

Local Power Spectra of the Earth's Atmosphere



Salah Kouhen

Linacre College

A thesis submitted for the degree of

Doctor of Philosophy

Supervised by

Professor Hannah M. Christensen, Professor David P. Marshall

Department of Physics, AOPP

University of Oxford

Michaelmas Term, 2024

Dedicated to Said and Jacqui Kouhen

Declaration

I declare that this thesis is entirely my own work and, except where stated, describes my own research.

Salah Kouhen

Linacre College

Acknowledgements

I am incredibly grateful to have had Professor Hannah Christensen as my supervisor. Her ability to ask the right question and quickly see the heart of the issue has been invaluable. I am also grateful for the freedom she gave me to explore interesting topics while also guiding me so that I didn't stray too far. Hannah has been incredibly supportive and understanding, which has been essential in navigating the past four years of a PhD and life. I hope that at least some fraction of her qualities have rubbed off on me.

I am also thankful to Professor David Marshall, who became my co-supervisor in my second year. David is an exceptionally friendly and supportive person who always has an entertaining anecdote, he also possesses the skill of seeing through to what is most interesting. I have benefitted enormously from his advice.

I am grateful to Tim Palmer, Kristen Strommen, Fenwick Cooper, Tim Woolings, Ray Pierrehumbert, and many colleagues at AOPP for engaging discussions and valuable insights.

I owe special thanks to Benjamin Storer and his group under Hussein Aluie at the University of Rochester for their collaboration. I feel deeply lucky to have chanced into a collaboration with Ben and having the opportunity to visit his group. Ben's technical skill always leaves me inspired, his friendly character made my visit to Rochester a pleasure and our meetings something I always look forward to.

I am thankful to NERC ([Grant NE/S007474/1]) and the Ruth and Nevill Mott Scholarship, whose funding made this PhD possible. I also acknowledge Linacre College for its welcoming atmosphere and excellent food. My thanks also go to the subdepartment of Physics, AOPP, for its student support, especially from Lucy Li, and computational support of Man-Suen Chan.

I am forever grateful for the love and support of my parents. To my mum, whose bottomless curiosity about the world and creativity led me to science and to my dad, whose work-ethic and tutoring in mathematics allowed me to arrive at a PhD. I love you both very much. I would also like to thank my auntie Katie, my Nan and my late grandfather for their love. My granddad's seemingly limitless knowledge of the world inspired me growing up, he is sorely missed.

To my friends, thank you for being there over the past four years. Thank you to my brother Stefan, to Samer, Ljubo, Tom and to my two comrades who moved with me from Warwick to Oxford, Aldi and James. I've been so lucky to meet each of you. I would also like to thank those friends I have met while at Oxford, thank you Gill, Andrew, Ruthie and Anupa. I hope to know you all for many years to come.

Abstract

This thesis is dedicated to the kinetic energy power spectrum of the Earth's atmosphere. This spectrum refers to the distribution of kinetic energy at different length scales of motion, from the energy in the largest swirling weather systems, to the smallest convective clouds. The spectrum is a multifaceted and deep topic, and it has earned academic attention for several reasons.

One reason is the inherent beauty the spectrum holds in many turbulent systems. Turbulence is a highly complex phenomenon that is composed of motion at a great range of scales; however, when viewed through the lens of a power spectrum, a simple power law often neatly describes the magnitude of the kinetic energy over a range of scales. This simplification of turbulence, which is an incredibly ubiquitous and challenging phenomenon in physics, invites us to investigate the spectral perspective further. In Part One of this thesis, we do this by applying novel coarse-graining techniques to extract spectra from the atmosphere. This technique allows us to bring a spectral perspective to spaces that were previously difficult to reach. For example, we are able to investigate how the distribution of energy changes at different locations on Earth, creating the first consistently generated map of mesoscale spectral slopes in the atmosphere. We also investigate how the spectrum changes under different conditions using a new method of "conditioned spectra"; some of these analyses are completely novel, such as the distribution of energy when mean sea level pressure is high or low or mesoscale spectral energy transfers are from large to small or small to large scales. Others are comparable to existing simulation or observational work, such as the spectrum corresponding to different levels of precipitation or different orographic heights. Spectra are often used to assess the smallest scale to which a simulation can represent the real world, and so a more nuanced understanding of local spectra could open the door to deeper and more meaningful comparisons between different models and between models and observations. In addition to spectra, our methodology can resolve local spectral fluxes. Spectral fluxes are central to parametrisations of unresolved scales in weather models, raising the possibility of evaluating and improving these parametrisations in the future through our methodology.

Another reason for investigating the spectrum is that it contains clues to a system's underlying dynamics. In Part Three, we present two-dimensional turbulence theory and review how details of the spectrum, such as the spectral slope or the direction of energy transfer, can indicate the underlying turbulent dynamics active within a particular range of scales. Mesoscale weather dynamics are much less understood than synoptic weather. It is still unclear how waves, turbulence and external forcings interact to produce the motion at these scales. Much of the progress towards understanding this topic has come from investigations into the spectrum. In Part One, we investigate how the direction of energy flux differs at different locations and under different atmospheric conditions. We also investigate the local partition of energy into rotational and divergent motion. From this, we are able to conclude that purely turbulent models of the atmosphere, which have become unpopular over the past 20 years, may be valid under certain conditions or in certain locations. We also support the prevailing notion that gravity waves are centrally important in energising the upper troposphere and lower stratospheric spectrum in the mesoscales. This work lays the foundations for a more detailed local breakdown of the mesoscale fluxes that could lead to a far deeper understanding of mesoscale processes than at present.

The last reason we will give for the academic attention the atmosphere's spectrum has attracted is the most

intriguing. The spectrum has a link with predictability. [Lorenz \(1969\)](#) showed that the power spectrum can be used in simple turbulent systems to predict the rate of error growth. In Part One, we explore local power spectra, a logical possibility is that if a link between the global spectrum and predictability exists, a link between local predictability and the local spectrum may exist. In Part Two, we provide background on the link between spectra and predictability. After giving this background, we extend the global theory of error growth in atmospheric turbulence by including the effects of the condensate, the shallowing of the spectrum that occurs at large scales. We show that including a realistic condensate significantly extends the predictability estimates returned by turbulence models (with the caveat that we ignore the anisotropy of the real atmosphere in our model). We then take the first steps towards a local model of error growth based on a cascade of error to larger scales by creating a toy model. This toy model illustrates how error would behave if it was related in a simple way to the local eddy turnover time. We point out this approach's limitations but find some broad consistencies with known spatial patterns in error growth, such as larger absolute errors in the midlatitudes and larger relative errors over orography. This work represents the start of investigations into the local relationships between error growth and the power spectrum, research which could be incredibly important for weather prediction and turbulence modelling generally. If a diagnostic or prognostic relationship between local spectra and error growth could be found, it may be of incredible utility, since representing the inevitable growth of error is the most challenging and important aspect of weather prediction.

Table of contents

Table of contents	xi
List of figures	xiv
I Spatially local mesoscale kinetic energy spectra	1
1 Spectra in the atmosphere, a background	2
1.1 Introduction to the thesis	2
1.1.1 Structure of the thesis	4
1.2 The global spectrum, the current picture	5
1.2.1 Definitions	5
1.2.2 Motivation for investigating the spectrum	6
1.2.3 The spectrum in analysis	7
1.2.4 The slope in analysis	11
1.2.5 Variations in time and location	12
1.3 Explaining the shape of the spectrum	13
1.3.1 Purely turbulent explanations	13
1.3.2 Wave explanations	15
1.4 Summary	16
2 Coarse-graining	17
2.1 Introduction	17
2.2 Extracting spectra via coarse-graining	17
2.3 Order, or when coarse-graining fails	21
2.3.1 The scaling behaviour of the coarse-grain spectrum	22
2.3.2 A bias towards steeper slopes	24
2.4 Local spectral slopes	25
2.5 Extension to vector fields	28
2.6 The coarse-grain energy flux	28
2.6.1 Deriving the kinetic energy budget for large scales	29
2.7 Conclusions and overview	33
3 Local mesoscale slopes in the atmosphere	34
3.1 Introduction	34
3.2 Mesoscale spectral slopes across the globe	36
3.2.1 200 hPa, the upper troposphere, lower stratosphere	36
3.2.2 600 hPa, the mid-troposphere	40
3.3 Convective and large-scale precipitation	45
3.4 Mean sea level pressure	47
3.5 Discussion and conclusions	48
4 Energy flux and spectral slope	50
4.1 Energy flux in the mesoscales	51
4.1.1 Flux in the stratosphere and upper troposphere: 200 hPa	51
4.1.2 Flux in the troposphere: 600 hPa	53
4.1.3 Orography, precipitation, pressure and flux	54
4.2 Spectra conditioned on spectral flux	57
4.3 Discussion and conclusions	59
5 A local divergent rotational decomposition	60
5.1 Introduction	60
5.2 Results	61

5.3	Conclusions	66
II	Predictability	68
6	Spectral error growth theory in 2D turbulence	69
6.1	Introduction	69
6.1.1	Key concepts	69
6.2	The Lorenz 1969 model	70
6.2.1	The model	70
6.2.2	The derivation of L69	71
6.2.3	Interpreting the L69 model	74
6.2.4	Limitations of the L69 model	75
6.2.5	A dimensional model for error growth	77
6.3	The effect of a condensate on L69	78
6.4	Discussion and conclusions	80
7	Towards a local nonlinear model for error growth	81
7.1	Introduction	81
7.2	Applying the dimensional model to the atmosphere	84
7.3	Summary and conclusions	88
III	Turbulence theory	90
8	Two-dimensional turbulence theory	91
8.1	Alpha models	91
8.2	Deriving the spectrum	92
8.2.1	Proving invariance of generalised enstrophy and generalised energy	93
8.2.2	Argument for upscale energy flux and downscale enstrophy flux	94
8.2.3	Finding the spectral slopes	97
8.3	Summary	100
9	Spectral slopes and error growth in simulated alpha models	101
9.1	Simulating alpha models psuedo-spectrally	101
9.2	Spectral slopes in alpha models	105
9.2.1	Speculation on extending dimensional arguments to Casimirs and nonlocal interactions	107
9.3	Error growth	110
9.3.1	Upscale vs upamplitude error growth	110
9.3.2	Twin-experiments in alpha models	110
9.4	Summary and conclusions	113
10	Conclusions	115
10.1	Summary of key results	115
10.1.1	Part One	115
10.1.2	Part Two	116
10.1.3	Part Three	117
10.2	Future work	118
10.2.1	Understanding mesoscale dynamics	118
10.2.2	Parametrisation improvement and model evaluation	119
10.2.3	Predictability	122
Appendix A	Fourier transform and structure function refresher	124
A.1	Fourier transform refresher	124
A.1.1	Introduction	124
A.1.2	The orthogonality of Fourier modes	124
A.1.3	Definitions	126
A.1.4	Plancherel's theorem	126

A.1.5	The power spectrum	127
A.1.6	The convolution theorem	127
A.1.7	Isotropy and the Hankel transform	128
A.2	Structure functions	128
A.2.1	Introduction	129
A.2.2	Relation to power spectra	129
A.2.3	Structure function blowup	130
A.2.4	Structure function scaling	131
Appendix B	SQG and alpha models derivation	132
B.1	The surface quasi-geostrophic (SQG) equation derivation	132
B.2	Showing SQG is an alpha model	134
Bibliography		136

List of figures

1.1	The spherical harmonic power spectrum at 25 pressure levels from the 2020 ECMWF analysis.	8
1.2	Best fit slopes for the spherical harmonic power spectrum of 2020 ECMWF analysis.	9
1.3	The Lindborg empirical fit to aircraft observations.....	10
1.4	Schematic of a 2D turbulence spectrum.	14
2.1	A schematic of the coarse-graining kernel, showing a basic kernel in real and spectral space.	19
2.2	A comparison of the spherical harmonic spectrum and the filtering spectrum.	24
2.3	A comparison of methods for aggregating local spectral slopes.	26
3.1	Global map of mesoscale spectral slope in analysis data at 200 hPa.....	36
3.2	ERA5 2020 average hourly precipitation and orography fields.	37
3.3	Global map of kinetic energy at 200 hPa.	37
3.4	Global map of mesoscale energy in analysis data at 200 hPa.	38
3.5	Orography conditioned spectra for 200 hPa.....	38
3.6	Precipitation conditioned spectra for 200 hPa.	39
3.7	Spectrum conditioned to exclude precipitation and orography for 200 hPa.....	40
3.8	Global map of mesoscale spectral slope in analysis data at 600 hPa.....	40
3.9	ERA5 2020 average hourly precipitation and orography fields.	41
3.10	Global map of kinetic energy at 600 hPa.	41
3.11	Global map of mesoscale energy in analysis data at 600 hPa.	42
3.12	Orography conditioned spectra for 600 hPa.....	43
3.13	Precipitation conditioned spectra for 600 hPa.	43
3.14	Spectrum conditioned to exclude precipitation and orography for 600 hPa.....	44
3.15	Large-scale and convective conditioned spectra for 200 hPa.....	45
3.16	Large-scale and convective conditioned spectra for 600 hPa.....	46
3.17	Mean sea level pressure conditioned spectra for 200 and 600 hPa.....	47
4.1	Purely rotational and divergent components of the flux.	50
4.2	Global map of mesoscale energy flux in analysis data at 200 hPa.....	52
4.3	Global map of mesoscale energy flux in analysis data at 600 hPa.....	53
4.4	Orography conditioned flux for 600 and 200 hPa.	54
4.5	Precipitation conditioned flux for 600 and 200 hPa.....	55
4.6	Large-scale and convective conditioned flux for 600 hPa.....	56
4.7	Dry and sea-level conditioned flux for 600 and 200 hPa.....	56
4.8	MSLP conditioned flux for 600 and 200 hPa.....	57
4.9	Flux conditioned flux for 600 and 200 hPa.	57
4.10	Flux conditioned spectra for 200 hPa.	58
4.11	Flux conditioned spectra for 600 hPa.	58
5.1	Rotational and divergent components of the global filtering spectrum.	61
5.2	Global map of total mesoscale divergent divided by rotational energy at 200 hPa.	62
5.3	Global map of total mesoscale divergent divided by rotational energy at 600 hPa.	63
5.4	Spectra conditioned on the ratio of divergent to rotational energy.....	64
5.5	Flux conditioned on the relative energy in divergent vs rotational energy.	65
6.1	Logarithmic value of $C(k, p)$ for different spectral slopes.....	75
6.2	L69 error growth predictions from Lindborg fit.	78
6.3	L69 error growth predictions from Lindborg fit with a condensate.	79
7.1	Dimensional error growth predictions from Lindborg fit.	83
7.2	Spectra fed to dimensional model.....	84
7.3	Dimensional spectra applied to ECMWF analysis.	85
7.4	RMSE in ECMWF IFS.	86

7.5	Local error estimate from dimensional model.	88
8.1	Schematic of an initial kinetic energy distribution.	96
8.2	Incorrect schematic of evolution of a kinetic energy distribution.	96
8.3	Correct schematic of evolution of a kinetic energy distribution.	97
8.4	Schematic of KE spectrum for alpha models.	99
9.1	Instantaneous alpha model KE spectra.	104
9.2	Evolution of alpha model spectral slopes.	105
9.3	Spectral slopes against theory in alpha models.	106
9.4	Temporal variability of spectral slopes in alpha models.	106
9.5	Casimir cascade predictions for spectral slopes.	108
9.6	Casimir nonlocal predictions for spectral slopes.	109
9.7	DKE fields for alpha models.	111
9.8	DKE spectra for alpha models.	112
9.9	Total error growth for alpha models.	113
10.1	Global map of seasonal mesoscale energy in analysis data at 600 hPa.	119
10.2	Illustrative example of using coarse-graining for model evaluation.	120
10.3	Schematic of effective resolution.	121

Part I

Spatially local mesoscale kinetic energy spectra

Chapter 1

Spectra in the atmosphere, a background

1.1 Introduction to the thesis

The spectral perspective

The motions of the atmosphere that determine our weather are complex. They are the result of interactions between a great range of scales, complicated external forcings, turbulent interactions between scales and propagating waves. Predicting its motion into the future well enough to have useful information about the weather may not even be possible more than two weeks in advance. And yet, there is a perspective that collapses a very similar kind of chaos into order.

