Hyperspectral Microwave Sensors for Atmospheric Science

Master's Thesis

Lukas Kluft

Universität Hamburg
MIN-Fakultät

Supervisors:
Prof. Dr. Stefan Bühler
Dr. Manfred Brath

Hamburg, February 2, 2017

Thema der Arbeit: Hyperspectral Microwave Sensors for Atmospheric Science

Abstract

Current microwave sensors are limited to a couple of channels. This restricts the amount of information in the measurement. A new technology called Kinetic Inductance Detectors (KIDs) is emerging. KIDs can be used to build low-noise hyperspectral microwave instruments. These instruments allow measurements with high spectral resolution.
In this thesis a concept for a hyperspectral microwave sensor is derived. A state-of-the-art radiative transfer model is used to simulate measured spectra. These simulations are used to perform humidity and temperature retrievals based on the Optimal Estimation Method (OEM).
The one-step linear retrieval solution is compared to a hyperspectral infrared sensor, similar to the Infrared Atmospheric Sounding Interferometer (IASI). Both instruments show comparable retrieval performance. Temperature profiles can be retrieved with a vertical resolution of roughly 1000 m. The resolution of water vapor soundings is around 1500 m for KID and 2000 m for IASI.
As a final check, an iterative retrieval is performed for the KID instrument. It can be shown that the retrieval provides robust results even for challenging atmospheric situations. The vertical resolution achieved by the iterative retrieval is poorer than the linear theory suggests. Perturbations smaller than 3–5 km are only partly detectable.
In summary, hyperspectral microwave sensors based on the KID technology are able to provide useful information for atmospheric science. In a next step the retrieval of cloud particles has to be evaluated. If the development remains constant an airborne prototype may be operable within the next years.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Remote sensing of atmospheric properties is a key element in numerical weather prediction as well as in climate research. In general, remote sensing can be performed in different spectral regions each having its own characteristics. Two important ones are the millimetre/submillimetre (microwave) and infrared ranges.

Infrared measurements offer a large amount of information as a lot of spectral lines can be observed. Instruments observing at a lot of different spectral channels are called hyperspectral instruments. One of them is the Infrared Atmospheric Sounding Interferometer (IASI) which is already in operational use and provides high resolution spectra of a broad part of the infrared spectral region. However infrared measurements are strongly sensitive to clouds (even thin cirrus) and atmospheric soundings can only be retrieved in clear-sky situations. This limitation constrains their use for continuous observations which are needed for numerical weather prediction. Also, a dry-bias is introduced to humidity measurements as clear-sky scenes are overrepresented (John et al., 2011).

Microwave instruments are less influenced by clouds and can observe water vapor and temperature even in moderately cloudy conditions. This advantage makes them particularly valuable for numerical weather prediction as cloudy regions are especially important for the future development of weather situations. Unfortunately, current microwave sensors lack the spectral resolution to compete with infrared sensors. AMSU-A, for example, offers only 15 discrete frequency channels (Rosenkranz, 2001). This limits its information content.

Kinetic Inductance Detectors (KIDs) combine the benefits of both spectral regimes. They allow hyperspectral microwave observations. The synergy of having a lot of channels that are only weakly affected by clouds offers the possibility of high resolution measurements even in cloudy conditions. This is a significant advantage compared to current technologies as it drastically increases the operating time in real measurement situations.

In this thesis the potential performance of a KID-based hyperspectral microwave sensor is investigated. For this purpose an instrument concept is developed, considering the latest technical benchmarks of the KID technology. A state-of-the-art radiative transfer model is used to simulate measurements for this sensor concept based on an atmospheric

data base. The results are compared with simulations for a hyperspectral infrared instrument, simulating IASI. An Optimal Estimation Method retrieval is used to yield retrieval quantities to benchmark the performance of both instruments.
In a final step the KID concept is used to simulate a real atmospheric retrieval. The setup is applied to perform an iterative retrieval using the Levenberg-Marquardt algorithm. The robustness of the simultaneous retrieval of the atmospheric quantities temperature, water vapor and ozone is examined. This is done using synthetic profiles as well as atmospheric soundings from weather model output. Finally, the results of those researches are used to decide, if a hyperspectral microwave sensor is able to provide comparable retrieval performance to current infrared sensors.
Chapter 2 introduces the field of satellite measurements and briefly describes the two sensor concepts, IASI and KID. Technical characteristics for both instruments are given and the measurement principle is explained. Chapter 3 provides an insight into the radiative transfer simulations that were done during this thesis. It is explained what knowledge and which assumptions are needed to properly simulate the emission and propagation of radiation in the atmosphere and how to simulate the behavior of an instrument. Also the atmospheric soundings used for the simulations are introduced. Chapter 4 describes the retrieval methods used and defines some fundamental retrieval properties. These are the fundamentals for the retrievals performed later, both, to compare the instrument concepts, and to retrieve atmospheric profiles. Chapter 5 compares the performances of the infrared and microwave sensor concepts. Chapter 6 describes the retrieval of actual atmospheric profiles for the microwave instrument. Chapter 7 gives an outlook on future developments of KID sensors and their applications for atmospheric sciences.

Chapter 2

Instruments

Atmospheric observations can be done using a variety of different measurement techniques. In remote sensing there are ground-based, airborne and satellite measurements, each of which offer individual benefits and downsides. The focus of this thesis is on satellite measurements, which provide a global coverage of atmospheric measurements.

Satellites can travel through space in different orbits. Geostationary satellites are located 35.786 km above the Earth's surface. At that distance the satellite needs exactly one day to fully orbit the earth. As the earth itself is rotating, the observation point on the surface is not moving. Geostationary orbits provide information of the same scene every time a measurement is taken. This way it is possible to observe temporal developments in the atmosphere. However geostationary satellites lack spatial resolution. In theory their spatial coverage is limited to the hemisphere that is pointing to them. This whole hemisphere is seen as a disk. In practice the regions closer to the rim, like the polar regions, are observed with a poor accuracy, if at all. To increase the coverage and spatial resolution satellites can be flown at a lower orbit, e.g. a polar-orbit. These are orbits a lot closer to the Earth (400–800 km). The shorter distance to the Earth directly increases the spatial resolution. The footprint is decreased and hence spatial coverage is reduced. But polar-orbiting satellites change their observation point quickly. So with increasing observation time (usually half a day) global coverage can be achieved.

Instruments have to fulfill strict requirements in order to be installed on a satellite. They have to work in the rough conditions of the space environment: Vacuum and sometimes large temperature changes, depending on insolation. Also repairs in space are not practical which requires the instruments to be reliable through out the whole mission (usually around 7 years).

Lastly all these requirements have to be compatible. The orbit, the measurement technique, and the quantities to be retrieved are highly interdependent.

This chapter briefly summarizes the concepts and technical characteristics of two measurement systems: Firstly, the infrared sensor IASI, which is already providing atmospheric temperature and humidity profiles, and second, a concept for a microwave instrument using a newly emerged technology, KID sensors.

2.1 Instrument noise

The main goal of this thesis is to intercompare different measurement techniques. To draw a proper comparison one needs well defined metrics to assess instrument performances. Measurement noise is one of the most important instrument characteristics. It is defined as the variability of the signal when observing an object at a constant temperature. There are several ways to define the noise of an instrument. In the following the noise-equivalent temperature difference $NE\Delta T$ is used:

$$NE\Delta T = \frac{T_{scene} + T_{inst}}{\sqrt{\Delta\nu\Delta t}} = \frac{T_{sys}}{\sqrt{\Delta\nu\Delta t}}. \tag{2.1}$$

This allows a straight-forward comparison of both instruments as this approach is also used by Clerbaux et al. (2009) to describe IASI. The $NE\Delta T$ is the standard deviation of the measured brightness temperature of a constant target. Brightness temperature is a unit for radiance, and the conventional unit for instruments measuring thermal radiation. T_{scene} is the noise temperature of the observed object (the Earth's atmosphere). T_{inst} is the noise temperature of the instrument. This means that even when observing the atmosphere with a perfect sensor noise would be present. Equation 2.1 shows that the sum of both values is the relevant quantity. Additionally, $\Delta\nu$ is the measured frequency interval (often referred to as the width of the instrument channel), and Δt is the integration time. A smaller frequency interval therefore increases the instrument noise, whereas a longer integration time t reduces the instrument noise.

2.2 Infrared Atmospheric Sounding Interferometer (IASI)

The Infrared Atmospheric Sounding Interferometer (IASI) is an instrument onboard the METOP polar-orbitting satellite. The following technical overview is obtained from Blumstein et al. (2004). IASI performs passive infrared remote sensing using a Fourier Transform Spectrometer. It operates in the 19 THz–83 THz (645 to 2760 cm^{-1}) spectral region. Although it is common to use wavenumbers when referring to the infrared spectral region, frequencies will be used here to simplify the comparison with the microwave. One difficulty of this thesis is to compare a conceptual instrument with an already existing one. To ensure a fair comparison, the infrared sensor setup is kept on a comparable conceptual level. The channels are evenly spread with a width of 7.5 GHz resulting in 8534 equidistant channels. The instrument noise is characterised for two spectral regions taking values between 0.1 and 0.2 K (Clerbaux et al., 2009). The retrieval with real IASI measurements would require a more sophisticated setup, but the chosen specifications are adequate to estimate IASI's performance in comparison to a conceptual microwave instrument.

Figure 2.1 gives an overview of IASI's observed spectra. The shown quantities result from radiative transfer simulations using the Atmospheric Radiative Transfer Simulator, ARTS (Section 3.1). The top panel shows the opacity for each relevant absorber. The

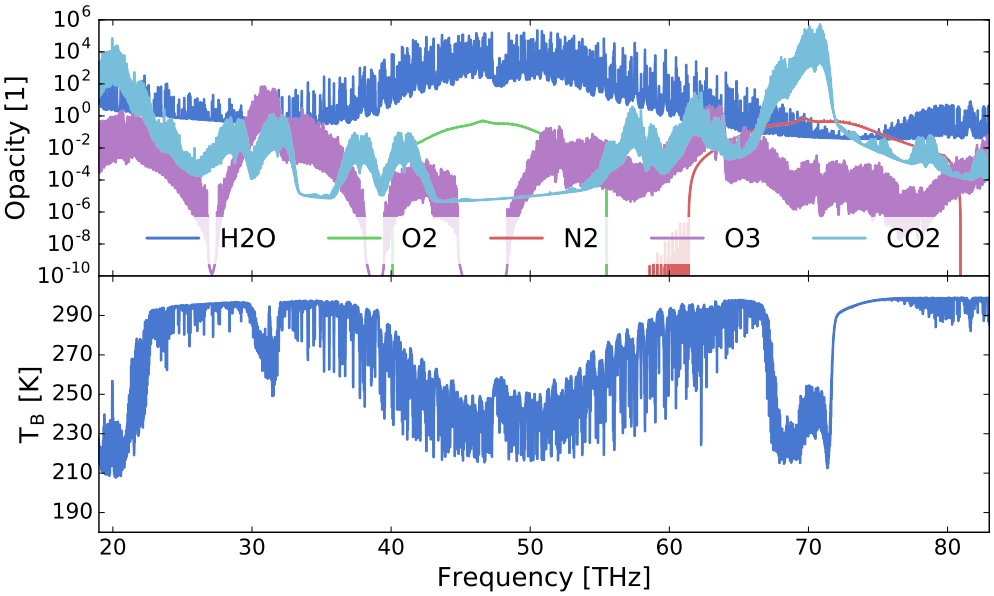

Figure 2.1: Zenith opacity (top) and brightness temperature (bottom) for an IASI-like instrument.

zenith opacity is defined as the path integral over the absorption coefficient α in each height z from the surface to the top of the atmosphere, in zenith direction:

$$\tau = \int_0^\infty \alpha(z)\, dz \tag{2.2}$$

A high opacity is related to a more opaque atmosphere. From the viewpoint of a satellite sensor the line of sight in an opaque atmosphere is shorter. Therefore the received radiation is emitted in high altitudes whereas in a transparent atmosphere the sensor can detect signals even close to the ground. For satellite measurements there is a rule of thumb that an opacity around 1 marks the transition between a transparent and an opaque atmosphere. Channels with opacities way smaller than 1 are called window or imaging channels, as they allow to look through the atmosphere. They provide measurements that can be interpreted as images of the observed quantity at the surface or total column quantities. Opaque channels ($\tau \gg 1$) are so called sounding channels. The radiation in these channels is emitted in different heights of the atmosphere. This offers information about the vertical distribution of the measured quantity.

Nearly the whole spectral region is dominated by water vapor absorption. Only carbon dioxide and ozone exceed water vapor around their absorption bands ($\approx$22 THz and $\approx$72 THz for CO_2 and $\approx$31 THz for O_3). The simulated brightness temperature spectrum (Figure 2.1 bottom panel) perfectly fits the expectations gained from the opacity. There are some window channels with a opacity far beneath 1 that can see the surface (290 K) and four major sounding bands ($\tau > 1$). These bands are caused by the main absorbers H_2O, CO_2 and O_3 and allow measurements up to the mesosphere ($\approx$230 K).

IASI is based on a Michelson interferometer, a variant of a Fourier Transform Spectrom-

eter. An interferometer splits the detected radiation into two beams and recombines them again. The optical interference between both beams contains useful information about the incoming radiation.
The first beam follows a path of constant length while the second path is adjustable. The difference between both paths is called Optical Path Difference (OPD). The detected energy of the recombined beams varies with the OPD. If the OPD is a multiple of the wavelength, both beams are in phase and the detected energy is at a maximum. If it is an odd multiple of half the wavelength the two beams cancel each other out, the resulting energy is zero.
The measured timeseries of radiation is called an interferogram. It can be shown that it is identical to the Fourier Transform of the incident radiation spectrum. Performing an Inverse Fourier Transform therefore provides the original spectrum. The Inverse Fourier Transform, besides calibration and correction of detection non-linearities, is done during the on-board processing.
IASI has to fulfill spectral requirements directly resulting from the observed gases. The on-board Michelson interferometer has a maximum OPD of ± 2 cm resulting in a spectral sampling of $0.25\,\mathrm{cm}^{-1}$. This is sufficient to resolve the CO_2 absorption lines around the $15\,\mu$m band.

