




CHALMERS
UNIVERSITY OF TECHNOLOGY



Precipitation estimation from the Arctic Weather Satellite: An initial study using probabilistic deep learning

Master's thesis in Engineering Mathematics and Computational Science

LUDVIG BERGSTRAND

DEPARTMENT OF ENVIRONMENTAL AND ENERGY SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2026

www.chalmers.se

MASTER'S THESIS 2026

**Precipitation estimation from the Arctic Weather
Satellite: An initial study using probabilistic deep
learning**

LUDVIG BERGSTRAND



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Environmental and Energy Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Precipitation estimation from the Arctic Weather Satellite:
An initial study using probabilistic deep learning
LUDVIG BERGSTRAND

© LUDVIG BERGSTRAND, 2026.

Supervisors: Eleanor May & Adrià Amell, Department of Environmental and Energy
Sciences

Examiner: Patrick Eriksson, Department of Environmental and Energy Sciences

Master's Thesis 2026
Department of Environmental and Energy Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in L^AT_EX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Abstract

Accurate precipitation information is important for understanding the hydrological cycle, improving weather forecasts and supporting hydrological applications. However, precipitation is difficult to observe with sufficient spatial and temporal coverage. Rain gauges provide only local measurements and ground-based radars are limited to regions where they are available. Satellite observations can complement these systems by providing precipitation information over larger and sparsely instrumented domains.

Passive microwave observations are widely used for precipitation estimation, and the newly launched Arctic Weather Satellite (AWS) offers a novel extension of this capability. AWS is a compact satellite carrying a 19-channel microwave radiometer and the first mission to include sub-millimetre channels around 325 GHz for this purpose. This thesis primarily investigates whether AWS observations can be used to retrieve surface precipitation rates using supervised machine learning. A secondary aim is to determine if these novel sub-millimetre channels have a positive effect on the retrieval performance. To achieve this, a dataset is constructed by matching AWS antenna temperatures with Multi-Radar Multi-Sensor (MRMS) radar precipitation estimates over the contiguous United States. Quantile regression neural networks are trained to predict conditional quantiles of the surface precipitation rate, providing both deterministic estimates and probabilistic estimates of retrieval uncertainty.

Results demonstrate that AWS antenna temperatures contain valuable information for surface precipitation retrieval. The trained model captures broad precipitation structures and achieves near-zero overall bias on the independent test set. Although the predictions do not fully capture the intensity of the precipitation the spatial mapping is consistent with the radar reference. An additional model trained without the sub-millimetre channels evaluates the specific impact of these sensors. Comparisons indicate that the sub-millimetre channels provide a modest but consistent improvement across the overall evaluation metrics. These findings serve as an initial assessment of AWS-based precipitation retrieval. This investigation is particularly relevant due to the upcoming launch of EPS-Sterna, a constellation of AWS satellites which could significantly enhance the temporal sampling of global precipitation observations. Further efforts could explore using other reference precipitation products

to expand the geographical domain, as well as incorporating longer data records as newer AWS observations become available.

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Eleanor May and Adrià Amell. Thank you for all the discussions, support and input throughout this project. It has been both educational and intriguing to study and write about this subject under your guidance. I would also like to thank Patrick Eriksson for his input and for acting as the examiner making the thesis possible.

Finally, I wish to thank everyone who has made my time at Chalmers educational and memorable.

Ludvig Bergstrand, Gothenburg, June 2026

Contents

1	Introduction	1
1.1	Remote sensing of precipitation	1
1.2	Research context	2
1.3	Aim and scope	3
2	Theory	5
2.1	Remote sensing of precipitation	5
2.1.1	Ground-based precipitation observations	5
2.1.2	Visible and infrared satellite observations	6
2.1.3	Spaceborne precipitation radar	7
2.1.4	Passive microwave satellite observations	7
2.2	Passive microwave radiometry	8
2.2.1	Brightness temperature and antenna temperature	8
2.2.2	Microwave channel sensitivity	9
2.2.3	Precipitation signatures	10
2.2.4	Radiative transfer and retrieval	11
2.3	Statistical retrieval theory	11
2.3.1	Forward and inverse problems	11
2.3.2	Bayesian retrieval formulation	12
2.3.3	Approaches to solving the inverse problem	13
2.4	Machine learning and quantile regression neural networks	13
2.4.1	Supervised learning and neural networks	13
2.4.2	Quantile regression	14
3	Data	16
3.1	Arctic Weather Satellite observations	16
3.1.1	Radiometer and channel groups	16
3.1.2	Scan geometry and spatial sampling	17
3.1.3	Level-1B observation product	17
3.2	Multi-Radar Multi-Sensor precipitation data	18
3.3	Additional comparison dataset	19

4	Method	21
4.1	Construction of the AWS–MRMS dataset	21
4.1.1	Study period and domain	21
4.1.2	AWS remapping and scan-position selection	21
4.1.3	Temporal and spatial matching with MRMS	22
4.1.4	Quality control and filtering	22
4.2	Feature and target construction	23
4.2.1	Input variables	23
4.2.2	Target variable	24
4.2.3	Log transformation of the target	25
4.2.4	Feature normalization	26
4.3	Train, validation and test split	26
4.4	Quantile regression neural network	27
4.5	Training procedure	28
4.6	Evaluation	28
4.6.1	Prediction summary	28
4.6.2	Deterministic metrics	29
4.6.3	Probabilistic diagnostics	29
5	Results	31
5.1	Training behaviour	31
5.2	Performance statistics on the test set	33
5.3	Comparison with MRMS reference swaths	38
5.4	Effect of the sub-millimetre channels	46
6	Discussion	48
6.1	Retrieval performance	48
6.2	Underestimation of high precipitation rates	49
6.3	Contribution of the sub-millimetre channels	50
6.4	Limitations of the evaluation	50
7	Conclusion	52

List of Figures

3.1	Example AWS observations for one antenna temperature channel over a 12-hour period. The dotted lines indicate the contiguous United States (CONUS).	18
4.1	Example AWS antenna temperature swaths of inputs for selected channels.	24
4.2	Distribution of MRMS surface precipitation rates in the full collocated AWS–MRMS dataset.	25
4.3	MRMS rain-rate distributions for the training, validation and test datasets.	27
5.1	Training and validation loss trajectories computed in log-transformed target space for the all-channel and reduced-channel QRNN models.	32
5.2	Probabilistic reliability diagrams across specific MRMS reference precipitation bins computed on the test set. The diagonal line represents perfect calibration while the yellow dashed line, which overlaps the diagonal line, represents the reliability for the training dataset in the log-transformed training space.	33
5.3	Example probabilistic prediction for an individual test sample. The predicted discrete quantiles are smoothed via interpolation to approximate the continuous cumulative distribution function (left) and then derive the probability density function (right).	34
5.4	Relative binned test-set metrics using posterior mean predictions normalized by the mean MRMS reference rain rate of each bin.	36
5.5	Observed and predicted rain-rate density distributions evaluated across the complete test dataset.	37
5.6	Conditional density plots of posterior mean predictions against MRMS reference values. The solid red line marks the conditional mean. The dashed line represents the ideal 1:1 identity line. The yellow regions define areas where the sample concentration drops to zero.	38

5.7	Swath comparison for the heavy-rain case. The top row presents absolute predictions, the MRMS reference target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.	40
5.8	Swath comparison for the moderate-rain case. The top row presents absolute predictions, the MRMS target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.	42
5.9	Swath comparison for the weak-rain case. The top row presents absolute predictions, the MRMS target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.	44
5.10	Spatial mapping of the predicted probability of exceeding specific precipitation thresholds for the heavy-rain case in figure 5.7 using the all-channel model.	45

List of Tables

3.1	Overview of the AWS radiometer channel groups. Footprint values are approximate nadir full width at half maximum (FWHM).	16
3.2	Main AWS Level-1B variables.	19
3.3	Main information from the MRMS Surface Precipitation Rate product.	19
4.1	Total number of samples in the training, validation and test datasets.	27
5.1	Test-set verification metrics using posterior mean predictions.	35

1 Introduction

1.1 Remote sensing of precipitation

The monitoring of precipitation is an important component for understanding the Earth's hydrological cycle and improving numerical weather prediction. Precipitation affects water availability, floods, droughts and the development of severe weather systems. Accurate observations are therefore important both for scientific studies of the climate system and for practical applications such as weather forecasting and hydrological modelling (Michaelides et al., 2009).

Despite its importance, precipitation remains difficult to measure accurately. Rain gauges provide direct measurements of precipitation at individual locations but only represent a small area around the measurement point. Ground-based weather radars provide broader spatial coverage and information about precipitation over larger regions. However, radar observations are limited to regions where radar networks are available and a dense radar network is required to cover a large area. This means that large parts of the world remain poorly observed by ground-based systems, especially over oceans, remote regions and sparsely instrumented land areas. Satellite observations can therefore complement ground-based systems by providing precipitation information over larger and more spatially uniform domains (Kidd et al., 2017; Kidd and Levizzani, 2011).

Satellite precipitation observations are commonly based on spaceborne radar, passive microwave and, to a lesser extent, visible and infrared observations. Spaceborne precipitation radars provide more direct information about precipitation structure but such instruments are only available on a limited number of satellites and have narrow swaths compared to passive sensors (Fabry, 2015). Visible and infrared sensors can provide frequent information about cloud systems, especially from geostationary orbit, but they mainly observe cloud-top properties. Passive microwave observations provide an important compromise because the measured radiation is affected by emission from liquid water and scattering by frozen hydrometeors inside clouds. For this reason, passive microwave observations are an important source of information for satellite precipitation retrievals where retrieval refers to the estimation of

precipitation from indirect satellite measurements (Stephens and Kummerow, 2007).

The Global Precipitation Measurement (GPM; Hou et al. 2014) mission is one of the main current satellite missions for precipitation observation. It combines measurements from the GPM Core Observatory which carries both a microwave imager and a precipitation radar, with observations from a constellation of passive microwave sensors. This makes it possible to provide global estimates of rain and snow. These satellites are flown close to earth, primarily at altitudes 800 km and below, and only observe a limited swath of the Earth at a time. As a result satellite precipitation products depend on the number of available sensors and their corresponding orbital sampling (Kidd and Levizzani, 2011).

The Arctic Weather Satellite (AWS), launched in 2024, is a small satellite carrying a 19-channel cross-track microwave radiometer measuring at frequencies that have been typically used for precipitation retrievals (Eriksson et al., 2025). AWS demonstrates a smaller and more cost-efficient concept for passive microwave observations and acts as the proto-flight mission for EPS-Sterna. EPS-Sterna is an approved constellation of polar-orbiting satellites based on the AWS concept. The first six satellites are planned for launch in 2029 and the constellation is expected to improve the temporal sampling of global microwave observations (EUMETSAT, 2026; SMHI, 2026).

The AWS radiometer differs from many previous sensors because it includes channels around 325 GHz, referred to as sub-millimetre channels since their wavelengths are smaller than 1 mm. These channels represent a new part of the spectrum for satellite meteorology and are sensitive to scattering by frozen hydrometeors (Eriksson et al., 2025; Duncan et al., 2025). They may therefore provide additional information for precipitation retrieval compared with lower-frequency microwave channels.

1.2 Research context

Passive microwave observations are already widely used for satellite precipitation retrieval. In the GPM processing system, the Goddard Profiling Algorithm (GPROF; Pfreundschuh et al. 2022a) retrieves precipitation from the passive microwave sensors of the GPM constellation, including the GPM Microwave Imager (GMI; Draper et al. 2015). This demonstrates that passive microwave radiances contain useful information about precipitation, even though the relation between the measured radiances and surface precipitation is indirect.

Machine-learning methods have recently been used to improve satellite precipitation retrievals and better exploit such indirect information. Several studies have applied deep learning to satellite observations to enhance the accuracy and spatial resolution of precipitation estimates (Sadeghi et al., 2022; Goroo et al., 2020).

A probabilistic formulation is particularly useful for precipitation retrieval because the same satellite observation can correspond to different surface precipitation rates. Quantile regression neural networks (QRNNs) address this by predicting several quantiles of the conditional distribution rather than only a single deterministic estimate. QRNNs have been shown to provide retrieval accuracy comparable to standard neural networks while also producing statistically consistent uncertainty estimates for non-Gaussian retrieval errors (Pfreundschuh et al., 2018). Neural-network-based versions of GPROF have shown improved retrieval performance for passive microwave precipitation retrievals while retaining the ability to provide probabilistic outputs (Pfreundschuh et al., 2022a). Related work has also applied probabilistic neural networks to precipitation retrieval from geostationary satellite observations (Pfreundschuh et al., 2022b; Amell et al., 2025). These specific studies provide the methodological background for the machine-learning approaches used in this thesis.

The new aspect in this thesis is the use of AWS observations for precipitation retrieval. AWS provides passive microwave observations in several sounding bands and includes sub-millimetre channels around 325 GHz (Eriksson et al., 2025). Previous studies of sub-millimetre microwave observations have mainly focused on their sensitivity to ice clouds and frozen hydrometeors. For AWS the sub-millimetre channels have been shown to add information in scenes with significant frozen water aloft (Duncan et al., 2025). Since many precipitating systems contain frozen hydrometeors above the melting layer these channels may provide information relevant for precipitation retrieval.

1.3 Aim and scope

The aim of this thesis is to investigate whether AWS observations can be used to retrieve surface precipitation rates using supervised machine learning. The retrieval models are trained and evaluated using collocations between AWS observations and precipitation estimates from the Multi-Radar Multi-Sensor (MRMS; (Zhang et al., 2016)) system. MRMS is a system that provides a radar-based precipitation product with high-resolution surface precipitation estimates over the contiguous United

States.

A second aim is to assess the contribution of the AWS sub-millimetre channels to the retrieval. This is done by comparing models trained with all available AWS channels to models where the sub-millimetre channels are excluded. The comparison is used to evaluate whether these channels provide additional information for retrieving surface precipitation from AWS observations.

The thesis focuses on data-driven precipitation retrieval using machine learning. The retrieval models use AWS antenna temperature observations together with viewing-geometry information. The main model type is the quantile regression neural network (QRNN) which is used to produce probabilistic precipitation retrievals.

