

# **Bayesian Methods for Climate Time-Series Forecasting**

Olivia Nguyen

Submitted in total fulfilment of the requirements  
of the degree of Doctor of Philosophy

March 2026

School of Mathematics and Statistics  
The University of Melbourne

*Unofficial template, not affiliated with or endorsed by The University of Melbourne. The author, committee and results in this sample are fictional. Delete this line before you submit.*

# Abstract

Seasonal climate forecasts guide decisions in agriculture, water management and emergency planning, and their usefulness depends on how honestly they express uncertainty. A forecast that says a coming season will be “slightly drier than average” is of little help to a farmer who needs to know the chance of a season dry enough to lose a crop. This thesis develops Bayesian methods for forecasting monthly rainfall and temperature from historical station records and large-scale climate indices, with the aim of producing forecasts whose probabilities can be taken at face value.

The first contribution is a hierarchical state-space model that shares information between nearby weather stations. Stations with short or broken records borrow strength from their neighbours through a prior whose correlation decays with distance, which matters in Australia, where many stations in the interior have decades of missing data. The second contribution is a Gaussian-process component that captures slowly varying trends, including the warming trend in temperature, without forcing them into a fixed parametric form. The third is an approximate inference scheme, based on a variational approximation with a Kalman smoother inside it, that scales to thousands of stations on ordinary hardware.

Forecasts are assessed with proper scoring rules and with the coverage of their prediction intervals, rather than with point accuracy alone, because a model can be accurate on average and still badly overconfident. Across 1,200 Australian stations and a ten-year verification period, the hierarchical model improves the continuous ranked probability score of three-month rainfall forecasts by 14% relative to a climatology baseline, and its 90% prediction intervals contain the observed value 89% of the time. Models fitted to each station independently are more accurate than climatology but overconfident, with coverage of only 84%.

The thesis concludes with a discussion of how probabilistic forecasts can be communicated to decision makers who are not statisticians, and of the limits of purely statistical models when the climate itself is changing and the past is an imperfect guide to the future.

# Declaration

This is to certify that:

- i. the thesis comprises only my original work towards the Doctor of Philosophy except where indicated in the Preface;
- ii. due acknowledgement has been made in the text to all other material used;
- iii. the thesis is fewer than 100 000 words in length, exclusive of tables, maps, bibliographies and appendices.

Olivia Nguyen, March 2026

# Preface

This thesis contains no published material. The station records used in Chapter 4 were obtained from a public national archive. The computations were performed on a university high-performance computing system. No third-party editorial assistance was used.

# Acknowledgements

I thank my supervisors for their generosity with ideas and time, my advisory committee for its guidance, and my fellow students in the School of Mathematics and Statistics for their company.

# Contents

|                                       |            |
|---------------------------------------|------------|
| <b>Abstract</b>                       | <b>ii</b>  |
| <b>Declaration</b>                    | <b>iii</b> |
| <b>Preface</b>                        | <b>iv</b>  |
| <b>Acknowledgements</b>               | <b>v</b>   |
| <b>1 Introduction</b>                 | <b>1</b>   |
| <b>2 Background</b>                   | <b>2</b>   |
| <b>3 A Hierarchical Spatial Model</b> | <b>3</b>   |
| <b>4 Results</b>                      | <b>4</b>   |
| <b>5 Conclusion</b>                   | <b>5</b>   |
| <b>A Variational Updates</b>          | <b>7</b>   |

# List of Figures

|     |                                                          |   |
|-----|----------------------------------------------------------|---|
| 3.1 | Hierarchical structure of the forecasting model. . . . . | 3 |
|-----|----------------------------------------------------------|---|

# List of Tables

|     |                                                               |   |
|-----|---------------------------------------------------------------|---|
| 4.1 | Three-month rainfall forecasts across 1,200 stations. . . . . | 4 |
|-----|---------------------------------------------------------------|---|

# Chapter 1

## Introduction

A farmer deciding what to plant in autumn cares less about the most likely rainfall than about the chance of a very dry season. Forecasts that report only a single value hide that information. Bayesian models make uncertainty explicit, and when they are well calibrated they support better decisions [1].

# Chapter 2

## Background

We model a station's monthly anomaly  $y_t$  with a latent state  $x_t$  that evolves over time,

$$y_t = Fx_t + \varepsilon_t, \quad x_t = Gx_{t-1} + \eta_t, \quad (2.1)$$

a dynamic linear model [3]. Gaussian processes provide flexible priors for smooth trends [2]. Probabilistic forecasts are compared with proper scoring rules such as the continuous ranked probability score [1].

# Chapter 3

## A Hierarchical Spatial Model

Stations share parameters through a hierarchical prior whose correlation decays with distance, so that sparsely observed stations borrow strength from their neighbours. Inference uses a variational approximation with a factorized posterior over stations.




Figure 3.1: Hierarchical structure of the forecasting model.

# Chapter 4

## Results

Table 4.1 compares three-month rainfall forecasts over a ten-year verification period.

Table 4.1: Three-month rainfall forecasts across 1,200 stations.

| Model                      | CRPS (relative to climatology) | 90% interval coverage |
|----------------------------|--------------------------------|-----------------------|
| Climatology                | 1.00                           | 0.90                  |
| Independent station models | 0.93                           | 0.84                  |
| Hierarchical model         | 0.86                           | 0.89                  |

# **Chapter 5**

## **Conclusion**

Sharing information across stations improves both the accuracy and the calibration of seasonal forecasts. Incorporating outputs of physical climate models as covariates is a natural extension.

# Bibliography

- [1] Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [2] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [3] Mike West and Jeff Harrison. *Bayesian Forecasting and Dynamic Models*. Springer, second edition, 1997.

# Appendix A

## Variational Updates

The variational posterior for each station's states is Gaussian and is updated with a Kalman smoother given the current hyperparameters, alternating with closed-form updates of the hyperparameters until the evidence lower bound changes by less than  $10^{-4}$ .