2.3 Kinetic Induction Detector (KID)

In the last years a new detector technology has emerged, Kinetic Inductance Detectors (KID). These cryogenic detectors are very sensitive and therefore used in a broad field of applications. The technology has only been around for a decade but has gone through a rapid development (Day et al., 2003; Doyle et al., 2008; Mauskopf et al., 2014; Baselmans et al., 2008; Bueno et al., 2014). Latest studies have shown that the technology is operable in space missions (Griffin et al., 2015). KIDs offer the technical possibility to build instruments containing large arrays of detectors. This way one can construct hyperspectral microwave sensors for meteorological use.
For this thesis a conceptual KID instrument is designed. Its specifications are chosen to provide high-resolution water vapor and temperature profiles. An overview of the instrument characteristics is given in Table 2.1. The individual spectral channel resolving power is 1000 for all channels. It is defined as the ratio of the center frequency ν and the channel width $\Delta\nu$. The resulting 2303 channels have an increasing width with higher frequencies. The $NeΔT$ was calculated with simulated brightness temperatures for an assumed 410 km orbit.
The values of the $NeΔT$ are between 0.007 and 0.025 K (see Figure 2.2). The relative channel width results in narrow channels for low frequencies. Therefore, channels at lower frequencies below 400 GHz are more noisy. Additionally, measured brightness temperatures are slightly higher in this spectral region. This further increases the $NeΔT$ (see Equation 2.1).
Figure 2.3 shows the results of radiative simulations for the instrument concept. The

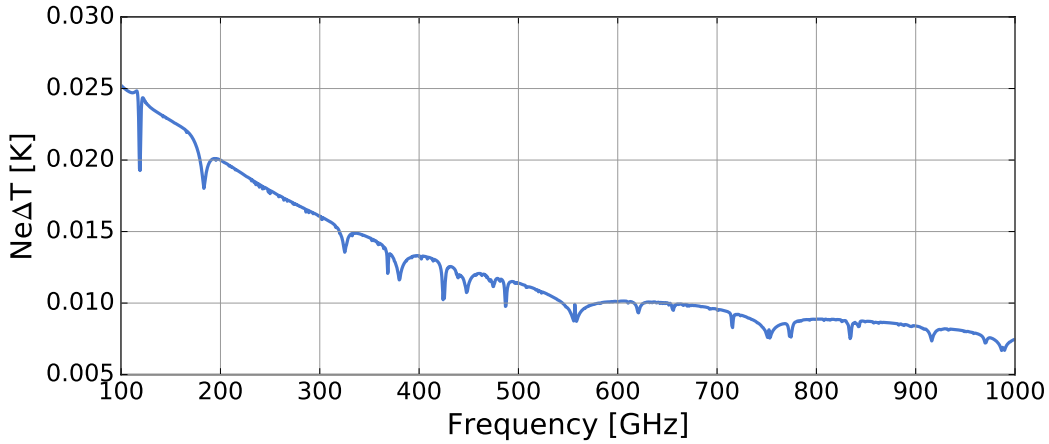

Figure 2.2: KID noise-equivalent temperature as a function of frequency.

calculations are based on a tropical atmosphere climatology (Section 3.2). The top panel shows the opacity for different absorbers. The whole microwave spectral region is dominated by the absorption due to water vapor. There is only one peak at 112 GHz driven by oxygen. Measurements in the microwave spectral region can be used for atmospheric sounding. The observed brightness temperatures are comparable to the infrared (290–230 K). This indicates, that both instruments are measuring signals from roughly the same heights.

Table 2.1: Characteristics of the KID instrument.

Orbit height	410 km
Instrument aperture	0.8 m
Channels	2303
Spectral range	100 GHz - 1 THz

2.3.1 Kinetic inductance

Some materials lose their electrical resistance when cooled beneath a critical temperature. This behavior is called superconductivity. It is caused by electrons forming so called Cooper Pairs, which consist of two electrons attracted to each other. This particle constellation is able to move through the material without resistance.

A KID is a superconducting resonator cooled to a sub-Kelvin temperature. These resonators are used as circuit elements which present an inductive load to a microwave stripline. When incident photons hit the detector they break the superconducting Cooper Pairs. This decreases the superconductivity which alters the resonant frequency of the detector. Changes in resonant frequency shift the detector frequency, amplitude and phase. This shift is proportional to the incoming photon flux and is therefore the actual measurement signal.

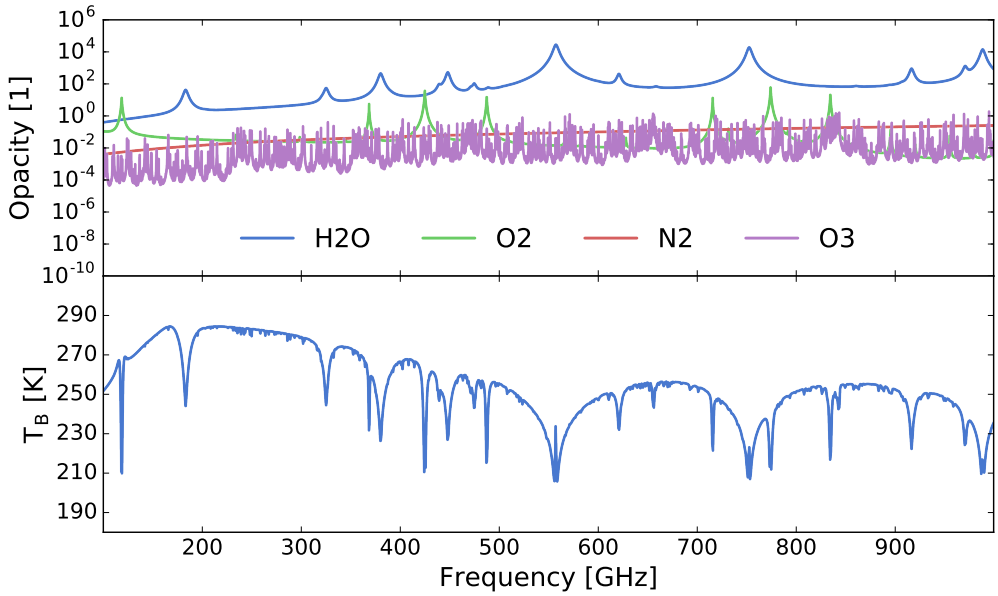

Figure 2.3: Opacity (top) and brightness temperature (bottom) for a KID instrument.

It is possible to tune each detector to a slightly different frequency. This means that the detector is more sensitive to photons with a specific energy. Large arrays of detectors tuned to different frequencies can be used as a spectrometer. This way hyperspectral microwave sensors can be built. These sensors have very little instrument noise and could therefore provide high resolution high accuracy measurements in the microwave region.

2.3.2 Instrument design

There are two different types of satellite instruments: pushbroom sensors and scanning sensors. Each of them offers different benefits and downsides.
A pushbroom consists of a line of pixels which scan the Earth simultaneously. Each of these pixels would be a spectrometer built with KID. The advantage of scanning multiple pixels at the same time is the increased integration time for each individual pixel. An increased integration time directly decreases the measurement noise as shown in Section 2.1. The technical downside of a pushbroom sensor is the limited amount of total detectors in the instrument. When using 10 detectors at a time, the number of frequencies covered by each detector has to be decreased by a factor of 10.
The second design is a scanning sensor where all detectors are used as one spectrometer and the spatial resolution is gained by spinning the whole instrument. This approach is more affected by noise as the integration time for each pixel decreases. It has to be studied whether the increased spectral resolution compensates the increasing measurement noise.

Chapter 3

Radiative transfer simulations

3.1 Atmospheric Radiative Transfer Simulator (ARTS)

The Atmospheric Radiative Transfer Simulator (ARTS) (Buehler et al., 2005; Eriksson et al., 2011) is a radiative transfer model for the millimeter and sub-millimeter spectral range. ARTS is implemented in a modular way which makes it very flexible to use.
ARTS is a line-by-line model. It is capable of performing radiative transfer simulations on arbitrary frequency grids. The absorption lines for the observed spectral range can be read from catalogues in various formats. For the calculations done in this thesis the latest version of the HiTran catalogue (Rothman et al., 2013) was used. The HiTran catalogue provides information to describe the shape of each line and its dependency on the atmospheric conditions. Those parameters make it possible to calculate the absorption coefficients at any given frequency and atmosphere.
The frequency ranges are chosen in different ways for both instruments. For the infrared a linear-spaced frequency grid with a spectral resolution of $0.25\,\mathrm{cm}^{-1}$ is used. This spacing was originally chosen to distinguish the rotational absorption lines of CO_2 (Blumstein et al., 2004). The linear spacing is a consequence of the used measurement technique (Michelson Interferometer, Section 2.2). When using a constant width the channels tend to become extremely narrow for high frequencies. For a spectrometer based on single sensors, this raises the technical issue of building a sensor that can resolve such small bands. Therefore a relative grid spacing is used. It is chosen to make the width of a channel relative to its center frequency constant:

$$\frac{\nu}{\Delta\nu} = 1000. \tag{3.1}$$

3.2 FASCOD atmospheres

Radiative transfer simulations need the current atmospheric state to simulate the propagation of the radiation inside the atmosphere. This atmospheric state consists of information on pressure, temperature and the mixing ratio of all relevant gases. The data has to be given for each level of interest to provide a decent image of the atmosphere. In

general, information on cloud particles (e.g. snow and cloud ice) are needed to properly describe the radiative transfer, especially scattering. But as the microwave sensor is only weakly affected by clouds, they are neglected for this study. All simulations are based on clear-sky atmospheres.
The FASCOD atmospheres are a commonly used set of atmospheric profiles published by Anderson et al. (1986). They represent climatologies at different locations and times e.g. subarctic-winter, midlatitude-summer or tropical. For this study a tropical atmosphere was used to ensure a strong temperature and water vapor signal.
As the used atmosphere has a huge impact on the results further studies have to prove that the found results can be obtained using different atmospheres. Nevertheless the FASCOD atmospheres represent a moderate atmospheric state far from extremes. It is reasonable to assume that the results are representative for arbitrary atmospheric profiles.

3.3 Backend channel response

Simulating instrument measurements with a transfer model introduces another level of complexity. After simulating the radiative transfer through the atmosphere one has to make assumptions on what a sensor would actually measure. This requires a detailed understanding of the detector concept to estimate how the different channels behave when detecting radiation.
In practice, radiances are calculated for a lot of frequencies inside the observed channel. Weights are used to combine the different radiances to estimate the observed measurement signal. Weighting functions can be of different forms and complexity. For this study a gaussian mean centered on the channel frequency was used. The width of a channel is characterised by the full width at half maximum of the gaussian. It is chosen to be relative to the spectral range so that channels at high frequencies are wider (see Figure 3.1).

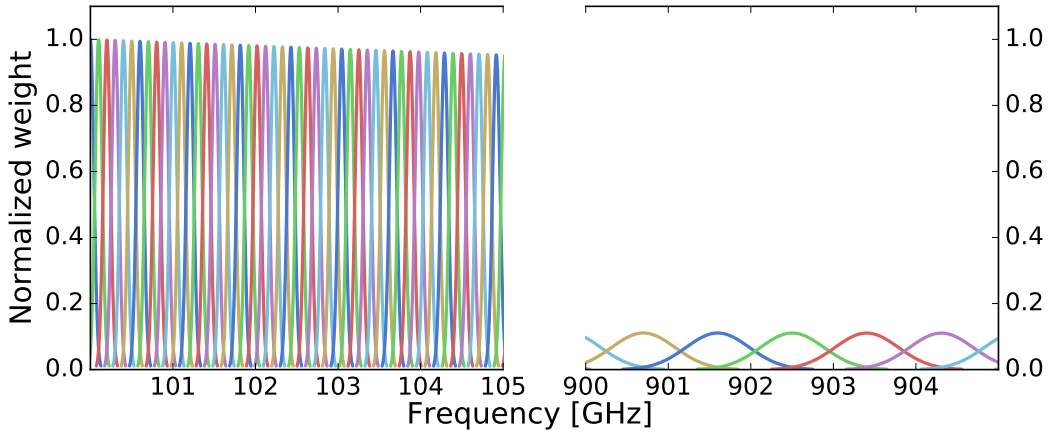

Figure 3.1: KID backend channel response function in different frequency regions.

Chapter 4

Linear retrieval theory

4.1 Optimal Estimation Method (OEM)

Remote sensing provides indirect measurements. Atmospheric properties are not measured directly but via their radiative signal. The measured signal has to be transferred into an actual information about the atmospheric conditions. This task is very challenging, as the problem is mathematically ill-posed. The radiation at the top of the atmosphere is a weighted mean over the radiation emitted in different atmospheric layers. A lot of different atmospheric states can be the solution for the measured spectrum. It is very difficult to find a reliable solution without further assumptions or information.

Inverse theory offers a variety of methods to retrieve atmospheric properties from measurements. For this thesis the Optimal Estimation Method (OEM) is used. The used equations are equivalent to the ones derived by Rodgers et al. (2000). The formal derivation of the methods and equations can be found there. But the basic principles will be described in this chapter.

OEM combines the measured signal and prior information on the atmospheric state to constrain the possible solutions. A Bayesian approach is used to find the optimal combination of both sources of information: For this purpose one needs to estimate how trustworthy this information is. When using a perfect instrument one would only need little prior information to constrain the solution. Without any measurements, a retrieval would rely on the prior information only. Real retrievals are a mixture of both situations depending on e.g. altitude, frequency range and observed quantity.

Inverse theory requires a detailed knowledge of the relation between the atmospheric state x and the measured signal y. In general, this is done using a forward model F that transfers from state to measurement space:

$$y = F(x). \tag{4.1}$$

The information coming from the forward model is a crucial part of the retrieval. It allows to evaluate how well an atmospheric state fits the measured signal.

4.1.1 Forward model

A forward model is used to simulate measured brightness temperatures emitted by a given atmosphere. The Atmospheric Radiative Transfer Simulator (ARTS, Section 3.1) is used as a full forward model. It is capable of solving the radiative transfer inside the atmosphere. Additionally, instrument characteristics are included to simulate the signal measured by an actual sensor. Radiative transfer is highly non-linear. A lot of processes are involved and interconnect, resulting in a complex system that is highly sensitive to changes in the atmospheric state. Therefore, the full forward model cannot be inverted directly. However, it is possible to do a linear approximation around an atmospheric state x_0:

$$y = K(x_0)(x - x_0) + y_0 + \varepsilon \tag{4.2}$$

with the Jacobian K and the measured spectrum for the linearization point y_0. Despite its non-linearity, it is assumed that the radiative transfer can be linearized locally with only a small error ε. This linearized equation can be inverted and allows to calculate an atmospheric state fitting the measurement. However, an accurate estimate of the Jacobian is needed. A more detailed discussion on this is given in the following subsection.

4.1.2 Jacobians

The matrix K, needed to transfer the current state x into the measurement space y, is called the Jacobian. It is the derivative of the measured brightness temperature with respect to changes in a state vector element:

$$K = \frac{\partial F(x)}{\partial x}. \tag{4.3}$$

The Jacobians are derived from radiative transfer simulations and need to be calculated for every quantity of interest. However, they are strongly dependent on the atmospheric state. For that reason it is necessary that the state x_0 does not differ significantly from the true atmospheric state x. Otherwise the linear approximation is no longer applicable. The Jacobians are calculated for different heights. This allows to determine heights that are most relevant for changes in the measured signal. These are the regions in which the retrieval is sensitive. As the Jacobians strongly differ for different frequencies, it is possible to retrieve atmospheric properties in different heights. Common instruments are limited to a couple of channels. These channels are chosen in such a way that a preferably wide range of heights is covered. Hyperspectral sensors cover a large range of frequencies. It is anticipated that the measurements provide height-resolved information on atmospheric conditions.

Figure 4.1 shows humidity and temperature Jacobians for a tropical atmosphere. The temperature Jacobians (Figure 4.1a) are straight-forward to understand. An increasing

temperature leads to a warmer signal in every channel. This behavior is true for all clear-sky cases. Figure 4.1b shows the water vapor Jacobians. The simulation was done assuming an ocean scene, resulting in a cold background for microwave detectors (surface emissivity $\varepsilon = 0.6$). This is the reason for the positive humidity Jacobians at frequencies beneath 150 GHz. Without any water vapor the atmosphere would be transparent and the cold ocean could be seen. Increasing the water vapor content makes the atmosphere more opaque. More radiation is emitted by the atmosphere and the cold background is masked. This leads to a warmer signal and therefore positive Jacobians. For higher frequencies the background is not relevant as the atmosphere is always opaque. The measurement signal in these sounding channels originates from higher altitudes. If the atmosphere gets more humid the layer that can be seen from the satellite rises as the sight distance decreases. The colder temperatures in high altitudes lead to a cooler signal and hence a negative humidity Jacobian.