The scope of the thesis is limited by the availability of collocations between AWS and MRMS. Since MRMS is only available over a limited geographical domain the results should not be interpreted as a global validation of AWS precipitation retrievals. The MRMS product also contains uncertainties and should not be treated as perfect truth. In addition, the study is affected by the early mission phase of AWS where calibration uncertainties, data availability and possible instrument issues may influence the results. The results should therefore be interpreted as an initial assessment of AWS-based precipitation retrieval over the contiguous United States.

The remainder of this thesis is structured as follows: Chapter 2 details the theory behind remote sensing and machine learning. Chapter 3 describes the data sources used. Chapter 4 outlines the methodology. Finally, Chapter 5 presents the results followed by a discussion and conclusions in Chapters 6 and 7.

2 Theory

This chapter introduces the concepts needed to understand the retrieval developed in this thesis. It first describes precipitation observing systems. It then introduces passive microwave radiometry and explains why retrieving precipitation from satellite radiances is an inverse problem. Finally, it describes quantile regression neural networks as a probabilistic retrieval method.

2.1 Remote sensing of precipitation

Precipitation can be observed using several types of measurement systems, including rain gauges, ground-based weather radars and satellite instruments. These systems differ in what they measure, their spatial and temporal coverage and how directly the measurements are related to precipitation rate. This is important because most precipitation datasets are not direct measurements of precipitation everywhere but estimates derived from observations through a retrieval or algorithm. Unless otherwise stated, the general concepts of precipitation remote sensing in this section are based on Michaelides et al. (2009), Stephens and Kummerow (2007) and Kidd and Levizzani (2011).

2.1.1 Ground-based precipitation observations

Rain gauges provide one of the most direct measurements of precipitation at the surface. They measure the amount of water collected at a specific location and are therefore often used as reference measurements for accumulated precipitation. However, a rain gauge represents only a very small area. Since precipitation can vary strongly over short distances, a single gauge does not necessarily describe the precipitation in the surrounding region. Gauge networks can reduce this problem but their coverage is uneven and they are often sparse in remote regions, mountainous areas and over the oceans (Kidd et al., 2017).

Ground-based weather radars provide spatially continuous information about precipitation over larger regions. A radar transmits microwave pulses and measures the signal backscattered by atmospheric water particles, such as rain drops, snowflakes,

ice crystals and graupel. These particles are collectively referred to as hydrometeors. The measured radar reflectivity is related to the amount, size and phase of hydrometeors but it is not itself a direct measurement of surface precipitation rate. A conversion from radar reflectivity to precipitation rate is therefore required. This conversion depends on assumptions about the microphysical properties of the precipitation (Fabry, 2015).

Radar-based precipitation estimates can be affected by several factors. Attenuation occurs when the radar signal is weakened as it passes through precipitation. Beam blocking occurs when terrain partly or fully obstructs the radar beam. Bright-band effects occur near the melting layer where melting snowflakes can produce enhanced radar reflectivity. In addition, the radar beam samples higher altitudes with increasing distance from the radar. Radar-based precipitation products are therefore valuable reference datasets but they should still be interpreted as estimates rather than perfect observations of surface precipitation (Fabry, 2015).

2.1.2 Visible and infrared satellite observations

Visible and infrared satellite observations are widely used to monitor cloud systems. Visible sensors measure reflected solar radiation and are therefore mainly available during daylight. Infrared sensors mainly measure thermal radiation emitted by the Earth–atmosphere system and can be used both during day and night. From geostationary orbit, visible and infrared sensors can provide frequent observations over large regions which makes them useful for tracking the development and movement of cloud systems.

The limitation for precipitation retrieval is that visible and infrared observations mainly provide information about cloud-top properties such as cloud reflectance and cloud-top temperature. These properties are related to precipitation only indirectly. For example, very cold cloud tops can indicate tall clouds that may produce heavy precipitation but they do not uniquely determine whether precipitation reaches the surface or how large the surface precipitation rate is. Retrievals based on visible and infrared observations therefore rely on statistical relationships between cloud-top properties and precipitation.

2.1.3 Spaceborne precipitation radar

Spaceborne precipitation radars are active instruments. They transmit microwave pulses toward the atmosphere and measure the signal backscattered by hydrometeors. The main advantage of spaceborne radar is that it provides direct measurements of the vertical structure of precipitation by measuring how long emitted microwave pulses take to return to the sensor. This makes radar observations highly valuable for developing and evaluating satellite precipitation retrievals (Hou et al., 2014). However, spaceborne precipitation radars typically have narrower swaths compared to passive microwave radiometers, which limits their spatial coverage and temporal sampling (Hou et al., 2014).

2.1.4 Passive microwave satellite observations

Passive microwave radiometers measure naturally emitted radiation from the Earth-atmosphere system. In contrast to radar, they do not transmit a signal. Instead they observe microwave radiances that are affected by the surface, atmospheric gases, cloud liquid water and precipitation-sized hydrometeors.

Passive microwave observations are important for precipitation retrieval because microwave radiation can interact with hydrometeors inside the cloud, not only with the cloud top. Liquid precipitation and cloud water can increase the measured radiance through emission, especially over radiometrically cold ocean surfaces. Frozen hydrometeors can reduce the measured radiance by scattering radiation away from the sensor. These emission and scattering signals provide information about precipitation and precipitating cloud systems.

The relation between passive microwave observations and surface precipitation is nevertheless indirect. The measured microwave signal depends not only on precipitation rate but also on the vertical distribution of hydrometeors, surface emissivity, atmospheric temperature and humidity, viewing geometry and sensor frequency. As a result, different atmospheric scenes can produce similar measured radiances. Passive microwave precipitation retrieval is therefore an inverse problem (further defined in Section 2.3.1) with substantial uncertainty. Because of their complementary strengths, passive microwave radiometers and spaceborne radars are often combined to create synergistic precipitation products.

2.2 Passive microwave radiometry

Unless otherwise stated, the physical principles of passive microwave radiometry described in this section are based on Rees (2013) while precipitation signatures are based on Stephens and Kummerow (2007) and Kidd and Levizzani (2011).

2.2.1 Brightness temperature and antenna temperature

A black body refers to an idealized physical body that perfectly absorbs and emits all electromagnetic radiation. Brightness temperature, T_B , is the temperature that a black body would need to have in order to emit the same measured spectral radiance at the observed frequency. It is therefore a radiance-equivalent temperature and should not be interpreted as the physical temperature of the surface, a cloud or a single atmospheric layer.

The relation between black-body radiance and physical temperature is given by Planck’s law. Expressed per unit frequency, the spectral radiance of a black body is

$$B_\nu(T) = \frac{2h\nu^3}{c^2} \frac{1}{\exp\left(\frac{h\nu}{kT}\right) - 1}, \quad (2.1)$$

where h is Planck’s constant, k is Boltzmann’s constant, c is the speed of light and ν is frequency. Brightness temperature is defined as the temperature that satisfies

$$I_\nu = B_\nu(T_B), \quad (2.2)$$

where I_ν is the spectral radiance.

At microwave frequencies and terrestrial temperatures $h\nu$ is much smaller than kT . Planck’s law can then be approximated by the Rayleigh–Jeans relation:

$$B_\nu(T) \approx \frac{2k\nu^2}{c^2} T. \quad (2.3)$$

In this limit, radiance is approximately proportional to temperature. Brightness temperature is therefore a convenient way to express radiance measurements, especially in microwave remote sensing.

A passive microwave radiometer receives microwave radiation through an antenna over a larger area. The antenna does not receive radiation from one exact direction

only but has an angular response called the antenna pattern which describes how strongly radiation from different directions contributes to the measured signal. In this sense, brightness temperature is a property of the object being observed whereas antenna temperature, T_A , is the angularly weighted measure of the radiation actually received by the instrument.

2.2.2 Microwave channel sensitivity

The information contained in a passive microwave channel depends strongly on frequency because atmospheric gases, hydrometeors and the surface interact differently with radiation at different wavelengths. At frequencies where gaseous absorption is weak, radiation from the surface and lower atmosphere can reach the satellite more easily. These frequency ranges are often called atmospheric window regions. Window channels can therefore contain strong contributions from the surface, lower atmosphere and hydrometeors.

At frequencies close to molecular absorption lines, gases absorb and emit radiation efficiently. An absorption line is a frequency range where a molecule interacts strongly with radiation. Near such a line, radiation from lower atmospheric layers is partly absorbed before reaching the satellite. The measured radiance is therefore dominated by emission from the layers from which radiation can still escape to the sensor. Channels closer to the centre of an absorption line are generally sensitive to higher atmospheric layers while channels farther from the line centre receive larger contributions from lower layers.

This principle is used for atmospheric sounding. Oxygen absorption bands are used for temperature sounding because oxygen is well mixed in the atmosphere. Its concentration is therefore comparatively well known, so variations in measured radiance near oxygen absorption features are mainly controlled by the temperature of the emitting atmospheric layers. Water-vapour absorption bands are used for humidity sounding because water vapour varies strongly in both space and height. Measurements near water-vapour absorption lines are therefore sensitive to the amount and vertical distribution of humidity, as well as to atmospheric temperature.

The altitude dependence of a channel is described by a weighting function. A weighting function shows how strongly different atmospheric layers contribute to the measured radiance. A channel should therefore not be interpreted as observing a single altitude. Each channel receives contributions from a range of heights with the largest

contribution coming from the layers where the weighting function is largest. For AWS, the simulated temperature and humidity sensitivities illustrate this altitude-dependent behavior across its specific sounding channels (Eriksson et al., 2025).

2.2.3 Precipitation signatures

Precipitation affects passive microwave measurements through emission, absorption and scattering with their relative importance depending on the frequency. Liquid water absorbs and emits microwave radiation. Rain drops and cloud liquid water can therefore increase the upwelling microwave radiance, especially at lower microwave frequencies. This emission signal is often easiest to detect over ocean where the microwave emissivity of the surface is relatively low and spatially homogeneous. The precipitation signal is then observed against a comparatively cold radiometric background.

Over land, the emission signal from precipitation is harder to isolate. Land surfaces usually have higher and more variable microwave emissivity than ocean surfaces. Vegetation, soil moisture, snow cover, surface temperature and topography can all affect the measured radiance. These surface effects can mask or resemble the radiometric signal from precipitation, especially for weak precipitation.

At higher microwave frequencies, scattering by frozen hydrometeors becomes more important because the radiation wavelength approaches the size of the ice particles. Ice particles, snow and graupel can scatter upwelling radiation away from the sensor which often reduces the measured antenna temperature in channels sensitive to ice scattering. The strength of this signal depends on the amount of ice, the particle size distribution, particle shape and the vertical structure of the cloud.

Liquid-water emission and ice scattering provide complementary information. Emission is more closely related to cloud liquid water and rain water in the lower atmosphere while scattering provides information about frozen hydrometeors above the melting layer. Since many precipitating systems contain ice above the melting layer, scattering-sensitive channels can contain information related to precipitation. The relation to surface precipitation is indirect and similar microwave signatures can occur for different vertical cloud structures and surface precipitation rates.

2.2.4 Radiative transfer and retrieval

The relation between the atmospheric state and the measured microwave radiance is described by radiative transfer. Radiative transfer accounts for how radiation is emitted, absorbed and scattered along the path from the surface and atmosphere to the satellite. In a non-scattering atmosphere, the measured upwelling radiation can be understood as the sum of attenuated surface emission and thermal emission from atmospheric layers along the path.

Using the Rayleigh–Jeans approximation, a simplified non-scattering solution of the radiative transfer equation can be written in terms of the brightness temperature T_B as

$$T_B(\tau_s) = T_B(0)e^{-\tau_s} + \int_0^{\tau_s} T(\tau')e^{-(\tau_s-\tau')} d\tau', \quad (2.4)$$

where $T_B(0)$ is the brightness temperature of the radiation entering the atmosphere at the lower boundary, $T(\tau')$ is the atmospheric temperature as a function of optical thickness and τ_s is the total optical thickness between the lower boundary and the sensor. Optical thickness is a measure of how much a given medium attenuates light. The first term represents radiation from the lower boundary attenuated by the atmosphere. The integral represents thermal emission from atmospheric layers also attenuated before reaching the sensor.

This expression is only a simplified non-scattering case. In cloudy and precipitating atmospheres scattering by hydrometeors must also be considered. The measured radiance then depends on surface emissivity, atmospheric temperature and humidity, viewing geometry and the vertical distribution and microphysical properties of hydrometeors. Radiative transfer defines the forward problem: given the surface and atmospheric state, it provides the radiance that would be measured by the satellite.

2.3 Statistical retrieval theory

The statistical retrieval framework outlined in this section is primarily based on the inverse methods described by Rodgers (2000) and Pfreunds Schuh et al. (2018).

2.3.1 Forward and inverse problems

A retrieval problem aims to estimate an unknown geophysical quantity from indirect observations. In atmospheric remote sensing the observations are related to the

atmospheric and surface state through radiative transfer and the properties of the observing instrument. This relation can be written as a forward model,

$$\mathbf{o} = F(\mathbf{s}) + \boldsymbol{\epsilon}, \quad (2.5)$$

where \mathbf{s} is the atmospheric and surface state, \mathbf{o} is the observation vector, F is the forward model and $\boldsymbol{\epsilon}$ represents measurement and modelling errors.

The forward problem is to compute the observations that would be measured for a known state. The retrieval problem is the inverse problem: the observations are known and the aim is to estimate a property of the state. In the context of this study, this property is the surface precipitation rate denoted by r .

The inverse problem is ill-posed because similar observations can originate from different atmospheric and surface states. This is further complicated by inherent measurement and modelling errors. For passive microwave precipitation retrieval, this means that surface precipitation cannot be determined uniquely from the measured radiances. A retrieval should therefore be interpreted as an estimate based on the information available in the observations rather than as a direct reconstruction of the true precipitation field.

