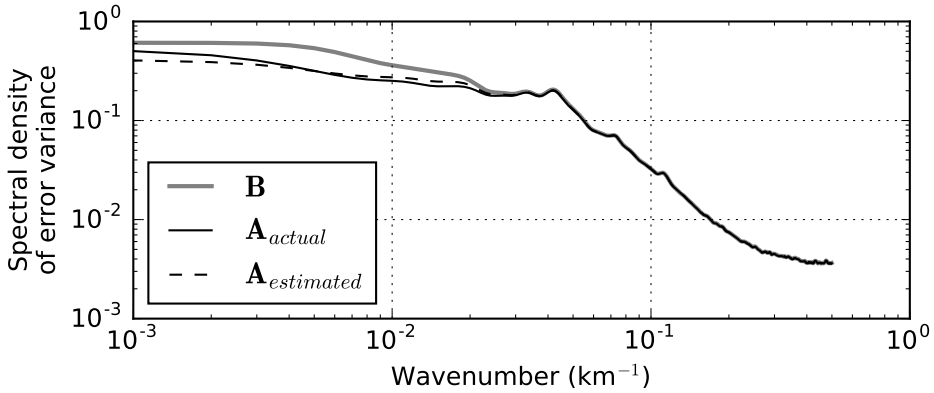

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

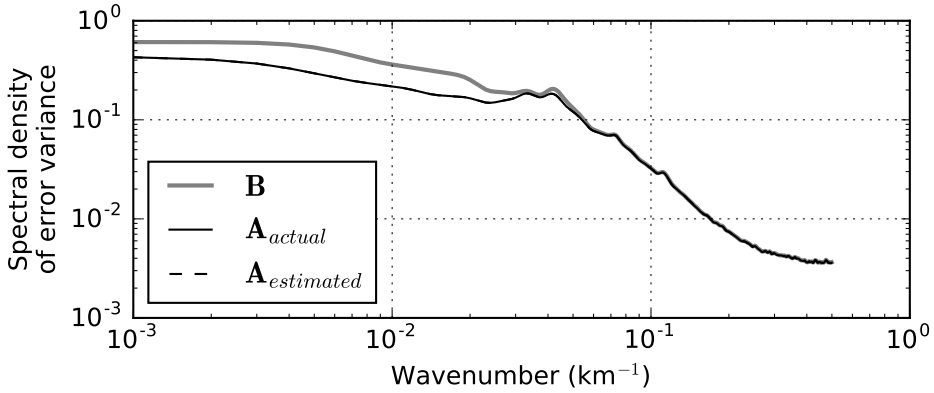

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.3: Spectral densities of thickness analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 150 km

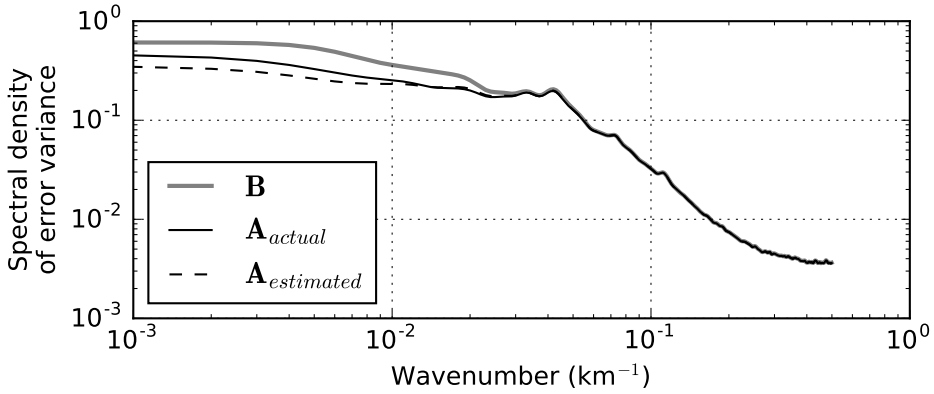

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

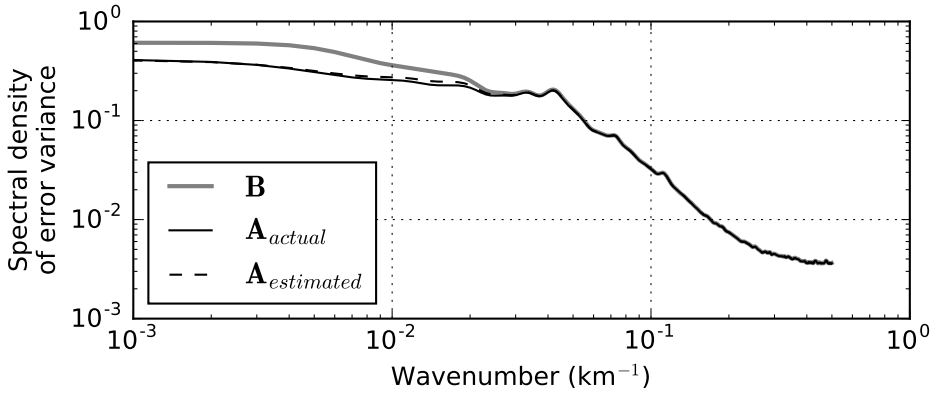

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

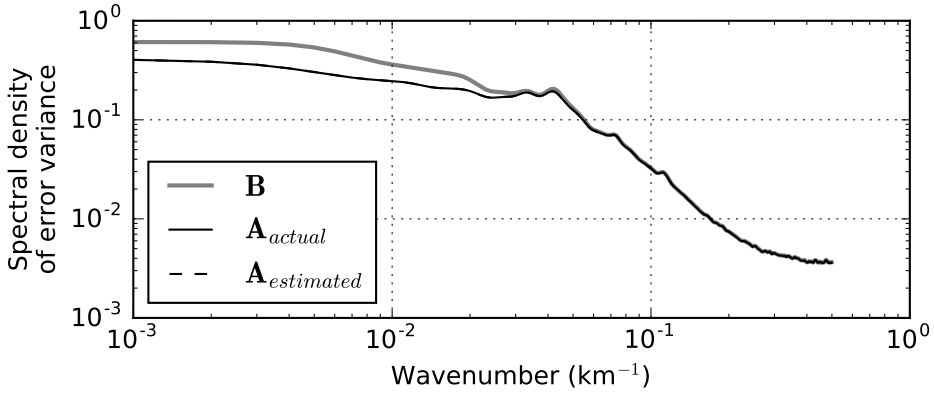

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.4: Spectral densities of thickness analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 50 km