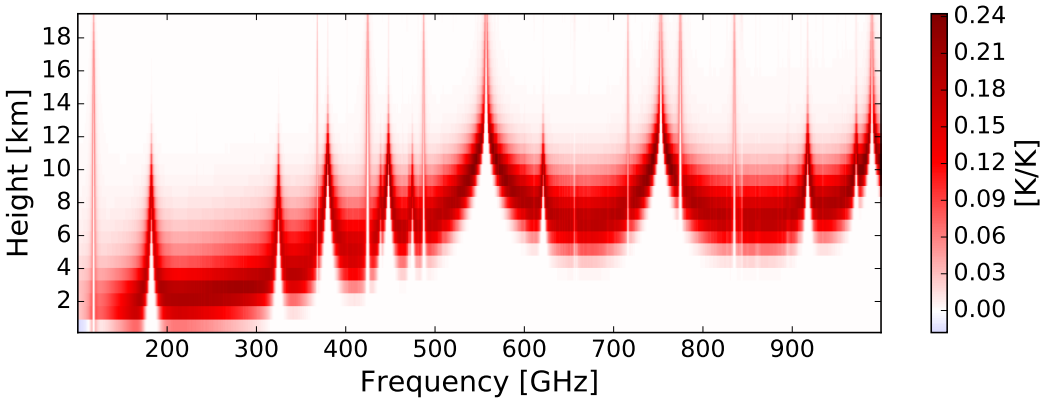

(a) Temperature Jacobians.

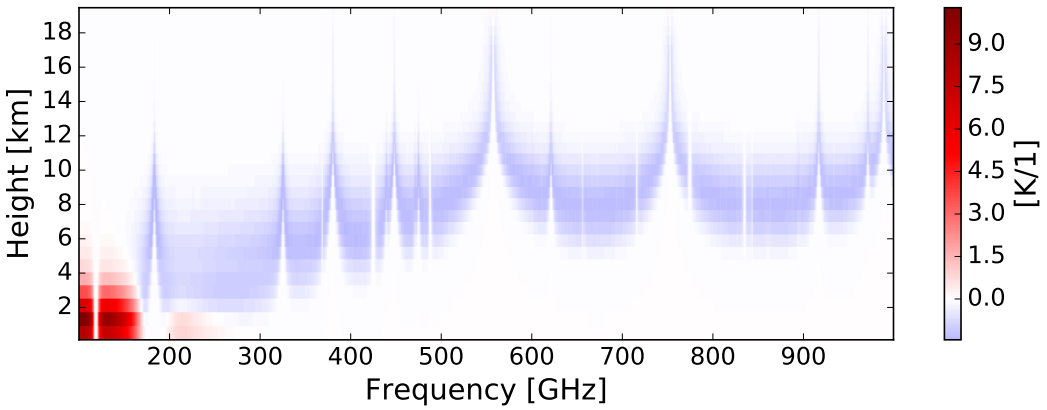

(b) Water vapor Jacobians.

Figure 4.1: Temperature and water vapor Jacobians for a KID sensor.

4.1.3 Linear solution

Using OEM, a solution for the linearized model can be found:

$$\hat{x} = x_a + \left(K^T S_y^{-1} K + S_a^{-1}\right)^{-1} K^T S_y^{-1}(y - K x_a) \tag{4.4}$$

with the retrieved atmospheric state $\hat{x}$, the linearization point x_a (a priori), the Jacobian K, the signal y and the covariance matrices for the measurement and the prior information, S_y and S_a, respectively.

Equation 4.4 consists of two terms. The first term is the state x_a around which the linearization is performed (a priori). The second term can be seen as a correction to adjust the spectra of x_a to be more consistent with the signal y. The spectral difference between the signal and the a priori according to the linearized forward model is given by $y - Kx_a$. It is weighted using the Jacobian K and the inverse measurement error covariance S_y^{-1}. For K close to zero or very high values of S_y, the whole second term becomes neglectable and the retrieval is equal to the a priori. This is the case for very inaccurate measurements or regions where no signal is emitted. The parenthesized expression at the beginning of the second term regulates which source of information to use. If the a priori knowledge is small (high values of S_a), Equation 4.4 simplifies to the linear solution

$$\hat{x} = x_a + K(y - Kx_a).$$

In reality one often finds a mixed case, where both sources of information are used for the retrieval. These are the cases where the measurement system is sensitive to the atmospheric signal and the a priori information is used to constrain the solution.

The linear solution is highly dependent on the linearization point being close to the true profile. As this is not the case for most real retrieval cases, the solutions tend to be very error-prone. There are iterative retrieval algorithms, that continuously update the linearization point. This improves the retrieval results significantly. A more detailed discussion of iterative retrieval algorithms and their results is given in Chapter 6.

Nevertheless, linear solutions are useful to explain the basic principles of OEM. Also they enable the evaluation of the general information content of a retrieval. For that reason a linear retrieval is used for the comparison of the two sensor concepts.

4.1.4 Averaging kernel

In a retrieval a true atmospheric state is indirectly transferred into a retrieved state through the measurement state. But it is possible to formulate a mapping, that directly transfers the true state into the retrieved state. This allows to evaluate which information from the true profile is actually used in the retrieval. The linear solution of the OEM allows to evaluate this relation between the retrieved profile and the actual atmospheric

state. This mapping is referred to as averaging kernel matrix A:

$$A = \left(K^T S_y^{-1} K + S_a^{-1}\right)^{-1} K^T S_y^{-1} K. \tag{4.5}$$

Simply spoken, it describes how the true atmosphere is projected into the retrieved state. Using A, the linear solution (Equation 4.4) can be written as

$$\hat{x} = A \cdot (x - x_a) + x_a. \tag{4.6}$$

For a perfect retrieval the averaging kernel matrix equals the identity matrix. This means, that every element of the true state vector is directly used as a retrieval result. It is possible to asses the performance of a retrieval by looking at the diagonal elements of the averaging kernel matrix. The retrieval works well for regions where the diagonal elements are close to 1.

Figure 4.2 shows the averaging kernel matrices for water vapor and ozone. The water vapor kernels (Figure 4.2a) show nearly perfect results for the troposphere. All values on the diagonal are close to 1 and there is barely any correlation between different altitudes. Above the tropopause the results weaken but the retrieval performance is still acceptable.

Averaging kernels for ozone (Figure 4.2b) are smaller in general. This directly indicates that the signal from water vapor is stronger and provides more information. But the ozone kernels perform well up to much higher altitudes around 50 km. The ozone layer is the main source for the detected ozone signal and therefore signals from high altitudes are received. The correlation between layers is high for all altitudes, so it is reasonable to assume that a retrieval would perform well but lack vertical resolution for these altitudes.

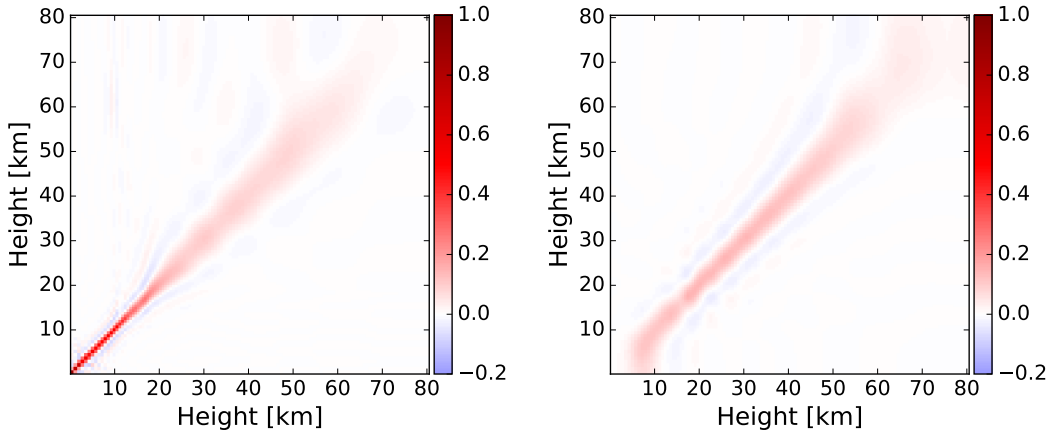

(a) Water vapor. (b) Ozone.

Figure 4.2: Averaging kernel matrices (color).

4.1.5 A priori covariances

The concept of OEM is based on combining a priori knowledge and measurements to get the best possible retrieval result. In order to evaluate which information should be used in a specific case, the quality of the different information needs to be quantified. For the measurement this is achieved by using the instrument error as well as the Jacobians. Together, both quantities give an insight on how much information is gained from the measurement. For the OEM to decide whether to weight the a priori information or the measurement stronger the quality of the a priori has to be specified. This is important and requires careful consideration of the reliability of the assumed climatology. In this thesis an approach by Schneider and Hase (2011) is followed to quantify the water vapor a priori covariances. In summary the method uses fixed covariances for given heights, in a way that the water vapor content in the troposphere is much more variable than in higher altitudes. Additionally a correlation length is assumed to specify the covariance between different height levels. The correlation length decreases with altitude. This behavior describes a high uncertainty within the troposphere. Figure 4.3a shows the resulting covariance matrix. It is smooth at all altitudes which is a desired feature because spurious peaks in the covariance matrix can lead to spurious peaks in the retrieval.

A mixed approach is chosen for the temperature a priori covariance. The covariance coefficients on the diagonal are calculated directly from a set of atmospheric profiles. These profiles are a subset of the dataset generated by Eresmaa et al. (2012). The authors used short-range forecasts produced by the European Centre for Medium-range Weather Forecasts (ECMWF) to create a dataset that describes the atmospheric composition. For this thesis the high resolution version (IFS137) with a refined grid of 137 vertical layers is used. The dataset is subdivided such that only tropical atmospheres are considered. The same exponential decay as for the water vapor is assumed. Figure 4.3b shows the temperature covariance matrix. In contrast to the water vapor the most uncertain region is located around the ozone layer in 30–50 km height. The matrix is not as smooth as for the water vapor but it does not show peaks, either.

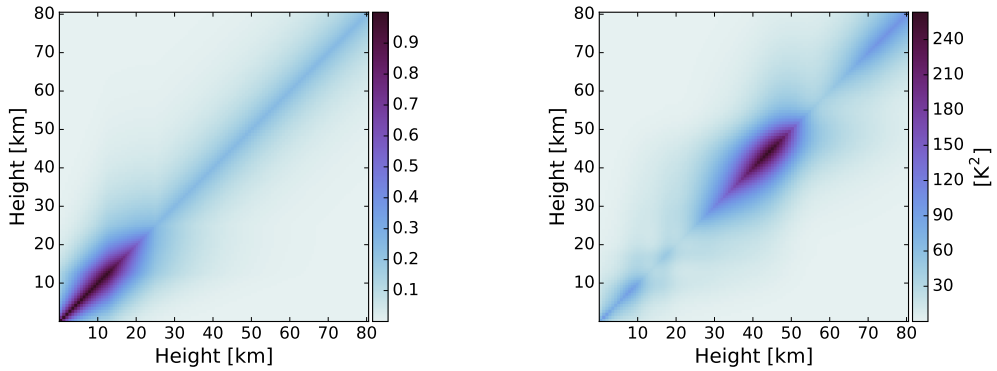

(a) Water vapor. (b) Temperature.

Figure 4.3: A priori covariance matrices.

4.2 Channel selection

The large number of channels offers potential benefits but poses difficulties to the implementation. Even with the newly developed KID technology (Section 2.3) it is most likely not possible to have sensors with thousands of channels. For that reason it is necessary to prove that the retrieval works with a decreased number of channels. This limitation is relevant for both possible instrument designs — a pushbroom or a scanning sensor.

Additionally, the retrieval of actual atmospheric profiles requires several iterations of simulations and post-processing. It is currently not practicable to apply an iterative approach to measurements with such a high spectral resolution. To reduce the computational power needed to perform the retrieval it is also helpful to reduce the number of observed channels.

Fortunately, the information content is not distributed evenly over the whole spectrum. So an intelligent selection of channels can provide a majority of the information. This selection can be done with an iterative approach. In a first step the information content for every single channel is identified and the most informative channel is chosen. Following this, the channel chosen is combined with the remaining channels and the information content for each channel pair is calculated. This way highly correlated channels are rejected because they do not contribute further information. The procedure is repeated to select the 100 most informative channels.

The reduction of degrees of freedom is used as a measure of the information content. It is defined as the trace of the averaging kernel matrix A:

$$\Delta dfs = tr(A). \tag{4.7}$$

Figure 4.4a shows the ranking of the 100 most informative channels. The axis of ordinates is not intuitive and needs some further explanation. The algorithm results in a sorted list of channels, starting with the channel that reduces the degree of freedom of the retrieval the most. All following channels add less information to the retrieval than this channel. Still the total amount of degrees of freedom decreases with each additional channel. The total reduction of the degrees of freedom is plotted when using all channels up to the observed one. The total reduction of degrees of freedom is represented by the space between the bar and the top of the plot. The y-axis is inverted to detect the information content of a single channel. The higher the bar the more information is contained in the measurement. It is obvious that channels located at the sides of dominant absorption lines are preferred (compare Figure 2.3). After all absorption lines are covered, the algorithm tends to accumulate channels around the most distinct absorption lines.

Figure 4.4b shows the reduction of degrees of freedom in dependence of the amount of channels used for the retrieval. The first 30 channels decrease the degrees of freedom linearly, which is the theoretical optimum. After this linear region the performance

reaches saturation very quickly. The 100 most informative channels reduce 40 degrees of freedom. When using all 2300 channels the reduction is 42. So a significant amount of the information is found in the selected channels.

The selection of channels allows more efficient simulations. It is possible to decrease the number of frequencies covered by the radiative transfer simulations by a factor of 20 without losing much information. A decreasing number of frequency channels proportionally reduces the computational power needed for the radiative transfer simulation. The profit during the Optimal Estimation retrieval is even higher as the frequency is used in higher orders.

These findings are highly dependent on the chosen atmosphere. A different atmosphere leads to differing Jacobians, which alters the distribution of information inside the measurement. Tests have shown that the information content decreases significantly when using the selected channels for other atmospheres. These results suggest that the concept is viable. It is possible to reduce the number of channels while keeping the retrieval performable. In order to get a robust channel selection statistics over several atmospheres have to be used.

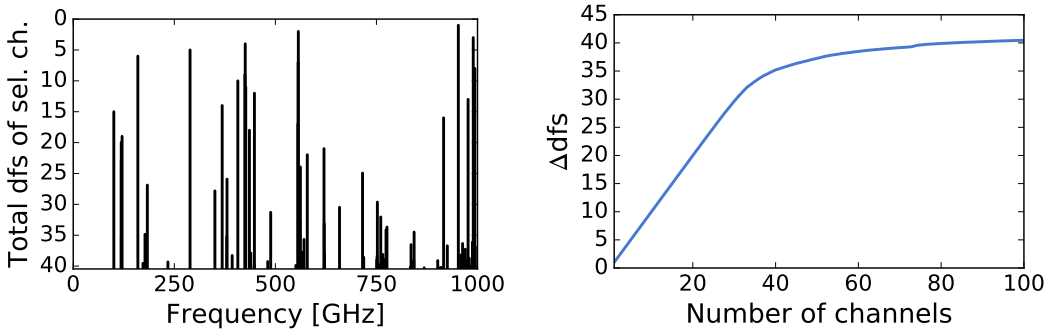

(a) Selected frequency channels and their rank regarding the reduction of degrees of freedom.

(b) Reduction of degrees of freedom in dependence of the number of selected channels.

Figure 4.4: Selection of most informative frequency channels.

4.3 Retrieval height levels

The number of retrieval levels is an important subject to deal with. Like the amount of channels the number of retrieval levels directly affects the computational costs of the simulation. But more levels also offer more information about the atmospheric state. Therefore it is a goal to reduce the number of levels as far as possible while still describing the atmosphere with a sufficient accuracy. Unfortunately the task is not as straight forward as it is for the frequency selection where it is possible to independently drop channels from the simulations. The task requires a reasonable trade-off between accuracy and computational costs.