2.3.2 Bayesian retrieval formulation

Bayesian retrieval theory provides a framework for representing uncertainty in ill-posed inverse problems. Instead of assigning one exact value to the retrieval target the Bayesian formulation describes the posterior distribution of the target given the observation. For a scalar precipitation target r and an observation vector \mathbf{o} , this can be written as

$$p(r \mid \mathbf{o}). \quad (2.6)$$

This distribution describes the precipitation rates that are consistent with the observation and with the prior information used in the retrieval.

Bayes' theorem relates the posterior distribution to the prior distribution and the likelihood,

$$p(r \mid \mathbf{o}) = \frac{p(\mathbf{o} \mid r)p(r)}{p(\mathbf{o})}. \quad (2.7)$$

Here $p(r)$ is the prior distribution of precipitation, $p(\mathbf{o} \mid r)$ is the likelihood of observing \mathbf{o} for a given precipitation rate and $p(\mathbf{o})$ is a normalizing factor.

In the context of this study, the primary target is surface precipitation, so this scalar formulation is used as a simplification, even though the observations mathematically depend on the full atmospheric state. A deterministic estimate such as a posterior mean or median, can be derived from the posterior distribution, but a single value does not represent the full retrieval uncertainty.

2.3.3 Approaches to solving the inverse problem

Retrieval methods can be understood as different ways of approximating the posterior distribution or summaries of it. Simulation-based methods use a physical forward model to generate observations from atmospheric states. Examples include optimal estimation and Monte Carlo methods. These approaches are useful because they are directly connected to the physics of the measurement and can be applied even when real collocations are limited. These methods can be computationally expensive when the forward model is complex, especially for cloudy or precipitating atmospheres where scattering is important.

Database-based retrievals use a collection of atmospheric states and corresponding observations which may be simulated or observed. A retrieval model can use the database entries to infer the retrieved quantity for a new observation. This type of approach has been used in passive microwave precipitation retrieval, including the Goddard Profiling Algorithm (GPROF) framework (Pfreundschuh et al., 2022a).

Machine-learning methods are not separate from these categories. They can be used with simulated databases or with empirical collocations. In both cases the trained model provides a computationally efficient mapping from observations to retrieval outputs.

2.4 Machine learning and quantile regression neural networks

2.4.1 Supervised learning and neural networks

Supervised learning is a machine-learning setting in which a model is trained from examples of inputs and corresponding target values. The purpose is to learn a mapping that can be applied to new inputs. For a regression problem the target consists of one or more continuous variables and the model output is a corresponding numerical

prediction.

A neural network represents this mapping using a sequence of layers. Each layer applies a linear transformation followed by a nonlinear activation function. A hidden layer l can be written as

$$\mathbf{h}^{(l)} = \sigma(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}), \quad (2.8)$$

where $\mathbf{h}^{(l-1)}$ is the output from the previous layer, $\mathbf{W}^{(l)}$ is a matrix of weights, $\mathbf{b}^{(l)}$ is a vector of biases and σ is the activation function. The weights and biases are the trainable parameters of the model.

The activation function introduces nonlinearity. Without nonlinear activation functions several layers would still be equivalent to one linear transformation. Activation functions therefore allow neural networks to represent nonlinear relationships between inputs and outputs.

During training the weights and biases are adjusted to minimize a loss function. The loss function measures the discrepancy between the model output and the target values in the training data. The choice of loss function determines which statistical property of the target distribution the model is encouraged to estimate (Goodfellow et al., 2016).

2.4.2 Quantile regression

Quantile regression is a regression method for estimating conditional quantiles of a target variable. For a target variable Y , an input vector \mathbf{x} and a quantile level $\tau \in (0, 1)$, the conditional quantile $q_\tau(\mathbf{x})$ is defined by

$$P(Y \leq q_\tau(\mathbf{x}) \mid \mathbf{x}) = \tau. \quad (2.9)$$

For a given input \mathbf{x} , the quantile level specifies the probability that the target variable lies below the corresponding conditional quantile.

A conditional quantile can be estimated by minimizing the quantile loss. For a target value y , a predicted conditional quantile \hat{q}_τ and a quantile level τ , the loss is

$$L_\tau(y, \hat{q}_\tau) = \begin{cases} \tau(y - \hat{q}_\tau), & y \geq \hat{q}_\tau, \\ (1 - \tau)(\hat{q}_\tau - y), & y < \hat{q}_\tau. \end{cases} \quad (2.10)$$

The loss is asymmetric. For an upper quantile, observations above the predicted quantile are penalized more strongly than observations below it. For a lower quantile, the opposite is true. This asymmetry makes the minimizer of the expected loss equal to the corresponding conditional quantile (Koenker and Bassett, 1978; Koenker, 2005).

Several conditional quantiles can be estimated simultaneously by minimizing the average loss over a set of quantile levels $\{\tau_j\}_{j=1}^M$,

$$L(y, \hat{\mathbf{q}}) = \frac{1}{M} \sum_{j=1}^M L_{\tau_j}(y, \hat{q}_{\tau_j}). \quad (2.11)$$

When a neural network is trained with this loss it is referred to as a quantile regression neural network (QRNN). The network outputs a vector of predicted conditional quantiles,

$$f(\mathbf{x}) = [\hat{q}_{\tau_1}(\mathbf{x}), \dots, \hat{q}_{\tau_M}(\mathbf{x})]^\top. \quad (2.12)$$

These predicted quantiles provide a discrete approximation of the conditional distribution of the target variable given the input (Pfreundschuh et al., 2018).

3 Data

3.1 Arctic Weather Satellite observations

This section describes the Arctic Weather Satellite (AWS) observations used in this thesis. The focus is on the radiometer channel groups, the scan geometry and the main variables contained in the Level-1B observation product.

3.1.1 Radiometer and channel groups

AWS carries a cross-track microwave sounder with four receiver chains covering several microwave frequency regions. The instrument was designed as a compact passive microwave sounder with a target orbit altitude of approximately 600 km. This altitude was chosen as a compromise between antenna-size requirements, ground coverage, launch cost and end-of-life deorbiting considerations (Eriksson et al., 2025).

The first receiver group contains channels in the 50–57 GHz oxygen absorption band and is used for temperature sounding. The second group contains the 89 GHz window channel. The third group combines channels around 166 GHz and 183 GHz, including a window-like channel and channels near the 183 GHz water-vapour absorption line. The fourth group contains the sub-millimetre channels around 325 GHz (Eriksson et al., 2025). Table 3.1 summarizes the AWS channel groups. The table is intended as an overview; detailed channel specifications are given by Eriksson et al. (2025).

The receiver chains are spatially separated in the instrument. As a result the foot-

Table 3.1: Overview of the AWS radiometer channel groups. Footprint values are approximate nadir full width at half maximum (FWHM).

Channel group	Channels	Frequency region	Approx. nadir FWHM
AWS1X	AWS11-AWS18	50-57 GHz	< 40 km
AWS2X	AWS21	89 GHz	< 20 km
AWS3X	AWS31-AWS36	166-183 GHz	< 15 km
AWS4X	AWS41-AWS44	325 GHz	< 15 km

prints of the channel groups are not fully co-located in the original Level-1B data because the feed-horns view the main reflector from slightly different directions (Eriksson et al., 2025).

3.1.2 Scan geometry and spatial sampling

AWS uses a rotating mirror to scan across the satellite track. The mirror rotates at a frequency of 0.84 Hz, corresponding to 50.4 rotations per minute. This gives an along-track distance between successive footprints of approximately 8.2 km (Eriksson et al., 2025).

During each rotation the instrument records Earth-view samples as well as views of cold space and the onboard calibration target. The Earth-view part of the scan contains 145 samples distributed approximately between scan angles of $\pm 55^\circ$ resulting in a swath width exceeding 2000 km. The sampling time is 2.5 ms and this sampling is retained in the Level-1B data (Eriksson et al., 2025).

The cross-track geometry means that the viewing angle changes across the scan. Near nadir, the line of sight is closest to vertical. Toward the swath edges, the viewing angle increases, leading to a longer atmospheric path length and a larger projected footprint on the surface. This scan-position dependence is an important property of the observations and is represented in the Level-1B viewing-geometry variables.

Figure 3.1 shows an example AWS swath for one antenna temperature channel. The figure illustrates the cross-track sampling of the instrument and the spatial structure of the observations across an overpass. The difference between land and ocean is visible, illustrating how this channel acts as a window channel.

3.1.3 Level-1B observation product

The AWS Level-1B product contains calibrated antenna temperatures together with geolocation, time and viewing-geometry information. Antenna temperatures are provided for the radiometer channels listed in Table 3.1. Latitude, longitude and observation time describe the location and time of each field of view while the satellite zenith angle describes the viewing geometry across the scan (Eriksson et al., 2025). Table 3.2 summarizes the main Level-1B variables of interest to this study.

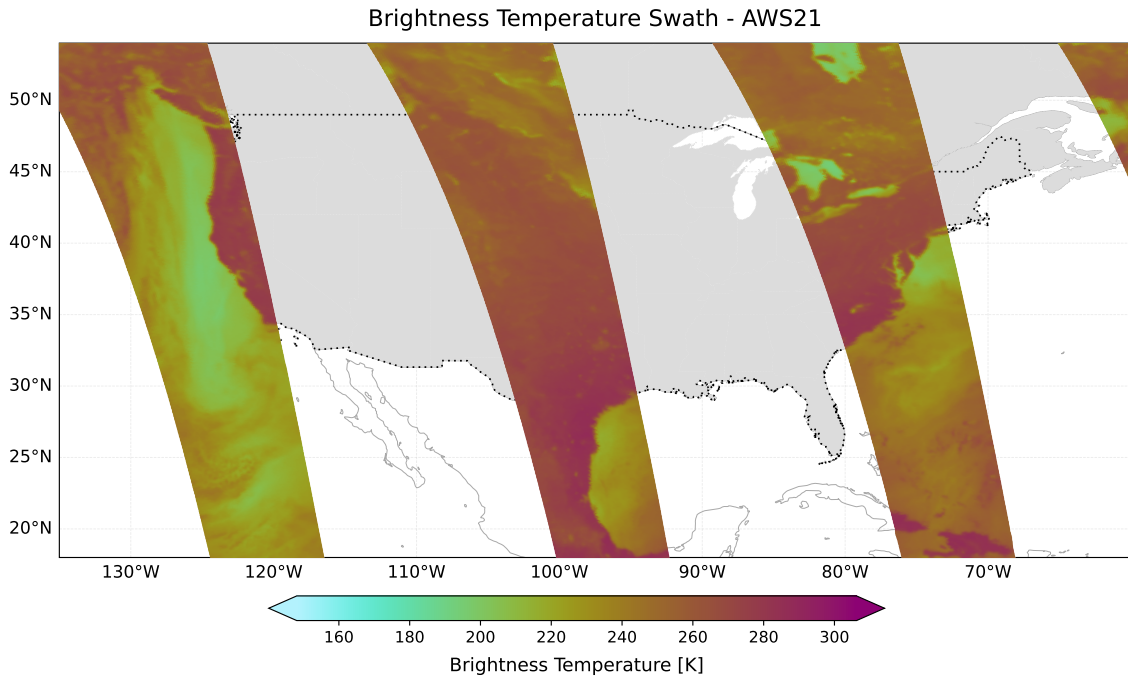


Figure 3.1: Example AWS observations for one antenna temperature channel over a 12-hour period. The dotted lines indicate the contiguous United States (CONUS).

3.2 Multi-Radar Multi-Sensor precipitation data

The reference precipitation data used is taken from the Multi-Radar Multi-Sensor (MRMS) system. MRMS is a radar-based and multi-sensor system that provides precipitation products using information from multiple weather radars together with additional observational and model data streams (Zhang et al., 2016). The MRMS product used is the Surface Precipitation Rate product which provides near-surface precipitation-rate estimates over the contiguous United States (NOAA, 2025).

The MRMS Surface Precipitation Rate product has a spatial resolution of approximately $1 \text{ km} \times 1 \text{ km}$ and a temporal resolution of 2 min. The precipitation rate is expressed in mm h^{-1} . Each MRMS field represents the estimated surface precipitation rate on a regular geographical grid at a specific valid time (NOAA, 2025). Table 3.3 summarizes the main information provided by the product.

Radar-based precipitation products contain uncertainties related to radar calibration,

Table 3.2: Main AWS Level-1B variables.

Variable	Description
Antenna temperatures	Calibrated radiometer observations for the AWS channels
Latitude and longitude	Geolocation of the fields of view
Observation time	Time of the AWS measurement
Satellite zenith angle	Viewing angle relative to the local vertical

Table 3.3: Main information from the MRMS Surface Precipitation Rate product.

Quantity	Description
Precipitation rate	Surface precipitation rate in mm h^{-1}
Latitude and longitude	Geographical position of the MRMS grid cells
Time	Time of the precipitation field

beam height, attenuation, beam blocking, bright-band effects and the conversion from radar reflectivity to precipitation rate. While no dataset provides perfect truth, ground-based radar is widely considered the most reliable option which justifies its use as the reference target. These uncertainties are relevant when MRMS is used as a reference dataset since differences between retrieved precipitation and MRMS can originate from uncertainties in either dataset (Zhang et al., 2016; Fabry, 2015).

3.3 Additional comparison dataset

The Integrated Multi-satellitE Retrievals for GPM (IMERG; Huffman et al. 2020) product is a global satellite precipitation product developed within the GPM mission. IMERG combines precipitation estimates from passive microwave sensors in the GPM constellation with infrared observations and rain gauge information to produce gridded precipitation estimates over large spatial domains.

The product is designed to improve the temporal sampling of satellite precipitation estimates by combining observations from multiple satellites. Passive microwave observations provide precipitation-sensitive information while infrared observations from geostationary satellites help fill temporal gaps between microwave overpasses. The Late Run IMERG product is available at 0.1° spatial resolution and 30 min

temporal resolution (Huffman et al., 2020). The Late Run version is used because the Final Run product was not available for the full evaluation period at the time of analysis.

4 Method

This chapter describes how the AWS and MRMS data are combined and used to train machine learning models for precipitation retrieval. The purpose of the method is to construct a supervised learning dataset where AWS observations are used as input and MRMS surface precipitation rates are used as reference targets.