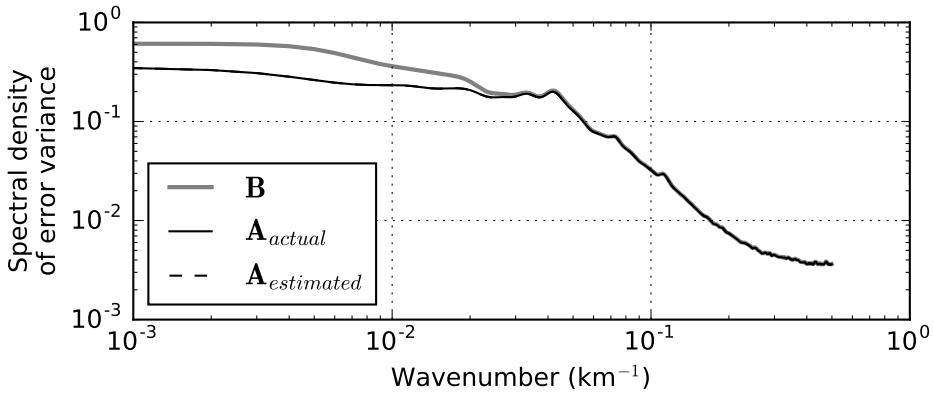

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

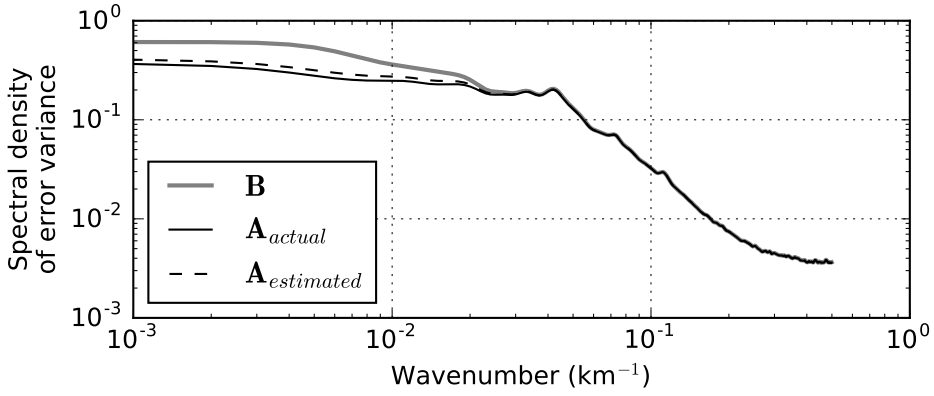

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

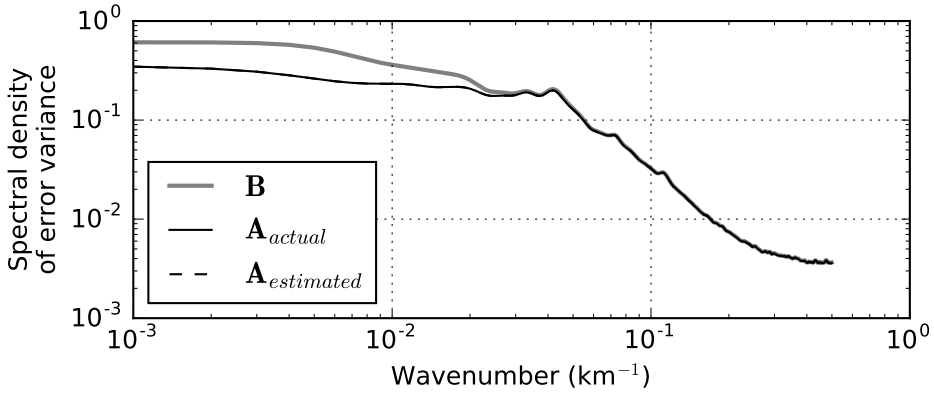

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.5: Spectral densities of thickness analysis error variances for the twin experiment scenario where the true observation error covariance matrix was diagonal