Fortunately, the linear algebra done during the retrieval is in theory not dependent on the number of levels used. For practical success, the a priori covariances need to be defined

in a way that takes varying numbers of levels into account. This is done following an approach by Schneider and Hase (2011) using well defined height dependencies for covariances and correlation lengths (Section 4.1.5).
Obviously, a coarser grid would waste some of the achievable resolution. A finer grid increases the computational a costs without improving the results. In practice one tries to match the internal resolution of the retrieval.
Figure 4.5 shows the vertical resolution for temperature and water vapor retrieval depending on the number of pressure levels. It can be seen that the adaptation of the a priori covariances works for the considered retrieval quantities. The vertical resolution is constant for 125 or more pressure levels. When choosing a coarser pressure grid, the vertical resolution is determined by the grid spacing. This is especially obvious for the run with 75 levels.
As a reasonable compromise 125 pressure levels are used for the instrument comparison. This makes it possible to determine differences in the resolution while limiting the number of pressure levels as far as possible. The iterative retrieval is run with 75 levels for performance reasons. For an iterative algorithm the run-time is more important than for a single run.

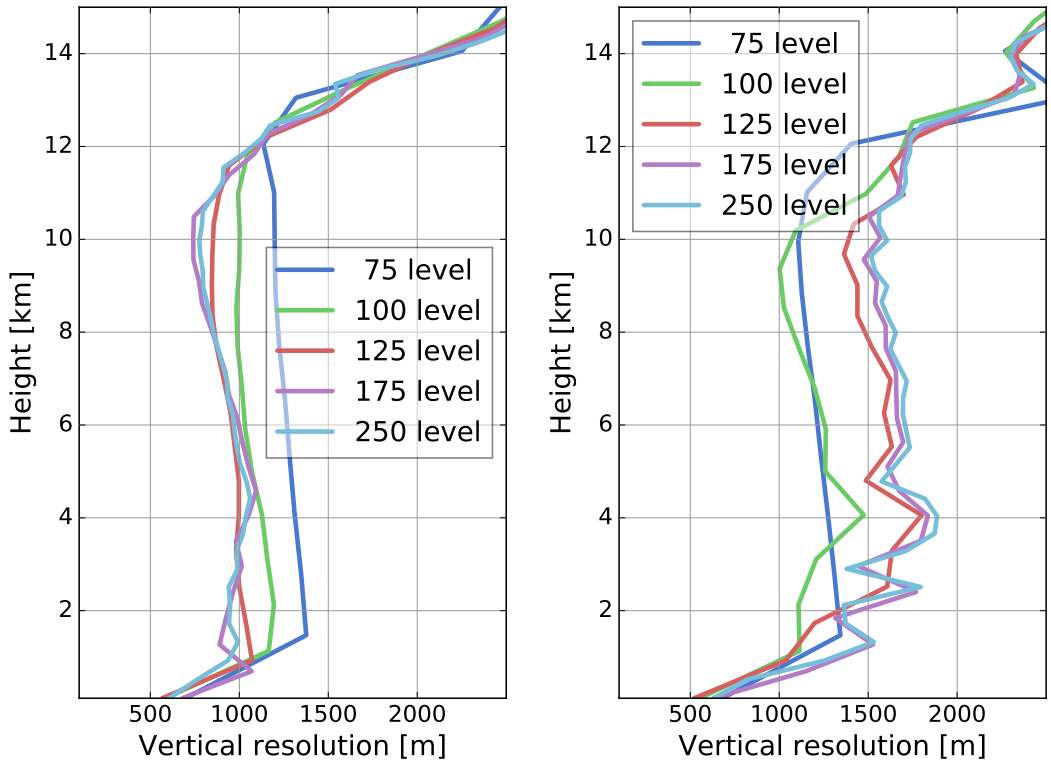

(a) Temperature.

(b) Water vapor.

Figure 4.5: Vertical resolution in dependence of the number of pressure levels.

Chapter 5

Comparing infrared and microwave retrieval

In the previous chapters all relevant concepts and assumptions were introduced. The sensor concepts in infrared and microwave were described (Chapter 2) and their benefits and downsides discussed. It has been shown how sensors can be simulated using a radiative transfer model, including relevant assumptions and settings (Chapter 3). Finally the basic concepts of an OEM retrieval were explained (Chapter 4). This framework provides the possibility to compare the general retrieval performance of both instruments. The quantities considered are the measurement response, the vertical resolution and the retrieval error, which allow a detailed estimate of the instrument sensitivity and the robustness of possible retrieval results.

5.1 Measurement response

The measurement response indicates the sensitivity of a retrieval in different heights. It is defined as the sum of the column kernels (Section 4.1.4) at each height level. Values close to 1 imply that the retrieval is sensitive to the changes in a specific height. If the measurement response is 0, no information is gained from the measured spectrum.
Figure 5.1 shows the averaging kernels and the corresponding measurement response for the KID instrument. The retrieval is sensitive to temperature up to 60 km (Figure 5.1a). A distinct change in the averaging kernels can be seen around the tropopause around 15 km. The averaging kernels in the troposphere are very narrow and peak at values around 0.8. This implies that these atmospheric layers can be retrieved individually with almost no correlation between neighboring layers. For higher layers (20–50 km) the averaging kernels get wider and overlap. Although the sensitivity is still high, an uncorrelated retrieval of values in different heights is no longer possible. The results for the water vapor retrieval show similar behavior (Figure 5.1b). A measurement response around 1 is achieved up to altitudes of approximately 50 km, indicating a high retrieval sensitivity towards water vapor. The averaging kernels are smaller than for temperature and peak around 0.5. Above the tropopause a significant overlap can be seen. The

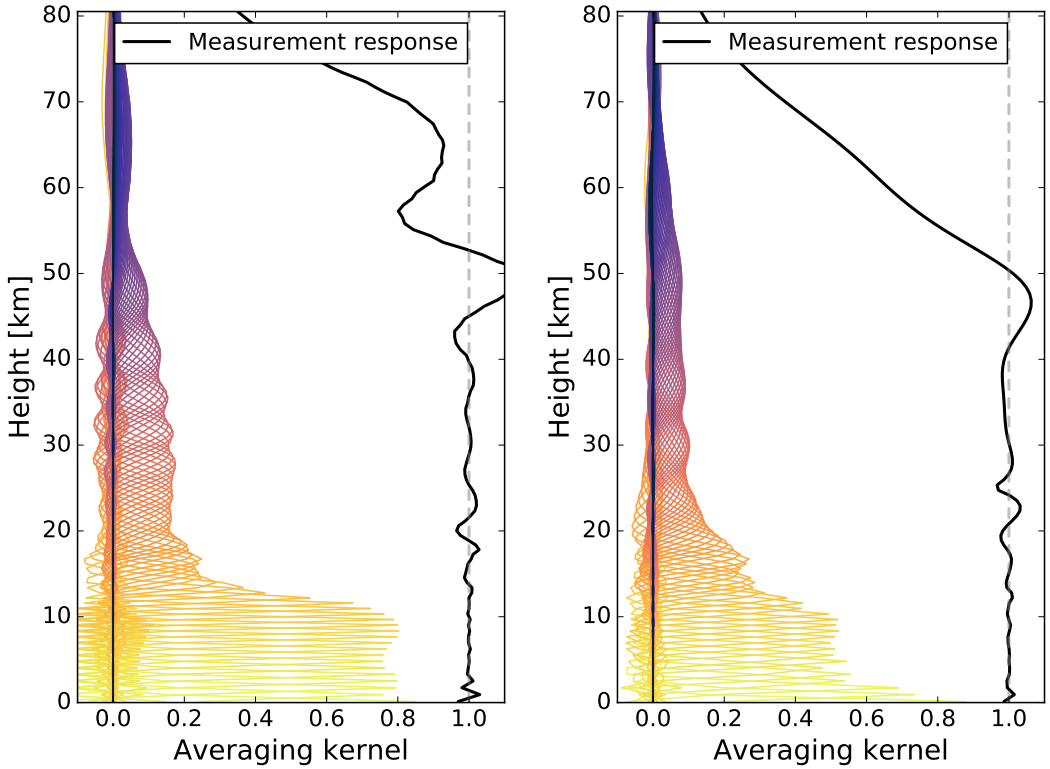

(a) Temperature.

(b) Water vapor.

Figure 5.1: Averaging kernels (color) and measurement response (black) for KID.

overlap is even stronger than for the temperature retrieval. This is caused by the low water vapor content in the upper atmosphere, which weakens the emission of longwave radiation in these layers.

Figure 5.2 shows the measurement response and the averaging kernels for the IASI-like simulations. The temperature retrieval (Figure 5.2a) for IASI and KID differs clearly in high altitudes (above 20 km). The measurement response is close to 1 for the whole observed altitude range (up to 80 km). This is a consequence of the different absorbers that can be observed in both spectral regions. While both sensors detect signals from water vapor, ozone and oxygen, IASI is also detecting some carbon-dioxide bands around 20 and 70 THz (compare Figure 2.1). As the CO_2 mixing ratio is virtually constant for all heights, these bands allow temperature retrieval nearly throughout the whole atmosphere. The water vapor retrieval (Figure 5.2b) is almost the same for IASI and KID. However the averaging kernels for IASI begin to overlap in lower altitudes and are slightly weaker. A possible explanation is the smaller frequency range of the KID instrument, which allows for more distinct observation of the H_2O bands (compare Figure 2.3).

In conclusion, both instruments are sensitive to water vapor and temperature over a wide altitude range (roughly up to 60 km). IASI is in advantage regarding temperature retrieval in the upper troposphere, as a simple result of the absorbers observed in the spectral region. KID is somewhat more sensitive to water vapor in the lower troposphere,

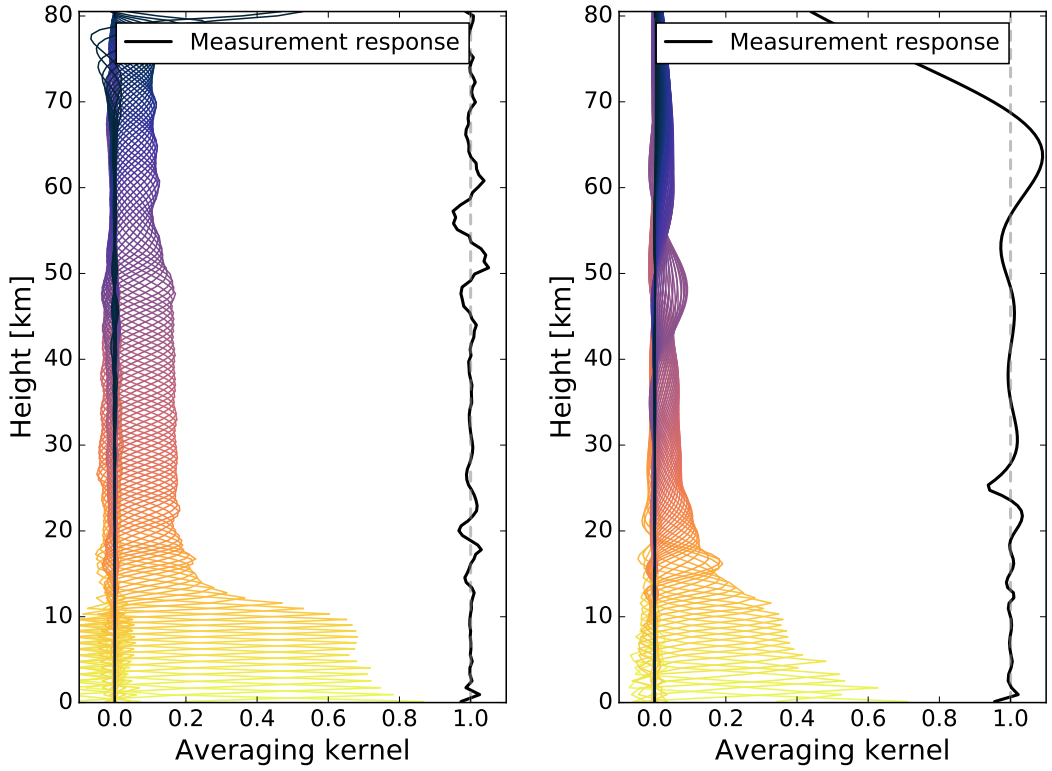

(a) Temperature.

(b) Water vapor.

Figure 5.2: Averaging kernels (color) and measurement response (black) for IASI.

possibly leading to more accurate retrievals in these layers.

5.2 Vertical resolution

The vertical resolution is used to practically describe the retrieval performance. It characterizes the size of the smallest feature in a profile that can still be resolved. The vertical resolution can be estimated from the averaging kernels of the retrieval. It is defined as the full-width at half maximum — the width of the kernel where its peak value is decreased to the half. As shown in the previous section the averaging kernels tend to be very narrow in the troposphere. This makes the calculation of the kernel width more complex. The problem is solved by using the two values above and below the half maximum. A linear interpolation of these values is used to find a good approximation of the pressure level where the half maximum is reached.

The vertical resolutions are shown in Figure 5.3. The results are presented per retrieval species to directly compare the performance of both instruments. Figure 5.3b shows the vertical resolution for the temperature retrieval. Both instruments perform equally well in the lowest 10 km. The achieved vertical resolution is around 1000 m, with KID performing slightly better. But the differences are small, especially when considering the conceptual state of the study. Larger differences can be seen when comparing the vertical resolution of water vapor (Figure 5.3b). Both instruments resolve vertical

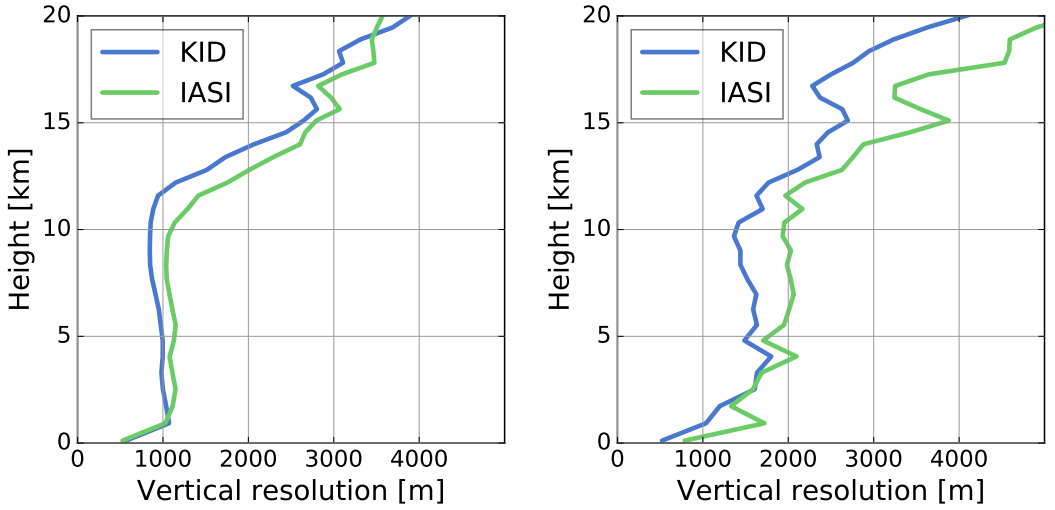

(a) Vertical resolution for temperature.

(b) Vertical resolution for water vapor.

Figure 5.3: Vertical resolution for KID and IASI.

structures of 1000–1500 m for heights below 5 km. In the upper troposphere, KID seems to appear to perform better. While the vertical resolution of the KID sensor is nearly constant around 1500 m, IASI only achieves a resolution of 2000 m. This behavior fits the findings in the previous section (5.1), where the averaging kernels tend to be less distinct for IASI in heights between 5 and 10 km. In general the examination of retrieval resolutions confirmed the positive impressions of KID. The linear OEM retrieval shows that a hyperspectral microwave sensor is able to achieve results comparable to IASI. For the troposphere the performance seems to be even better.

It has to be said that the smallest resolvable feature in a real retrieval setup is probably larger than this (see Chapter 6). But as the full-width at half maximum is used as an estimate of the resolution for both instruments, it is sufficient for a comparison.