Two-dimensional turbulence shares much of the flavour of the chaotic motion of the atmosphere. Suppose we provide kinetic energy to a thin film of fluid at a particular lengthscale, then a curious thing will happen. Vortices will merge and grow until they reach the outer scale, while filaments of vorticity will stretch and fold until they are dissipated. This complex motion becomes remarkably simple when a spectral perspective is taken. If the magnitude of the kinetic energy at each scale is plotted, the energy at scales smaller than the one we injected energy into will go like L^3 while energy at larger scales will go like $L^{5/3}$. A large chunk of information has been simplified to two numbers, $5/3$ and 3 .

In our atmosphere, specifically in the upper troposphere and lower stratosphere, the spectrum reveals a similarly beautiful simplification. In observations, the picture is flipped compared to that in the 2D case described above. At large scales there is a 3 power law and at small scales there is a $5/3$ power law. The simplicity of this perspective, on a topic that is defined by complexity, invites us to explore the spectral perspective further. The gap here is between a simple spectral perspective and a complex physical space perspective, full of specific mechanisms and structures.

To bridge this gap, in this thesis, we employ a novel method for extracting local power spectra from the atmosphere. Coarse-graining involves successively blurring the velocity fields of the atmosphere to infer the energy at each scale and location. In this way, we are able to probe how the spectrum changes in different locations and how it is affected by conditions such as precipitation or sea level pressure. A local spectral perspective may be particularly useful for model validation. Spectra are already used to investigate how well a model can resolve different scales, in the future, they may be used to compare model skill in representing convection or orographic forcing. Our methodology also enables us to find the spectral fluxes at different locations. This could be incredibly useful for creating or improving parametrisations — particularly parametrisations of the mesoscales for low-resolution weather or climate models.

In addition to local spectral fluxes, coarse-graining gives us the ability to compare the amount of energy in rotational and divergent spectra locally. These two pieces of information are useful tools for inferring the mechanisms driving mesoscale dynamics, which are less understood than synoptic weather dynamics. The discovery that the average flux of energy in the mesoscales was from large to small scales in the early 2000s severely reduced the popularity of several theories of the mesoscales where energy from small-scale forcing, such as convection, played a leading role (see Section Three of this chapter). This included theories where

two-dimensional Navier-Stokes turbulence, forced at small scales, could explain the dynamics. The picture changed again in the mid-2010s when it was recognised that motion in the upper troposphere and lower stratosphere, the most intensely studied portion of the atmospheric spectrum due to the availability of aircraft observations, was dominated by divergent rather than rotational motion. Simple two-dimensional turbulence theories involve rotational, rather than divergent, flow. In this thesis, we find via a local analysis that both of these conditions are state and location-dependent. This opens the possibility that, although still not good explanations for the average shape of the global power spectrum, these turbulent arguments could explain the spectrum locally, at particular times or locations. That is, two-dimensional turbulent dynamics could dominate the mesoscale motion under some conditions.

Beyond diagnostics

In 1969, Edward Lorenz ([Lorenz, 1969](#)) made a seminal contribution to turbulence theory. By building a model for error growth in turbulence, he concluded that depending on the power law followed by the fluid's kinetic energy spectrum, a fluid may have a finite horizon of predictability. This is a stunningly counterintuitive result when first encountered. In a deterministic system, no matter how precisely one measures the initial conditions, there is a time beyond which one can never predict the motion. In the atmosphere, for the error to grow to the scales where what could be considered a weather prediction is no longer possible, is thought to take around two weeks. To escape the range where this result holds, it is currently thought we would need to measure the atmosphere to a resolution of centimetres, tens of thousands of times a second, at which point the motions of humans and animals would also need to be predicted. Weather prediction beyond a month is not just hard, it may be impossible.

Lorenz's results are specific to a flavour of two-dimensional turbulence. However, given that motions in the atmosphere are more complex than this simple model, the theory likely provides an upper, not lower, bound on predictability. If we take these results seriously, the central objective of weather prediction is not to produce higher and higher resolution predictions for the weather but to find the exact envelope of possibilities that diverge from our starting point. The forecasts for the weather do not suddenly go wrong after two weeks, error grows over time. In Lorenz's model, smaller scales saturate, that is, we lose all ability to predict them, before the largest scales. In this way, scales smaller than 400 kilometres may be completely independent of the initial conditions after a few days, while we retain some ability to predict the largest weather scales for almost two weeks. For this reason, if we discovered an alien weather-predicting machine, its output would probably be a distribution. It would show how an initial distribution, the initial condition with its uncertainty, evolved until it contained no useful information other than that inferable from external forcings.

This is an incredibly difficult goal, but progress has been made towards it. Data assimilation is the science of determining the current atmospheric state's distribution by combining observational and simulation data. Forecasts also evolved from deterministic to probabilistic in the early 90s with the invention of ensemble forecasting ([Toth and Kalnay, 1993](#); [Palmer et al., 1993](#)). One may wonder if it is possible to simulate the distribution rather than solving the dynamics directly. Surprisingly, the answer is yes. The evolution of the distribution is governed by the Liouville equation. However, despite the fact that it is possible to write down the equations, the approach is hamstrung by the curse of dimensionality. Instead of a single simulation, we are attempting to solve a continuum of them. This is especially problematic because instead of evolving smoothly, where initially close points remain close, the distribution will be stretched across a high dimensional strange attractor of possible states. Another approach is needed.

Rather than approaching the problem from this theoretical perspective, where progress appears impossible, it may be more fruitful to return to the more empirical perspective of approaches like ensemble forecasting. As Lorenz established, there is a link between error growth and the power spectrum in turbulent systems. An avenue of investigation is then to what extent this link holds in the atmosphere and whether short-term growth in error can be linked to local power spectra obtained from coarse-graining. Local predictions of error growth are not new. If one linearises the governing equations and calculates local Lyapunov exponents, the local error growth for short times can be approximated. However, the limitation of this line and other linear stability analysis methods, is they are linear theories. That is, they cannot capture the complexity of fundamentally nonlinear turbulence. In particular, that means they cannot represent the saturation of error, and so are only genuinely equipped for studying the growth of small errors over short intervals.

Turbulence is a ubiquitous problem in physics. From weather prediction to aircraft design to magnetic confinement fusion, its complexity is felt in a vast range of fields. It is also a beautiful topic. In Part Two, we discuss the first steps towards local error models based on local spectra, and in Part Three, we introduce investigations on alpha models, which could be a convenient setting for future investigations into error growth in turbulence.

1.1.1 Structure of the thesis

We now give an overview of the structure of the thesis. We have organised the research into three parts:

Part I: Spatially local mesoscale kinetic energy spectra

Chapters one to five are focused on local mesoscale kinetic energy spectra. The first chapter provides an overview of the global distribution of energy in the atmosphere as it is currently understood, including a colour plot of slope against scale and altitude, which is a novel way to present this data. We then summarise the key theories proposed to explain the global average spectrum. In the second chapter, we explain the recently developed coarse-graining methodology used to produce local spectra in this thesis, including details of how we extended this method to produce local spectral slopes and conditioned spectra.

Chapter Three presents the first consistent local maps of kinetic energy spectral slope, the patterns we observe lend support to the theory that vertically propagating gravity waves are the dominant source of energy for the spectrum in the lower stratosphere. Some of the work in this chapter has been published in [Kouhen et al. \(2024\)](#). We display spectra conditioned on precipitation (including partition into convective and large-scale sources), orographic height and mean sea level pressure (MSLP). Chapter Four shows some of the first spatial maps of spectral kinetic energy flux through the atmospheric mesoscales. We also show the first fluxes conditioned on precipitation, orography and MSLP, as well as spectra conditioned on spectral flux. Finally, in Chapter Five we examine the first local divergent and rotational kinetic energy spectra decomposition. This gives us insight into where divergent or rotational motion dominates at a particular pressure level. As a result of our investigation, we conclude that two-dimensional turbulence could still explain mesoscale dynamics under certain conditions, such as strongly rotational motion.

Part II: Predictability

Chapters Six and Seven are concerned with predictability. In Chapter Six, we start by giving an overview of the atmospheric predictability literature and then extend this body of work by showing, for the first time, the significant effect including a more realistic condensate to simple turbulence predictability models has. Including the condensate (the shallowing of the spectral slope at large scales) extends the any-skill

predictability horizon significantly, more than doubling it. It is generally assumed in seasonal forecasting that skill originates from features with longer timescales than the atmosphere, such as sea surface temperatures or features of the atmosphere with longer timescales due to complex mechanisms, such as the quasi-biennial oscillation. However, if our result holds in more realistic settings (for instance, when extended to a beta plane), it could imply that seasonal forecasts can be extended by better representing large-scale features in initial conditions, even if the processes involved are purely turbulent. In Chapter Seven, we propose a simple local nonlinear turbulent error growth model for the atmosphere. This represents a first step towards investigating the potential of local spectra for predictability quantification. We discuss our toy model's successes and shortcomings in predicting error growth patterns observed in forecasts.

Part III: Turbulence theory

The final part of the thesis is concerned with spectra in two-dimensional turbulence. In Chapter Eight, we give an overview of two-dimensional turbulence theory, deriving the spectral slopes of a family of two-dimensional turbulence models (alpha models), including the -3 slope observed in the synoptic atmosphere. In Chapter Nine, we show the results of error growth experiments with alpha models, showing how the relationship between alpha and the spectral slope changes past an alpha of two and how the flavour of error growth switches from upscale to upamplitude at the same transition point. These models can serve as testbeds for future investigations into error growth in the atmosphere.

In the final chapter, we conclude the thesis, summarising our findings and giving directions for future research.

1.2 The global spectrum, the current picture

This thesis is broadly concerned with the distribution of kinetic energy among the different length scales present in the atmosphere, the power spectrum. We will start by reiterating the motivation to study this distribution, present the current understanding of the spectrum in the atmosphere and then give an overview of the explanations for why the shape we observe prevails.

Before we can begin, a few concepts need to be introduced.

1.2.1 Definitions

Power spectra and spectral slopes

The kinetic energy power spectrum is used to quantify the distribution of energy across scales in a flow. For example, in the atmosphere, we may suppose an intense rain storm is associated with more energetic motion on scales of 1 to 100 km, so if we were to evaluate the kinetic energy power spectrum over a region with such a storm, we should see greater values at the small scales than in a more quiet region.

Generally, in turbulent flows, we observe a power-law-like relationship between length scale and the spectral energy density at that scale. That is $E(L) \propto L^{-\beta}$, where L is the length scale of interest and $E(L)$ is the spectral density of energy at that length scale. The negative of the power in this power law, β , is often called the spectral slope.

The reason for this terminology is the power spectrum is usually written in terms of wavenumber and plotted on a log-log axis. Wavenumber is $k = 2\pi/L$, so $E(k) \propto k^\beta$, taking the log of both sides, we have

$\log E(k) = \beta \log k + C$, so on a log-log axis we have a linear relationship with a slope of β . Hence, the slope represents the distribution of energy among scales, with a shallower slope indicating relatively more energy in small-scale motion and a steeper slope implying relatively more energy in large-scale motion.

There are various methods for calculating a flow's power spectrum. The most common methods, Fourier analysis and structure functions, are reviewed in Appendix A. In this thesis, we will evaluate local spectra using a coarse-graining methodology detailed in Chapter Two.

Atmospheric mesoscales

The mesoscales (lit., “middle scales”) are those smaller than the synoptic scales of weather systems but larger than individual convection cells. There are a few possible definitions of the mesoscales; we could define it in terms of the dominant processes, but to avoid a delineation that is time and altitude-dependent, we will define the mesoscales to be between 10 and 400 km.

IFS and ECMWF analysis

In this work, we will analyse data from the European Centre for Medium-Range Weather Forecasting, ECMWF. ECMWF assimilates observations into an operational model to make forecasts for global weather. To make these forecasts, initial conditions representing the best guess of the current atmospheric state must be produced. These initial conditions are called the ECMWF analysis and represent the combination of observations and the expected current state according to the ECMWF weather model, the integrated forecasting system, IFS.

The analysis should not be confused for reanalysis. Reanalysis is the recalculation of atmospheric conditions using a consistent model (as opposed to the often upgraded model used to produce analysis) and is often at a lower resolution than analysis. The ECMWF analysis for 2020 is at a high resolution of $1/12^\circ$ or approximately 9 km. By contrast, the most recent ECMWF reanalysis, ERA5, is at $1/4^\circ$ or 31 km ([Hersbach et al., 2020](#)).

It should also be stressed that observations and analysis are different things. Analysis is a combination of simulation and observational data. When we refer to observations in this thesis, we exclusively refer to aircraft, radiosonde, satellite or other direct observational datasets.

1.2.2 Motivation for investigating the spectrum

Four practical motivations are typically given for understanding the spectrum—predictability, parametrisation, model evaluation and understanding mesoscale dynamics.