Spectral Densities of Analysis Error Variances – Ice Concentration

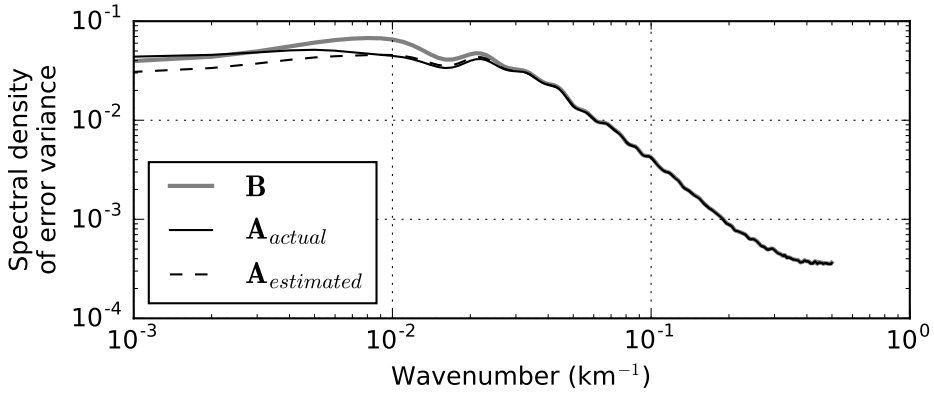

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

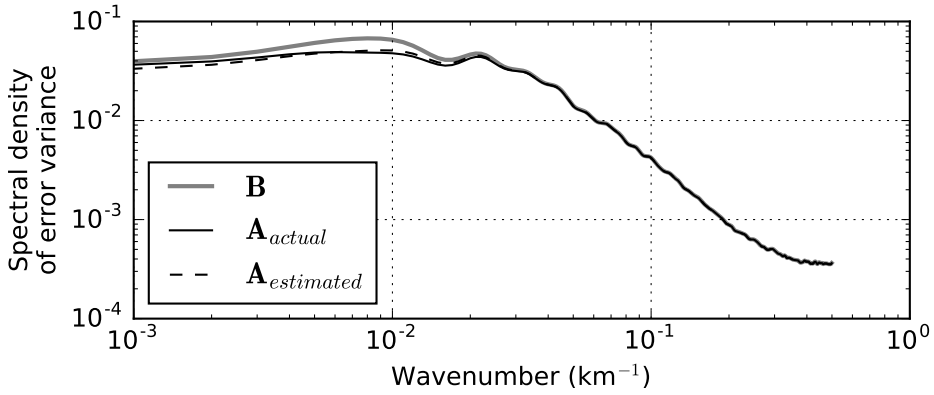

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

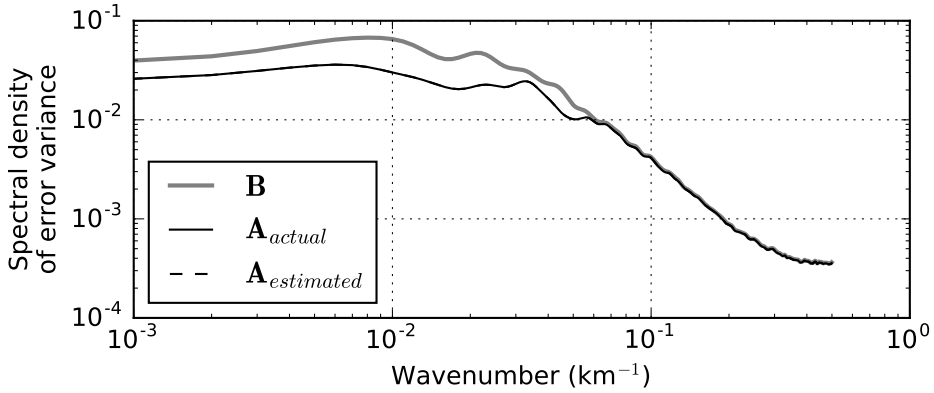

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.6: Spectral densities of concentration analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 150 km

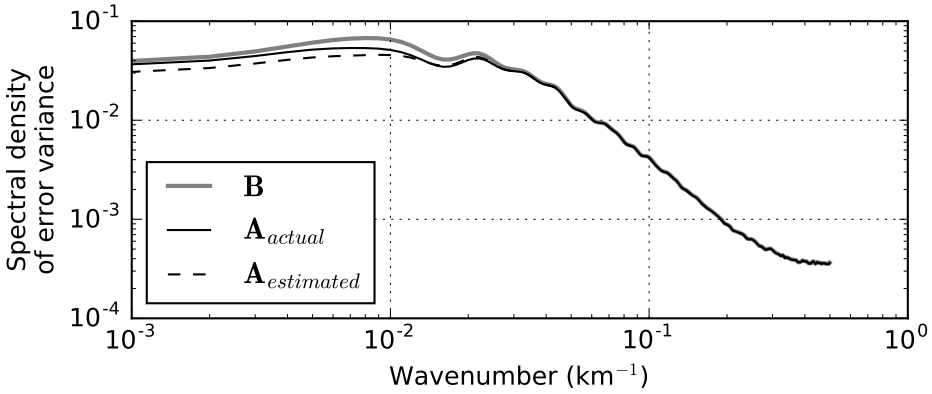

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

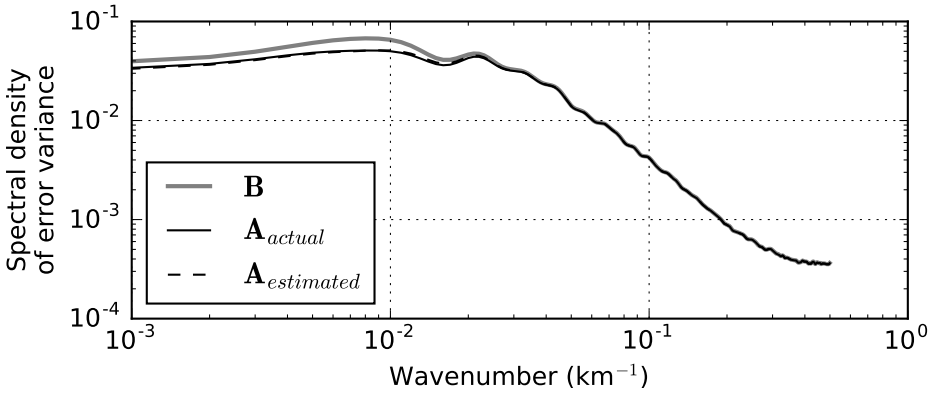

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

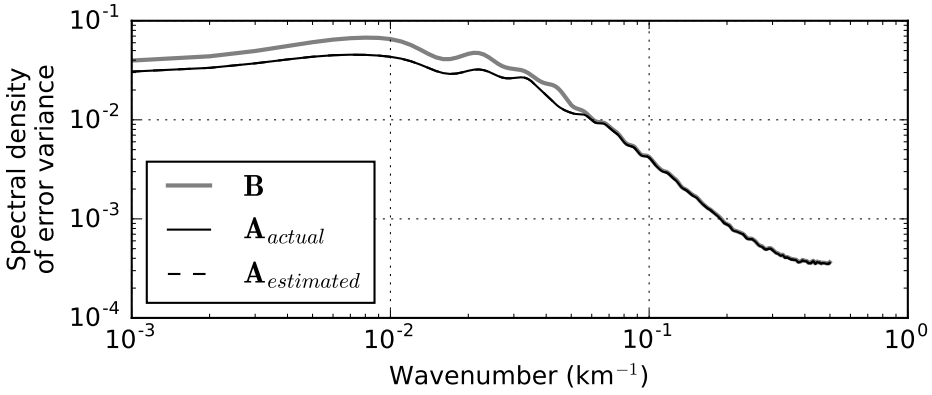

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.7: Spectral densities of concentration analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 50 km