5.3 Retrieval error

When comparing different retrieval methods it is necessary to compare the retrieval error. The retrieval error quantifies the uncertainty of a retrieved atmospheric state. OEM allows to directly estimate the retrieval error from the averaging kernel matrix and the assumed error covariances. Following Rodgers et al. (2000) the retrieval error S is described through:

$$S^{-1} = K^T S_y^{-1} K + S_a^{-1}. \tag{5.1}$$

The first term combines the Jacobian K with the noise of each channel S_y. Noise in channels with a strong Jacobian yields to a strong retrieval error, whereas noisy channels with a weak Jacobian do not affect the result that much. For the extreme case of a zero Jacobian, the error is obviously equal to the a priori covariance. The Jacobians

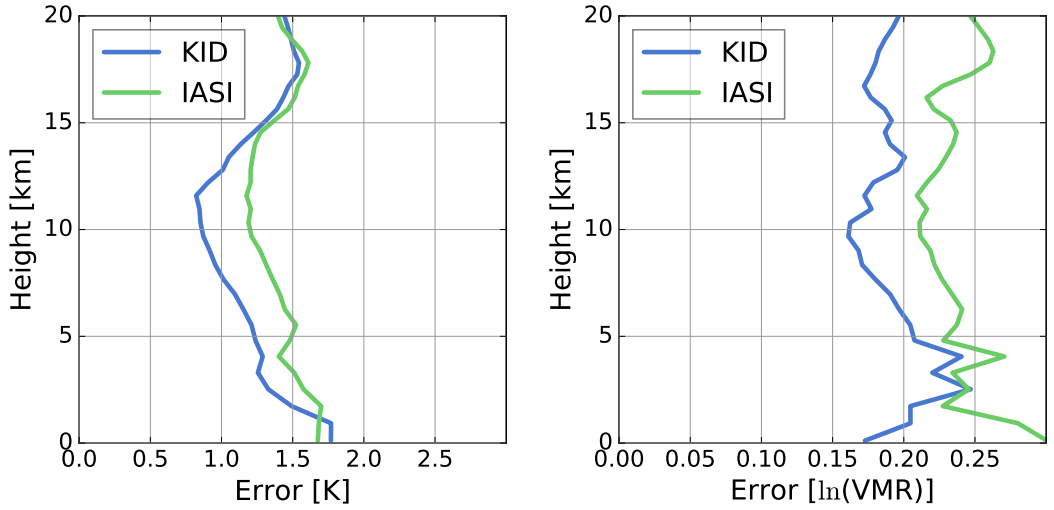

(a) Retrieval error for temperature.

(b) Retrieval error for water vapor.

Figure 5.4: Retrieval error for KID and IASI.

allow to also determine the height-dependency of the error. Therefore the retrieval performance can be evaluated as a function of height (e.g. troposphere for water vapor and stratosphere for ozone).

Figure 5.4 shows the retrieval error for the temperature and water vapor retrievals of both instruments. The temperature retrieval error is similar in the lowest atmospheric layer (Figure 5.4a). Both retrievals show an error of 1.5 K. The differences grow for higher altitudes. In 5–12 km height, KID's error decreases to roughly 0.7 K. Whereas IASI only improves slightly to an error of 1.2 K.

Regarding water vapor, KID performs slightly better. The retrieval error in logarithmic scale is between 0.16 and 0.22 up to 10 km (Figure 5.4b). The error of the IASI-like infrared sensor shows a similar height-dependency but is larger, reaching values between 0.21 and 0.3.

The retrievals performed with both instruments are prone to errors in the same order of magnitude. KID performs slightly better within the troposphere (below 15 km). An interesting behavior can be seen close to the ground. While the error for IASI increases significantly, the error for KID decreases to 0.17. The reason for this difference is the cold background in the microwave simulation ($\varepsilon = 0.6$), which leads to a distinct signal. In reality, the emissivity is not perfectly known. The real error is expected to be larger.

In summary, the overall performance of the KID concept is comparable with the IASI-like instrument. The measurement responses of both sensors indicate high sensitivity to temperature and water vapor up to high altitudes around 50 km. Temperature can be retrieved with a vertical resolution around 1000 m. The water vapor retrieval resolves structures of 1000 to 2000 m.

Chapter 6

Non-linear iterative retrieval

The goal of a retrieval is to translate a set of measurements into information about a specific quantity. OEM can be used to achieve this goal. It uses a priori information to estimate the radiative properties (namely the Jacobians). The a priori state does not represent the real atmosphere perfectly. For some conditions it may differ a lot from the real state. Therefore the calculated Jacobians are only a rough estimate of the dependency of emitted radiation and the atmospheric state. This limits the possible retrieval performance significantly. To improve the quality of the radiative transfer an iterative approach can be used. During each iteration not only the retrieved profile but also the assumed Jacobians are updated. These iterative algorithms are computationally expensive, because the radiative transfer as well as the linear algebra for the OEM has to be done over and over again. For that reason, the non-linear iterative retrieval is performed for the KID sensor only.

6.1 Cost function

In order to evaluate the different iteration steps it is necessary to define proper metrics. Like this it is possible to prevent the retrieval from developing into non-physical solutions. This metric is called a cost function. The cost function is also used to determine if the retrieval has converged into a stable state. It is defined in a way to penalize deviations from the measured spectrum y as well as differences to the a priori state x_a:

$$\chi^2(x) = [x - x_a]S_a^{-1}[x - x_a] + [F(x) - y]S_y^{-1}[F(x) - y]. \tag{6.1}$$

The weighting of the different terms is dependent on the measurement covariance S_y and the a priori covariance matrix S_a. If the measurement noise is low and the knowledge of the a priori state is limited, it is reasonable to trust the measurement. In this case the differences between the measured spectrum and the calculated spectra are weighted more than differences to the a priori state. Using χ^2 one can quantify the quality of a retrieved state compared to the measurement and the chosen a priori state. When starting with x_a obviously only the difference to the measured spectra is non-zero. So

the aim of the retrieval is to reduce this term while keeping the deviation to x_a moderate. Mathematically speaking we are looking for a minimum in the cost function. Unfortunately this cannot be done analytically as the evaluation of the cost function depends on the forward model. But there are different methods of finding a function's minimum. The approach used in this study is the so called Levenberg-Marquardt algorithm. The basic concept and a sketch of its implementation is given in the following section.

6.2 Levenberg-Marquardt algorithm

The Levenberg-Marquardt algorithm is an iterative approach to find the minimum of the cost function. The solution is found by performing several iterations i of an OEM retrieval (Rodgers et al., 2000). Hereby, the next step x_{i+1} is dependent on the current solution x_i:

$$x_{i+1} = x_i + \left[(1+\gamma)S_a^{-1} + K_i^T S_y^{-1} K_i\right]^{-1} \left\{K_i^T S_y^{-1}[y - F(x_i)] - S_a^{-1}[x_i - x_a]\right\}. \quad (6.2)$$

The algorithm includes a damping parameter γ which controls the step size done by the retrieval in each iteration. For high γ values the algorithm follows the steepest descent. This leads to robust results but is very slow. The algorithm needs a vast amount of iterations to converge to a solution. If γ is zero the Levenberg-Marquardt algorithm is equal to the linear OEM solution. The idea is to adjust γ after each iteration to get the benefits of both methods. The retrieval is started with a very high γ (order of 10^9), so that the solution is found by following the steepest descent. After each iteration the cost function is evaluated to see, if the retrieval fits the measurement better. If this is the case, γ is decreased which leads to an increasing step size. Simply put, when the results are good the retrieval gets more scope for development in the next step. If the cost function increased, γ is also increased. This prevents the solution from "exploding". As soon as the retrieved state further deviates from the measurement, it is discarded and the course is adjusted to closer follow the steepest descent. The procedure is as follows:

1. Compute χ^2.
2. Choose a value for γ.
3. Solve the linear equation (Eq. 6.2).
4. Run the forward model and evaluate χ^2.
5. If χ^2 increased, increase γ and go back to 3.
6. If χ^2 decreased, decrease γ and update the retrieval solution.

A practical problem for each iterative algorithm is to find or define a criterion when to stop further iterations. The straight-forward solution is to define a threshold for the

cost function. This is not practical because the starting value of χ^2 can vary in order of magnitudes depending on the measured spectrum; and so does the final value of χ^2 for the converged solution. A more robust criterion is to use the relative change in the cost function. If χ^2 decreases less than 1‰ the solution is considered to be converged and further iterations are skipped. A more detailed analysis of the convergence behavior of different runs is given at the end of this chapter (Subsection 6.3.3).

6.3 Retrieval results

To gain a better understanding of the retrieval behavior, runs with different atmospheric profiles were performed. Afterwards several retrieval metrics were compared to find special characteristics and maybe link them to features in the specific atmosphere.
For that reason, forward simulations were calculated to get simulated measurements of a KID sensor. The atmosphere sample consists of synthetic profiles and randomly selected Eresmaa profiles. Synthetic profiles give the opportunity to investigate certain features, like the vertical resolution, without other side-effects. The retrieval of an actual (model) profile is a final check to see how the retrieval performs in a realistic setup. Artificial noise is added to the simulated spectra to make the setup more realistic. The noise is normally distributed and the mean noise for each channel is equal to the assumed $Ne\Delta T$ (Section 2.1).
The most relevant quantity is the retrieved atmospheric profile ("Retrieval"). It is shown beside the used a priori ("A Priori") and the true profile ("Truth"). The retrieved profile is presented with an uncertainty range. This allows an estimate of the value range in agreement with the retrieval. In addition the averaging kernel as well as the measurement response are investigated. This is done to ensure that the sensitivity of the retrieval is constant during all iterations. A final relevant measure is the so called kernel smoothing. It shows how the retrieval would look like for a perfect measurement without any measurement error. It therefore allows to visualize the smallest features possible to detect.
All these measures are calculated in the state space. To fully understand a retrieval it is necessary to also investigate the measurement space. Therefore the simulated spectra for the measured signal, the a priori state and the final retrieval result are shown. As the retrieval tends to show very good agreement in spectral space, it is often hard to distinguish the retrieved spectrum from the signal. This is why the residuum between both quantities is calculated, too. It is shown alongside the noise-equivalent temperature (Section 2.1). This gives the opportunity to decide whether remaining differences are caused by limited measurement accuracy, or if the solution found is the optimal estimation although some error in the spectral signal is left.
Together all these quantities allow a detailed analysis of the retrieval performance. The results are expected to be strongly dependent on the atmospheric state. The regarded atmospheres were chosen to cover a wide range of features to get a proper estimate of this dependency. Both investigations are described in detail in the following section.

6.3.1 Synthetic sine wave retrieval

Synthetic profiles are an easy way to investigate certain aspects of the retrieval performance alone. They offer the possibility to only change some parts of an atmospheric situation separately to see how the retrieval responds to this adjustment. In the following, synthetic profiles are used to see how sensitive the retrieval is towards structures that span over different vertical spaces.

The setup is simplified as far as possible by setting all retrieval quantities (water vapor, temperature and ozone) to the a priori values. Afterwards an artificial sine wave with different wavelengths (5 and 10 km) is added to the water vapor profile, keeping the other quantities fixed. The retrieval is still run in "full-mode", meaning that all three quantities are retrieved and the retrieval is in principle free to adjust the values of temperature and ozone. Albeit the usage of a sine wave is a hard test for the retrieval, as for every small positive difference in the profile there is an equalizing negative one in one of the next layers. This leads to a reduction in the spectral signal, as the radiation at top of the atmosphere is always a weighted mean over several atmospheric layers. Therefore an added sine wave may be harder to detect than simple perturbations in single layers.

Figure 6.1 shows the retrieval results for the synthetic profile with an added sine wave of 10 km wavelength. This test case simulates a good-natured retrieval situation. The true profile is close to the used a priori state and there are only very smooth structures to be retrieved. The measurement response indicates that the retrieval is sensitive up to heights of around 50 km. Indeed the vertical resolution is not constant over the whole height range. The averaging kernels peak in the troposphere with a sudden drop at the tropopause (at approximately 12 km). The retrieval solution represents the true profile and is in agreement with the promising linear retrieval metrics (Chapter 4). For higher altitudes the retrieval lacks the vertical resolution to resolve the sine wave. In heights between 15 and 50 km the retrieved solution has adjusted from the a priori to slightly higher values.

The robust retrieval results in the state space can also be verified by looking into the measurement space. Figure 6.2 shows the simulated brightness temperature spectrum for the signal, the a priori and the retrieved profile. All spectra are similar for a broad frequency range, especially for high frequencies above 300 GHz. This is a direct result of initializing the temperature and ozone profiles with the true values. Therefore the simulated spectrum around the oxygen lines is already in good agreement. But there are differences around the water vapor spectral lines (e.g. 183 GHz). In these regions signal and a priori have differences of up to 10 K. The retrieval eliminates these differences nearly completely. The signal and the simulated spectra for the retrieval solution can hardly be distinguished. Figure 6.3 shows the residuum between signal and retrieved solution in spectral space. The residuum is below 0.05 K for most of the observed frequencies. The retrieval reduced the differences in the spectra to an order of magnitude comparable to the assumed instrument noise (grey area). This is the theoretical optimum, as the retrieval does not consider information smaller than the noise.

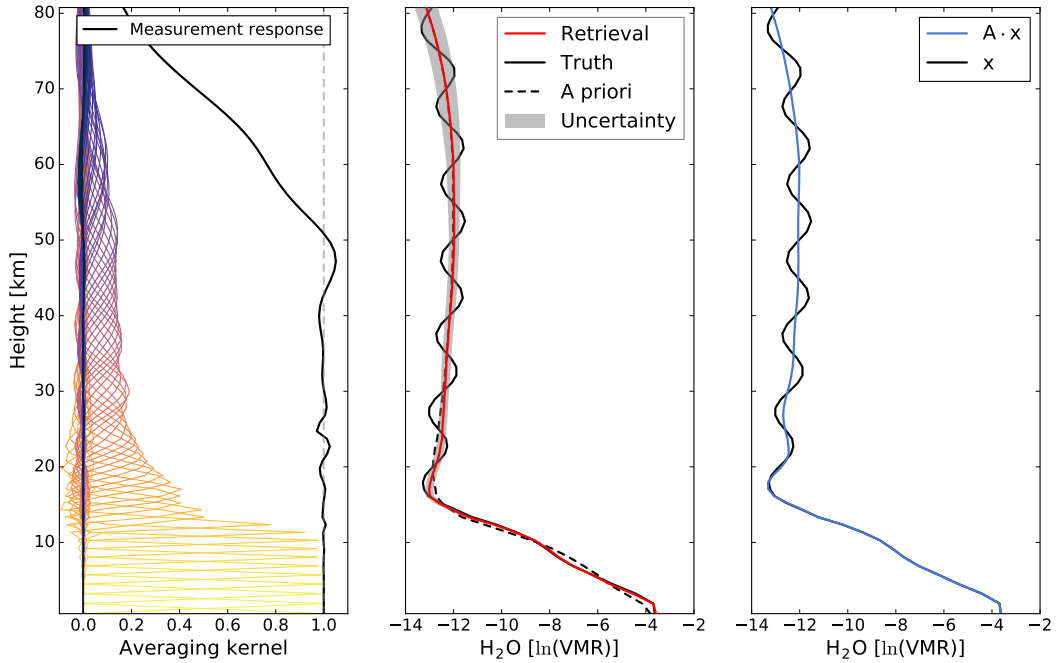

Figure 6.1: Synthetic water vapor retrieval for KID. The added sine wave has a wavelength of 10 km.