Predictability

The kinetic energy power spectrum can be linked to the growth of error in turbulent systems. This was first shown in [Lorenz \(1969\)](#) using a turbulence closure model, an argument we show in a didactic manner in Chapter Six. The logic behind such a connection is that if the error is initially concentrated at small scales, it may grow by infecting the scale of motion twice as large and then twice as large again and so on. The timescale for each infection to take place may then be proportional to the eddy turnover time at the relevant scale. This eddy turnover time will be directly related to the kinetic energy at that scale. We explain this dimensional argument more completely later in the thesis, but it is sufficient to see for now that such a connection is plausible. It is well known that errors in weather predictions grow inhomogeneously and are state-dependent. In this thesis, we investigate how the kinetic energy - and so implicitly, the eddy turnover

times - vary under different conditions and locations. This may have a correspondence with error growth in some scenarios, this is discussed further in Chapter Seven.

Parametrisation

Any simulation of the atmosphere will not resolve all the scales of motion. As such, the effect of these smaller scales on the flow must be parametrised. One clear instance where our methodology, which can reveal local fluxes of energy, could be useful, is parametrisations of the flux of energy into and out of the small scales (e.g. [Shutts \(2005\)](#)). In Chapter Four, we show the first local atmospheric flux maps in the mesoscales produced with high resolution analysis data; these can likely be useful in informing the design of future parametrisation schemes.

Model evaluation

The spectrum is often used to evaluate the performance of a model. For example, through intercomparison of model spectra ([Li et al., 2023](#)), comparison against observations ([Dueben et al., 2020](#)) or through the concept of effective resolution ([Skamarock, 2004](#); [Bolgiani et al., 2022](#)). The local and conditioned spectra presented in this thesis in Chapter Three could be used to deepen these comparisons. For instance, instead of comparing global spectra between models, spectra under different specific conditions, such as orographic height or convective precipitation, could be compared. This may reveal unnoticed biases between models; for instance, such a method could reveal an under-representation of mesoscale energy in regions of intense convective precipitation.

Mesoscale dynamics

Possibly, the most important use of spectra has been for understanding how the mesoscales in the atmosphere operate. This has been a long process, some landmarks being ([Nastrom et al., 1984](#); [Cho and Lindborg, 2001](#); [Callies et al., 2014](#)). Further details are given in this chapter's section on explanations for the spectrum. In Chapter Three, we show local maps of spectral slope, which support the explanation that spectra in the lower stratosphere are energised primarily by vertically propagating gravity waves. In Chapters Four and Five, by studying local energy fluxes and the Helmholtz decomposition of our analysis data, we show that the discarded two-dimensional turbulence theory for the mesoscales could still be a good explanation for mesoscale motion under certain conditions.

1.2.3 The spectrum in analysis

Figure 1.1 shows the spherical harmonic power spectrum (the spherical counterpart to the Cartesian Fourier spectrum reviewed in Appendix A) for the ECMWF 2020 analysis. Darker shades of blue show pressure levels higher in the atmosphere, more details are given shortly. It is typical for power spectra to be plotted with a wavenumber on the x-axis, and so length-scale decreases from left to right in the figure. There are several features of interest.

At the largest scales, the spectrum is relatively flat, this is called the condensate. This term refers to a region of the spectrum where the accumulation of energy (from an inverse transfer) results in a flattened spectrum. The nomenclature is borrowed from quantum mechanics; in a Bose-Einstein condensate, many particles can accumulate in a single quantum state. Similarly, in this case, much of the energy is accumulating in a few Fourier modes ([Kraichnan, 1967](#)). In the atmosphere, the condensate is associated with the Rhines scale; changing planetary vorticity with latitude leads to anisotropy in the flow, causing the emergence of

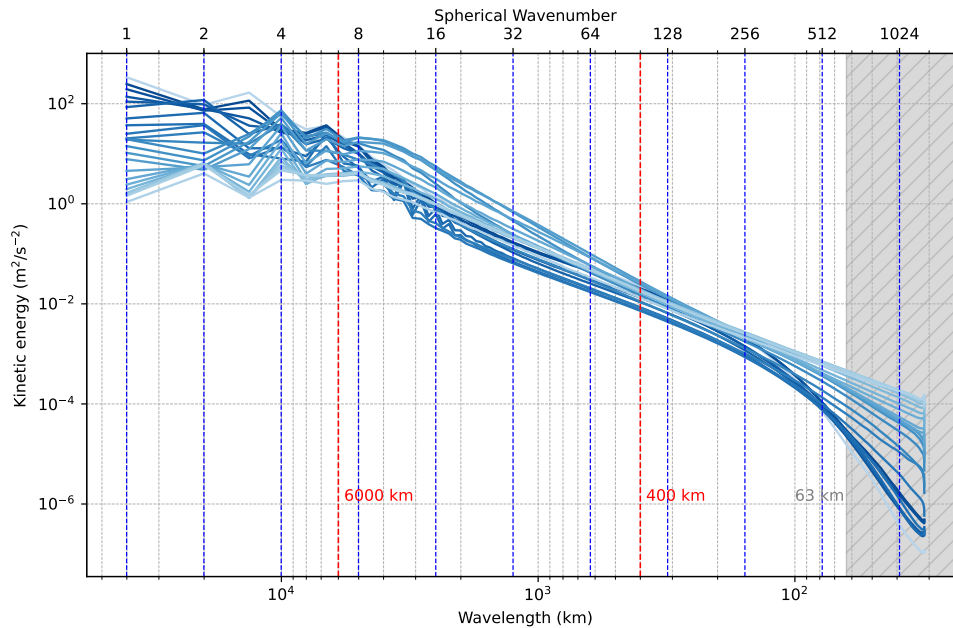


Figure 1.1: The spherical harmonic power spectrum at 25 pressure levels from the 2020 ECMWF analysis. Darker shades of blue show pressure levels higher in the atmosphere. This figure shows the general shape of the spectrum in the atmosphere. There is a large-scale flattening called the condensate, a range below this that is smoother and pressure level dependent and a steepening of the spectrum at the smallest scales in many cases. This figure demonstrates the differences in the magnitude of the global spectrum with altitude. Spectral slopes calculated over the spectral ranges denoted by blue dashed lines are shown in Figure 1.2. Red dashed lines denote 6000 km and 400 km, the boundaries of the synoptic scales and the mesoscales. The greyed-out region shows the estimated effective resolution using the heuristic presented in [Skamarock \(2004\)](#).

jets. Spectrally, the energy in the meridional component of the wind drops precipitously with scales greater than the Rhines scale, while the zonal component continues to rise. Overall, this results in a flatter portion of the spectrum at the largest scales ([Rhines, 1975](#)).

Between 6000 km and 400 km, the synoptic scales, the spectrum steepens. How steep the spectrum is here and the amount of energy it contains is heavily dependent on the pressure level. In the mesoscales the behaviour is again heavily pressure level dependent. After providing a little more background, we will discuss these features in detail.

In the upper troposphere and lower stratosphere (pressure levels between 300 and 100 hPa), observations suggest the spectrum follows a -3 spectral slope in the synoptic scales and a $-5/3$ slope in the mesoscales ([Nastrom et al., 1984](#); [Lindborg, 1999](#)) (see Figure 1.3 and later discussion). For this reason, it is common to plot reference lines with these gradients on spectral plots like Figure 1.1. We will resist this temptation and instead show the spectral slopes the spectra actually follow over power two wavenumber octaves denoted by blue dashed lines in Figure 1.2. Observations of the spectrum in the troposphere are sparse ([Frehlich and Sharman, 2010](#)), and I know of no observational spectra in the upper stratosphere. There is no reason to impose a -3 or $-5/3$ slope on the data. Since the shape of the spectra at different pressure levels is shown in Figure 1.2, the purpose of Figure 1.1 is to gain comfort with the shape of spectra and understand the relative magnitude of spectra at different pressure levels.

Since the IFS model is spectral, spherical harmonic coefficients are the native format variables are stored in, so producing the global spectra shown in Figure 1.1 is relatively straightforward. In 2020, IFS operated on

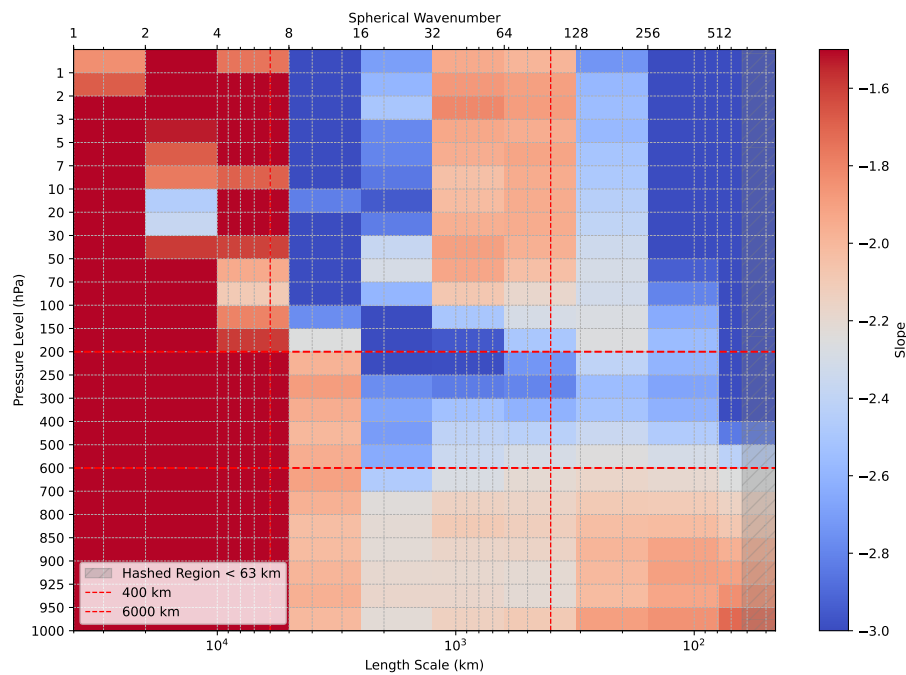


Figure 1.2: Best fit slopes for the spherical harmonic power spectrum of 2020 ECMWF analysis. These are calculated by finding lines of best fit to the spectrum shown in Figure 1.1 over power two wavenumber intervals denoted by dashed blue lines in that figure. Vertical dashed red lines show 6000 and 400 km, the boundaries of the synoptic scales and the mesoscales. Horizontal dashed lines show the two pressure levels we focus on in Chapters 3, 4 and 5 of this thesis, 200 and 600 hPa. We are able to see that a mesoscale slope of $-5/3$ and a synoptic slope of -3 are oversimplistic descriptions of the data. In the troposphere, at altitudes below 600 hPa, the spectrum never steepens to a value comparable to -3 . In the upper troposphere and lower stratosphere, this description holds better in that the synoptic scales are bluer. The blue mesoscales likely reflect underrepresented divergent motion in the IFS analysis since, in observations, the spectrum is much closer to $-5/3$ here. At lower pressure levels, the transition from a steep to shallow slope occurs at larger scales, as noted in Burgess et al. (2013). Due to the shallowness of the condensate, it appears as a red block on the left of the figure. The estimated effective resolution is shown hashed here as in Figure 1.1.

137 vertical levels ranging from 1000 to 0.01 hPa. This data has been interpolated onto 25 pressure levels between 1 and 1000 hPa. The power spectra of these 25 levels are plotted with darker shades of blue for higher altitudes.

Small scale steepening

A common heuristic for the effective resolution of an atmospheric simulation is $7\Delta x$, that is, only scales greater than seven times larger than the grid size are likely to be well resolved (Skamarock, 2004). In Figure 1.1 we have greyed out scales smaller than 63 km (7×9 km). It is interesting to note that this effective resolution was originally defined in terms of when the spectral slope begins to steepen. Visually, we can see that this steepening occurs at different scales at different altitudes, and at the lowest altitudes, it does not occur at all. It is customary to assume that steepening before $7\Delta x$ is due to numerical dissipation. However, given the shallow spectra at lower altitudes, the low energy at high altitudes is probably due to a lack of vertically propagating gravity waves, perhaps due to their absorption by critical layers (Malardel and Wedi, 2016; Avalos, 2024). We will provide more background on explanations for the spectrum's shape later in this chapter. The key point is that identifying the length scale at which the spectrum deviates from a $-5/3$ slope as the effective resolution is flawed for assessing the upper stratosphere. This is because the mesoscale

spectrum at these altitudes is not known from observations, this analysis data is our best estimate for the atmospheric state there. Hence, if we wish to evaluate the effective resolution at these higher altitudes, we should instead define the spectra of some suitable higher-resolution simulation or dataset to be our truth and then define the scale at which the spectra deviate from that truth to be the effective resolution. This extended concept of effective resolution is outlined further in the Future Work section of Chapter Ten.

The canonical spectrum

We show the classical mesoscale delineation line of 400 km and the synoptic scale of 6000 km. It is visually made more clear in Figure 1.2 that the shape of the spectrum in the upper troposphere and lower stratosphere (UTLS) does undergo a change around 400 km. Focussing on 200 hPa, the coloured blocks directly above the upper dashed red line in Figure 1.2, we can see shallower white slopes in the mesoscales transition into deeper blue slopes in the synoptic scales. However, it is also clear from Figure 1.1 that this change in shape is not universal to all altitudes.

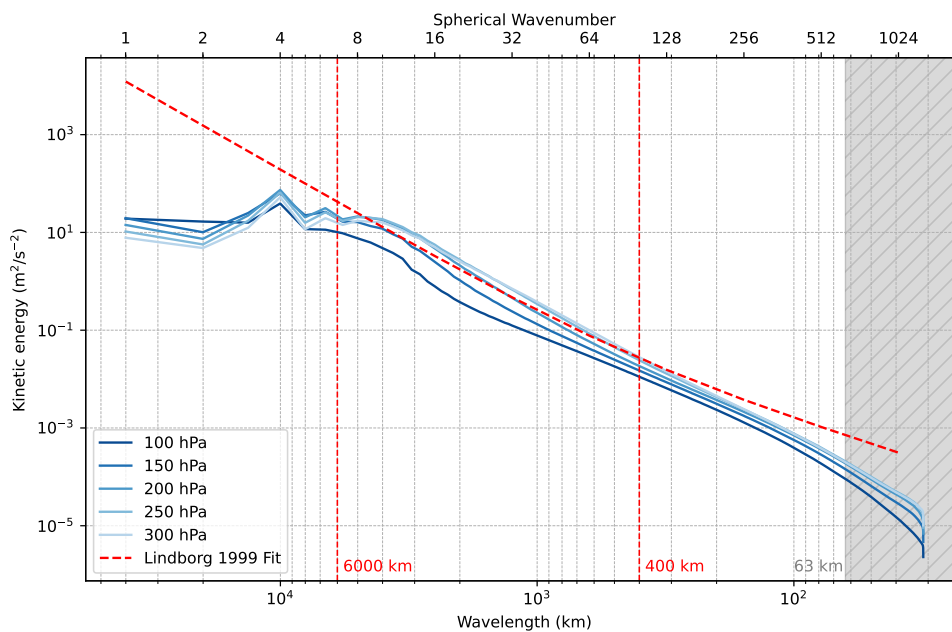


Figure 1.3: This figure shows the empirical fit to aircraft observations presented in Lindborg (1999). This was derived from a separate dataset but agrees well with the Nastrom et al. (1984) aircraft observations. The empirical fit follows a -3 spectral slope that transitions to a $-5/3$ slope in the mesoscales. Flight tracks at the very longest scales are nonexistent in the dataset, leading to a large deviation between the fit and the condensate in analysis. The analysis spectrum at 100, 150, 200, 250 and 300 hPa are plotted since these are the only pressure levels that could be expected to be comparable to these observations. This is because the cruising altitude of the commercial airliners used to gather the data lies within this range of pressure levels. We see that mesoscale shallowing is underrepresented in the analysis but note that Figure 1.2 shows us that some level of shallowing does occur at each pressure level as the length scale decreases, although after this there is a severe steepening which we speculated on when discussing Figure 1.1.