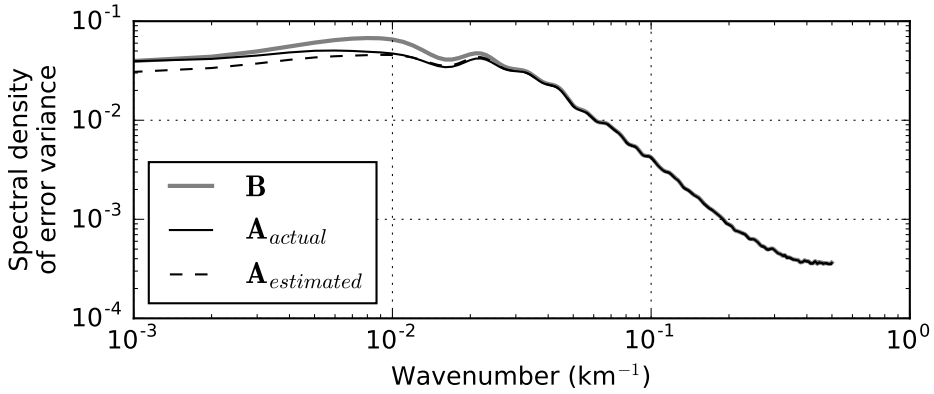

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

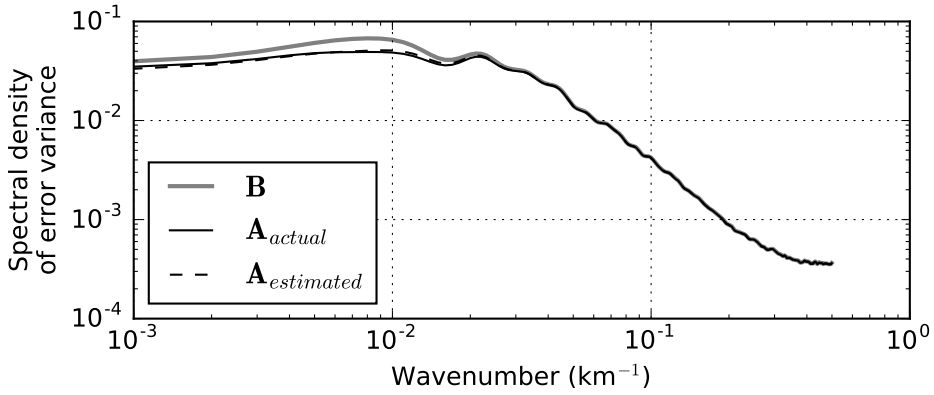

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

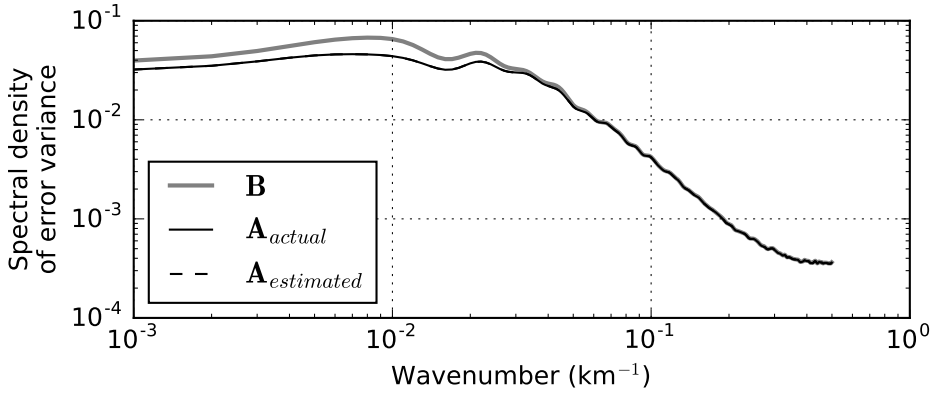

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.8: Spectral densities of concentration analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 150 km