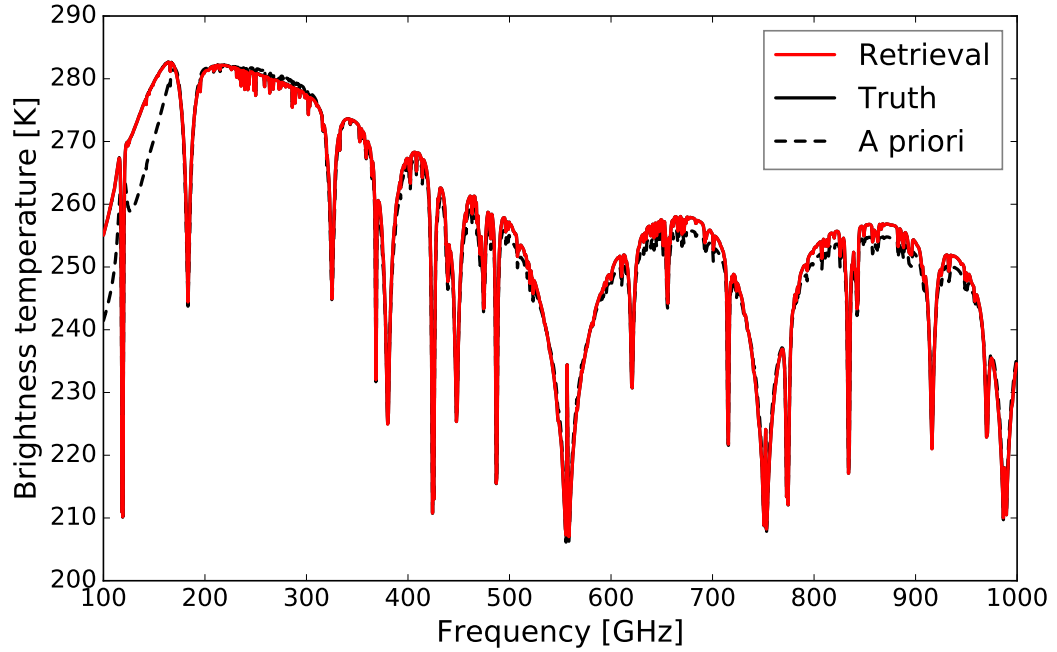

Figure 6.2: Simulated brightness temperatures for KID. The observed profile was superimposed by a sine wave of 10 km wavelength.

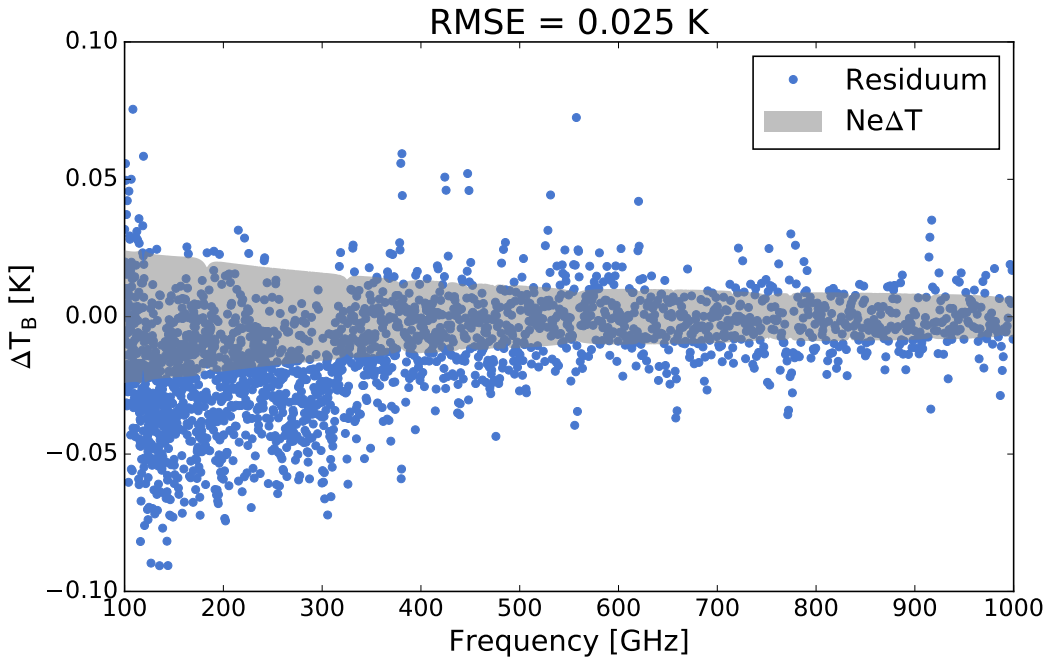

Figure 6.3: Residuum between the simulated measurement and the spectrum of the retrieval result. The observed profile was superimposed by a sine wave of 10 km wavelength. The grey area shows an estimate of the instrument noise.

These results lead to the conclusion that the retrieval works well for smooth atmospheric profiles. This result indicates the general capability of the retrieval as a whole. As shown in the previous chapters, an iterative retrieval algorithm consists of a variety of different parts, all of which have to be chosen consistently to work reliably. The capability to retrieve a simple profile gives confidence in the chosen setup and lays the foundation of further, more demanding test cases.

The next step is to decrease the wavelength of the added sine wave to 5 km. This leads to finer structures in the profile to be retrieved and therefore a more challenging retrieval. The structures are comparable to these found in the troposphere. Temperature and ozone are again set to the true values, but free retrieval quantities. Figure 6.4 shows the retrieval overview. The measurement response as well as the averaging kernels are very similar to the results of the previous run. This is due to the fact that the simulations for both retrieval results show nearly identical Jacobians. The information coming from different atmospheric layers are the same in both runs. More interesting is the resolution of the finer structures in the actual profile. The retrieval is very sensitive in the lower atmosphere, hence the water vapor content close to the surface is retrieved correctly. But even fine structures in higher layers can be resolved. Especially the small inversion at around 5 km height is retrieved accurately. However some small differences can be seen in all altitudes. The retrieval and the true profile do not match as well as for the previous run. This indicates that the structures get too fine to add a significant signal to the measurement. The practical retrieval resolution seems to be reached.

Figure 6.5 shows the spectra for the retrieval, the true profile and the a priori. The a

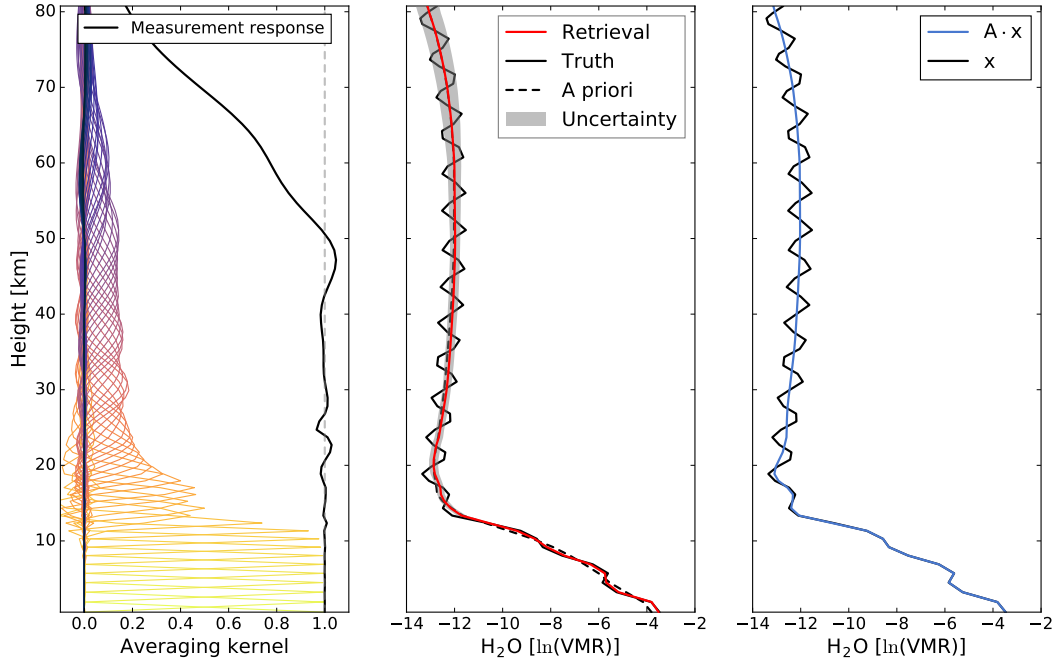

Figure 6.4: Synthetic water vapor retrieval for KID. The added sine wave has a wavelength of 5 km.

priori spectra shows differences around the water vapor spectral lines as already known from the previous run. The spectra agree even more for higher frequencies, where the three lines can hardly be distinguished. This indicates that small disturbances do not add much signal to the measured spectrum.

A look at the remaining residuum (Figure 6.6) strengthens this impression. The simulated spectrum for the retrieval result is in even better agreement with the measurement than for the smooth profile. The spectral differences closely follow the assumed measurement noise. Therefore the retrieval does not see differences to minimize any more.

In summary, the study of retrieval results for synthetic profiles has lead to important results. The retrieval suggests general consistency which strengthens the confidence in retrieval solutions. Initializing retrieval quantities with the true values leads to the expected behavior, that these quantities are not changed by the retrieval. This is not shown in plots, but the retrieved profiles for temperature and ozone were constant for all iterations during both retrieval setups. Also, the information on the retrieval sensitivity calculated from the measurement response agrees with the observed results. The retrieval of water vapor profiles is sensitive up to 50 km, with a decent vertical resolution limited to the troposphere. In addition the two case studies allow a more practical estimate of the vertical resolution than in the previous chapters. The usage of the full-width at half maximum of the averaging kernels is only as theoretical concept to quantify the general correlation between different retrieval layers. In practice, structures slightly smaller than 5 km seem to be the detection limit. Smaller structures do not add enough information to the measured signal.

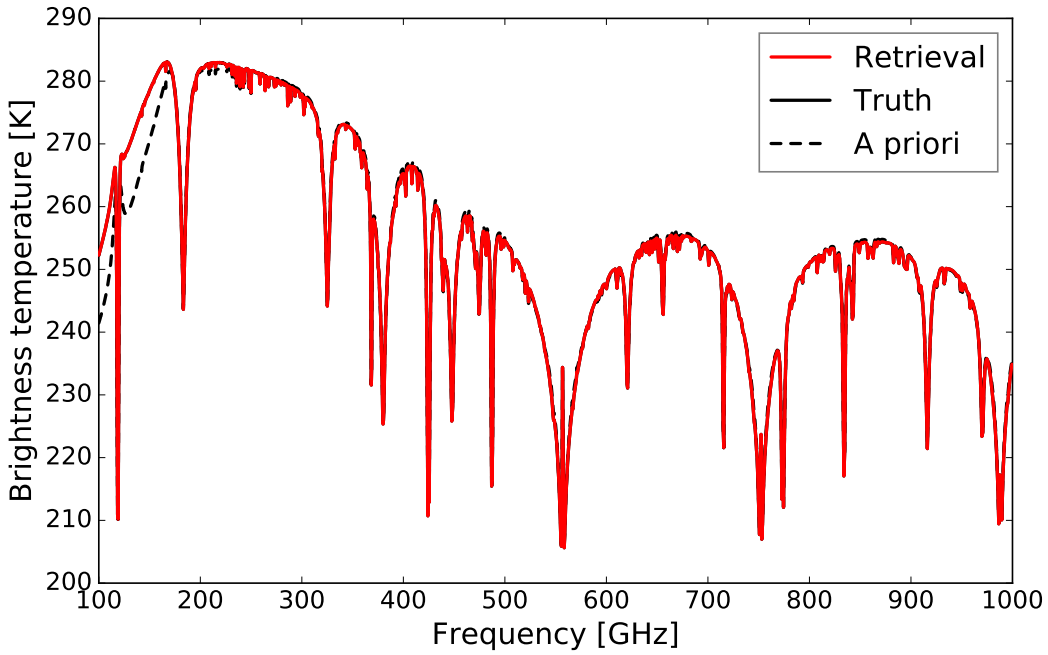

Figure 6.5: Simulated brightness temperatures for KID. The observed profile was superimposed by a sine wave of 5 km wavelength.

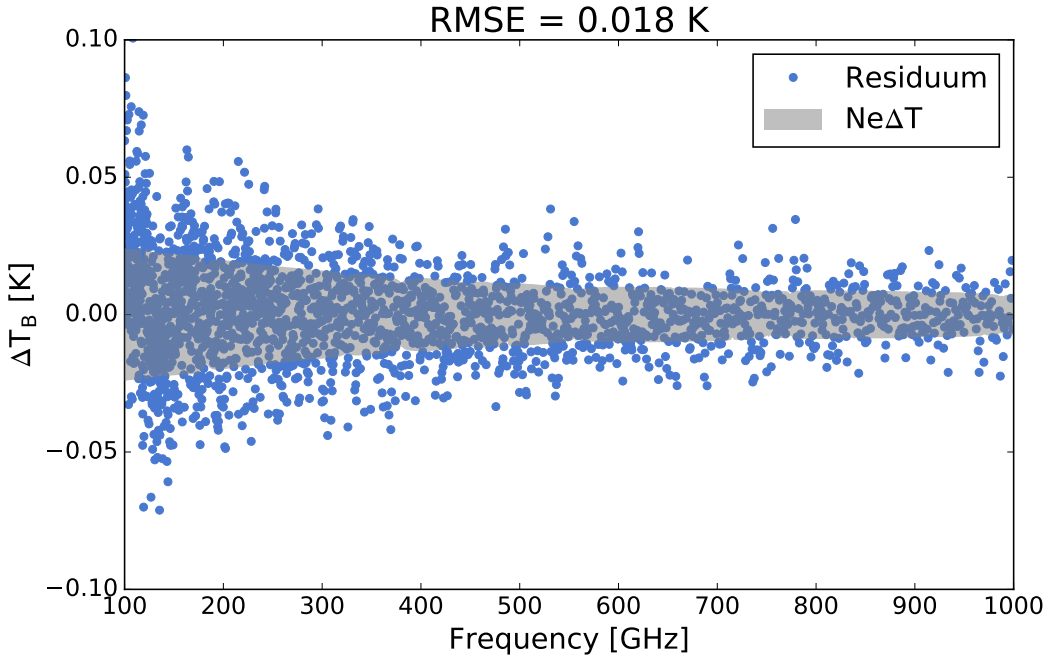

Figure 6.6: Residuum between the simulated measurement and the spectrum of the retrieval result. The observed profile was superimposed by a sine wave of 5 km wavelength. The grey area shows an estimate of the instrument noise.

6.3.2 Retrieval of atmospheric soundings

As a final test case, some real atmospheric soundings were retrieved. The Eresmaa et al. (2012) dataset is again the basis for the simulations. A random subset of 10 profiles was used to simulate measurements. ARTS was used as a forward model to simulate the received measurement signal. Afterwards the iterative retrieval algorithm was used to try to retrieve the atmospheric state. This opens the possibility to test the retrieval in a realistic setup while still being able to compare the results with a known state. The retrieval is again run in its full mode — with water vapor, temperature and ozone being retrieved. In this test case all retrieval quantities show significant differences to the a priori. The algorithm has to adjust all species to decrease the discrepancies in spectral space. This introduces some additional complexity with respect to the discussion of the retrieval results. Errors in the retrieval of one species induce errors in the other retrieval quantities. So while trying to evaluate the retrieval performance for each species, we will have to keep the others in mind to explain certain features.

Most of the randomly chosen soundings are smooth and close to the a priori. The iterative retrieval works well for these cases and yields results comparable to those of the synthetic soundings. A detailed discussion of these retrievals is omitted, because they do not give new insights into the retrieval performance. Instead the retrieval with the most relevant features will be discussed. This allows to investigate the retrieval response to special atmospheric conditions. Furthermore, some internal coherences of the iterative retrieval can be understood.