It is here that I must introduce the typical or canonical understanding of the atmospheric spectrum. In Figure 1.3 we have plotted an empirical fit to observations of the spectrum from aircraft campaigns (Lindborg, 1999). The fit shows that in the mesoscales, a $-5/3$ slope prevails while in the synoptic scales, there is a -3 slope. This canonical view of the spectrum became widely accepted due to the work of Nastrom et al. (1984).

Nastrom et al. (1984) is a landmark analysis of commercial airline wind speed flight data obtained through

the Global Atmospheric Sampling Program (GASP). Since commercial airliners move much more quickly than almost any atmospheric phenomena, there was no need to employ a frozen turbulence Taylor hypothesis¹, as had complicated measurements of the spectrum made via radiosondes or fixed devices (see discussion in [Gage \(1979\)](#); [Nastrom et al. \(1984\)](#)). This made their analysis decisive. Their analysis measured a range of scales between 10,000 and 1 km. They found a clear $-5/3$ slope below around 400 km that transitioned to a -3 slope at larger scales. However, it is key to note that these measurements are all at the cruising altitude of commercial aircraft, which means between 300 and 200 hPa. They are also concentrated in the northern hemisphere extratropics. The results have been confirmed with another similar aircraft dataset in [Lindborg \(1999\)](#), the empirical fit to these observations, which matches well with those in [Nastrom et al. \(1984\)](#), is plotted in Figure 1.3.

By analysing structure functions of wind measurements from the global rawinsonde network and from Aircraft Communications, Addressing, and Reporting System/Aircraft Meteorological Data Relay (ACARS/AM-DAR) data, [Frehlich and Sharman \(2010\)](#) extended the analysis to altitudes as low as 700 hPa. Their analysis shows that the lowest altitude does not undergo a transition and remains shallow until the condensate, as it does in the ECMWF analysis (Figures 1.1 and 1.2).

1.2.4 The slope in analysis

Inspired by the results of [Nastrom et al. \(1984\)](#) and [Lindborg \(1999\)](#), it has been common for studies to assume a $-5/3$ slope must prevail below a certain scale to be replaced by one closer to -3 at larger scales ([Frehlich and Sharman, 2010](#); [Dueben et al., 2020](#)). This has led to simply providing a $-5/3$ and -3 reference slope near spectra to become standard practice when presenting results from analysis or simulations. Considering that some theories predict a -2 slope in the mesoscales ([Kafiabad et al., 2019](#)), which is not easily discernable from $-5/3$ by eye, this practice is less than ideal. In Figure 1.2, we show the best-fit slopes for power two octaves of wavenumber. These octaves are shown in Figure 1.1 as vertical dashed blue lines. This is a novel way to present the atmospheric spectrum. We will now describe the features of this figure.

With the new representation of the spectrum in Figure 1.2, some features become immediately apparent to us. First, focusing on the UTLS (200 hPa has been picked out with a red horizontal line) one sees that here the mesoscale to synoptic transition to a steeper slope occurs at a smaller scale than anywhere else in the atmosphere. This was first noted in [Burgess et al. \(2013\)](#). It is also clear that although the analysis does undergo a shallowing in the mesoscales, this is still a steeper slope than is seen in observations. That IFS tends to produce slightly steeper slopes than observations is known, see for example [Stephan et al. \(2022\)](#). One can also note that the scale at which the condensate begins to become important is smaller than at higher altitudes but seems quite consistent throughout the troposphere and lower stratosphere.

In the lower troposphere, we see that although the significant steepening seen at higher altitudes is absent, there is still a steepening at larger scales, which seems to be consistent with the observations of [Frehlich and Sharman \(2010\)](#) (although, in their dataset, there is a strong latitudinal effect). In the upper atmosphere (pressure levels below 200 hPa), we see a shallow red region starting from several thousand kilometres and ending at around 400 km. At these altitudes, the steep blue interval has reduced to about a decade of wavenumbers compared to lower altitudes. In the mesoscales, the slope at these altitudes steepens

¹The Taylor hypothesis is the assumption that turbulent structures are fixed or frozen into the background flow. This allows one to draw an equivalence between measuring turbulence at a given location at an interval of one second, and measuring the turbulence at the same instant 10 meters apart, given the background flow speed is 10 m/s. Hence, with this assumption, one can convert temporal into spatial spectra. The validity of this approach in the atmosphere was contested, although it seems the spectral slopes measured in this manner broadly align with the later $-5/3$ results.

sharply, perhaps due to less gravity wave activity as we speculated earlier based on the analysis of [Malardel and Wedi \(2016\)](#) and [Avalos \(2024\)](#). Bearing in mind a bias towards steeper slopes and given the sparsity of spectrum observations of the upper atmosphere, Figure 1.2 provides a qualitative overview of our current knowledge of the shape of the atmosphere's kinetic energy power spectrum.

Limitations of the ECMWF Integrated Forecasting System

Several limitations of the ECMWF IFS must be mentioned. First, the model is hydrostatic. This means it does not represent some high-frequency gravity waves. This is the least serious of the limitations for our upper mesoscale focus, since the general consensus is that non-hydrostatic effects only become dominant when horizontal resolutions exceed 10 km ([Wedi et al., 2010](#)). The two more serious limitations are that the model uses a semi-implicit, semi-Lagrangian time-stepping scheme and that gravity waves are reduced in the data assimilation. A semi-Lagrangian time-stepping scheme trades numerical stability for altered gravity waves with reduced phase speeds and amplitudes. In particular, see the reduced gravity wave activity in [Preusse et al. \(2014\)](#). In addition, to not degrade forecast skill, the initial state found via data assimilation has deliberately low levels of unbalanced gravity-wave energy. In the 4D-var data assimilation system employed in 2020, this is achieved using a digital filter initialization (DFI). The details of this scheme for controlling the amount of gravity waves can be found in the documentation ([Yessad, 2019](#)). However, despite this, we found that spectra calculated for a forecast lead time of ten days and those calculated using the analysis produced qualitatively similar versions of Figures 1.1 and 1.2, demonstrating that data assimilation was not dominating the spectral characteristics of the analysis data.

1.2.5 Variations in time and location

Having encountered the shape of the average global spectrum at different altitudes as it is currently understood, we focus on local spectra. Local spectra, for example, the spectra measured over tall orography or in the presence of convective precipitation, can teach us about the way phenomena interact with the flow, help us improve parametrisations (for example, by quantifying how orography redistributes kinetic energy to different scales) and may teach us about how these phenomena affect the predictability of the weather. We begin with the current understanding of how the spectrum changes with location.

Location variations

Latitudinal variations have been noted in observations ([Frehlich and Sharman, 2010](#)), full physics simulations ([Judt, 2018](#)) and in reanalysis ([Bolgiani et al., 2022](#)). In general, mesoscale slopes appear to be shallower in the tropics and spectra less energetic there. The first attempt to measure spectra for each longitude and latitude point was made in [Wang and Sardeshmukh \(2021\)](#). Unfortunately, this used a spherical harmonic technique which produces inconsistent results and ERA5, which does not reproduce the mesoscale shallowing likely due to insufficient resolution. In this thesis coarse-graining and ECMWF analysis are used to show significant spatial variation in mesoscale slopes, particularly corresponding to orography and precipitation (see Chapter Three). We also show that the shallowing in the tropics is associated with precipitation and that shallow slopes do not extend to dry regions in the tropics.

Orographic variations

Variations of the spectrum with orography have been noted in observations, in particular great increases in energy in the UTLS over the Rockies are noted in [Nastrom et al. \(1987\)](#). Removing orography from models was seen to have an extreme effect on the spectrum ([Malardel and Wedi, 2016](#)). Similar variations can

be seen in reanalysis (Wang and Sardeshmukh, 2021). In this thesis, we show that this is also the case in analysis.

Variations with precipitations

To my knowledge, there is no analysis of direct observational data comparing spectra for precipitating and dry conditions. In a high-resolution local model over Germany, Selz et al. (2019) found that spectra calculated along lines that intersected precipitating regions generally had a shallower mesoscale slope than those that were dry. In Chapter Three, this finding is replicated but by using a coarse-graining technique outlined in Chapter Two.

Variations with pressure

Vonich and Hakim (2018) analysed many aircraft missions and found that higher-category storms exhibited steeper mesoscale spectra. By conditioning on mean sea level pressure in Chapter Three, we show that this result is replicated in the ECMWF analysis.

1.3 Explaining the shape of the spectrum

We now give an overview of the various theories for explaining why shallow or steep spectral slopes should occur at all in the atmosphere. Explaining the atmospheric spectrum's shape has been an active area of research for over fifty years. The most promising current theories for the mesoscales involve the vertical propagation of gravity waves in the upper atmosphere and the interaction between waves and turbulence in the lower atmosphere. However, many interesting ideas have been presented over the years.

In Chapter Three, we will show there is a great deal of location dependence for the shape of the spectrum. For many of the following ideas, even if their relevance in explaining the global mesoscale spectrum has been debunked, they may still be a dominant explanation at a particular time or location. By understanding these explanations, we can infer when particular mesoscale dynamics may be present. In fact, in Chapters Four and Five, we show that it is likely that under conditions of strongly rotational flow, turbulent explanations can emerge as dominant for explaining measured spectral slopes.

We can split explanations for the atmospheric spectrum into purely turbulent and wave explanations.

1.3.1 Purely turbulent explanations

-3 as a direct 2D turbulence enstrophy cascade

Turbulence and spectral slopes became deeply linked when Kolmogorov (1941) showed that in homogenous isotropic three-dimensional turbulence, the power spectrum should display a $-5/3$ slope according to dimensional arguments. If one assumes that there is a constant spectral flux of kinetic energy through a range of scales then one can show that a $-5/3$ spectral slope must present itself. In two-dimensional turbulence, it was found that kinetic energy moves to larger, not smaller, scales through the merger of vortices, and it is enstrophy (or the local rotation of the fluid) that moves to smaller scales. Extending Kolmogorov (1941) to two-dimensional turbulence, Kraichnan (1967) predicted a -3 slope for scales smaller than an energy injection and a $-5/3$ slope at larger scales. This spectrum is shown in schematic form in Figure 1.4. This theory is covered in detail in Part Three of this thesis.

The -3 in Figure 1.4 is widely understood to be related to the steep slope in the synoptic atmosphere.

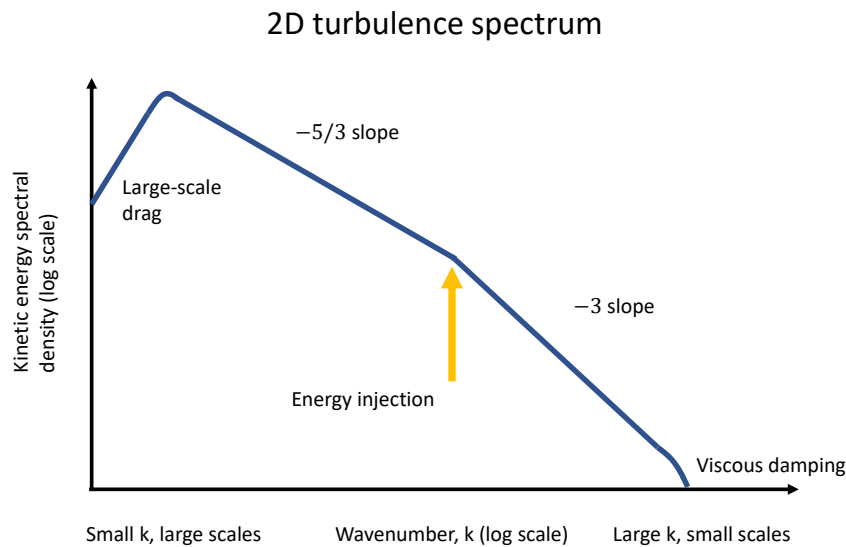


Figure 1.4: A schematic for the kinetic energy power spectrum in two-dimensional turbulence. At smaller scales than energy injection, an enstrophy direct cascade inertial range with a -3 spectral slope forms. This -3 slope extends until viscous effects become important at the Kolmogorov scale. At scales larger than the injection of energy, the merger of vortices leads to an inverse energy cascade inertial range with a slope of $-5/3$. This range ends when the effects of large-scale drag become important. Such a drag must be present as the flow is assumed to be in equilibrium. Details of the theoretical derivation of this picture are found in Chapter Eight.

Charney (1971) showed this by adapting the arguments of Kraichnan (1967) to the quasi-geostrophic equations. The quasi-geostrophic equations are the result of accounting for the next higher order terms in geostrophic balance, allowing for the flow to evolve in time. These equations are a good model for synoptic scale flow in the atmosphere. The Charney (1971) prediction of a -3 slope is explained in more detail in Chapter Eight. His explanation is still generally accepted.

It has been noted that in simpler models without the nonlinearities required for a cascade, a -3 slope can still emerge (O’Gorman and Schneider, 2007). However, since O’Gorman and Schneider (2007) fails to explain how a -3 slope would emerge for the rotational component of the flow (Callies et al., 2014), the utility of the Charney (1971) explanation remains.

-3 as an observational error (Aside)

The proposition that the -3 slope in the atmosphere is an artefact of how the observational data was processed is sometimes cited as illustrative of how controversial the spectrum can be. That the -3 slope was an issue of interpreting aircraft data was proposed by Lovejoy et al. (2009), specifically that the aircraft follow not paths of constant height but rather of constant pressure and that this led to an erroneous -3 regime. However, this proposal contradicts observational and model-based evidence (Frehlich and Sharman, 2010; Skamarock et al., 2014).

-5/3 from a direct 3D turbulent cascade

At the smallest scales in the atmosphere, no one contests that a $-5/3$ spectral slope should appear as a result of a direct cascade of energy from large to small scales, as in three-dimensional turbulence theory (Kolmogorov, 1941). A dimensional argument showing how a $-5/3$ slope arises in this context is given in Part

Three. This has never been proposed as an explanation for the shallow mesoscales, however. This is because the atmosphere is a stratified and rotating fluid. These factors make the atmosphere only modelable as 3D turbulence at scales as small as 100 meters, with the exception of highly energetic events such as storms, where scales could perhaps reach as large as 10 km (Skamarock, 2004; Gage, 1979).