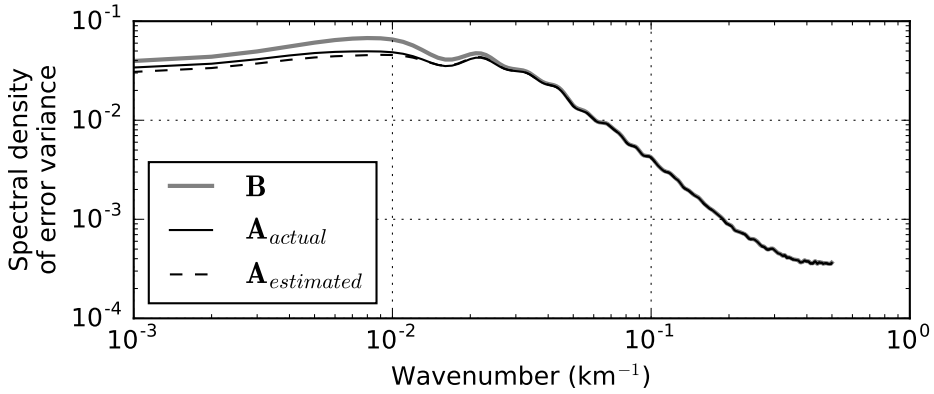

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

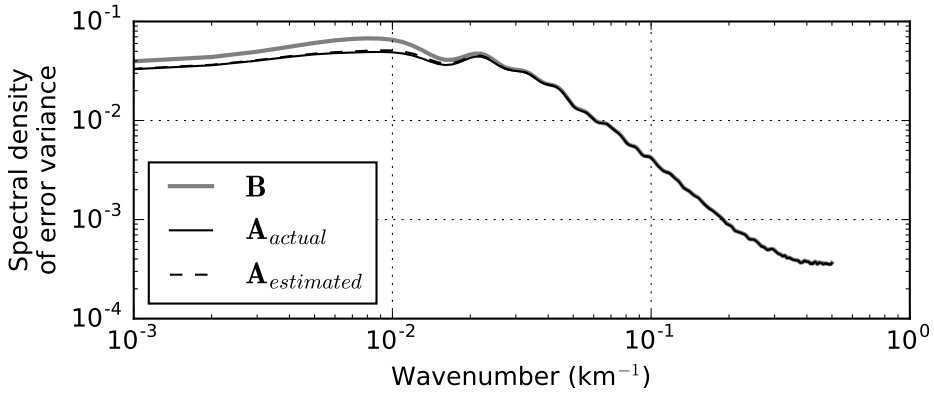

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

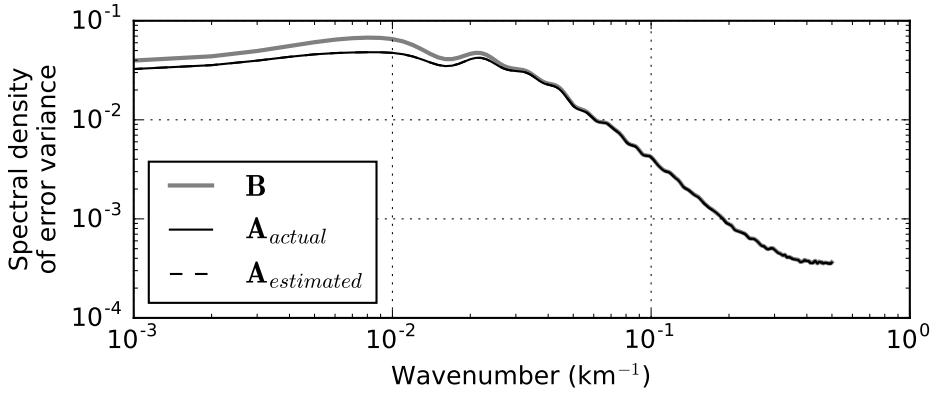

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.9: Spectral densities of concentration analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 50 km

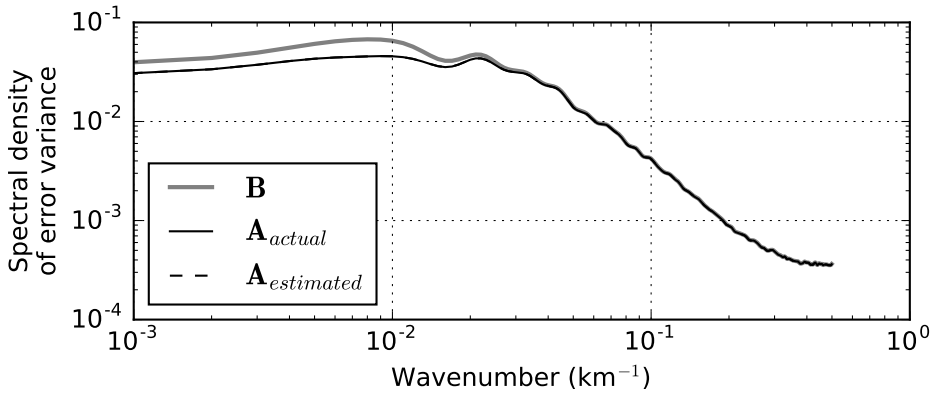

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

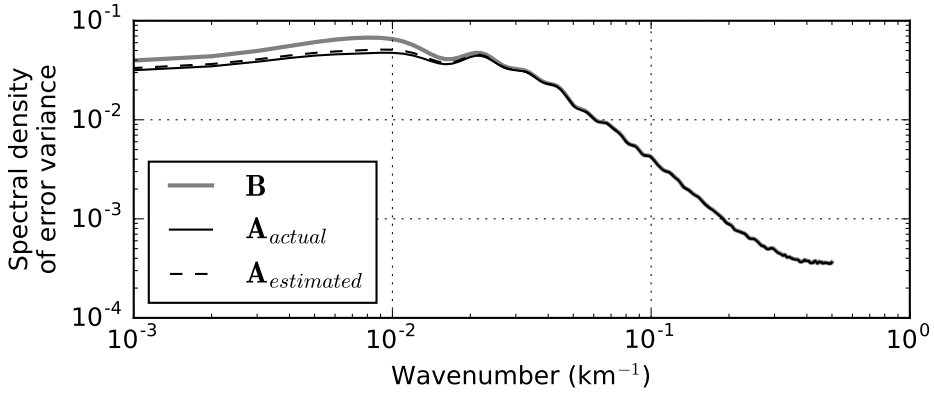

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

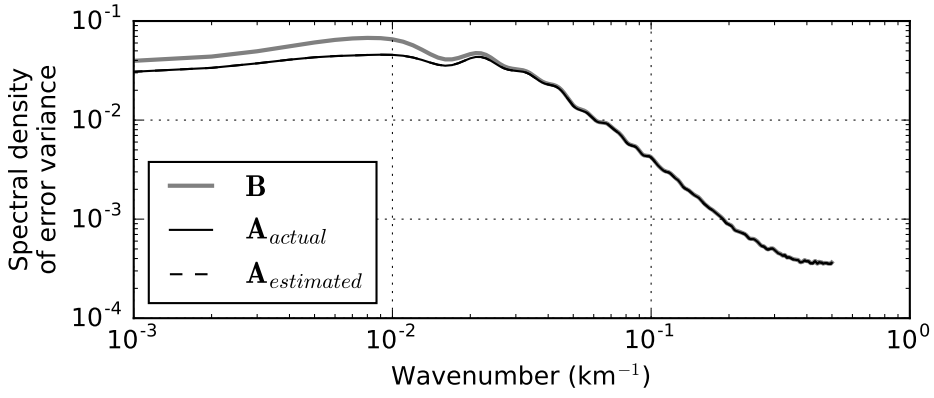

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.10: Spectral densities of concentration analysis error variances for the twin experiment scenario where the true observation error covariance matrix was diagonal