The chosen sounding is very challenging as it shows vast differences to the a priori. The tropopause is only weakly distinct. Also, the temperature profile is warmer around the tropopause than in an average climatology. This gives the chance to inspect some interesting features including, but not limited to, the interactions of different retrieval species.

Figure 6.7 shows the overview plot for the water vapor retrieval. The averaging kernels show a similar behavior to the previous section. The retrieval is sensitive for regions below 15 km altitude. Therefore the retrieved solution nearly perfectly matches the water vapor mixing ratio close to the ground. There are significant deviations around the tropopause. Above 15 km the solution found underestimates the true water vapor content. Especially the strong feature located at the tropopause ($\approx$ 15 km) leads to errors in lower atmospheric layers. This oscillating propagation of deviations is a direct result of the radiative transfer inside the atmosphere. If the signal in one layer is too high, this can be balanced with a decreased signal in a neighboring layer. The question of interest is, what causes the problems above the troposphere. On that score a look on the temperature retrieval (Figure 6.8) gives an explanation. As already mentioned in the first paragraph the temperature profile is unusual as it has a weakly distinct tropopause. This can be seen in a warm tropopause temperature and small vertical temperature gradients. The vast differences between the true temperature profile and the chosen a priori state result in significant deviations in spectral space (Figure 6.9). The simulated

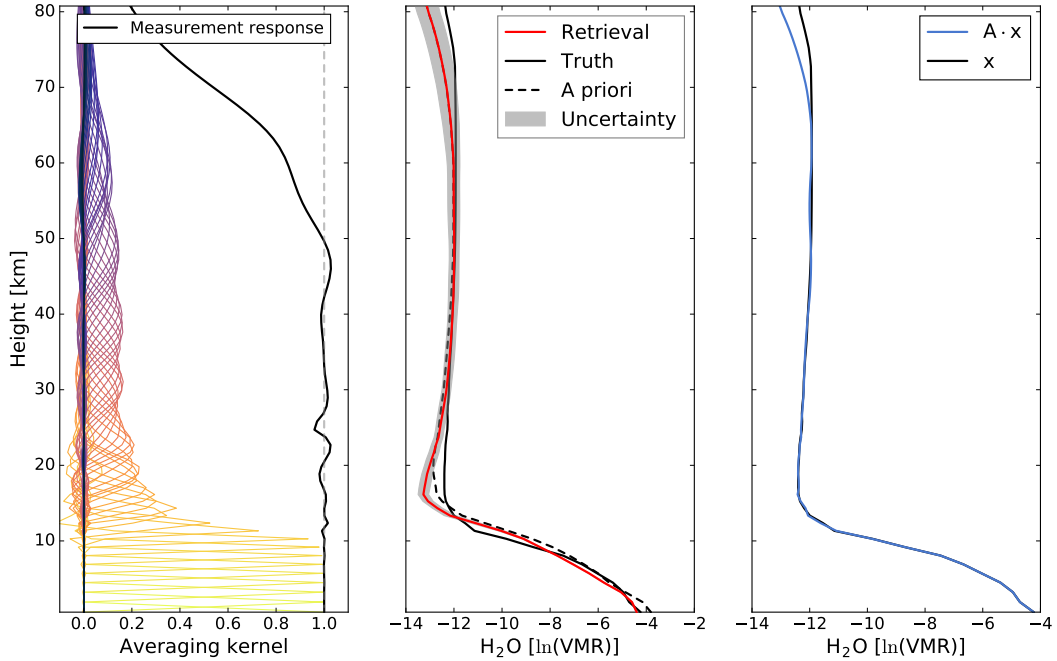

Figure 6.7: Water vapor retrieval for an Eresmaa profile.

measurement and the a priori spectrum differ by several Kelvin in the whole spectral range. Changes in atmospheric temperature lead to analogues changes in the observed brightness temperature. In the end the retrieval result is in good agreement with the true temperature profile. The differences of around 10 K in the beginning are decreased significantly. Also the shape of the temperature profile follows the true temperature profile. But there are some features that cannot be found in the retrieved solution, most noticeably the differences around the tropopause. The retrieved state is off by 2–3 K and the vertical gradients are slightly off. These discrepancies become amplified in the water vapor retrieval. The algorithm tries to draw out the remaining spectral residuum caused by the wrong temperature profile by adjusting the water vapor profile. Even though this works in spectral space, the retrieved profiles can be off. This development is not penalized by the cost function as the improvement in spectral space is weighted more than the discrepancies to the a priori.

Finally the ozone retrieval in Figure 6.10 is considered. The averaging kernels show a sensitivity of the retrieval over the whole altitude range. But the vertical resolution is coarser than for the other species. This can be seen in the retrieved ozone profile. The solution found roughly follows the shape of the true profile, but it is not capable of resolving finer structures. It seems that the retrieval is only useful to estimate the integrated ozone content. Nevertheless the retrieved profile allows to estimate the distribution of atmospheric ozone in general.

To wrap up the case study of an actual retrieval, the spectral state is examined. Figure 6.9 shows the simulated brightness temperature spectrum. As already mentioned, the a priori and the simulated measurement differ in an order of up to 10–15 K in the whole simulated frequency range. The reason for the general bias especially at higher

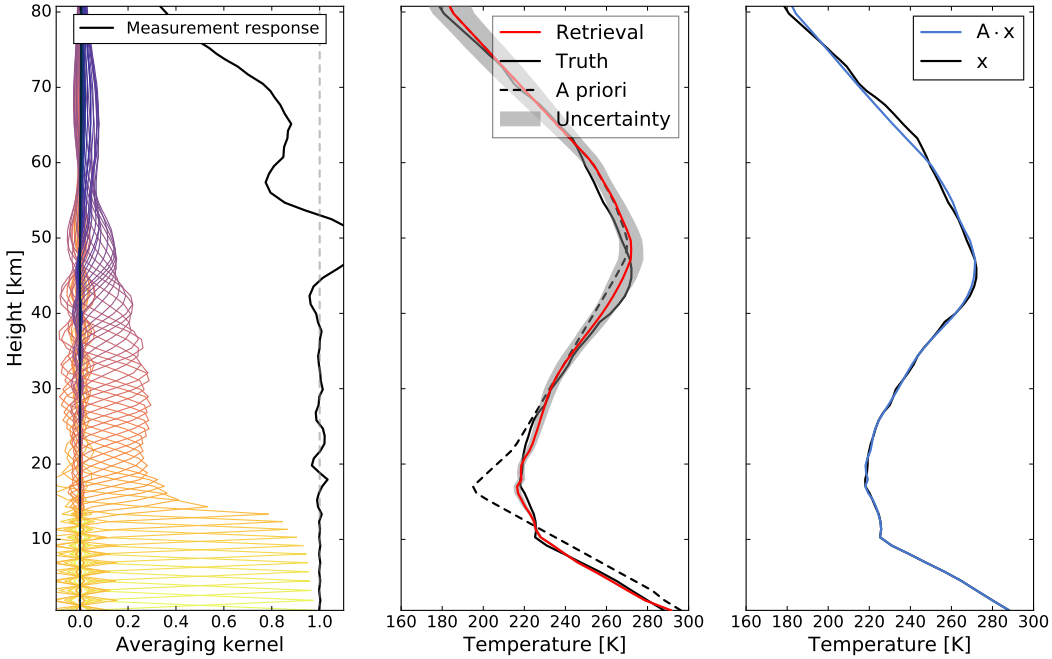

Figure 6.8: Temperature retrieval for an Eresmaa profile.

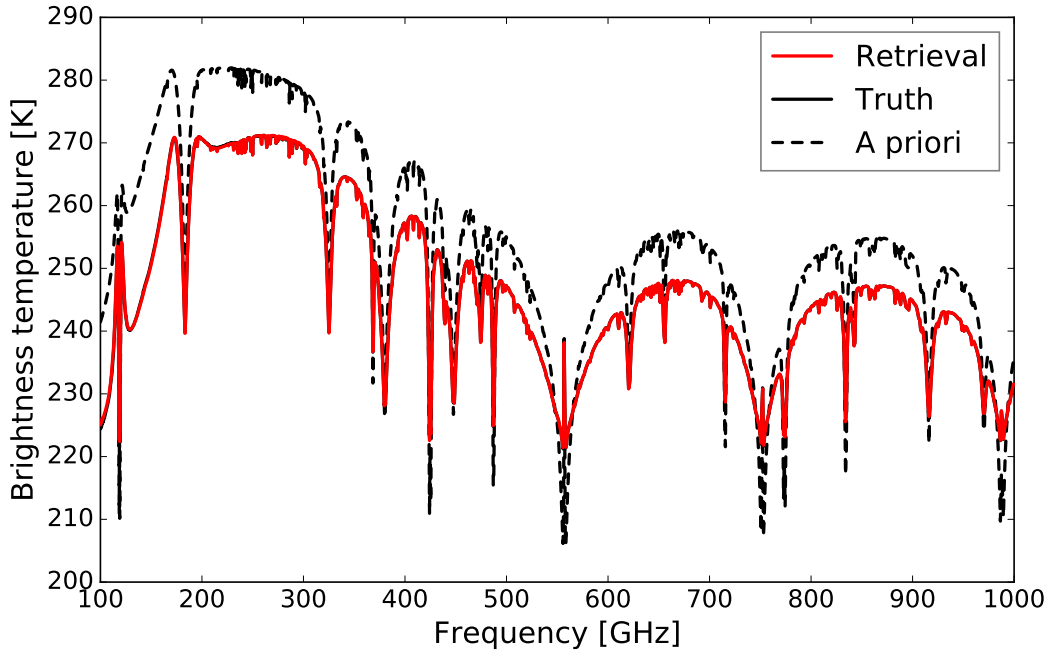

Figure 6.9: Simulated brightness temperature spectrum for an Eresmaa profile.

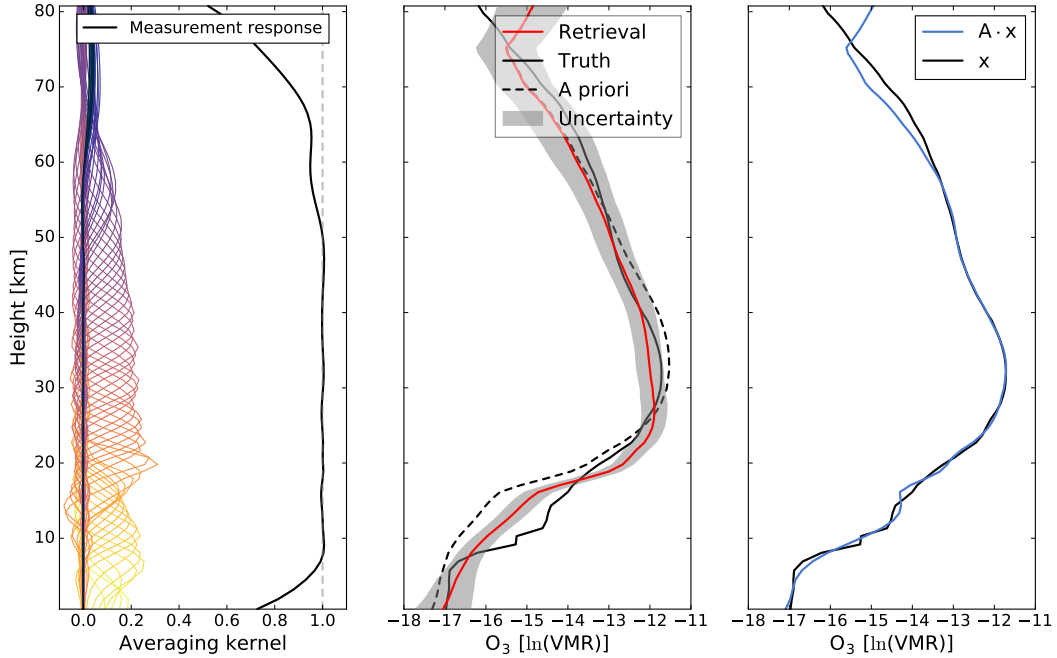

Figure 6.10: Ozone retrieval for an Eresmaa profile.

frequencies is the difference in the temperature profile, which can be seen directly in the spectrum. In addition the behavior around the H_2O spectral lines is different from the a priori state. Therefore it is a success for the retrieval algorithm to minimize the residuum between the retrieval spectrum and the signal down to 1 K. A closer look at the residuum (Figure 6.11) shows that the remaining error is below 0.2 K for frequencies above 300 GHz. In general, frequencies dominated by oxygen absorption agree very well in the spectral space. These are regions mainly driven by the temperature profile. The differences in the final iteration start to look randomly distributed but are 2–3 times the size of the assumed measurement noise. The results for lower frequencies look worse at first sight. At frequencies around the 183 GHz line, the residuum is over 0.6 K. This is due to the insufficiencies in the retrieval of the water vapor profile around the tropopause. In general the retrieval solution is still good, keeping in mind the challenging atmosphere. It was far off the a priori and had some unusual features, that proved themselves hard to be retrieved. Therefore it is a success that the retrieval algorithm was capable of finding solutions that both, followed the shape of the true profile and also stayed stable during the whole iterative process. Moreover the differences in the retrieval seem to be acceptable for most real applications. The temperature profile is showing differences of around 2 K in the lower 30 km, which is a good result. The water vapor retrieval is more difficult to evaluate. For the lower troposphere the differences in log-space are below 0.2, which is sufficient, regarding the influence of the temperature retrieval. Around 15 km the error is over 1 (over 100 %), but the absolute values in this altitude are of several orders of magnitude lower than close to the surface. These regions nearly do not affect the spectral signal at all, which makes them hard to retrieve, although they are very interesting for atmospheric sciences.

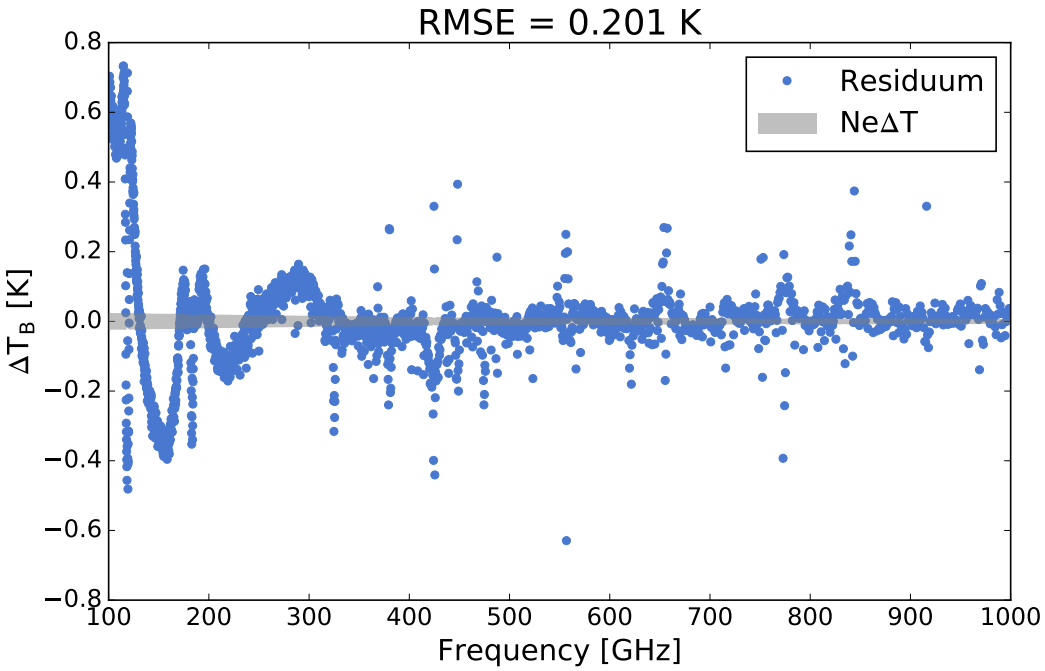

Figure 6.11: Residuum between the simulated measurement and the spectrum of the retrieval result (Eresmaa).