-5/3 as a 2D inverse cascade

The most popular early theory for the mesoscale -5/3 slope was that it was due to an inverse cascade of energy as in Kraichnan (1967). This was proposed by Gage (1979), who gave small-scale convection at around 1 km as the source of the small-scale energy. Lilly (1983) clarified how convective systems may inject energy by analogy with wake turbulence in 3D stratified fluids. Since the theory of Kraichnan (1967) maintained that energy and enstrophy could not cascade through the other's inertial ranges, Gage (1979) concluded that there must be a spectral sink somewhere between 400 and 1000 km in the atmosphere. That no mechanism for such a sink seemed possible was a major detraction for the theory until Lilly (1989) showed in a turbulence closure model that two inertial ranges could coexist at the same scales with little interaction between them, a result confirmed in simulations by Maltrud and Vallis (1991). However, this theory fell out of favour when it was shown that the transfer of energy in observations of the mesoscales was from large to small, not small to large scales (Cho and Lindborg, 2001). It will be shown that, by creating spatial maps of spectral kinetic energy flux, a downscale transfer does not prevail in the mesoscales everywhere on Earth (Chapter Four) and the possibility that in certain situations, an inverse cascade exists can be resurrected.

-5/3 as a 2D direct cascade

Although the 2D Navier-Stokes equations produce a -3 slope at scales smaller than energy injection, there are other two-dimensional fluid equations relevant to the atmosphere. One of these is the surface quasi-geostrophic equations which produce a -5/3 slope below an injection of energy. This was proposed to explain the mesoscale slope by Tulloch and Smith (2006, 2009). However, these equations are horizontally divergence-free and so fail to explain the observational results of Callies et al. (2014), which showed the UTLS spectrum to be dominated by divergent motion - waves rather than turbulence - in the mesoscales. This essentially kills any two-dimensional turbulent explanation for the UTLS spectrum since these theories are horizontally divergence-free. In Chapter Five, we discuss how this situation is location-dependent, and show that under conditions of strongly rotational flow, the two-dimensional turbulence phenomenology may be present.

1.3.2 Wave explanations

Shallow slopes due to direct forcing from gravity waves

That gravity waves could be responsible for the atmospheric mesoscale shallowing was first proposed in Dewan (1979) and expanded upon by VanZandt (1982). The idea is that features such as orography, convection and instabilities generate gravity waves at a range of scales, energising the mesoscales and so causing the shallowing we observe.

In the form of vertically propagating gravity waves, this is currently the dominant theory for the mesoscale shallowing in the UTLS (Callies et al., 2014; Malardel and Wedi, 2016; Avalos, 2024). As opposed to a turbulent triad cascade as is present in the previous theories or the need for any kind of transfer of energy between horizontal scales, this explanation posits that energy is delivered directly to the shallower scales

from gravity wave activity lower in the atmosphere. These can be generated by orography, fronts, storms or other features. In Chapter Three, we note that our map of local spectral slopes at 200 hPa is similar to maps of gravity wave activity, lending support to this theory.

Shallow slopes due to a cascade of gravity waves

Rather than the result of forcing at multiple scales, some theories hold a cascade of waves to smaller scales to be centrally important. [Lelong and Riley \(1991\)](#) showed that triad interactions could catalyse energy transfer between wave modes, i.e. that gravity waves could undergo a cascade within a turbulent flow. This was expanded on by [Bartello \(1995\)](#). Such transfers do not seem to be dominant in analysis or simulation in the UTLS, but may be central in the troposphere ([Malardel and Wedi, 2016](#); [Avalos, 2024](#)).

Shallow slopes due to a stratified turbulent cascade that also produces gravity waves

In stratified turbulence, depending on the Froude number and possibly the presence of rotation, turbulent cascades that resemble either two-dimensional or three-dimensional turbulence can occur along with the generation of gravity waves. This has been proposed as a mechanism for the mesoscale shallowing ([Lindborg, 2006](#); [Bartello and Tobias, 2013](#)). This explanation has the disadvantage that simulations with vertical resolution much less than those required by the theory have reproduced the $-5/3$ spectral slope ([Skamarock et al., 2014](#)). However, as the resolution of simulations continues to increase, such as those in [Dueben et al. \(2020\)](#), the mechanism or nature of the spectral slopes in the mesoscales may shift due to the influence of stratified turbulence. It would be interesting to revisit the analysis performed in this thesis for a global model that resolves stratified turbulence, although since it is estimated that a resolution on the order of 100 meters is required for such a simulation ([Lindborg, 2006](#)), this should not be possible for 10-15 years, assuming improvements in computational science continue at their current pace.

-2 due to the scattering of gravity waves by slower turbulent features

[Kafiabad et al. \(2019\)](#) showed that gravity waves scattered by stationary turbulent features produce a -2 spectral slope. This has been extended to slowly moving turbulence in [Cox et al. \(2023\)](#) where it was shown to still hold.

Shallowing due to direct imprint through unspecified turbulent or wave means

Related to the idea of a mixed generation of gravity waves, one can consider that forcing on the atmosphere is not localised to a particular wavenumber but rather spread out across a range of scales, often obeying a power law. An example is orography, which obeys a $-5/3$ spectrum ([Skamarock, 2004](#)). When the spectrum of orography was better represented, [Malardel and Wedi \(2016\)](#) found that a shallower spectrum was observed.

A mix of explanations

It has been conjectured that a mix of mechanisms is responsible for the shallowing ([Selz et al., 2019](#)). Since the energy-weighted sum of local spectra generated by coarse-graining is the global spectrum, the methodology we present in this thesis is uniquely positioned to help investigate the mix of mechanisms hypothesis. In Chapter Three, we will explore the effects on the total spectrum in the analysis of excluding spectra from points of high orography and precipitation, thus quantifying their contribution to the global spectrum.

1.4 Summary

In this Chapter, we have summarised the current understanding of the distribution of energy among scales in the atmosphere. We have reviewed the observational evidence and show slopes found in the ECMWF analysis for 2020 using a novel visualisation of the data (Figure 1.2). We emphasised that the Nastrom and Gage spectrum, with its mesoscale slope of $-5/3$ and its steep synoptic slope of -3 with a transition at a wavelength of around 400 km is actually specific to the upper troposphere and lower stratosphere. The values of the shallow and steep slope and where one transitions into the other, can vary with altitude, location and the prevailing atmospheric conditions.

We then presented an overview of the theories that have been proposed over the past 55 years for the shape of the spectrum, including the reasons why solely turbulent theories fell out of favour. For explaining the shallowing in the upper atmosphere, the current frontrunner is vertically propagating gravity waves. For the troposphere, some level of turbulence-wave interaction is assumed to be responsible. However, given the variability in the spectrum in space and time explored in this thesis, we have emphasised the possibility that many of these explanations can coexist under different conditions or locations.

Chapter 2

Coarse-graining

2.1 Introduction

This chapter introduces a method for obtaining local power spectra in the atmosphere. As we saw in Chapter One, the power spectrum of the atmosphere varies at different locations and under different conditions. Therefore, a consistent way to probe the spectrum while retaining spatial information is valuable for investigating analysis or simulation data. In this thesis, we will employ coarse-graining to achieve this end. The reader may be familiar with wavelets, which serve the same function. Wavelets are a subclass of coarse-graining (Sadek and Aluie, 2018). For the full background on spherical vector coarse-graining, one should refer to Storer et al. (2022); Buzzicotti et al. (2023); Storer et al. (2023); Aluie et al. (2018); Aluie (2019) and the references therein.

In this chapter, we will give a general overview of the method, including how it was extended to local spectral slopes and conditioned spectra for the work in Chapters Three, Four and Five. We will explain the benefits and limitations of the method, particularly the concept of order and how this prevents the method from evaluating spectral slopes outside of a certain range and the bias towards steeper spectral slopes exhibited by the method. We have endeavoured to present concepts as clearly as possible, filling in many steps left unexplained in the places this material is presented in the literature. In particular, we will explain why a Helmholtz decomposition is necessary more concisely than is currently available in the literature and show a full derivation of the local spectral flux.

The reader is assumed to be familiar with the Fourier transform and its extension to the sphere, the spherical harmonic transform. A refresher on Fourier techniques can be found in Appendix A. If the reader wishes to use the coarse-graining methodology described in this chapter on a dataset of their own, they are invited to use the open-source software FlowSieve (Storer and Aluie, 2023).

2.2 Extracting spectra via coarse-graining

This section mainly follows Sadek and Aluie (2018).

Let us consider a turbulent 2D velocity field. Our goal will be to extract the spectrum from this via coarse-graining. This is possible since when we coarse-grain (read locally average or smooth) at a particular scale, we discard energy from motions at smaller scales. For example, suppose there was a wave of wavelength 10 meters in the flow, if we were to coarse-grain with a kernel - or weighting function for the local average - that had a length scale of 100 meters, then we would completely blur out the 10 meter wave. Thus, the kinetic energy associated with that wave would be removed.

Generally, coarse-graining at a scale ℓ removes the energy associated with motion at scales smaller than ℓ . Therefore, if E_ℓ is the kinetic energy of the flow coarse-grained at scale ℓ , then this is the cumulative energy of all scales larger than ℓ , so the energy at scale ℓ should be attainable by taking a derivative, $-\frac{\partial}{\partial \ell} E_\ell$.

We will now formalise these ideas.

The definition of coarse-graining

Restricting ourselves to cartesian space of dimension d , \mathbb{R}^d , the coarse-grained velocity is given by

$$\bar{\mathbf{u}}_\ell(\mathbf{x}) = \int_{\mathbb{R}^d} G_\ell(\mathbf{r}) \mathbf{u}(\mathbf{x} + \mathbf{r}) \, d\mathbf{r}, \quad (2.2.0.1)$$

where G_ℓ is the coarse-graining kernel of scale ℓ . This is a convolution, $\bar{\mathbf{u}}_\ell(\mathbf{x}) = (G_\ell * \mathbf{u})(\mathbf{x})$. We are locally averaging \mathbf{u} according to the kernel, or weighting function, G_ℓ .

Convolution and local averaging

The definition of the convolution is

$$f * g(x) = \int_{-\infty}^{\infty} f(y)g(x - y) \, dy. \quad (2.2.0.2)$$

If one examines this integral and pictures what happens as y varies for a given x , one sees that this is locally averaging using a weighting function $g(-x)$, which is the mirror image of g . We can see this if $x = 0$ as

$$f * g(0) = \int_{-\infty}^{\infty} f(y)g(-y) \, dy, \quad (2.2.0.3)$$

which is the average of $f(x)$ weighted by $g(-x)$. Since we will only consider radially symmetric kernels, we can think of convolution and local averaging as synonyms.

Properties of the kernel

Our kernel at a particular scale will be obtained by rescaling a mother kernel G . In particular,

$$G_\ell(\mathbf{r}) := \ell^{-d} G\left(\frac{\mathbf{r}}{\ell}\right). \quad (2.2.0.4)$$

We will impose a few properties on our mother kernel that its rescaled children will inherit. It will be normalised so that

$$\int_{\mathbb{R}^d} G(\mathbf{r}) \, d^d \mathbf{r} = 1, \quad (2.2.0.5)$$

we assume it decays faster than any polynomial for large \mathbf{r} , which will ensure the existence of derivatives of its Fourier transform, and we will assume that its main support is within a region of size unity (such that the rescaled kernels have support mainly in a region of size ℓ). We will also assume that the kernel is radially symmetric.

Power spectrum of filtered fields

According to the strategy we outlined earlier for finding the spectrum, our first step must be to find the kinetic energy spectrum of the filtered flow. We have

$$\widehat{\mathbf{u}}_\ell(\mathbf{x}) = (G_\ell * \widehat{\mathbf{u}})(\mathbf{x}) \quad \text{definition, hat denotes Fourier transform} \quad (2.2.0.6)$$

$$= \widehat{G_\ell(\mathbf{x})\mathbf{u}(\mathbf{x})} \quad \text{convolution theorem (Appendix A)} \quad (2.2.0.7)$$

$$= \widehat{G}(\ell\mathbf{k})\widehat{\mathbf{u}}(\mathbf{k}). \quad \text{scaling property of FT} \quad (2.2.0.8)$$

Therefore, the power spectrum of the filtered fields are

$$E_\ell(\mathbf{k}) = |\widehat{G}(\ell\mathbf{k})|^2 E(\mathbf{k}). \quad (2.2.0.9)$$

To understand the filtered spectrum, consider the kernel (left) and its Fourier transform (right) shown in Figure 2.1. For small wavenumbers, the Fourier transform of the kernel will be close to one (this is precisely what makes it a good weighting function), so for large-scales, $E_\ell(\mathbf{k})$ matches $E(\mathbf{k})$. For small scales, the value is close to zero, and the energy at these scales is removed. This is the property of local averaging mentioned in the introduction, that it removes the energy associated with small-scale features.

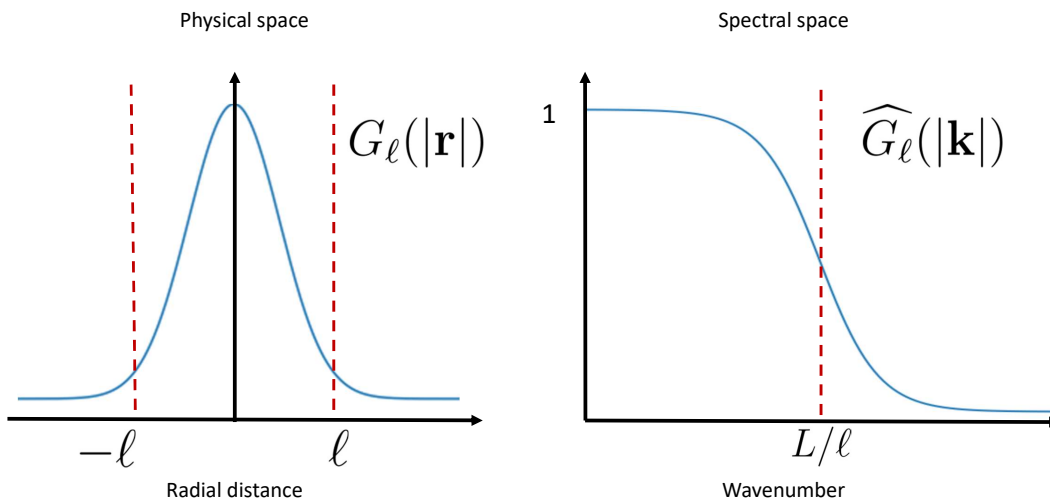


Figure 2.1: A schematic of the coarse-graining kernel, showing a basic kernel in real and spectral space. Real space is shown on the left, spectral space on the right. The filtering length-scale is given by ℓ , in physical space the main support is within a radius ℓ of the origin. In spectral space, this corresponds to the main support being for wavenumbers less than the wavenumber corresponding to ℓ , L/ℓ where L is the outer scale of the system.