The workflow consists of three main parts. First, AWS observations are remapped, filtered and collocated with MRMS precipitation fields, as presented in Section 4.1. Second, the collocated data are transformed into input features and target values suitable for neural-network training, as described in Section 4.2. Third, Quantile Regression Neural Networks are trained and evaluated using deterministic metrics, probabilistic diagnostics and qualitative visualizations, which is outlined in Sections 4.3 to 4.6. The final target variable in all experiments is the MRMS Surface Precipitation Rate.

4.1 Construction of the AWS–MRMS dataset

4.1.1 Study period and domain

The dataset is constructed from AWS observations and MRMS precipitation fields covering the period from 1 January 2025 to 31 December 2025. It should be noted that observations from the first three months of the year are expected to be of lower quality. Only AWS observations overlapping the MRMS domain are considered. Since MRMS is available over the contiguous United States the resulting dataset is limited to this geographical region.

4.1.2 AWS remapping and scan-position selection

The AWS channel groups are not perfectly co-located in the original Level-1B data, as described in Chapter 3. To construct one common input vector for each field of view, the channel groups are remapped to a common field-of-view geometry before collocation with MRMS. The AWS3X group, containing the channels around 166–183 GHz, is used as the reference geometry. The other AWS channel groups are

remapped to the AWS3X field-of-view grid.

After remapping the selected antenna temperature channels share a common set of field-of-view locations. The latitude and longitude coordinates of the AWS3X grid are then used for the spatial collocation with MRMS.

Of the 145 total cross-track measurements, only the 80 scan positions closest to nadir are retained. The outer scan positions are excluded because the viewing angle, projected footprint size and surface projection differ most strongly near the swath edges. This reduces the influence of extreme scan-angle effects while retaining the central part of the cross-track scan.

4.1.3 Temporal and spatial matching with MRMS

MRMS Surface Precipitation Rate fields are available every 2 min. For each AWS observation the MRMS field closest in time is selected. The maximum time difference to the nearest MRMS field is therefore approximately 1 min. This remaining time mismatch is a source of uncertainty, especially for rapidly evolving or fast-moving precipitation systems.

The MRMS product has a native grid spacing of approximately 1 km which is finer than the AWS footprint size. The AWS3X footprint is approximately 10 km at nadir and increases away from nadir. To reduce the scale mismatch between MRMS and AWS, the MRMS precipitation field is spatially smoothed using a 5×5 grid-cell moving average window corresponding to an area of approximately 5×5 km while retaining the native 1 km grid spacing.

After averaging each AWS field of view is matched to the nearest MRMS grid point in the averaged precipitation field. The resulting MRMS value is used as the reference surface precipitation rate for the corresponding AWS observation. This does not explicitly model the exact spatial weighting of the AWS antenna footprint but it reduces the mismatch between the native MRMS grid scale and the satellite observation scale.

4.1.4 Quality control and filtering

Samples containing missing or invalid values are removed before training. Invalid measurements include not-a-number entries, infinite values or unreadable scans. This

includes samples with missing AWS antenna temperatures, missing viewing-geometry information or missing MRMS precipitation values.

The same filtering rules are applied before constructing the training, validation and test datasets. After collocation and filtering the dataset contains approximately 4.18×10^7 samples.

4.2 Feature and target construction

4.2.1 Input variables

Each remaining sample represents a single collocated point. These individual samples contain a common AWS field of view, the selected AWS antenna temperatures, the corresponding viewing-geometry information and one matched MRMS precipitation-rate value. For each collocated sample, the input variables consist of AWS antenna temperatures and satellite zenith-angle information. The antenna temperature inputs depend on the model configuration. The full-channel model uses all 19 AWS antenna temperature channels. The reduced-channel model excludes the four sub-millimetre channels around 325 GHz.

Satellite zenith angle is included to represent the viewing geometry across the scan. One zenith-angle variable is retained for each AWS channel group. The two input configurations are therefore

$$\mathbf{x}_{\text{full}} = (T_{A,1}, \dots, T_{A,19}, \theta_1, \dots, \theta_4), \quad (4.1)$$

and

$$\mathbf{x}_{\text{reduced}} = (T_{A,1}, \dots, T_{A,15}, \theta_1, \dots, \theta_4), \quad (4.2)$$

where $T_{A,i}$ denotes an AWS antenna temperature channel and θ_j denotes the satellite zenith angle for AWS channel group j . The full-channel model has 23 input features while the model without the sub-millimetre channels has 19 input features. Figure 4.1 shows an example of the AWS antenna temperature input fields for selected channels.

Latitude, longitude and observation time are used for collocation but are not used as model input features in the main retrieval models. This choice focuses the retrieval on radiometric and viewing-geometry information rather than on explicit location-specific or time-specific climatological information.

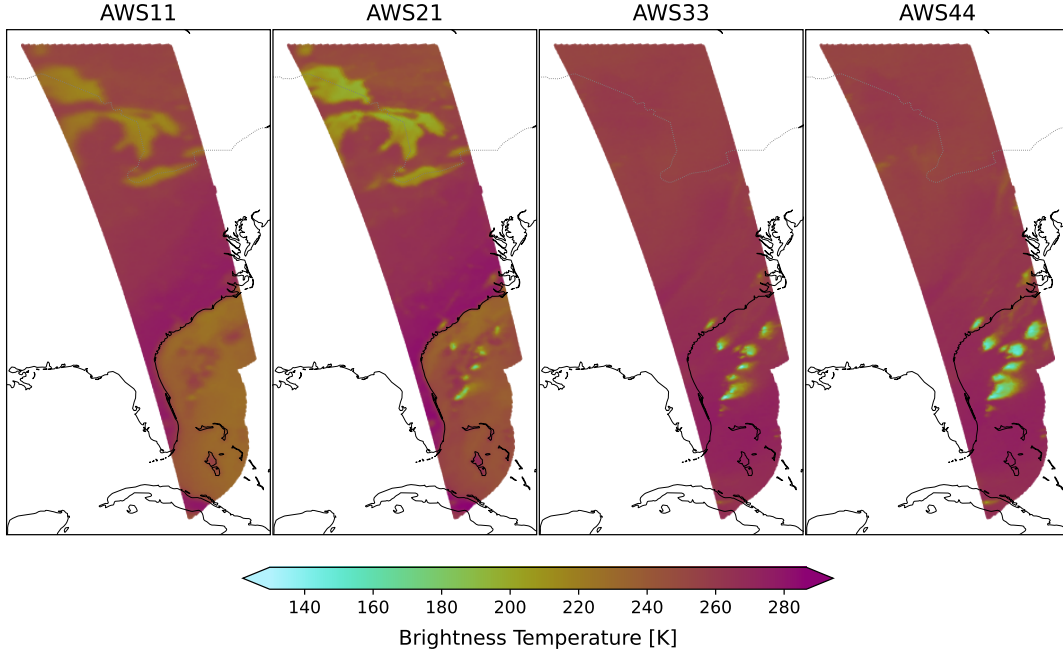


Figure 4.1: Example AWS antenna temperature swaths of inputs for selected channels.

4.2.2 Target variable

The target variable is the MRMS Surface Precipitation Rate matched to each AWS observation. It is denoted by

$$y = R_{\text{MRMS}}, \quad (4.3)$$

where R_{MRMS} is the MRMS surface precipitation rate in mm h^{-1} .

The target variable is non-negative and strongly skewed with many zero or near-zero precipitation samples and comparatively few high-intensity precipitation samples. This imbalance affects both model training and the interpretation of evaluation metrics.

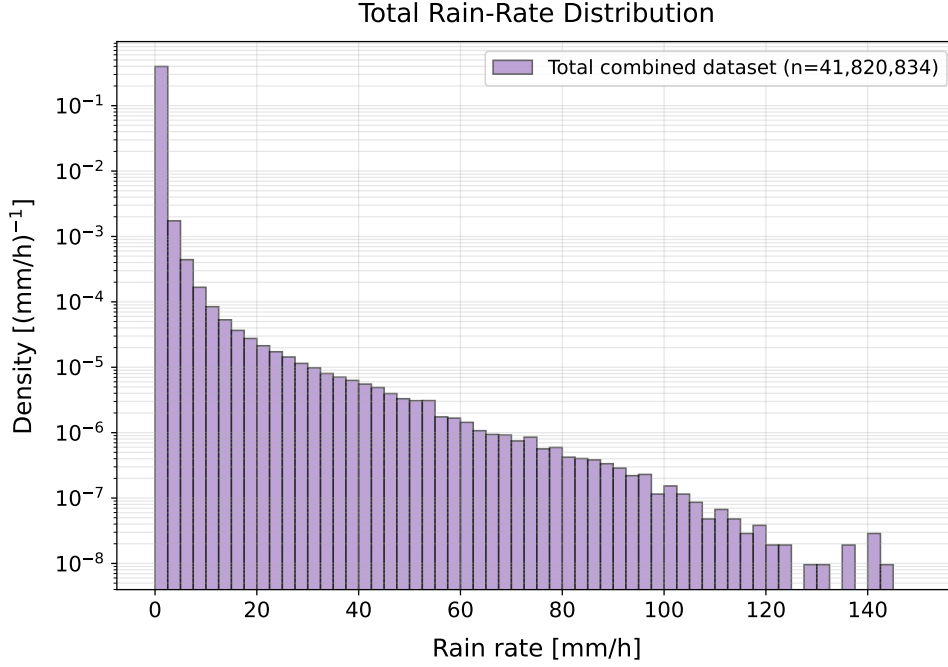


Figure 4.2: Distribution of MRMS surface precipitation rates in the full collocated AWS–MRMS dataset.

4.2.3 Log transformation of the target

The MRMS rain-rate distribution contains many zero values and has a large dynamic range. To reduce this dynamic range the QRNNs are trained in a log-transformed target space. Let y denote the MRMS surface precipitation rate in mm h^{-1} . The transformed target $z = \phi(y)$ is defined as

$$z = \phi(y) = \begin{cases} \log y, & y > 0, \\ u, & y = 0, \end{cases} \quad (4.4)$$

where

$$u \sim U(\log 10^{-5}, \log 10^{-4}). \quad (4.5)$$

Non-zero precipitation values are transformed using the natural logarithm. Zero precipitation values are assigned a random value in log space corresponding to a rain rate between 10^{-5} and 10^{-4} mm h^{-1} . This avoids undefined logarithms for

zero precipitation while keeping no-rain samples close to zero rain rate after back-transformation.

Predicted quantiles are transformed back to rain-rate units using the inverse transformation

$$\hat{y}_\tau = \phi^{-1}(\hat{z}_\tau) = \begin{cases} \exp(\hat{z}_\tau), & \hat{z}_\tau > \log 10^{-4}, \\ 0, & \hat{z}_\tau \leq \log 10^{-4}. \end{cases} \quad (4.6)$$

This thresholding maps very small predicted rain rates to zero precipitation.

4.2.4 Feature normalization

The input features have different numerical ranges. Antenna temperatures are given in Kelvin while zenith angles are angular variables. To improve the stability of neural-network training each input feature is standardized. Each input feature is standardized using statistics computed from the training set:

$$x_i^* = \frac{x_i - \mu_i}{\sigma_i}, \quad (4.7)$$

where μ_i and σ_i are the training-set mean and standard deviation of feature i . The same normalization parameters are applied to the validation and test sets.

4.3 Train, validation and test split

The collocated dataset is split by calendar day. A day-based split is used instead of a random sample split to reduce temporal leakage between training and evaluation datasets. If nearby observations from the same weather system were randomly split, very similar samples could appear in both training and evaluation sets.

For each month observations from the 2nd, 10th, 18th and 26th day are used for validation. Observations from the 11th and 21st day are used for testing. All remaining days are used for training. This split keeps complete days separated between the datasets while sampling validation and test cases throughout the year. The validation set is used for model selection and monitoring during training. The test set is kept separate and used for final evaluation. Figure 4.3 shows the MRMS rain-rate distributions for the training, validation and test datasets. Table 4.1 summarizes the total number of samples in each dataset.

Table 4.1: Total number of samples in the training, validation and test datasets.

Dataset	Number of samples
Training	33,470,572
Validation	5,483,584
Test	2,866,678




Figure 4.3: MRMS rain-rate distributions for the training, validation and test datasets.

4.4 Quantile regression neural network

The QRNN predicts selected conditional quantiles of the log-transformed MRMS rain-rate target. The quantile levels are

$$\tau = [0.01, 0.05, 0.15, 0.25, 0.35, 0.45, 0.50, 0.55, 0.65, 0.75, 0.85, 0.95, 0.99]. \quad (4.8)$$

For an input vector \mathbf{x} , the network output is

$$f(\mathbf{x}) = [\hat{z}_{\tau_1}(\mathbf{x}), \dots, \hat{z}_{\tau_Q}(\mathbf{x})]^\top, \quad (4.9)$$

where $\hat{z}_{\tau_j}(\mathbf{x})$ is the predicted conditional quantile of the transformed target z at quantile level τ_j . The output dimension is therefore $Q = 13$.

The QRNN models are fully connected feed-forward neural networks. The architecture used for the main experiments consists of four hidden layers with 256 neurons in each layer. The rectified linear unit (ReLU) activation function is applied after each hidden layer:

$$\text{ReLU}(x) = \max(0, x). \quad (4.10)$$

The quantile loss (Equation 2.10) is applied to the log-transformed target z . During training the loss is averaged over the predicted quantile levels and over the samples in each mini-batch as in equation (2.11). Both models use the same target transformation, normalization strategy, network architecture, loss function and evaluation procedure. The comparison therefore isolates the effect of including or excluding the AWS sub-millimetre channels.

4.5 Training procedure

Training is performed using mini-batch gradient-based optimization with a batch size of 1024. The model parameters are optimized using AdamW (Loshchilov and Hutter, 2019) with an initial learning rate of 10^{-4} . A cosine learning-rate scheduler is used during training to gradually reduce the learning rate over the training epoch.

The models are trained for a maximum of 20 epochs. After each epoch the model is evaluated on the validation set. The model state with the lowest validation loss is saved and used for later evaluation. Training is performed on a GPU because GPU acceleration makes mini-batch optimization substantially more efficient.