Spectral Densities of Analysis Error Variances – Ice Velocity

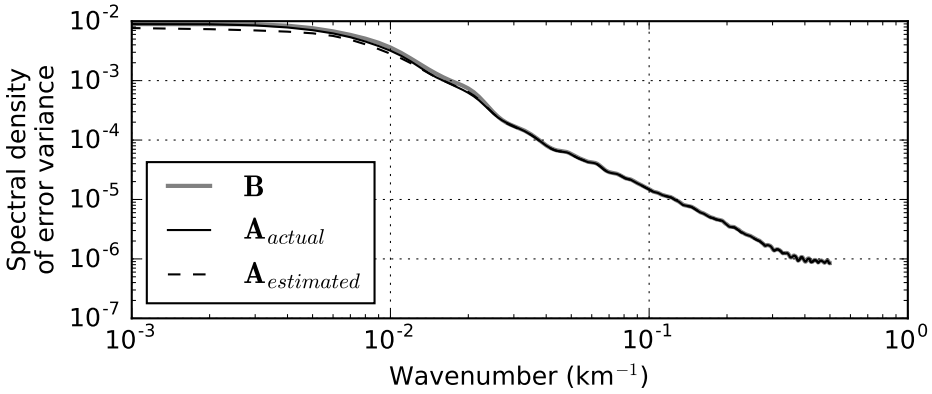

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

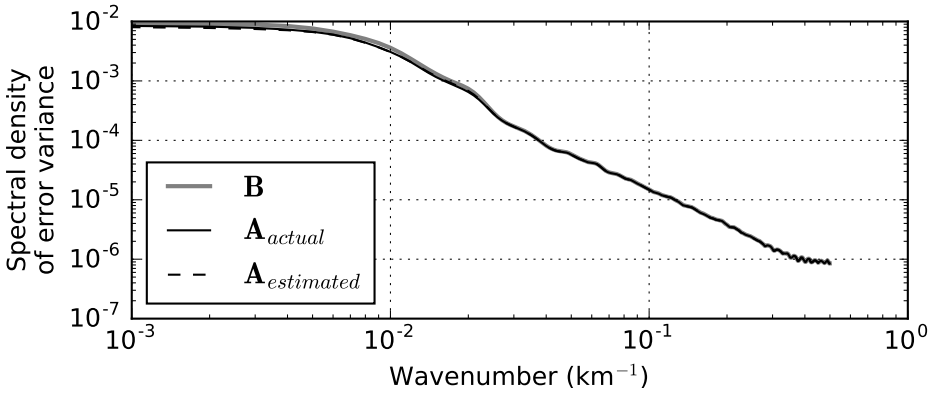

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

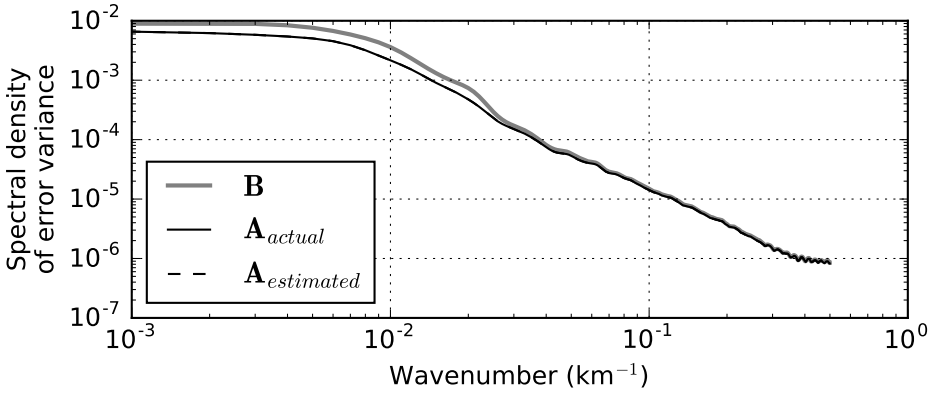

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.11: Spectral densities of velocity analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 150 km

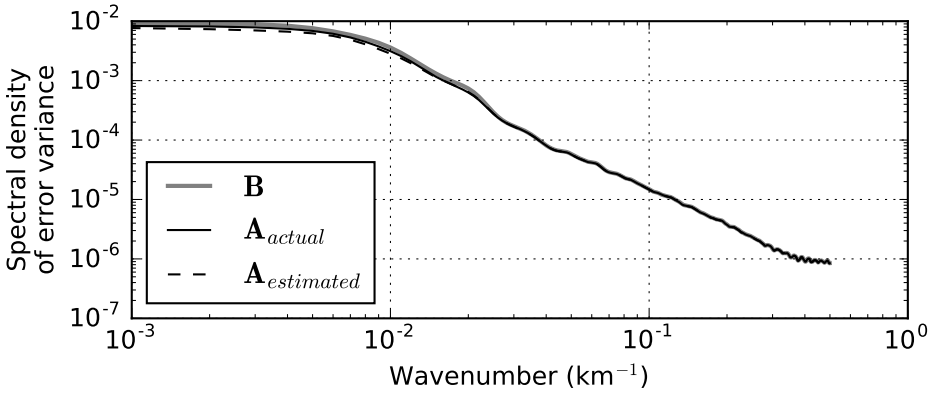

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

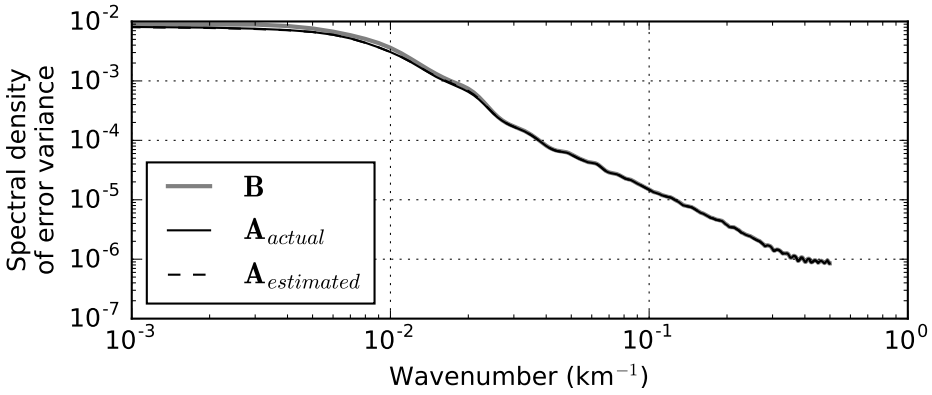

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

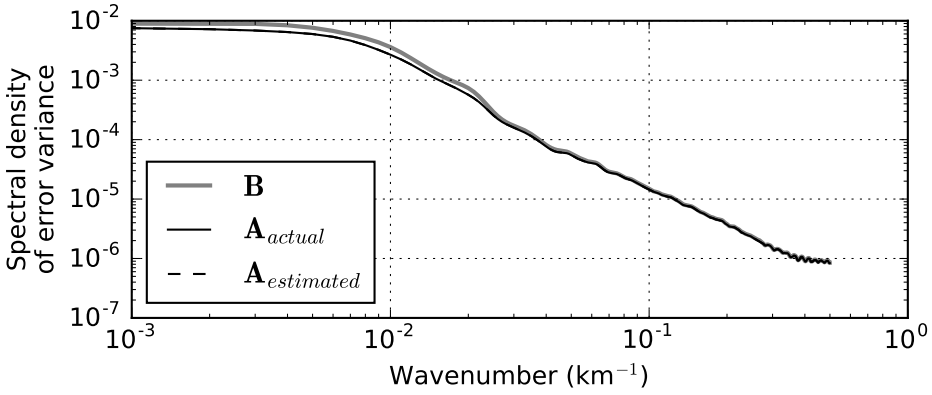

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.12: Spectral densities of velocity analysis error variances for the twin experiment scenario where the true observation error covariance matrix had a Gaussian correlation structure and a decorrelation length of 50 km