Coarse-grain kinetic energy vs coarse-graining the kinetic energy

An easy misunderstanding is that coarse-graining the kinetic energy of a flow directly and finding the kinetic energy of the coarse-grained velocity field will give equivalent results. This is not the case. The reason is kinetic energy is a nonlinear function of the velocity.

Consider a simple 1-dimensional flow with $u = \sin(Ax) + \sin(Bx)$. If we coarse-grain or Fourier transform we will have large coefficients around wavenumbers $k = A$ and $k = B$. When we square this to get the power spectrum (Appendix A) we will still have energy at only these two scales. However, the following formula shows that the kinetic energy will have more length-scales than the underlying flow, leading to spurious lengthscales if we coarse-grain this field:

$$u^2 = 2 \sin(Ax) \sin(Bx) + \dots = \cos((A - B)x) - \cos((A + B)x) + \dots \quad (2.2.0.10)$$

We must, therefore, coarse-grain the velocity fields and then calculate the kinetic energy from these to find the kinetic energy at the different length scales of the flow.

The coarse-grain spectrum

We are now in a position to define the coarse-grain power spectrum. We define the coarse-grain/filter spectrum as

$$\bar{E}(k_\ell) := \frac{d}{dk_\ell} \frac{\langle |\bar{\mathbf{u}}_\ell(\mathbf{x})|^2 \rangle}{2}. \quad (2.2.0.11)$$

The angle brackets denote a spatial average and k_ℓ refers to the wavenumber associated with lengthscale ℓ , which is $k_\ell = L/\ell$ if L is the outer scale. As we have just shown, the total kinetic energy of the velocity coarse-grained at scale ℓ contains almost all the energy up to that scale and almost none of the energy after it. Thus, if we coarse-grain at two different scales, ℓ_1 and ℓ_2 , we can subtract their total kinetic energies to deduce the kinetic energy contained in the scales between the two length scales. Taken to the logical conclusion, this means we can reconstruct the entire power spectrum by taking infinitesimal steps in coarse-graining scale and subtracting the resulting total energies, which is precisely how Formula 2.2.0.11 operates.

The coarse-grain spectrum as a function of the Fourier spectrum

We can write down the relationship between the spectrum we have just defined and the Fourier spectrum explicitly. Restricting ourselves to one dimension for simplicity, we have

$$\bar{E}(k_\ell) = \frac{d}{dk_\ell} \frac{\langle |\bar{u}_\ell(x)|^2 \rangle}{2} \quad \text{definition} \quad (2.2.0.12)$$

$$= \frac{d}{dk_\ell} \int_0^\infty \left| \widehat{G}\left(\frac{L}{k_\ell}k\right) \right|^2 E(k) dk \quad \text{integral of spectrum is energy, } \ell = L/k_\ell \quad (2.2.0.13)$$

$$= \int_0^\infty \frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell}k\right) \right|^2 E(k) dk. \quad (2.2.0.14)$$

Note that in [Sadek and Aluie \(2018\)](#) the factor L is omitted.

It requires a little thought to see what $\frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell} k\right) \right|^2$ looks like. First, we are taking the derivative with respect to k_ℓ , that is the cutoff scale in Figure 2.1, not the x-axis k in that figure. Now, suppose we increase the cutoff scale infinitesimally by δ in that figure. For scales much smaller or greater than k_ℓ the kernel does not change much, it is still very close to one or zero. For scales close to the increased cutoff, the value of \widehat{G}_ℓ will increase, increasing most at $k_\ell + \delta$. Hence, $\frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell} k\right) \right|^2$ will actually be a bump centered on k_ℓ . This is the key concept to understand the coarse-grain spectrum: the coarse grain spectrum is the local average of the Fourier spectrum in spectral space.

Furthermore, as k_ℓ increases, $\frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell} k\right) \right|^2$ is dilating. That is, we are averaging over a wider weighting function. There is a stronger blur on spectral information from small-scales than large-scales. Note that this may not be apparent on a traditional plot of the power spectrum since it is standard to use a log-log axis, meaning despite there being a greater blur, the resolution may appear sharper for higher wavenumbers.

2.3 Order, or when coarse-graining fails

In the previous section, we defined a coarse-graining spectrum and then discovered that this spectrum amounted to locally averaging the Fourier spectrum in spectral space.

We are interested in turbulent flow fields that usually follow a power law-like structure. Suppose the spectral slope is very steep. Then there is much more energy at the large than the small scales. If we blur this, at some point, the energy that is leaked into the small scales from the larger ones will be greater than the energy actually at those scales; they will be swamped out. At that point, we will measure an anomalously shallow slope due to this blurring effect, and our coarse-graining spectrum will be misleading.

This blurring will always be present to some extent but we now quantify when a slope cannot be replicated by the coarse-graining spectrum and how we can alleviate the problem by choosing a kernel of sufficiently high “order”. We will find that if one uses the most basic kernels, such as the top hat or Gaussian, spectral slopes steeper than -3 and perhaps shallower than -1 are not reproducible. In practice the lower limit is far more of an issue than the upper limit. For the same underlying reasons, this is also the case when extracting the spectrum using structure functions (Appendix A).

After discussing order, we will explain how, even within the range of slopes we can resolve, there will be a bias towards steeper slopes in the filter spectrum than in the Fourier spectrum.

The order of a kernel

The order of a kernel is how many vanishing moments it has. In particular, the order p of a kernel G is the largest p such that (again remaining in 1-dimension for simplicity)

$$\int_{-\infty}^{\infty} G(x)x^n dx = 0 \quad \forall n \leq p. \quad (2.3.0.1)$$

Here n is a non-zero positive integer.

Note that our kernels are always even functions. Since the symmetric integral about zero of any even function multiplied by an odd function is always zero, the order of our kernel will always be odd as any odd moment of the kernel is zero.

2.3.1 The scaling behaviour of the coarse-grain spectrum

Now let's find how the coarse-grain spectrum scales with wavenumber in the limit of large wavenumber. We have

$$\bar{E}(k_\ell) = \int_0^\infty \frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell}k\right) \right|^2 E(k) dk \quad \text{Equation 2.2.0.14}$$

(2.3.1.1)

$$= \underbrace{\int_0^{k_\alpha} \frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell}k\right) \right|^2 E(k) dk}_{\text{Term 1}} + C \underbrace{\int_{k_\alpha}^\infty \frac{d}{dk_\ell} \left| \widehat{G}\left(\frac{L}{k_\ell}k\right) \right|^2}_{\text{Term 2}} k^\beta dk. \quad \text{splitting integral, } E(k) = Ck^\beta$$

(2.3.1.2)

We will now find the scaling behaviour of Term 1 and Term 2.

Term 2

From this point, we define the scale $s = L \frac{k}{k_\ell}$. Substituting s into Term 2, we have

$$\text{Term 2} = Ck_\ell^\beta \int_{L \frac{k_\alpha}{k_\ell}}^\infty \frac{d}{ds} \left| \widehat{G}(s) \right|^2 \left(\frac{s}{L}\right)^{\beta+1} ds \quad (2.3.1.3)$$

$$\propto k_\ell^\beta. \quad (2.3.1.4)$$

Hence, if the integral converges and Term 2 is dominant for large wavenumbers, then $\bar{E}(k_\ell) \propto k_\ell^\beta$. That is the coarse-grain slope is the same as the Fourier slope for large wavenumbers. We will discuss convergence a little more after we look at Term 1.

Term 1

In Term 1, we see the importance of the kernel's order. Our strategy will be to Taylor expand the Fourier transform of the kernel around wavenumber zero. In particular,

$$\widehat{G}^{(n)}(k)|_{k=0} = \int_{-\infty}^\infty G(x)x^n dx \quad (2.3.1.5)$$

and so derivatives up to the order p of the kernel will be zero. The Taylor expansion is

$$\widehat{G}(k) = \widehat{G}(0) + k\widehat{G}^{(1)}(0) + k^2 \frac{\widehat{G}^{(2)}(0)}{2!} + \dots \quad \text{Taylor expansion} \quad (2.3.1.6)$$

$$= 1 + k^{p+1} \left(\frac{\widehat{G}^{(p+1)}(0)}{(p+1)!} + k^2 \frac{\widehat{G}^{(p+3)}(0)}{(p+3)!} + \dots \right) \quad \text{normalisation, order of kernel, kernel is even} \quad (2.3.1.7)$$

$$= 1 + k^{p+1} \phi(k). \quad \text{this defines } \phi(k) \quad (2.3.1.8)$$

Now we examine the term $\frac{d}{dk_\ell} \left| \widehat{G} \left(\frac{L}{k_\ell} k \right) \right|^2$. We have

$$\frac{d}{dk_\ell} \left| \widehat{G} \left(\frac{L}{k_\ell} k \right) \right|^2 = -\frac{s}{k_\ell} \frac{d}{ds} |\widehat{G}(s)|^2 \quad \text{substitution} \quad (2.3.1.9)$$

$$= -\frac{2s}{k_\ell} \widehat{G}(s) \frac{d}{ds} \widehat{G}(s) \quad \text{chain rule} \quad (2.3.1.10)$$

$$= -\frac{2s}{k_\ell} (1 + s^{p+1} \phi(s)) ((p+1)s^p \phi(s) + s^{p+1} \phi'(s)) \quad \text{Taylor expansion, derivative} \quad (2.3.1.11)$$

$$= -\frac{2s}{k_\ell} ((p+1)s^p \phi(s) + s^{p+1} \phi'(s)) \quad \text{expanding brackets} \quad (2.3.1.12)$$

$$+ (p+1)s^{2p+1} \phi(s)^2 + s^{2p+2} \phi(s) \phi'(s).$$

Since s goes like k_ℓ^{-1} while ϕ is a constant to leading order in s and ϕ' goes at least like s^2 , we have four terms that go like $k_\ell^{-(p+2)}$, $k_\ell^{-(p+5)}$, $k_\ell^{-(2p+3)}$ and $k_\ell^{-(2p+5)}$ respectively.

Hence, examining Equation 2.3.1.2, we can see that for large k_ℓ ,

$$\text{Term 1} \propto k_\ell^{-(p+2)}. \quad (2.3.1.13)$$

Steep slope failure

We can therefore conclude that we can only resolve a slope β if $\beta > -(p+2)$, as otherwise, it will be drowned out by energy smeared from the larger to the smaller scales.

In particular, let us consider the most standard kernels we may use, the top hat and Gaussian. These are both order one. We can see this since, as even functions, the integral of their product with x must be zero. However, as functions that are positive everywhere, the integral of their product with x^2 cannot be zero. Hence these common kernels cannot reproduce slopes steeper than -3.

Shallow slope failure

Although [Sadek and Aluie \(2018\)](#) concerns itself only with the boundary of resolving steep spectral slopes, it seems to the author that a similar boundary should exist for shallow or positively sloped spectra. This should be the case since if the spectral slope is sufficiently shallow, the integral in Equation 2.3.1.3 should not converge.

We will not endeavour to prove this conjecture, but since the integrand is just a local average of k^β it stands to reason that as k_ℓ tends to infinity,

$$\int_{k_\alpha}^{\infty} \frac{d}{dk_\ell} \left| \widehat{G} \left(\frac{L}{k_\ell} k \right) \right|^2 k^\beta dk \quad (2.3.1.14)$$

should have the same convergence properties as

$$\int_{k_a}^{\infty} k^\beta dk. \quad (2.3.1.15)$$

That is, the integral will not converge if $\beta > -1$.

Slope limits in practice

In practice, the steep slope constraint is the stronger of the two limits. If one generates a field with a spectral slope of -5 and then calculates the coarse-grain spectral slope using a kernel of order one, a -3 slope will be calculated. However, the shallow limit is less of an issue. Since the divergence of the integral arises from arbitrarily large values of k , one is still able to compute the filter spectrum and may see slopes quite consistent with those from Fourier analysis.

2.3.2 A bias towards steeper slopes

The above results were concerned only with slopes at the limit of infinite wavenumber. In practice, we see important biases when using the coarse-graining method at more prosaic wavenumbers, which we now present.

Consider a steep spectral slope being coarse-grained using a kernel of sufficiently high order to resolve it. Assuming the derivative of the kernel with respect to the wavenumber cutoff is symmetric in spectral space, the locally averaged value at a particular wavenumber is going to be dominated by the values at smaller wavenumbers. This implies greater values of spectral energy density at those scales than in the Fourier spectrum. However, we know from the previous analysis that at the large wavenumber limit, the slope must approach that of the Fourier spectrum. Hence, the slope must steepen at these intermediate scales.

Figure 2.2 compares the spherical harmonic and filtering spectrum for the 2020 ECMWF analysis. In our analysis, we resolved the filter spectrum up to a filtering scale of 1000 km. We can see the relative steepening of the spectral slope compared to the Fourier spectrum. Also, note that there is no reason for the filter scale and the corresponding wavenumber to coincide exactly.

Overall, however, Figure 2.2 demonstrates a good agreement between the two methods, showing that general trends in slope are indeed reproduced, although there is a slight steepening bias present. The resolved mesoscales in the operational analysis are indicated with red dashed lines.

It may be interesting to reapproach the spectral slope analysis in a more spectrally local manner, such as by approximating $\frac{d}{dk_\ell} \left| \widehat{G} \left(\frac{L}{k_\ell} k \right) \right|^2$ as a Gaussian. This may allow us to quantify the steepening seen in Figure 2.2. We will leave this for future work.