4.6 Evaluation

4.6.1 Prediction summary

For deterministic metrics a posterior-mean estimate is derived from the predicted quantiles. Let $\hat{Q}(\tau_j)$ denote the predicted rain-rate quantile at quantile level τ_j after back-transformation to mm h^{-1} . The posterior mean is approximated by numerical

integration of the predicted quantile function:

$$\hat{y}_{\text{mean}} \approx \frac{1}{\tau_Q - \tau_1} \sum_{j=1}^{Q-1} \frac{\hat{Q}(\tau_j) + \hat{Q}(\tau_{j+1})}{2} (\tau_{j+1} - \tau_j). \quad (4.11)$$

This approximation uses the available predicted quantile range from $\tau_1 = 0.01$ to $\tau_Q = 0.99$.

4.6.2 Deterministic metrics

Deterministic performance is evaluated by comparing the point estimate \hat{y}_i with the MRMS reference value y_i . The mean squared error is

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2. \quad (4.12)$$

The root mean squared error is

$$\text{RMSE} = \sqrt{\text{MSE}}. \quad (4.13)$$

The mean absolute error is

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i|. \quad (4.14)$$

The bias is

$$\text{Bias} = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i). \quad (4.15)$$

The Pearson correlation coefficient is

$$r = \frac{\sum_{i=1}^N (\hat{y}_i - \bar{\hat{y}})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (\hat{y}_i - \bar{\hat{y}})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}}. \quad (4.16)$$

4.6.3 Probabilistic diagnostics

The probabilistic predictions are evaluated using the continuous ranked probability score (CRPS). For a predictive distribution F and an observation y , CRPS can be written in terms of the quantile function $Q(\tau)$ as

$$\text{CRPS}(F, y) = 2 \int_0^1 L_\tau(y, Q(\tau)) d\tau, \quad (4.17)$$

where L_τ is the quantile loss defined in equation (2.10). Since the QRNN predicts a finite set of quantiles this integral is approximated numerically using the predicted quantile levels. The mean CRPS across all samples is reported as the final bulk metric.

Calibration is evaluated by comparing predicted quantile levels with observed frequencies. For a well-calibrated quantile prediction approximately a fraction τ of the observations should fall below the predicted τ -quantile. Reliability diagrams are used to visualize this relationship.

5 Results

This chapter presents the training behaviour and retrieval results. The final retrieval performance is evaluated on the independent test set. Unless otherwise stated, all metrics and plots are computed after transforming the predicted quantiles back to physical rain-rate units. The posterior mean is used as the deterministic precipitation estimate. The all-channel model is treated as the main retrieval model while the reduced-channel model is included in the figures and tables to support the assessment of the sub-millimetre channels in Section 5.4.

5.1 Training behaviour

The learning curves of the models provide an initial diagnostic of training stability. As shown in Figure 5.1, the validation curves decline rapidly during the first few epochs. The minimum validation loss was achieved at epoch 9 for the all-channel model and epoch 11 for the reduced-channel model. Beyond these epochs, while the training losses continue to decrease, the validation curves exhibit a slight upward trend. This behavior indicates mild overfitting which justifies selecting these specific earlier checkpoints. The models are stopped at these epochs and the resulting network weights are used for all subsequent evaluation on the test set.

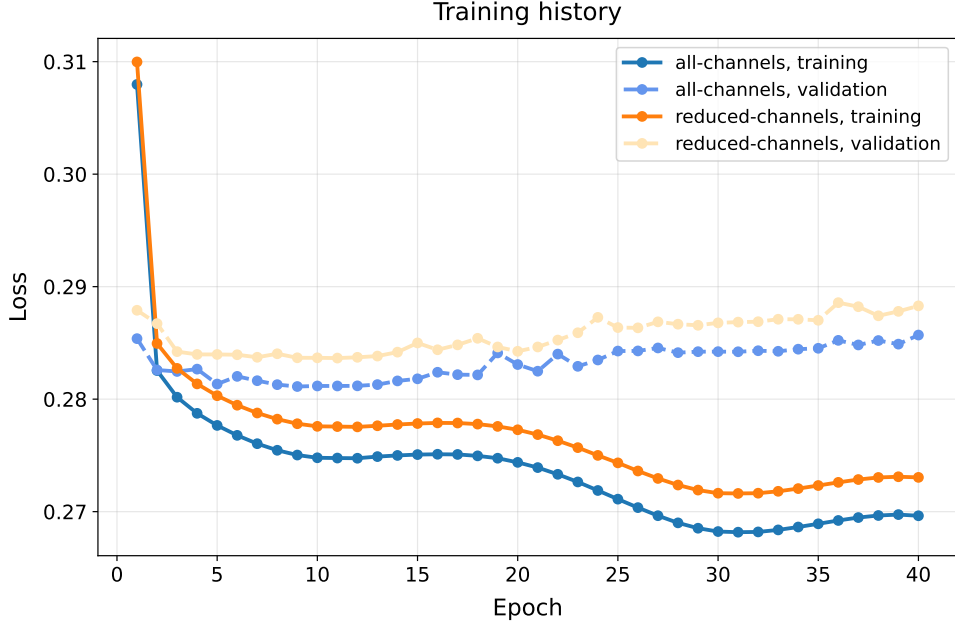


Figure 5.1: Training and validation loss trajectories computed in log-transformed target space for the all-channel and reduced-channel QRNN models.

The calibration of the predicted probability intervals is evaluated using the reliability diagrams in Figure 5.2. Reliability diagrams are computed by calculating the empirical frequency of observations that fall below each predicted quantile level. The dashed yellow line represents the reliability of the training dataset computed in the log-transformed training space. It tracks close to the perfect-calibration diagonal, verifying that the model is well calibrated across the overall dataset. When back-transformed to physical units and separated into individual bins the reliability curves dip below the perfect calibration line. This indicates that the model’s predicted quantiles are systematically lower than the actual observed rain rates. The deviation in the lowest rain-rate bin is a mathematical artifact of the zero-bound threshold, as true no-rain observations will inherently fall below or equal to any non-negative predicted quantile.

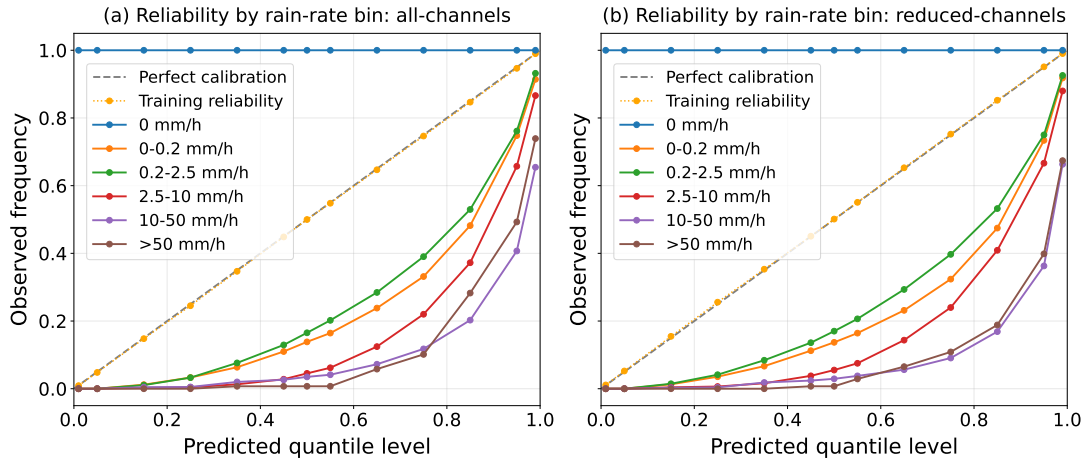


Figure 5.2: Probabilistic reliability diagrams across specific MRMS reference precipitation bins computed on the test set. The diagonal line represents perfect calibration while the yellow dashed line, which overlaps the diagonal line, represents the reliability for the training dataset in the log-transformed training space.

5.2 Performance statistics on the test set

Using the predicted quantiles we can approximate the continuous cumulative distribution function (CDF) and from that CDF we can derive the probability density function (PDF). Figure 5.3 displays this for one individual test sample. In this instance, the calculated posterior mean tracks near the true reference value while the spread of the predictive distribution expresses the physical uncertainty associated with the antenna temperature inputs.

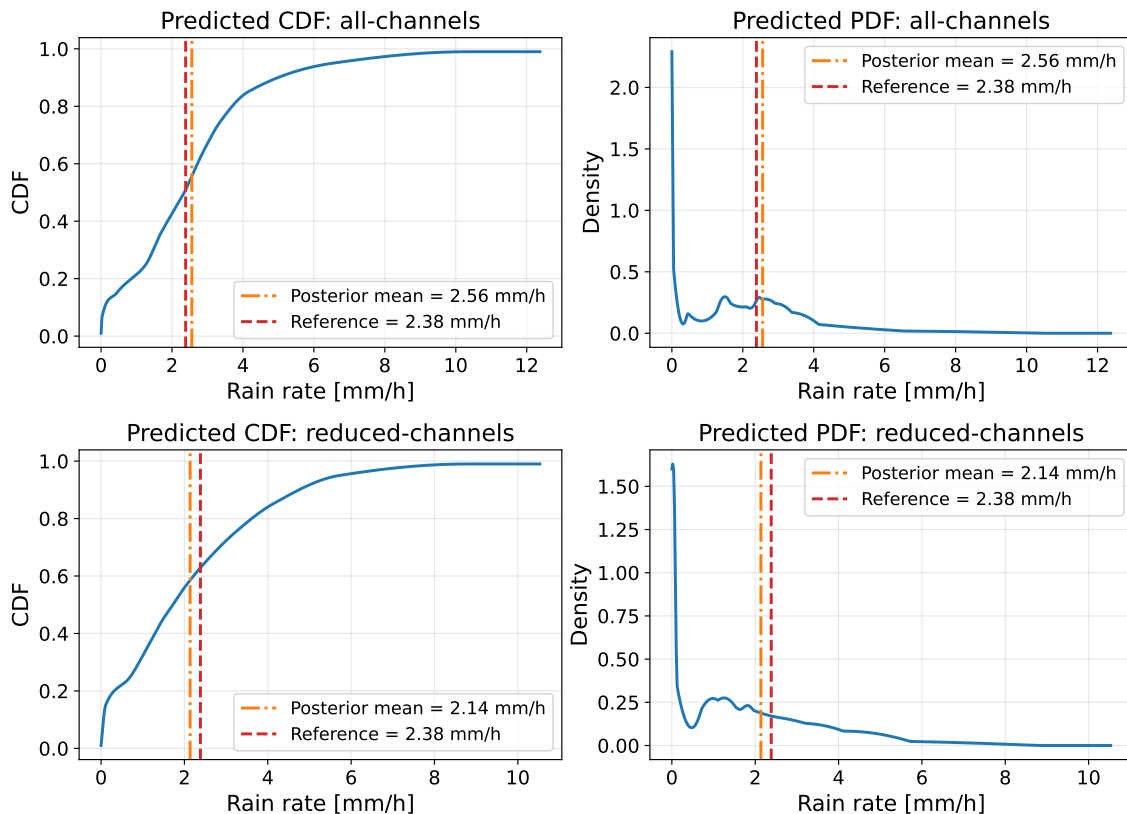


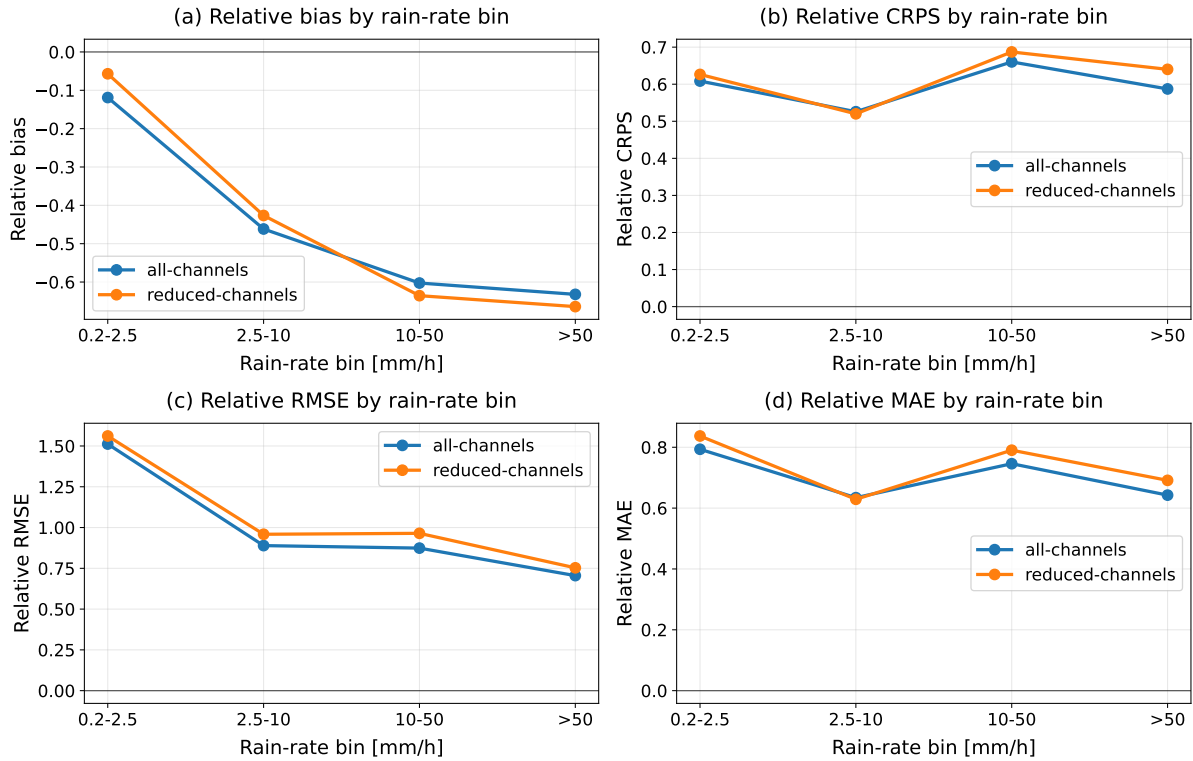
Figure 5.3: Example probabilistic prediction for an individual test sample. The predicted discrete quantiles are smoothed via interpolation to approximate the continuous cumulative distribution function (left) and then derive the probability density function (right).