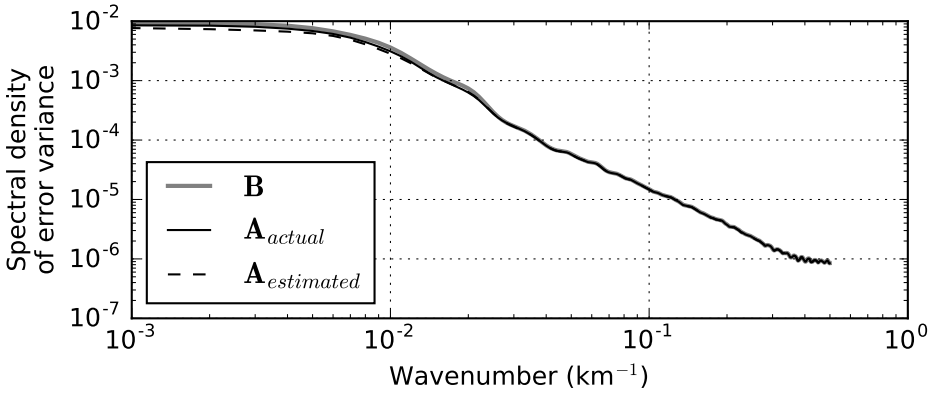

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

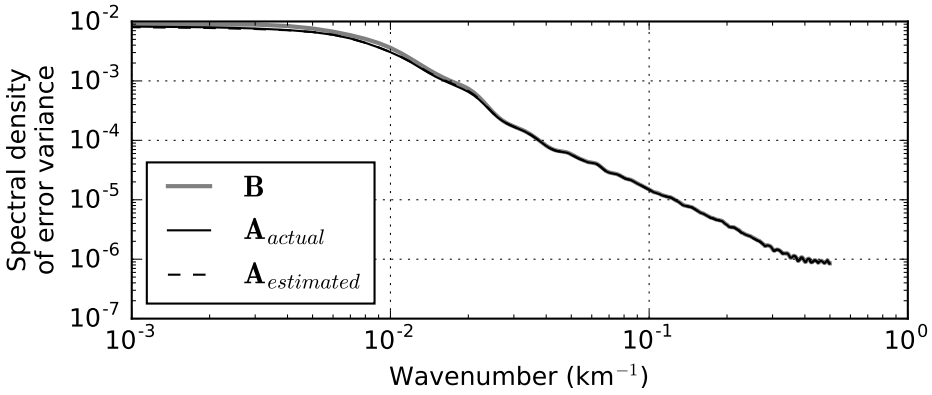

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

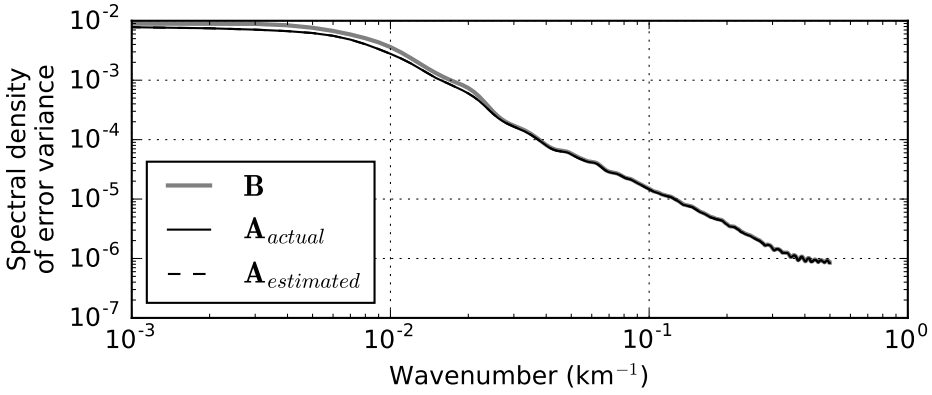

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.13: Spectral densities of velocity analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 150 km