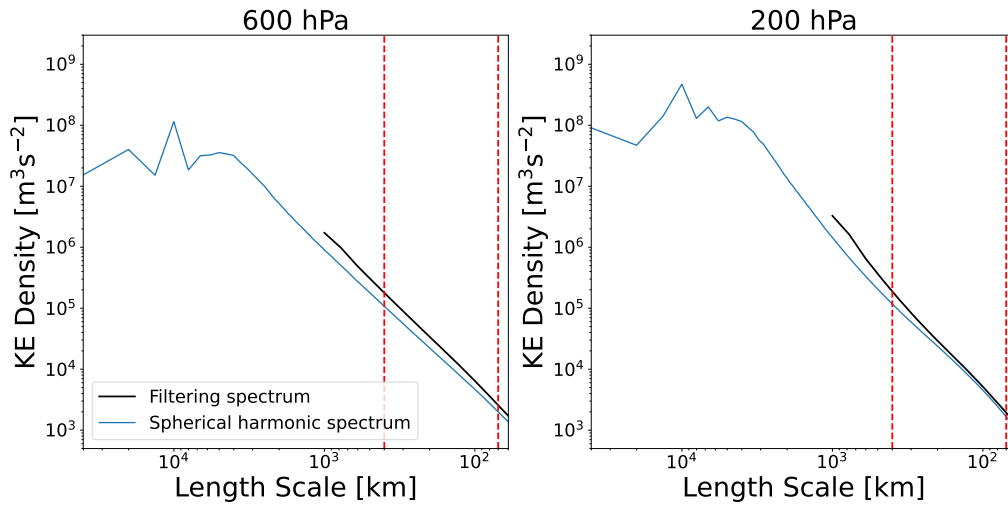


Figure 2.2: The average spherical harmonic spectrum for 2020 compared with the average filtering spectrum for 600 hPa (left) and 200 hPa (right). The filtering spectrum is consistently steeper than the spherical harmonic spectrum due to the blurring of large-scale energy into the smaller scales by coarse-graining. Since the coarse-graining technique is energy preserving, if extended to the largest scales, the filtering spectrum would have the same total area under the curve as the spherical harmonic spectrum.

2.4 Local spectral slopes

We may extend this coarse-grain spectrum to a local coarse-grain spectrum. To do so we simply avoid taking a spatial average and retain spatial information

$$\bar{E}(k_\ell, \mathbf{x}, t) := \frac{\partial}{\partial k_\ell} \frac{|\bar{\mathbf{u}}_\ell(\mathbf{x}, t)|^2}{2}. \quad (2.4.0.1)$$

Conditioned spectra

If we aggregate spectra for points that satisfy particular conditions, we have what we define as a “conditioned spectra”. This means that the conditioned spectrum represents the average spectrum centred on a point that satisfies a particular condition.

Contribution statement:

I proposed this method and calculated all the conditioned spectra in this thesis by postprocessing fields I generated by running FlowSieve.

Conditioned fluxes

Although we have not yet introduced local fluxes, it is clear that a similar procedure can be used to obtain conditioned fluxes.

Contribution statement:

I proposed this method and calculated all the conditioned fluxes in this work by postprocessing fields I generated by running FlowSieve.

Local spectral slopes

Given that we have now defined a power spectrum at each time and location, we can find the spectral slope of that spectrum. This is found with the formula

$$\text{Spectral Slope} = \frac{\partial}{\partial \log k_\ell} (\log \overline{E}) (k_\ell, \mathbf{x}, t) \quad (2.4.0.2)$$

$$= k \frac{\frac{\partial^2}{\partial k_\ell^2} |\overline{\mathbf{u}}_\ell|}{\frac{\partial}{\partial k_\ell} |\overline{\mathbf{u}}_\ell|}. \quad (2.4.0.3)$$

Contribution statement:

I proposed this method. I computed all spectral slopes of conditioned spectra in this thesis by post-processing fields I generated with Flowsieve. Spatial maps of spectral slope in this thesis were generated by averaging and postprocessing spectral slopes directly output by Flowsieve. My method for generating slopes was improved and implemented into the code by the Flowsieve team.

Aggregating the spectral slope

There are several ways we could aggregate slopes calculated in the above manner. For instance, we could simply take the mean of their values, use an energy-weighted mean, use the median, or use an energy-weighted median. These four possibilities are shown in Figure 2.3 aggregating in time as a function of spatial location.

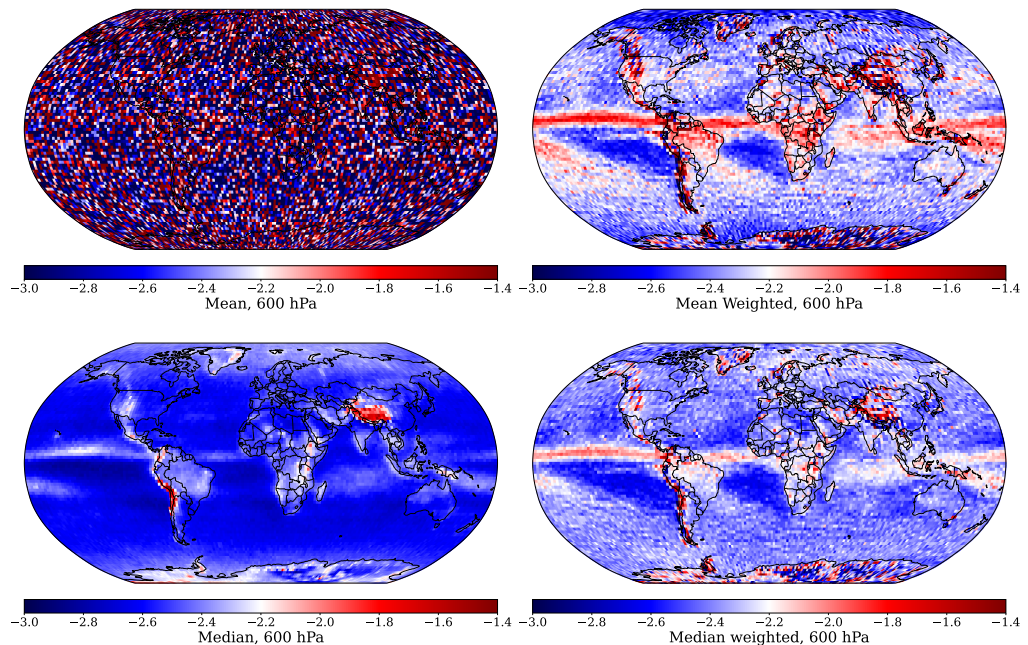


Figure 2.3: Four methods for aggregating local spectral slopes are demonstrated using the 2020 ECMWF analysis at 600 hPa. This is a result presented in Chapter Three. The top left panel shows the result of taking a mean in time of the spectral slopes calculated at each location and time in ECMWF analysis. The top right panel shows the energy-weighted mean (the method used in this thesis). The bottom left panel shows the median. The bottom right panel shows the energy-weighted median.

The raw mean is equivalent to measuring the spectral slope at each location and then averaging those

slopes. Due to low-energy events with extreme spectral slopes, the mean is incredibly noisy. These outliers can be dealt with by weighting according to energy or taking a median. Taking a mean weighted by energy is equivalent to calculating the average spectrum at a location of interest for the period of interest and then finding the slope of that spectrum. Precisely, we weigh each slope at each scale by the corresponding kinetic energy density at that scale. Figure 2.3 is then produced by finding the average of these weighted slopes over the mesoscales. This is the preferred method that we follow.

The median produces a qualitatively similar figure but with a bias towards steeper spectral slopes when compared to the energy-weighted slopes. This means that shallower slopes are produced by more intermittent high mesoscale energy events. The quantitative differences are reduced when the energy-weighted median is used. The exception to the quantitative picture being equivalent is the poles. Here, it is likely that the steeper slopes seen in the weighted mean are due to more intermittent events with high energy on larger scales.

Negative energy

When considered globally, the coarse-grained energy will only decrease with increasing filter length-scale, as we exclude the energy in more scales. However, locally this does not need to be the case. If, in a particular location, there is very little energy, then when blurring over a larger area, the energy associated with that point goes up. In this scenario, the derivative will be negative, and so we will assign a negative energy to that scale at that location. This is, of course, an artefact of the technique.

This can often occur if one considers a particular time slice and location in the atmosphere, but when averaged sufficiently, either spatially or temporally, negative energy becomes atypical in our application. Partially in the interests of computational memory requirements but also to aid in sufficiently averaging the spectral information, means were taken over spatial areas by coarsening from a $1/12^\circ \times 1/12^\circ$ to $2^\circ \times 2^\circ$ resolution (for spectral slopes, these means are energy weighted).

Our kernel

The kernel we use in this work is third order. That means it can resolve slopes as steep as -5. In particular, our coarse-graining operation of a scalar on a sphere (as we will discuss in the next section, this is the only coarse-graining procedure we need to explicitly define) is given by

$$\bar{f}(\mathbf{x}, t) = \int_{\mathbf{y} \in \Omega} f(\mathbf{y}, t) G(\gamma(\mathbf{x}, \mathbf{y})) dA, \quad (2.4.0.4)$$

where f is a scalar field on the sphere Ω , coarse-grained f is denoted by \bar{f} and $\gamma(\mathbf{x}, \mathbf{y})$ is the geodesic distance on the sphere with G the filtering kernel.

That kernel is defined by

$$G(\delta) = \frac{1}{A} \left[\frac{1}{2} (1 - \tanh(10D(\delta))) - c \cdot \exp(-(2D(\delta))^2) \right], \quad (2.4.0.5)$$

where A, c are normalization factors and $D(\delta) = \delta/(\ell/2) - 1$. This is the ‘‘high order’’ kernel implemented in FlowSieve (Storer and Aluie, 2023).

A note on an alternative approach

In [Wang and Sardeshmukh \(2021\)](#) local maps of spectral slopes are presented. The method employed in that study is to sum up spherical harmonics with length scales corresponding to the mesoscales, i.e. those with degrees corresponding to wavelengths within a particular range. They do this for two ranges, one corresponding to between 400 and 625 km and another to between 250 and 400 km. The authors treated these fields as local kinetic energy at these scales. They then assumed that this energy followed a power law relationship at each location and plotted that coefficient for the mesoscales by interpolating between the two points.

The fundamental problem with this approach is spherical harmonics are not spatially local. The situation is entirely analogous to Fourier analysis, to represent a bump localised at the origin with compact support, waves must be introduced at every location in the domain. There is no guarantee that the energy in the spherical harmonics at a particular range of scales corresponds to the local mesoscale energy there. This is seen in the noisy slope values presented in their study.

Another drawback was that they only performed the slope map analysis on ERA5, ERAI, and a low-resolution GFS and only at the 200 hPa pressure level. These datasets are not high enough resolution to resolve the mesoscale slope correctly ([Skamarock, 2004](#)), and indeed, it is known that the mesoscale shallowing is not well resolved in ERA5 ([Li et al., 2023](#)).

2.5 Extension to vector fields

In the next section, we will derive the coarse-grain spectral flux. The property that coarse-graining as an operation on vectors commutes with derivatives is key to that derivation.

This is simple for coarse-graining in Cartesian space, following from a basic property of convolutions. However, the situation on the sphere is more complex. [Aluie \(2019\)](#) explicitly defines a coarse-graining operation that convolves with the differential operators on a sphere in terms of the so-called Edmonds basis. However, in practice, this definition is never used in FlowSieve, which implements this spherical coarse-graining procedure ([Storer and Aluie, 2023](#)). The reason for this will soon become clear.

Extending derivative commuting coarse-graining to vectors

On the sphere, coarse-graining scalar fields is already a well-defined operation. That is, simply local averaging using geodesic distance (Equations 2.4.0.4 and 2.4.0.5).

Therefore, if a coarse-graining operation on vectors exists on the sphere that commutes with derivatives, we can already calculate its action on any vector field. This is because if we Helmholtz decompose \mathbf{u} into $\mathbf{u} = \nabla\Phi + \nabla^\perp\Psi$, we can filter the *scalars* Ψ and Φ to get $\bar{\mathbf{u}} := \nabla\bar{\Phi} + \nabla^\perp\bar{\Psi}$.

Hence, we can calculate coarse-grained vector fields by Helmholtz decomposing and then performing a scalar coarse-graining operation on each component.

Proving such an operation exists

We will simply point out here that this means the constructive proof in [Aluie \(2019\)](#) can be greatly simplified since all that is required in practice is proof of existence. We leave this simplification for future work.

Coarse-graining tensors

In the course of finding the local spectral fluxes in the next section we will be required to coarse-grain tensors. This requires a similar procedure to the Helmholtz decomposition previously outlined. However, unlike in the case of vectors, it was found that the expense of this decomposition did not justify the increased accuracy it provided to results (private communication, Benjamin Storer) and so in the implementation in FlowSieve at the time of writing, a shortcut is taken. Consult the documentation for further details on this point (Storer and Aluie, 2023).

2.6 The coarse-grain energy flux

Spectral fluxes can be a powerful analysis tool for designing parametrisations and deducing the dynamics of mesoscale weather. For instance, one could diagnose if an unknown two-dimensional fluid exhibited a $-5/3$ due to classical two-dimensional turbulence or surface quasi-geostrophic dynamics by analysing the direction of energy flux. As shown in Chapter Eight, classical two-dimensional turbulence is associated with an inverse cascade of kinetic energy to large scales, while surface quasi-geostrophic dynamics are associated with a direct cascade of energy to small scales. We now derive a formula for calculating the local spectral flux.

2.6.1 Deriving the kinetic energy budget for large scales

This section mostly presents material found in Aluie et al. (2018).

That our coarse-graining operation commutes with derivatives is useful because it allows us to manipulate the governing partial differential equations directly. This allows us to find a term Π_ℓ , the transfer of energy to smaller scales at each time and location in our flow.

We arrive at this term by first coarse-graining the Navier-Stokes equations. This results in a term that resembles a Reynolds stress, representing the nonlinear turbulent effects of small scales on larger scales.

Then, we will manipulate this coarsened Navier-Stokes equation into a kinetic energy budget for the large scales. That “subscale stress” term will transform into a term that tells us the flux of energy from smaller to larger scales at each point in the flow, Π_ℓ .

Coarse-graining the governing equations

Suppose our 3D equations of motion are the rotating Boussinesq equations

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho_0} \nabla p - f \wedge \mathbf{u} + \nu \Delta \mathbf{u} + \frac{\rho}{\rho_0} \mathbf{g} + \mathbf{F} \quad (2.6.1.1)$$

where f is the Coriolis frequency, ν is the viscosity, ρ_0 is the reference density, and \mathbf{F} is some forcing term.

We can filter this equation so that τ_ℓ represents applying a coarse-graining filter at scale ℓ . This means $\bar{\mathbf{u}}_\ell$ is the smoothed velocity with scales smaller than ℓ removed.

Since we have shown τ_ℓ commutes with derivatives, most terms remain the same in the coarse-grained equations of motion. For example, $\nu \Delta \mathbf{u}$ becomes $\nu \Delta \bar{\mathbf{u}}_\ell$. The exception is the nonlinear advection term.