Table 5.1 summarizes the statistical performance across the entire test set. The baseline all-channel model exhibits a global absolute bias close to zero (-0.002 mm h^{-1}) indicating that positive and negative errors largely cancel out over the complete domain. The MAE is low (0.067 mm h^{-1}) reflecting a small average absolute error that is heavily dominated by the highly frequent background no-rain conditions. The RMSE is comparatively larger (0.655 mm h^{-1}) reflecting its mathematical sensitivity to localized high-intensity errors. A Pearson correlation (Corr.) of 0.587 shows that the model captures broad physical trends but a portion of the precipitation variance remains uncaptured by the posterior mean prediction.

Table 5.1: Test-set verification metrics using posterior mean predictions.

Model	RMSE [mm h ⁻¹]	MAE [mm h ⁻¹]	Bias [mm h ⁻¹]	Corr.	CRPS [mm h ⁻¹]
All-channel	0.655	0.067	-0.002	0.587	0.040
Reduced-channel	0.708	0.071	0.002	0.525	0.041

To separate the dominance of no-rain and low-rain samples from the active precipitating regimes, Figure 5.4 presents relative performance metrics across discrete MRMS reference rain-rate intervals. The lowest rain-rate interval is excluded from these subplots due to its large relative scaling effects, a mathematical consequence where denominators near zero amplify small absolute differences. In the higher bins the relative bias trends systematically downward, falling below -0.6 for the highest categories. This structural decline indicates an underestimation of higher precipitation intensities. The relative RMSE generally decreases across the higher rain-rate bins. The relative CRPS and MAE reach their lowest values in the moderate 2.5–10 mm h⁻¹ interval, rise in the 10–50 mm h⁻¹ bin and then decrease again for the most extreme precipitation.



First bin excluded from plots: 0-0.2 mm/h
 all-channels: bias=17.698, CRPS=5.172, RMSE=136.189, MAE=18.397, n=2765455
 reduced-channels: bias=19.821, CRPS=5.664, RMSE=144.757, MAE=20.531, n=2765455

Figure 5.4: Relative binned test-set metrics using posterior mean predictions normalized by the mean MRMS reference rain rate of each bin.

Evaluating how closely the models mirror the empirical population demographics provides an indicator of retrieval consistency. Figure 5.5 compares the total continuous rain-rate density of the true MRMS test reference against the predictions on the same test dataset. The model effectively captures the high-density, low rain-rate core of the distribution. However, the predicted density curves fall below the test reference above 10 mm h^{-1} , producing fewer high-intensity precipitation estimates than are present in the MRMS reference. The physical and statistical reasons for this underestimation are explored further in Section 6.2.

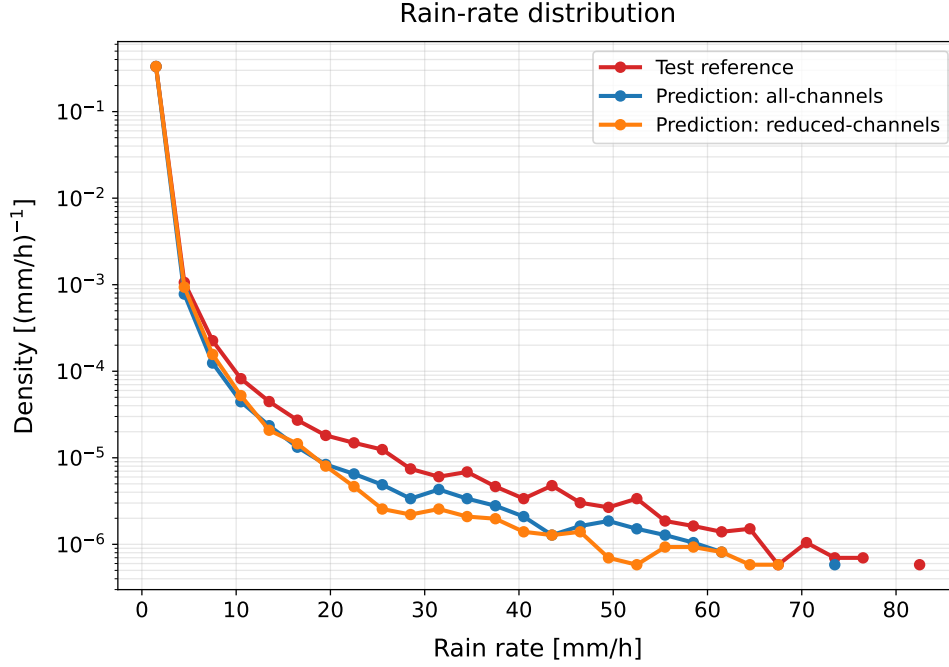


Figure 5.5: Observed and predicted rain-rate density distributions evaluated across the complete test dataset.

Figure 5.6 expands this evaluation by plotting the conditional density of the posterior mean predictions against the continuous MRMS reference field. A high concentration of points is visible at very low rain rates. In this region below 1.0 mm h^{-1} , the conditional mean changes weakly and forms a plateau. Through the intermediate regimes between 1.0 and 10 mm h^{-1} , the conditional mean tracks the identity line closely. However, for higher rain rates above 10 mm h^{-1} , the highest density of predicted values flattens out, causing the conditional mean to diverge from the ideal identity line. This behaviour is further discussed in Sections 6.1 and 6.2.

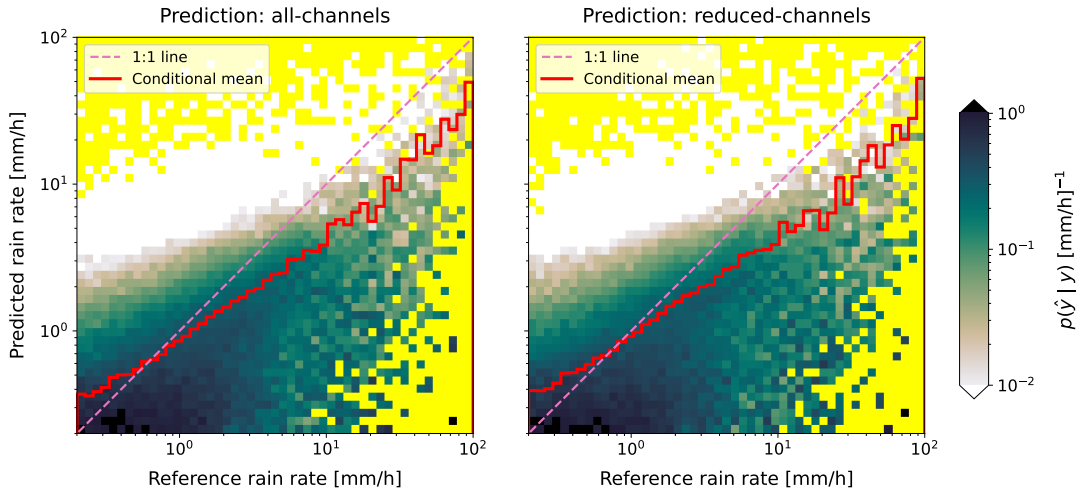


Figure 5.6: Conditional density plots of posterior mean predictions against MRMS reference values. The solid red line marks the conditional mean. The dashed line represents the ideal 1:1 identity line. The yellow regions define areas where the sample concentration drops to zero.

5.3 Comparison with MRMS reference swaths

Selected AWS overpasses are used to evaluate the spatial structure of the retrievals. The three examples represent different precipitation regimes: a heavy-rain case, a moderate-rain case, and a weak-rain case. These qualitative diagnostics examine how the predictions compare with the collocated MRMS reference and the operational gridded Late Run IMERG product. IMERG provides an independent comparison source which is particularly useful for assessing scenes where the retrieval disagrees with MRMS.

Figure 5.7 displays a heavy-rain case characterized by an active precipitation front. The model reproduces the general precipitation pattern but exhibits distinct regional differences. In the upper part of the swath the model generally underpredicts the rain rate as demonstrated in Figure 5.7(e). This corresponds to the antenna temperature inputs in Figure 5.7(h)–(k) where clear radiometric precipitation signatures are largely missing. The lower part of the swath exhibits a strong and clearly visible signal in the antenna temperatures. In this active southern region, the difference map in Figure 5.7(e) reveals a mixture of overestimation and underestimation. Additionally,

the MRMS target and the IMERG product differ from each other, especially in the upper portion of the swath. The MRMS target and the IMERG product differ from each other in the upper portion of the swath where IMERG seems to underestimate the spatial extent of the system. The AWS retrieval aligns more closely with the broader structure of the MRMS radar reference, demonstrating an advantage over IMERG in capturing the full scale of the precipitation.

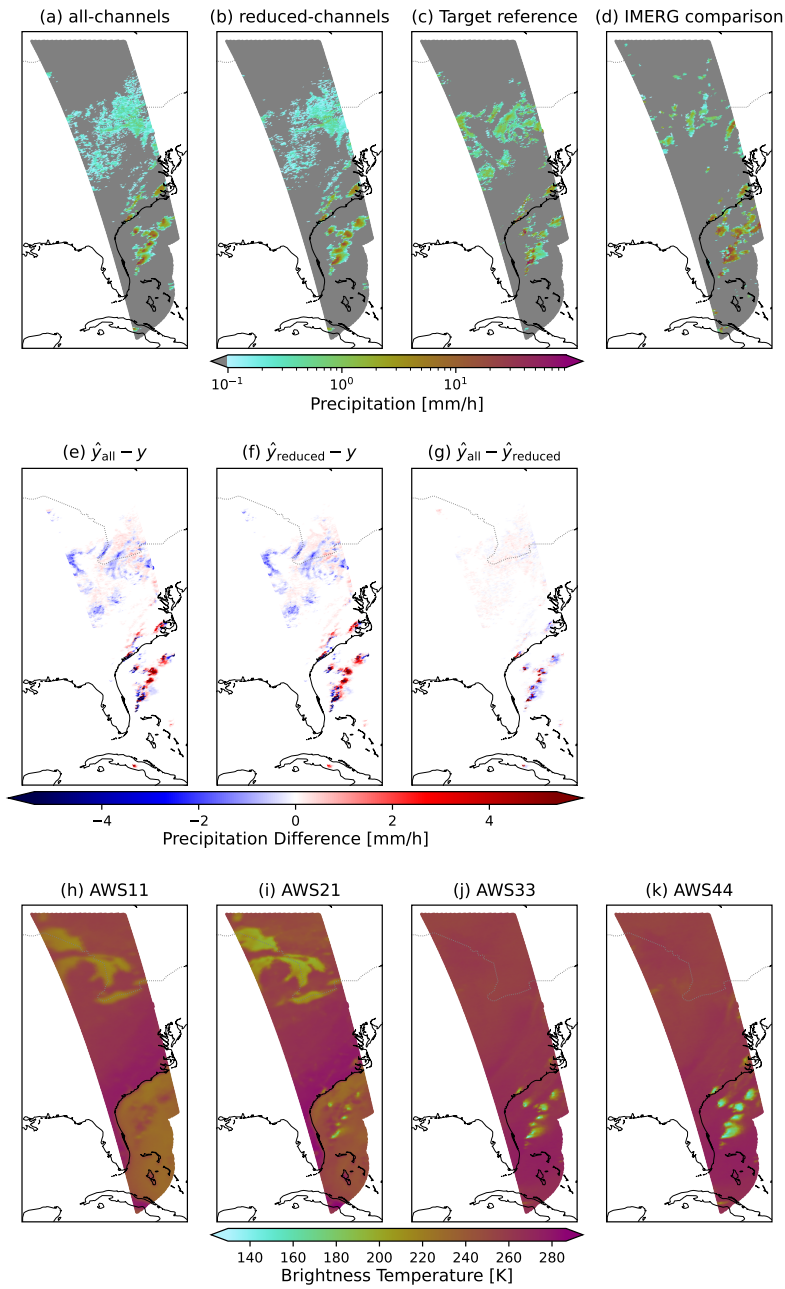


Figure 5.7: Swath comparison for the heavy-rain case. The top row presents absolute predictions, the MRMS reference target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.

Figure 5.8 displays a moderate-rain case captured over a structured precipitating system. The retrieval captures the broad spatial pattern of the precipitation but produces a smoother, more rounded field compared to the MRMS reference. This structure closely mirrors the patterns observed in the antenna temperature inputs with a particularly clear signal visible in the sub-millimetre channel shown in Figure 5.8(k). When examining the difference map in Figure 5.8(e), a mixture of both overestimation and underestimation is present across the precipitating system. Furthermore, the MRMS target and the IMERG product exhibit quite similar spatial patterns in this scene although the IMERG field appears somewhat noisier and coarser due to its lower spatial resolution.