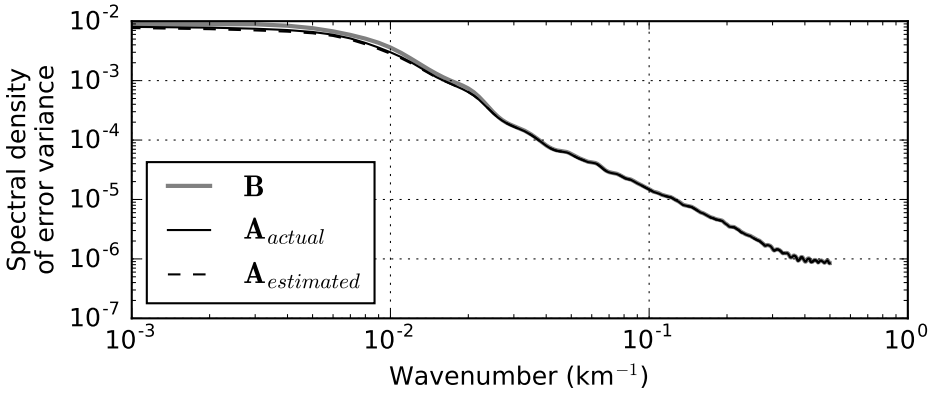

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

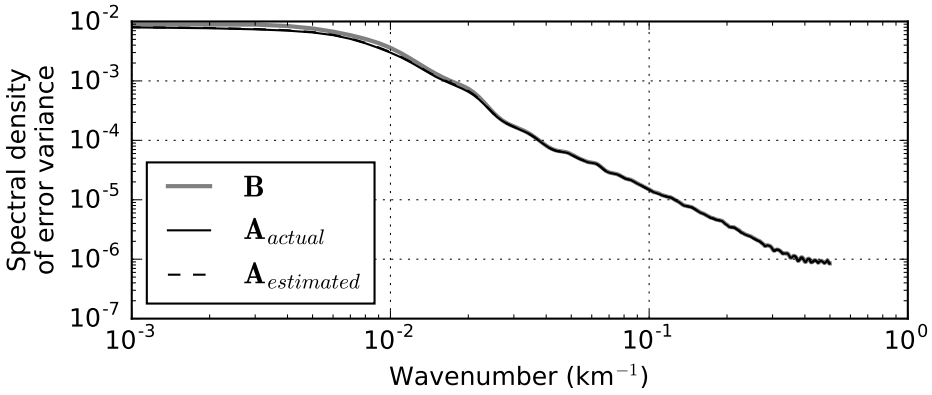

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

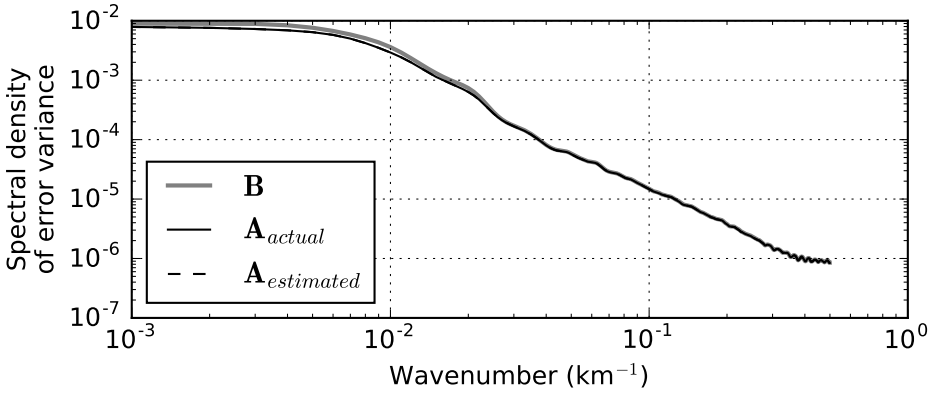

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.14: Spectral densities of velocity analysis error variances for the twin experiment scenario where the true observation error covariance matrix had an exponential correlation structure and a decorrelation length of 50 km

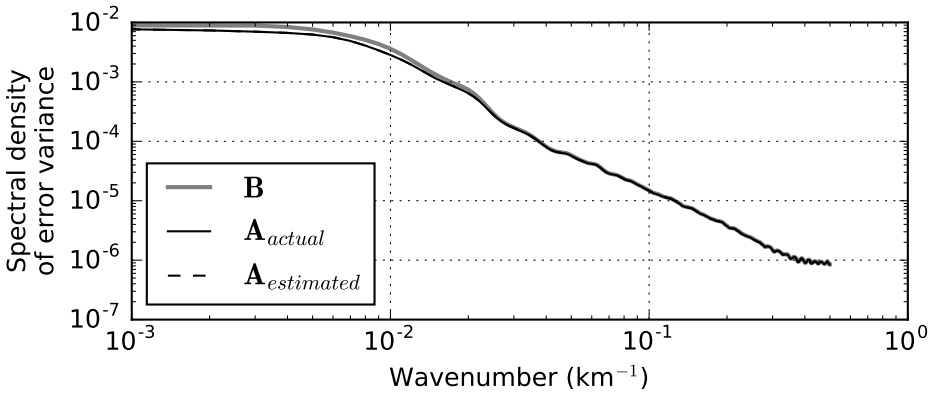

(a) $\mathbf{R}_{est} = \mathbf{R}_{diag}$

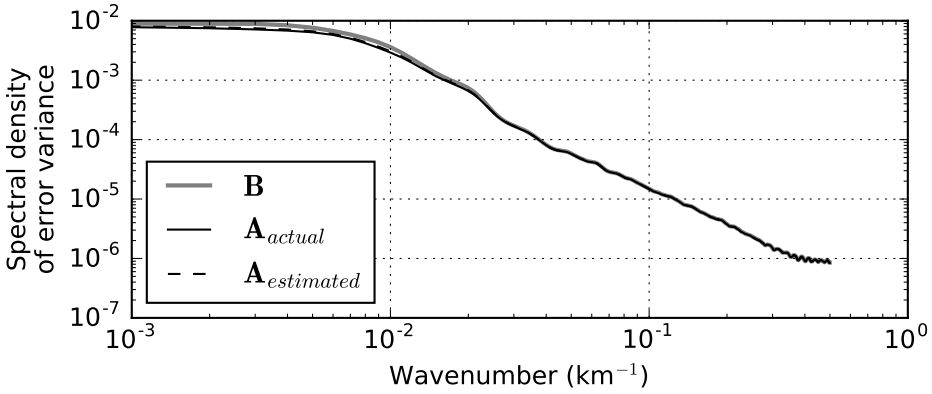

(b) $\mathbf{R}_{est} = \mathbf{R}_{diag} * 2.4$

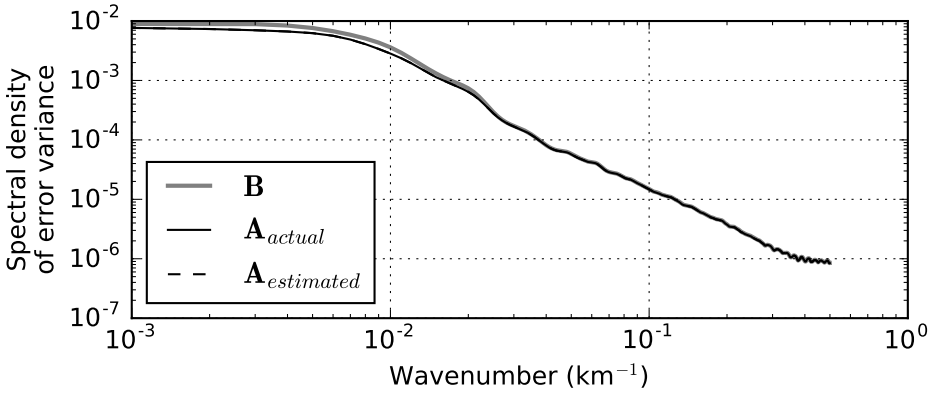

(c) $\mathbf{R}_{est} = \mathbf{R}_{true}$

Figure A.15: Spectral densities of velocity analysis error variances for the twin experiment scenario where the true observation error covariance matrix was diagonal