6.3.3 Retrieval convergence

It is important to understand how a retrieval converges in different situations. This information is mandatory for a retrieval in operational use as retrieval solutions have to be found as fast as possible. Reasonable knowledge of convergence behavior is the key to reduce the retrieval runtime. But it is also interesting from a pure scientific point of view. To understand how fast a retrieval adjusts to specific atmospheric situations allows to roughly estimate the sensitivity towards features that can be seen in the profile. The change in the value of the cost function (Equation 6.1) can be used as a proxy for the retrieval convergence. The closer the retrieved solution gets to the truth, the smaller the change done in each iteration gets. Another possibility is to use the change of the profile in state space. This solution seems more intuitive and robust, but one has to define statistical quantities to evaluate the change of all retrieval quantities at all levels using just one scalar. This would lead to further complexity in the retrieval.

The value of the cost function varies several orders of magnitude for different atmospheres. To compare the convergence behavior it is necessary to normalize χ^2. This is done by dividing all values of the cost function by the value for the a priori state. Hereby the value range is limited to be between 0 and 1. This makes the form of the decay for different atmospheric situations comparable.

Figure 6.12 shows the value of χ^2 during the iterative retrieval of different atmospheric situations. Dashed lines mark the iterations where the convergence criterion is fulfilled and further iterations are skipped. For most of the profiles the number of iterations needed to find a converged solution is between 15 and 30 iterations. A significant

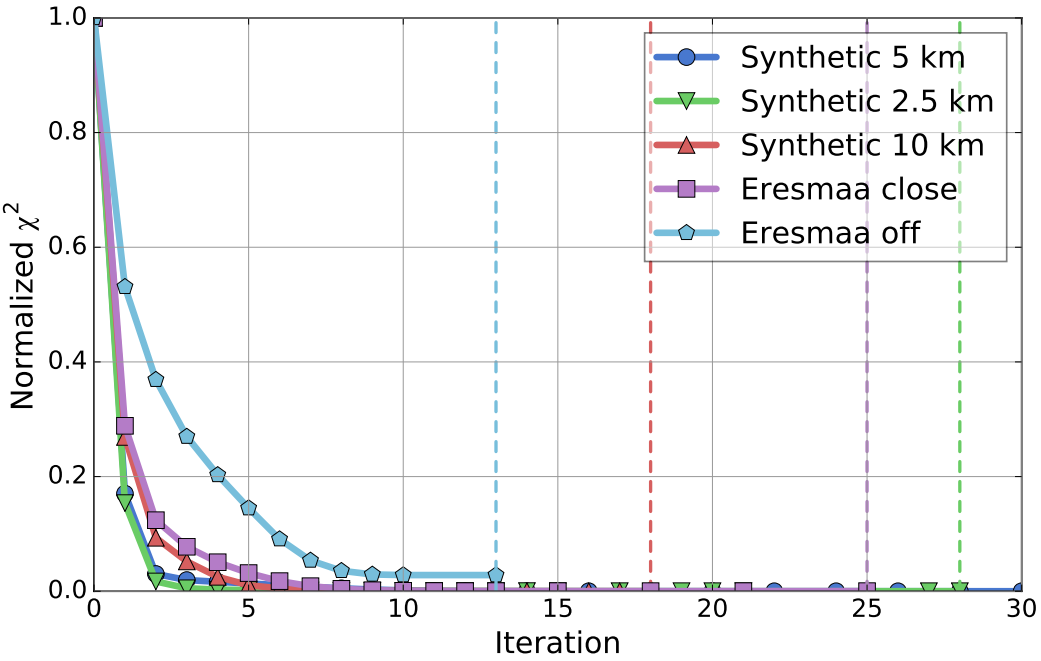

Figure 6.12: Normalized cost function χ^2 during several iterations for different atmospheres. The retrieved atmospheres consist of three profiles with artificial sinusoidal perturbations (2.5, 5 and 10 km wavelength) and two Eresmaa profiles (one close to the "A Priori" and one vastly differing).

outlier is the synthetic profile with an artificial sine wave of 5 km period. It takes 99 iterations to fulfill the convergence criterion. The solution improves with each further step. But this improvement is only visible in spectral space.

One way to prevent such long runs in operational use could be to limit the maximum number of iterations even more. To finally address this problem one has to investigate if the retrieval changes significantly in state space, or if further iterations are not needed to retrieve profiles with the required accuracy.

Another relevant result is the form of the decay for different atmospheric profiles. Figure 6.12 shows that the decreasing rate during the first iterations varies a lot. The two synthetic profiles with short sine waves have a strong decay in in the first few iterations. This can be explained by the fact that these sine waves do not add much difference in the measurement space (Figure 6.5). Even small retrieval steps during the first iterations nearly disperse the spectral difference. This fits the fact that the Eresmaa profile which is far off the a priori state has the slowest decay of all retrievals. It is completely off in spectral space during the first iterations and therefore needs more time to develop towards the true profile. All remaining retrievals lie between these extremes. They all encourage the very intuitive idea, that a retrieval converges faster, if the a priori and the true state are similar.

6.4 Robustness of the retrieval

Although the settings for the retrieval setup were chosen thoughtfully, there is still a high chance of inadequacies. The iterative retrieval algorithm is dependent on a lot of different parameters and assumptions. Starting with the instrument design and the consequent instrument characteristics, through to the assessment of known a priori information. All assumptions concerning the instrument are error-prone, because of its conceptual state. The final channel number, the channel frequencies and even the instrument noise can be different for an actual prototype. These changes could lead to differences in the retrieval performance. It is important to be sure about the possible consequences to obtain confidence in the robustness of the results. Another important part of the algorithm altogether is the a priori information. As described in Chapter 4 the goal of the retrieval is to find a solution, that fits the measurement but agrees with an assumed a priori knowledge. This knowledge and its presumed certainty have a strong influence on the retrieval. To obtain confidence in the presented results, retrieval runs with a slightly changed setup were performed.

At first the noise-equivalent temperature was altered. Two runs, one with double and one with halved instrument noise were performed. Especially the run with increased noise was important to ensure that the retrieval performance is stable, even if a future prototype is more noisy than currently assumed. The results have not shown a significant difference to the control run. A detailed presentation of the retrieval quantities is therefore omitted. It seems like the sheer amount of channels offers robust information. So even with a higher noise in each channel, the information gained from the measurement is still decent.

In a second study the a priori covariance matrices were decreased and increased. The expected behavior was a retrieval result, that sticks more or less to the chosen a priori state. A decreased covariance had hardly any influence on the retrieval outcome. So even when setting up a run with a high certainty in the a priori information, the retrieval yields an unchanged solution. This, again, suggests that the information content of the retrieval is stable and the results are robust. The more interesting result was found when increasing the a priori covariance, which means less trust in the a priori. For the chosen setup halving the values of the covariance matrix lead to exaggerating solutions in the first iteration. To sum up, the chosen a priori covariance seems to express the a priori knowledge in a reliable way. The retrieval is still free enough to develop towards the true atmospheric state, but is kept from diverging.

Chapter 7

Conclusions and outlook

The goal of this thesis was to derive a concept for a hyperspectral microwave sensor and to investigate its value for atmospheric science. The sensor is based on Kinetic Inductance Detectors and could provide several thousand channels in the microwave spectral range.

In a first study the instrument was compared to an idealized infrared instrument. The infrared sensor is a modified version of IASI. An OEM retrieval based on state of the art radiative transfer calculations using ARTS were used to compare basic performance characteristics of both instruments. It has been shown that the sensor concept is capable of achieving comparable results as the IASI-like instrument. Linear retrieval results have shown vertical resolutions around 1000 m for temperature and 1500 m for water vapor within the troposphere. In higher altitudes the resolution decreases, but the retrieval is still sensitive.

To further investigate the performance during a real retrieval, a couple of test retrievals were performed. KID is capable of resolving water vapor profiles with structures of around 3–5 km vertical range. The simultaneous retrieval of several retrieval quantities (water vapor, temperature, ozone) even works for extreme cases.

Summarizing, a hyperspectral microwave sensor based on KID technology would provide a very useful tool for atmospheric science. High resolution profile retrievals could be performed even in cloudy conditions. This is an improvement for climate sciences in general. Atmospheric soundings with high vertical resolution in cloudy areas would be a benefit for numerical weather prediction. The development of cyclones depends highly on these conditions. The increased spatial coverage would be beneficial for climate research as it could help to improve the knowledge on the global humidity distribution. Currently these statistics are biased towards dry values, because clear-sky measurements are overrepresented (John et al., 2011).

Nevertheless there are several open issues that require further investigation. One of the most important ones is to examine the possible retrieval of cloud particles. Besides its usage to retrieve water vapor and temperature profiles during cloudy conditions, the future sensor could be used to retrieve quantities of the cloud itself (e.g. snow, ice and liquid water content). In order to achieve this goal the radiative transfer has to be

adapted to include clouds. ARTS is capable of doing this, so future studies could be based on the work done during this thesis.
The reduction of frequencies is a necessary step to make retrieval calculations operable in practice. This thesis showed, as a proof of concept, that it is possible to retrieve reliable results even when only using 100 selected frequency channels. But those selected frequencies have to be chosen in a more elaborate way, to be appropriate for a wide range of atmospheric conditions. It has been shown that a selection based on only one atmospheric profile is not robust. A statistical approach using a set of sample atmospheres seems to be the best solution. The brute-force solution used in this thesis has to be adapted.
It is necessary that further studies keep track of the latest developments regarding the KID technology. Operating temperature, instrument noise and possible channel numbers are essential for the instrument design. A close cooperation with scientists developing this technology is mandatory.
Although these are a lot of issues to solve, all these tasks are manageable. The results of this thesis are very promising and encourage the further scientific investigation of hyper-spectral microwave sensors. If those studies lead to positive results and the technology fulfills the required benchmarks, it seems possible to develop a prototype for airborne missions within the next years.

Bibliography

Anderson, G. P., Clough, S., Kneizys, F., Chetwynd, J., and Shettle, E. P. (1986). AFGL atmospheric constituent profiles (0.120 km). Technical report, DTIC Document. 18

Baselmans, J., Yates, S., Barends, R., Lankwarden, Y., Gao, J., Hoevers, H., and Klapwijk, T. (2008). Noise and sensitivity of aluminum kinetic inductance detectors for sub-mm astronomy. *Journal of Low Temperature Physics*, 151(1-2):524–529. 14

Blumstein, D., Chalon, G., Carlier, T., Buil, C., Hebert, P., Maciaszek, T., Ponce, G., Phulpin, T., Tournier, B., Simeoni, D., et al. (2004). IASI instrument: Technical overview and measured performances. In *Optical Science and Technology, the SPIE 49th Annual Meeting*, pages 196–207. International Society for Optics and Photonics. 12, 17

Buehler, S., Eriksson, P., Kuhn, T., Von Engeln, A., and Verdes, C. (2005). ARTS, the atmospheric radiative transfer simulator. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 91(1):65–93. 17

Bueno, J., Coumou, P., Zheng, G., de Visser, P., Klapwijk, T., Driessen, E., Doyle, S., and Baselmans, J. A. (2014). Anomalous response of superconducting titanium nitride resonators to terahertz radiation. *Applied Physics Letters*, 105(19):192601. 14

Clerbaux, C., Boynard, A., Clarisse, L., George, M., Hadji-Lazaro, J., Herbin, H., Hurtmans, D., Pommier, M., Razavi, A., Turquety, S., et al. (2009). Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder. *Atmospheric Chemistry and Physics*, 9(16):6041–6054. 12

Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A., and Zmuidzinas, J. (2003). A broadband superconducting detector suitable for use in large arrays. *Nature*, 425(6960):817–821. 14

Doyle, S., Mauskopf, P., Naylon, J., Porch, A., and Duncombe, C. (2008). Lumped element kinetic inductance detectors. *Journal of Low Temperature Physics*, 151(1-2):530–536. 14

Eresmaa, R., Benedetti, A., and McNally, A. P. (2012). Diverse profile database of aerosol and trace gas concentrations from the Monitoring Atmospheric Composition

and Climate short-range forecasts. Technical report, NWP SAF Report No. NWPSAF-EC-TR-015, 12 p. 24, 43

Eriksson, P., Buehler, S., Davis, C., Emde, C., and Lemke, O. (2011). ARTS, the atmospheric radiative transfer simulator, version 2 . *Journal of Quantitative Spectroscopy and Radiative Transfer*, 112(10):1551 – 1558. 17

Griffin, M., Baselmans, J., Baryshev, A., Doyle, S., Grim, M., Hargrave, P., Klapwijk, T., Martin-Pintado, J., Monfardini, A., Neto, A., et al. (2015). SPACEKIDS: Kinetic inductance detector arrays for space applications. In *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2015 40th International Conference on*, pages 1–2. IEEE. 14

John, V. O., Holl, G., Allan, R. P., Buehler, S. A., Parker, D. E., and Soden, B. J. (2011). Clear-sky biases in satellite infrared estimates of upper tropospheric humidity and its trends. *Journal of Geophysical Research: Atmospheres*, 116(D14). 9, 51

Mauskopf, P., Doyle, S., Barry, P., Rowe, S., Bidead, A., Ade, P., Tucker, C., Castillo, E., Monfardini, A., Goupy, J., et al. (2014). Photon-noise limited performance in aluminum LEKIDs. *Journal of Low Temperature Physics*, 176(3-4):545–552. 14

Rodgers, C. D. et al. (2000). *Inverse methods for atmospheric sounding: theory and practice*, volume 2. World scientific Singapore. 19, 32, 36

Rosenkranz, P. W. (2001). Retrieval of temperature and moisture profiles from AMSU-A and AMSU-B measurements. *Geoscience and Remote Sensing, IEEE Transactions on*, 39(11):2429–2435. 9

Rothman, L., Gordon, I., Babikov, Y., Barbe, A., Benner, D. C., Bernath, P., Birk, M., Bizzocchi, L., Boudon, V., Brown, L., Campargue, A., Chance, K., Cohen, E., Coudert, L., Devi, V., Drouin, B., Fayt, A., Flaud, J.-M., Gamache, R., Harrison, J., Hartmann, J.-M., Hill, C., Hodges, J., Jacquemart, D., Jolly, A., Lamouroux, J., Roy, R. L., Li, G., Long, D., Lyulin, O., Mackie, C., Massie, S., Mikhailenko, S., Müller, H., Naumenko, O., Nikitin, A., Orphal, J., Perevalov, V., Perrin, A., Polovtseva, E., Richard, C., Smith, M., Starikova, E., Sung, K., Tashkun, S., Tennyson, J., Toon, G., Tyuterev, V., and Wagner, G. (2013). The HITRAN2012 molecular spectroscopic database . *Journal of Quantitative Spectroscopy and Radiative Transfer*, 130:4 – 50. HITRAN2012 special issue. 17

Schneider, M. and Hase, F. (2011). Optimal estimation of tropospheric H_2O and δD with IASI/METOP. *Atmospheric Chemistry and Physics*, 11(21):11207–11220. 24, 27

Versicherung an Eides statt

Hiermit versichere ich an Eides statt, dass ich die vorliegende Arbeit im Studiengang Master of Science Meteorologie selbstständig verfasst und keine anderen als die angegebenen Hilfsmittel — insbesondere keine im Quellenverzeichnis nicht benannten Internet-Quellen — benutzt habe. Alle Stellen, die wörtlich oder sinngemäß aus Veröffentlichungen entnommen wurden, sind als solche kenntlich gemacht. Ich versichere weiterhin, dass ich die Arbeit vorher nicht in einem anderen Prüfungsverfahren eingereicht habe und die eingereichte schriftliche Fassung der auf dem elektronischen Speichermedium entspricht.

Hamburg, den 2. Februar 2017 Lukas Kluft