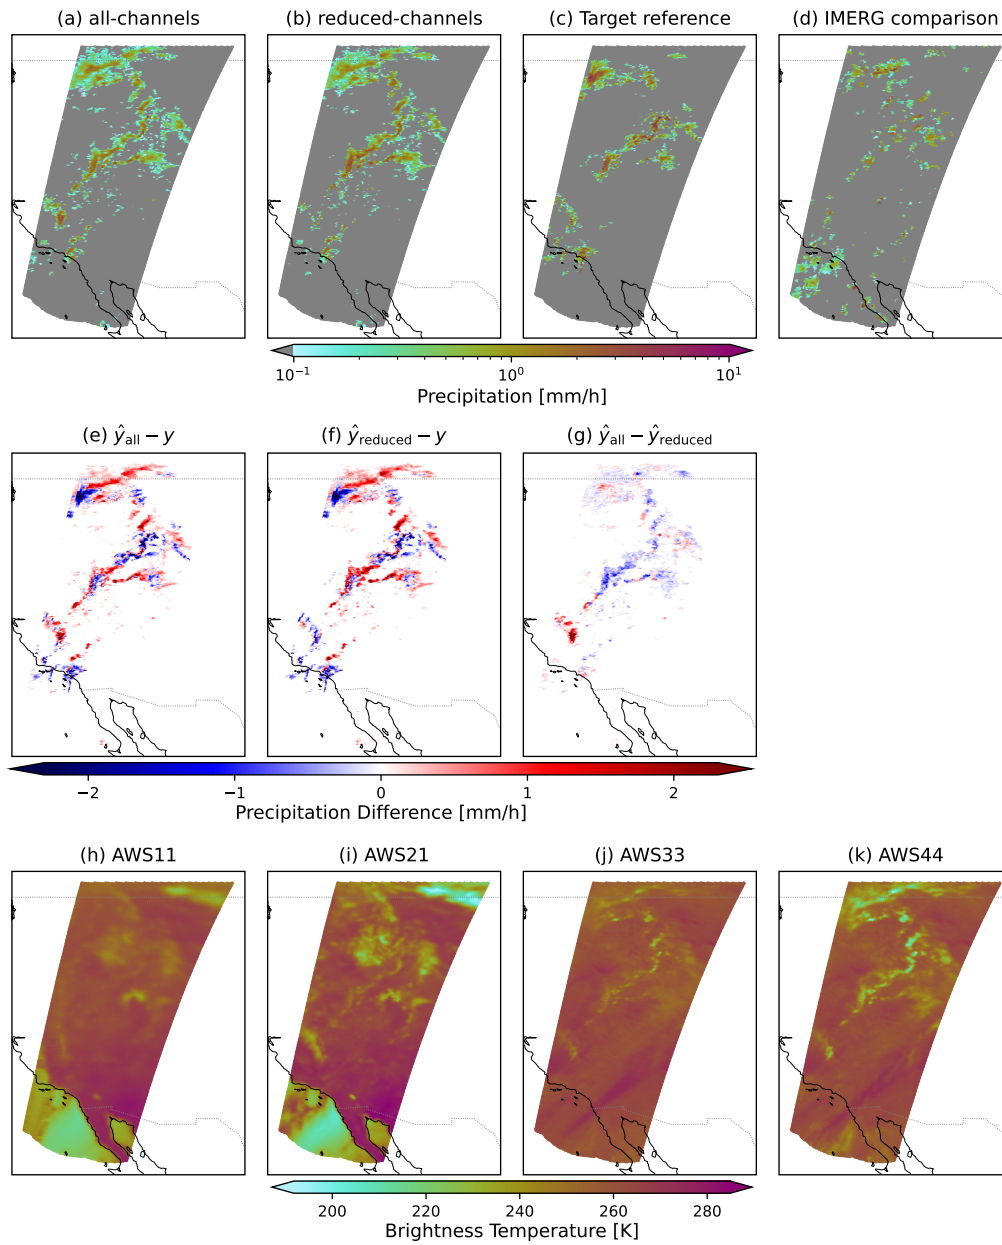


Figure 5.8: Swath comparison for the moderate-rain case. The top row presents absolute predictions, the MRMS target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.

Figure 5.9 displays a weak-rain case where the vast majority of the scene contains no precipitation. The model correctly identifies the general location of the light rain in the upper part of the swath but spreads it over a wider area than the MRMS target. As shown in Figure 5.9(h)–(k) there is no clearly visible radiometric signal in the antenna temperatures. This lack of a distinct signature is expected for such light precipitation. Furthermore, it is notable that the IMERG comparison product does not detect rain in this scene. In this specific instance, the AWS retrieval outperforms the operational IMERG product by successfully identifying the presence of light precipitation that IMERG misses. This illustrates both the capability of the QRNN to extract subtle radiometric signals and the inherent difficulty of evaluating weak precipitation when operational products disagree.

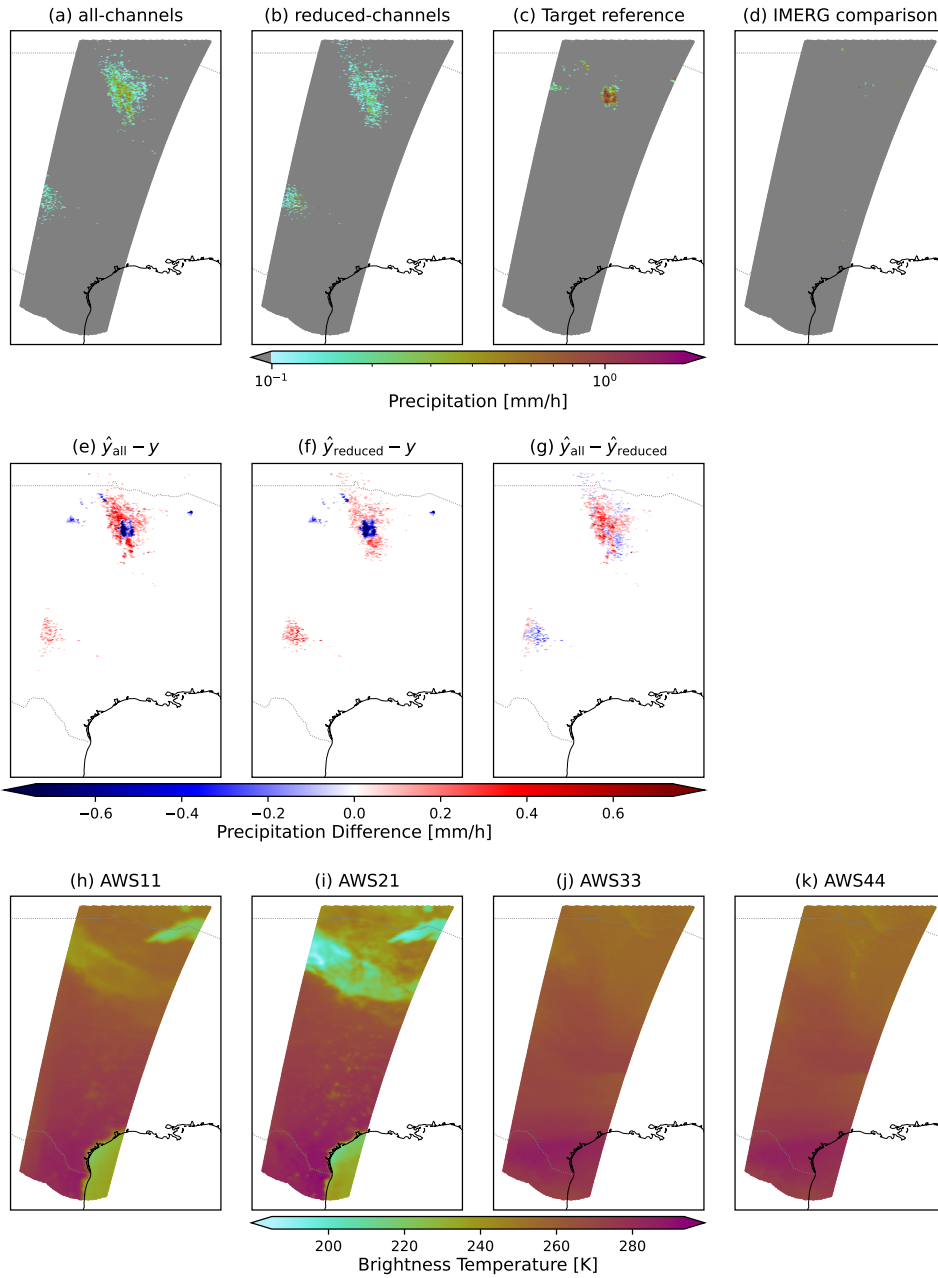


Figure 5.9: Swath comparison for the weak-rain case. The top row presents absolute predictions, the MRMS target and the Late Run IMERG field. The middle row displays absolute differences, and the bottom row displays the AWS antenna temperature inputs.

Overall the swath examples show that the retrievals capture the general spatial structure of the MRMS reference but with smoother fields and reduced local intensity contrasts which will be discussed in Section 6.2. To complement these spatial overpasses, Figure 5.10 visualizes the predicted probability of exceeding specific precipitation thresholds. This is shown solely for the same example as the heavy-rain case.

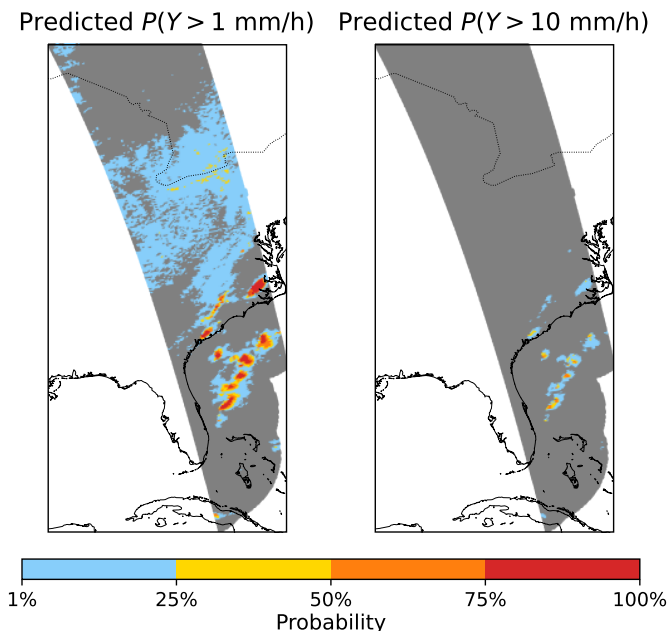


Figure 5.10: Spatial mapping of the predicted probability of exceeding specific precipitation thresholds for the heavy-rain case in figure 5.7 using the all-channel model.

The probability maps align closely with the structural characteristics of the precipitation regimes. The 1 mm h^{-1} map identifies the expansive boundaries of the precipitating cloud mass. The 10 mm h^{-1} threshold concentrates the highest probability values directly over the active cores. When comparing this to the heavy-rain case, the regions with a high probability of exceeding 10 mm h^{-1} align clearly with the highest intensity precipitation features in the MRMS reference. This demonstrates that even when the deterministic posterior mean does not capture the exact intensity of rain, the probabilistic output successfully captures the behaviour of the precipitation.

5.4 Effect of the sub-millimetre channels

The impact of the sub-millimetre channels is quantified by comparing the all-channel model against the reduced-channel configuration. A systematic improvement is visible across the bulk validation metrics detailed in Table 5.1. Including the sub-millimetre channels drops the RMSE from 0.708 to 0.655 mm h⁻¹, reduces the global CRPS from 0.041 to 0.040 mm h⁻¹ and increases the linear correlation coefficient from 0.525 to 0.587. This baseline improvement is supported by the binned metrics in Figure 5.4 where the all-channel model maintains a lower relative RMSE across all active precipitation categories. The relative MAE is also lower in most bins, with the slight exception of the 2.5–10 mm h⁻¹ interval.

The density distributions in Figure 5.5 show that the all-channel density tracks slightly closer to the test reference above 10 mm h⁻¹ indicating a marginally better recovery of the heavy tail. However, when evaluating the conditional density (Figure 5.6) and the probabilistic calibration curves (Figure 5.2) both models exhibit similar behaviour. The conditional mean plateaus and the reliability curves overlap closely between the two models, showing that the sub-millimetre channels do not produce obvious differences in the baseline probabilistic calibration or the conditional response.

However, the spatial difference fields across the overpasses reveal that the sub-millimetre channels provide useful localized updates. As shown in Figure 5.7(g) for the heavy-rain case, distinct structural adjustments occur in the active southern region. Furthermore, the all-channel model predicts slightly heavier precipitation across the upper part of the swath.

In the moderate-rain case shown in Figure 5.8(g), the differences form a scattered mix of positive and negative updates along the core structural paths of the precipitating system. This demonstrates that the sub-millimetre channels introduce localized intensity shifts to the spatial geometry of the front.

In the weak-rain case, the difference map in Figure 5.9(g) shows that the all-channel model predicts slightly heavier precipitation in the upper part of the swath where the actual rain is located. It predicts lower rain rates across the rain-free regions in the lower half. This suggests that the sub-millimetre channels help bound both higher and lower rain areas more accurately.

These metrics, distributions, and spatial comparisons indicate that the sub-millimetre channels provide useful additional information, contributing to an overall improvement of the precipitation retrieval.

6 Discussion

This chapter discusses the main findings from the AWS precipitation retrieval. The focus is on the retrieval performance, the underestimation of high precipitation rates, the contribution of the sub-millimetre channels and the main limitations of the evaluation framework.

6.1 Retrieval performance

The results show that AWS antenna temperatures contain valuable information for surface precipitation retrieval. The models successfully capture the general precipitation regimes and the spatial structure of precipitation systems across the selected swath examples. This indicates that the QRNN learns a meaningful physical relationship between the satellite observations and the MRMS surface precipitation rates. This success aligns with previous research demonstrating that passive microwave observations can provide highly effective surface precipitation retrievals (Pfreundschuh et al., 2022a; Kidd and Levizzani, 2011).

When evaluating the spatial swaths, structural differences between the MRMS target and the Late Run IMERG product can be seen. IMERG provides an independent comparison source. This is valuable for scenes where the retrieval disagrees with MRMS, helping to identify if the scene is inherently difficult to retrieve. As shown in Figure 5.9 for the weak-rain case, IMERG can miss lighter events. Furthermore, in the heavy-rain case shown in Figure 5.7, the precipitation in the upper part of the swath is not as prevalent in IMERG compared to the MRMS target. These disagreements might indicate an ambiguous radiometric scene rather than just a model error.

When evaluating the conditional density (Figure 5.6) a plateau is observed below 1.0 mm h^{-1} . This horizontal structure indicates a loss of model sensitivity to low physical rain rates. As the true target shifts across these near-zero values the model predictions remain locked. This is likely related to the difficulty of differentiating weak precipitation inputs from background surface emissivity variations and the inherent instrument noise floor of the satellite sensors. Passive microwave retrievals

frequently struggle with this thresholding effect over land as the emission signal of liquid water is often masked by the variable surface emissivity, specifically when there is very little rain (Kidd and Levizzani, 2011; Stephens and Kummerow, 2007).

6.2 Underestimation of high precipitation rates

A limitation of the retrieval is the underestimation of stronger precipitation. This underestimation could stem from a combination of data sparsity, algorithmic properties, and sensor resolution limits.

First, there is a strong imbalance in the precipitation target. Most samples correspond to no rain or weak precipitation while intense precipitation is rare. For example, bins exceeding 50 mm h^{-1} represent only a tiny fraction of the total dataset population. This data sparsity makes it difficult for the network to confidently predict extreme values which is reflected in the reliability diagrams where the predicted upper quantiles are systematically too low for heavy rain cases.

Second, a major contributing factor to the smoothed peak intensities is the mathematical use of the posterior mean as the point estimate. Since the network predicts a highly skewed, zero-bounded probability distribution, taking the expected value is highly sensitive to this skewness. It naturally suppresses the extreme peak values in the upper tail of the full distribution. This forces the deterministic estimate to be inherently conservative.

Furthermore, the spatial scale difference between AWS and MRMS introduces an unavoidable smoothing effect. While the MRMS data is kept on a fine 1 km grid, it was smoothed prior to collocation. This smoothing ensures that when using the nearest point method, errors from spatial mismatch are not as severe. However, the AWS microwave sounder integrates radiances over a much broader area ($\sim 10 \text{ km}$ at nadir). The combination of this large physical footprint and the target smoothing naturally acts as a spatial filter that rounds off high-frequency peaks.

Finally, this underestimation of high precipitation rates is a recognized challenge in passive microwave remote sensing. Previous studies evaluating operational satellite algorithms, such as GPROF, have also noted difficulties in capturing the full intensity of extreme precipitation peaks (Pfreundschuh et al., 2022a; Kidd and Levizzani, 2011). The ambiguous relationship between upwelling microwave radiances and surface precipitation means that highly intense, localized precipitation often do

not produce a proportionally distinct signal compared to moderate rainfall, leading to more conservative retrieval estimates (Stephens and Kummerow, 2007).

6.3 Contribution of the sub-millimetre channels

The comparison between the all-channel and reduced-channel models suggests that the sub-millimetre channels provide useful additional information for the retrieval. The all-channel model performs slightly better in the overall metrics and consistently lowers relative errors across most active precipitation-intensity bins.

This contribution is visually confirmed in the swath comparisons from Section 5.3. The absolute difference maps demonstrate that deviations between the two models are highly localized. Rather than applying a uniform shift, the sub-millimetre channels adjust values locally. For example, they introduce structural changes directly inside the active southern region for the heavy-rain case. In the weak-rain case, these channels help bound the system by predicting slightly more rain where precipitation exists and less rain in the dry regions.

This result is physically plausible. The sub-millimetre channels are expected to be more sensitive to scattering by frozen hydrometeors than lower-frequency microwave channels (Eriksson et al., 2025; Duncan et al., 2025). Information from these channels can therefore help the model better identify existing cloud structures related to precipitation and also help it more confidently identify clear-sky regions with no rain.

The improvement should not be overinterpreted. Both models show similar broad behaviour and both underestimate heavy precipitation. The added value of the sub-millimetre channels should therefore be understood as part of a multi-channel retrieval rather than as a complete solution on their own.

6.4 Limitations of the evaluation

The evaluation is limited by the use of MRMS as the reference product. While MRMS provides a highly valuable, high-resolution reference dataset over the United States, it can still contain uncertainties related to radar calibration, attenuation, beam height and the conversion from radar reflectivity to surface precipitation rate (Zhang et al., 2016). Differences between the AWS retrieval and MRMS are math-

ematically treated as retrieval errors, although a portion of these differences may originate from uncertainties in the MRMS product itself.

In the upper part of the swath for the heavy-rain case (Figure 5.7), the precipitation structures present in the MRMS reference lack obvious corresponding signals in the raw AWS antenna temperature inputs as shown in Figure 5.7(h)–(k). This could indicate a limitation in the information content of the AWS observations themselves. Passive microwave antenna temperatures can saturate or become highly ambiguous over intense precipitation (Stephens and Kummerow, 2007). The AWS data may not contain enough specific information to determine the exact level of precipitation. Without clear, distinct signals in the input data, the network inevitably struggles to reconstruct localized intensity peaks.

Additionally, this study is inherently affected by the early mission phase of the newly launched AWS satellite. Initial calibration uncertainties, data availability constraints and potential early instrument noise may influence the results. The size of the dataset is also constrained by the limited temporal availability of AWS data. The collocation procedure introduces another layer of uncertainty. Each AWS observation is matched to the nearest MRMS field in time and to a spatially averaged MRMS neighbourhood. Temporal mismatch can be particularly impactful for fast-moving or rapidly evolving weather systems.

Finally, the geographical scope is a limitation. The dataset is restricted to observations collocated with MRMS over the contiguous United States. Other regions may contain fundamentally different surface types, background emissivities and atmospheric conditions.

7 Conclusion

This thesis investigated precipitation estimation from the Arctic Weather Satellite (AWS) using quantile regression neural networks. The retrieval was trained and evaluated using AWS antenna temperatures collocated with MRMS surface precipitation rates. The main aim was to assess whether AWS observations contain useful information for precipitation retrieval and to investigate the specific contribution of the sub-millimetre channels.

The results show that AWS antenna temperatures contain valuable information related to surface precipitation. The retrievals effectively capture the general precipitation regimes and the spatial structure of precipitation systems across the domain. However, for heavier rain, the model does not achieve the same intensity as the radar reference.

To isolate the impact of the sub-millimetre channels, a model using all available frequency channels was compared against a baseline model without the sub-millimetre channels. This comparison confirms that the sub-millimetre channels provide useful additional information. The full model performs consistently better across the bulk verification criteria, slightly reducing errors and increasing correlation. The spatial analysis shows that these channels help refine the retrieved fields by adjusting values locally rather than applying a uniform shift.

This evaluation serves as an initial regional assessment. While MRMS provides a useful high-resolution reference it can still contain uncertainties and these results do not establish global AWS retrieval performance. Future work could evaluate the retrieval over longer multi-year time periods and expand to different geographical regions by using the GPM mission. The retrieval methodology could be improved by utilizing images as input, enabling the model to retrieve information from a region of points.

Overall this thesis provides an initial demonstration that AWS observations can be used for probabilistic precipitation retrieval. The results show that the retrieval captures meaningful structural trends and that the sub-millimetre channels contribute useful, targeted information to the multi-channel passive microwave framework.

Bibliography

- Amell, A., Hee, L., Pfreundschuh, S., and Eriksson, P. (2025). Precipitation retrievals from the arctic weather satellite using quantile regression neural networks. *Journal of Geophysical Research: Atmospheres*. DOI: <https://doi.org/10.1029/2025JD044595>.
- Draper, D. W., Newell, D. A., Wentz, F. J., Krimchansky, S., and Skofronick-Jackson, G. M. (2015). The global precipitation measurement (GPM) microwave imager (GMI): Instrument overview and early on-orbit performance. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(7):3452–3462. DOI: <https://doi.org/10.1109/JSTARS.2015.2403303>.
- Duncan, D. I., Bormann, N., Fielding, M. D., Geer, A. J., McEvoy, P., May, E., and Eriksson, P. (2025). A first global view of sub-millimetre radiances from the arctic weather satellite. Preprint, ESS Open Archive. DOI: <https://doi.org/10.22541/essoar.176279324.43846104/v1>.
- Eriksson, P., Emrich, A., Kempe, K., Riesbeck, J., Aljarosha, A., Auriacombe, O., Kugelberg, J., Hekma, E., Albers, R., Murk, A., Pedersen, S. M., John, L., Stake, J., McEvoy, P., Rydberg, B., Dybbroe, A., Thoss, A., Canestri, A., Accadia, C., Colucci, P., Gherardi, D., and Kangas, V. (2025). The arctic weather satellite radiometer. *Atmospheric Measurement Techniques*, 18:4709–4729. DOI: <https://doi.org/10.5194/amt-18-4709-2025>.
- EUMETSAT (2026). Europe backs transformative polar satellite constellation. Accessed 2026-05-16. URL: <https://www.eumetsat.int/europe-backs-transformative-polar-satellite-constellation-2026>.
- Fabry, F. (2015). *Radar Meteorology: Principles and Practice*. Cambridge University Press, Cambridge. DOI: <https://doi.org/10.1017/CB09781107707405>.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA. URL: <https://www.deeplearningbook.org/>.
- Gorooh, V. A., Kalia, S., Nguyen, P., Hsu, K., and Sorooshian, S. (2020). Deep neural network cloud-type classification (deep-CTC) model and its application

- in evaluating PERSIANN-CCS. *Journal of Hydrometeorology*, 21(12):2715–2727. DOI: <https://doi.org/10.1175/JHM-D-20-0065.1>.
- Hou, A. Y., Kakar, R. K., Neeck, S., Azarbarzin, A. A., Kummerow, C. D., Kojima, M., Oki, R., Nakamura, K., and Iguchi, T. (2014). The global precipitation measurement mission. *Bulletin of the American Meteorological Society*, 95(5):701–722. DOI: <https://doi.org/10.1175/BAMS-D-13-00164.1>.
- Huffman, G. J., Bolvin, D. T., Braithwaite, D., Hsu, K., Joyce, R., Kidd, C., Nelkin, E. J., Sorooshian, S., Tan, J., and Xie, P. (2020). Integrated multi-satellite retrievals for the global precipitation measurement (GPM) mission (IMERG). *Satellite Precipitation Measurement*, pages 343–353. DOI: https://doi.org/10.1007/978-3-030-24568-9_19.
- Kidd, C., Becker, A., Huffman, G. J., Muller, C. L., Joe, P., Skofronick-Jackson, G., and Kirschbaum, D. B. (2017). So, how much of the Earth’s surface is covered by rain gauges? *Bulletin of the American Meteorological Society*, 98(1):69–78. DOI: <https://doi.org/10.1175/BAMS-D-14-00283.1>.
- Kidd, C. and Levizzani, V. (2011). Status of satellite precipitation retrievals. *Hydrology and Earth System Sciences*, 15:1109–1116. DOI: <https://doi.org/10.5194/hess-15-1109-2011>.
- Koenker, R. (2005). *Quantile Regression*, volume 38 of *Econometric Society Monographs*. Cambridge University Press, Cambridge. DOI: <https://doi.org/10.1017/CB09780511754098>.
- Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50. DOI: <https://doi.org/10.2307/1913643>.
- Loshchilov, I. and Hutter, F. (2019). Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*. URL: <https://arxiv.org/abs/1711.05101>.
- Michaelides, S., Levizzani, V., Anagnostou, E., Bauer, P., Kasparis, T., and Lane, J. E. (2009). Precipitation: Measurement, remote sensing, climatology and modeling. *Atmospheric Research*, 94(4):512–533. DOI: <https://doi.org/10.1016/j.atmosres.2009.08.017>.
- NOAA (2025). Surface precipitation rate (spr) - version 12.2. Accessed 2026-05-17. URL: <https://vlab.noaa.gov/web/wdtd/-/surface-precipitation-rate-spr-1>.

- Pfreundschuh, S., Brown, P. J., Kummerow, C. D., Eriksson, P., and Norrestad, T. (2022a). GPROF-NN: a neural-network-based implementation of the Goddard Profiling Algorithm. *Atmospheric Measurement Techniques*, 15:5033–5060. DOI: <https://doi.org/10.5194/amt-15-5033-2022>.
- Pfreundschuh, S., Eriksson, P., Duncan, D., Rydberg, B., Håkansson, N., and Thoss, A. (2018). A neural network approach to estimating a posteriori distributions of Bayesian retrieval problems. *Atmospheric Measurement Techniques*, 11(8):4627–4643. DOI: <https://doi.org/10.5194/amt-11-4627-2018>.
- Pfreundschuh, S., Ingemarsson, I., Eriksson, P., Vila, D. A., and Calheiros, A. J. P. (2022b). An improved near-real-time precipitation retrieval for Brazil. *Atmospheric Measurement Techniques*, 15:6907–6933. DOI: <https://doi.org/10.5194/amt-15-6907-2022>.
- Rees, W. G. (2013). *Physical Principles of Remote Sensing*. Cambridge University Press, Cambridge, 3 edition. DOI: <https://doi.org/10.1017/CB09781139017411>.
- Rodgers, C. D. (2000). *Inverse Methods for Atmospheric Sounding: Theory and Practice*, volume 2 of *Series on Atmospheric, Oceanic and Planetary Physics*. World Scientific, Singapore. DOI: <https://doi.org/10.1142/3171>.
- Sadeghi, M., Nguyen, P., Hsu, K., and Sorooshian, S. (2022). Deep neural network high spatiotemporal resolution precipitation estimation (Deep-STEP) using passive microwave and infrared data. *Journal of Hydrometeorology*, 23(4):597–617. DOI: <https://doi.org/10.1175/JHM-D-21-0194.1>.
- SMHI (2026). Dags att förverkliga det nya satellitprogrammet EPS-Sterna. Accessed 2026-05-16. URL: <https://www.smhi.se/nyheter/nyheter/2026-01-29-dags-att-forverkliga-det-nya-satellitprogrammet-eps-sterna>.
- Stephens, G. L. and Kummerow, C. D. (2007). The remote sensing of clouds and precipitation from space: A review. *Journal of the Atmospheric Sciences*, 64(11):3742–3765. DOI: <https://doi.org/10.1175/2006JAS2375.1>.
- Zhang, J., Howard, K., Langston, C., Kaney, B., Qi, Y., Tang, L., Grams, H., Wang, Y., Cocks, S., Martinaitis, S., Arthur, A., Cooper, K., Brogden, J., and Kitzmiller, D. (2016). Multi-Radar Multi-Sensor (MRMS) quantitative precipitation estimation: Initial operating capabilities. *Bulletin of the American Meteorological Society*, 97(4):621–637. DOI: <https://doi.org/10.1175/BAMS-D-14-00174.1>.