Manipulating the advective derivative

It is convenient to write the advection term $(\mathbf{u} \cdot \nabla)\mathbf{u}$ as $\nabla \cdot (\mathbf{u}\mathbf{u})$, since it is then straightforward for us to apply our differentiation commuting coarse-graining to this. The product $\mathbf{u}\mathbf{u}$ is the outer product. This vector calculus identity is easier to see in Einstein notation. Then we have $\mathbf{u}\mathbf{u} = u_i u_j$ so

$$\begin{aligned}\nabla \cdot (\mathbf{u}\mathbf{u}) &= \partial_j (u_i u_j) \\ &= u_j \partial_j u_i + u_i \partial_j u_j \\ &= (\mathbf{u} \cdot \nabla)\mathbf{u} + \mathbf{u}(\nabla \cdot \mathbf{u}).\end{aligned}$$

We are concerned with incompressible flow (in 3D) since our speeds are generally much less than the speed of sound, so

$$\nabla \cdot (\mathbf{u}\mathbf{u}) = (\mathbf{u} \cdot \nabla)\mathbf{u}.$$

Since we want to write Equation 2.6.1.1 for the large-scale component only, we add and subtract $\overline{\mathbf{u}_\ell} \cdot \nabla \overline{\mathbf{u}_\ell}$ from the equation. Using our outer product form we can see that

$$\overline{(\mathbf{u} \cdot \nabla)\mathbf{u}_\ell} - (\overline{\mathbf{u}_\ell} \cdot \nabla)\overline{\mathbf{u}_\ell} = \nabla \cdot (\overline{\mathbf{u}\mathbf{u}_\ell} - \overline{\mathbf{u}_\ell} \overline{\mathbf{u}_\ell}).$$

Defining our sub-scale stress $\overline{\tau}_\ell$ to be

$$\overline{\tau}_\ell(\mathbf{u}, \mathbf{u}) := \overline{\mathbf{u}\mathbf{u}_\ell} - \overline{\mathbf{u}_\ell} \overline{\mathbf{u}_\ell},$$

we have

$$\overline{(\mathbf{u} \cdot \nabla)\mathbf{u}_\ell} - (\overline{\mathbf{u}_\ell} \cdot \nabla)\overline{\mathbf{u}_\ell} = \nabla \cdot \overline{\tau}_\ell(\mathbf{u}, \mathbf{u}).$$

Hence our coarse-grained equations of motion are

$$\frac{\partial \overline{\mathbf{u}_\ell}}{\partial t} + \overline{\mathbf{u}_\ell} \cdot \nabla \overline{\mathbf{u}_\ell} = -\nabla \cdot \overline{\tau}_\ell(\mathbf{u}, \mathbf{u}) - \frac{1}{\rho_0} \nabla \overline{p}_\ell - f \wedge \overline{\mathbf{u}_\ell} + \nu \Delta \overline{\mathbf{u}_\ell} + \frac{\overline{\rho}_\ell}{\rho_0} \mathbf{g} + \overline{\mathbf{F}}_\ell, \quad (2.6.1.2)$$

$$\nabla \cdot \overline{\mathbf{u}_\ell} = 0. \quad (2.6.1.3)$$

Deriving the Kinetic Energy (KE) budget for large scales

The large-scale kinetic energy budget is

$$\frac{\partial}{\partial t} \left(\rho_0 \frac{|\overline{\mathbf{u}_\ell}|^2}{2} \right)^{(G)} + \nabla \cdot \mathbf{J}_\ell = -\overline{\Pi}_\ell^{(R)} - \rho_0 \overline{\mathbf{u}_\ell} \cdot \nabla \overline{p}_\ell^{(G)} + \rho_0 \overline{\mathbf{u}_\ell} \cdot \overline{\mathbf{F}}_\ell^{(G)} + \overline{\mathbf{u}_\ell} \cdot \mathbf{g}^{(G)} + \overline{\mathbf{u}_\ell} \cdot \mathbf{f}_\ell^{(G)}, \quad (2.6.1.4)$$

where

$$\mathbf{J}_\ell = \rho_0 \frac{|\overline{\mathbf{u}_\ell}|^2}{2} \overline{\mathbf{u}_\ell}^{(B)} + \frac{\overline{p}_\ell}{\rho_0} \overline{\mathbf{u}_\ell}^{(G)} - \rho_0 \nu \nabla \frac{|\overline{\mathbf{u}_\ell}|^2}{2}^{(G)} - \rho_0 \overline{\mathbf{u}_\ell} \cdot \overline{\tau}_\ell(\mathbf{u}, \mathbf{u})^{(R)},$$

and

$$\Pi_\ell^{(R)} = -\rho_0 \overline{S_\ell}^{(R)} : \overline{\tau}_\ell(\mathbf{u}, \mathbf{u}).$$

$\overline{S_\ell}$ is the large scale strain tensor,

$$\overline{S_\ell}^{(R)} = \frac{1}{2} (\nabla \overline{\mathbf{u}}_\ell + \nabla \overline{\mathbf{u}}_\ell^T).$$

Equation 2.6.1.4 is derived by performing $\rho_0 \overline{\mathbf{u}}_\ell \cdot$ (Equation 2.6.1.2). The derivation requires some vector calculus; we start by showing how the blue term (B) is derived using Einstein notation and then indicate how to get the green terms (G) similarly. Finally, we show how the most important term (R), the subscale flux Π_ℓ , is found.

Deriving the blue term

We show that

$$\rho_0 \overline{\mathbf{u}}_\ell \cdot (\overline{\mathbf{u}}_\ell \cdot \nabla) \overline{\mathbf{u}}_\ell = \nabla \cdot \left(\rho_0 \overline{\mathbf{u}}_\ell \frac{|\overline{\mathbf{u}}_\ell|^2}{2} \right).$$

This follows by considering the following in Einstein notation:

$$\begin{aligned} \mathbf{u} (\mathbf{u} \cdot \nabla) \mathbf{u} &= u_i u_j \partial_j u_i \\ &= \partial_j \left(u_j \frac{u_i u_i}{2} \right) - \frac{u_i u_i}{2} \partial_j u_j \\ &= \nabla \cdot \left(\mathbf{u} \frac{|\mathbf{u}|^2}{2} \right) - \frac{|\mathbf{u}|^2}{2} \nabla \cdot \mathbf{u} \\ &= \nabla \cdot \left(\mathbf{u} \frac{|\mathbf{u}|^2}{2} \right). \end{aligned}$$

Deriving the green terms

The green terms follow in the same way by deriving the following equalities:

$$\mathbf{u} \cdot \partial_t \mathbf{u} = \partial_t \frac{|\mathbf{u}|^2}{2}, \quad (2.6.1.5)$$

$$\mathbf{u} \cdot \nabla p = \nabla \cdot (p \mathbf{u}), \quad (2.6.1.6)$$

$$\mathbf{u} \cdot (f \wedge \mathbf{u}) = 0, \quad (2.6.1.7)$$

$$\mathbf{u} \Delta \mathbf{u} = \nabla \cdot \left(\nabla \frac{|\mathbf{u}|^2}{2} \right) - (\nabla \mathbf{u})^2. \quad (2.6.1.8)$$

Deriving the red terms

The red terms in equation 2.6.1.4 both come from $\rho_0 \overline{\mathbf{u}}_\ell \cdot \nabla \cdot \overline{\tau}_\ell(\mathbf{u}, \mathbf{u})$. In Einstein notation, we examine

$$u_j \partial_j \tau_{ij} = \partial_j (u_j \tau_{ij}) - \partial_i u_j \tau_{ij}.$$

We can separate the tensor $\partial_i u_j$ into a symmetric and antisymmetric part. The reason this is useful is that the contraction of a symmetric and an antisymmetric matrix is zero, and $\tau_{ij} = \overline{u_i u_j} - \overline{u_i} \overline{u_j}$ is symmetric.

Symmetric/Antisymmetric Contraction

If A is symmetric and B is antisymmetric, then

$$A_{ij} = A_{ji} \quad \text{and} \quad B_{ij} = -B_{ji}.$$

Since $A_{ij}B_{ij} = A_{ji}B_{ji}$ (we are simply summing over all indices, so the result is the same), and by symmetry $A_{ij}B_{ij} = -A_{ji}B_{ji}$, we have $A_{ij}B_{ij} = -A_{ij}B_{ij}$. Hence

$$A_{ij}B_{ij} = 0.$$

We can decompose $\partial_i u_j$ into a symmetric strain tensor

$$S = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$$

and an antisymmetric rotation tensor

$$\Omega = \frac{1}{2}(\partial_i u_j - \partial_j u_i).$$

This simplifies the equality to

$$\begin{aligned} u_j \partial_j \tau_{ij} &= \partial_j (u_j \tau_{ij}) - S_{ij} \tau_{ij} \\ &= \nabla \cdot (\mathbf{u} \cdot \boldsymbol{\tau}) - S : \boldsymbol{\tau}. \end{aligned} \tag{2.6.1.9}$$

Here $:$ is the contraction, so $S : \boldsymbol{\tau} = S_{ij} \tau_{ij}$.

We are then faced with an interesting question. The divergence part of Equation 2.6.1.9 can be taken either into the transport term on the left of Equation 2.6.1.4, or we could consider it part of the subscale flux since it is due to the subscale stress. As discussed in [Aluie et al. \(2018\)](#), several different definitions of the subscale flux exist in the literature that differ by the divergence of some term, essentially, by what choice the author made at this stage. In other words, there is a gauge freedom in the definition of Π_ℓ . Thankfully, we can rule out definitions that do not have Galilean invariance: the energy flux from small scales in the fluid should not depend on the velocity of an observer.

If we move to a reference frame moving at speed \mathbf{U} , then we have

$$\nabla \cdot (\mathbf{u} \cdot \boldsymbol{\tau}) \Rightarrow \nabla \cdot (\mathbf{U} \cdot \boldsymbol{\tau}) + \nabla \cdot (\mathbf{u} \cdot \boldsymbol{\tau}),$$

so the first term in 2.6.1.9 is not Galilean invariant. Since $\boldsymbol{\tau}$ and S are Galilean invariant (simply replace $\mathbf{u} \Rightarrow \mathbf{U} + \mathbf{u}$), the contraction $S : \boldsymbol{\tau}$ is also invariant. Hence we define

$$\boxed{\Pi_\ell := -\rho_0 \overline{S}_\ell : \overline{\boldsymbol{\tau}}_\ell(\mathbf{u}, \mathbf{u})}.$$

A note on flux and order

It is unclear how the spectral blurring inherent in the coarse-graining approach, which led to our consideration of order and resolvable spectral slopes, affects the spectral fluxes we derive. It is nontrivial to say how the term $\overline{S}_\ell : \overline{\tau}_\ell$ would change, since formally one may expect the magnitude of \overline{S}_ℓ to be underestimated and the value of $\overline{\tau}_\ell$ to be perhaps overestimated (since one roughly represents the energy in the large scales, the other the small scales and we assume the blurring results in some of the energy in the large scales appearing in the small scales). It should be possible to evaluate the bias, if it exists, by comparing the value of Π_ℓ found from coarse-graining with that found through a direct Fourier method as in [Capet et al. \(2008\)](#). We leave this investigation for future research.

Breaking down Π_ℓ

There are three velocities that appear in Π_ℓ : two from $\overline{\tau}_\ell(\mathbf{u}, \mathbf{u})$ and one from \overline{S}_ℓ . If we perform a Helmholtz decomposition of \mathbf{u} so $\mathbf{u} = \mathbf{D} + \mathbf{R}$, or a divergent and rotational component, then the subscale flux can be decomposed into

$$\begin{aligned} \Pi_\ell(\mathbf{u}, \mathbf{u}, \mathbf{u}) &= \Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R}) + \Pi_\ell(\mathbf{D}, \mathbf{D}, \mathbf{D}) \\ &\quad + \Pi_\ell(\mathbf{R}, \mathbf{D}, \mathbf{D}) + \Pi_\ell(\mathbf{D}, \mathbf{R}, \mathbf{D}) + \Pi_\ell(\mathbf{D}, \mathbf{D}, \mathbf{R}) \\ &\quad + \Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{D}) + \Pi_\ell(\mathbf{R}, \mathbf{D}, \mathbf{R}) + \Pi_\ell(\mathbf{D}, \mathbf{R}, \mathbf{R}). \end{aligned} \quad (2.6.1.10)$$

The purely rotational flux $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$ must account for the 2D turbulence upscale energy transfer. This is because the 2D turbulence theory is divergence-free, and so must be entirely rotational. Thus, if present, the associated transfers must be between only the rotational terms.

The purely divergent flux should account for much of the downscale energy transfer from wave-wave interactions. A particularly large term in many contexts is $\Pi_\ell(\mathbf{D}, \mathbf{R}, \mathbf{R})$, which accounts for the large-scale compression or expansion of vortices. In our geophysical context, most of the energy is contained in the rotational part of the flow. In general, this means the more rotational parts appear the larger the corresponding flux.

In the oceans, as shown in [Storer et al. \(2023\)](#), for scales greater than a few thousand kilometres, energy is moved between scales by the general circulation ($\Pi_\ell(\mathbf{D}, \mathbf{R}, \mathbf{R})$). For example, the horizontally divergent Ekman flow near the equator stretches vortices meridionally, causing the length scale of vortices to increase and so energy to be transferred to larger scales. This is called the ‘‘piston effect’’ in [Storer et al. \(2023\)](#). For smaller scales, down to the Rossby radius where the Coriolis force ceases to dominate (this ranges from 10 km at the poles to 100 km as the equator is approached), the transfer is dominated by $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$. There is a strong upscale transfer of energy here, which is interpreted to be 2D turbulence. At scales smaller than the Rossby radius, a downscale flux is seen. However, the resolution of the data prevented interpretation of this as 3D turbulence, wave dynamics or another mechanism ([Storer et al., 2023](#)).

In the atmosphere, preliminary results indicate that similar dynamics exist. Due to the larger scales in the atmosphere (due to the greater velocities), the large-scale circulation is less dominant than in the oceans, [private communication with Benjamin Storer]. However, the atmospheric mesoscales have not yet been probed in a similar way, again due to the resolution of the data analysed. In Chapter Four, we begin the

study of fluxes in the ECMWF analysis and show the first analysis of a relationship between the local energy spectrum and local energy flux.

2.7 Conclusions and overview

In this chapter, we have given an overview of the coarse-graining approach for obtaining local spectra, including an appraisal of the strengths and weaknesses of the approach.

We have shown how the coarse-graining spectrum is calculated, including the key intuition that it is a local averaging in spectral space of the Fourier spectrum. We have examined when this coarse-graining spectrum fails to replicate the slopes observed in the Fourier spectrum due to insufficient order. We have explained why a Helmholtz decomposition is employed to calculate coarse-grained vector fields and shown how to derive a local spectral flux.

In the course of this chapter we have made some new contributions to the discussion on this methodology. We extended the technique to examine local spectral slopes, discussing the merits of using an energy-weighted slope for our application. We have introduced conditioned spectra and fluxes for the first time. We have pointed out that not only a steep slope limit, but also a shallow slope limit exists for the range of slopes the method can resolve, conjecturing that this lower limit is $\beta = -1$. In addition, we have commented on the possibility of simplifying the proof of a derivative commuting coarse-graining operation found in [Aluie \(2019\)](#).

In the rest of Part One, we will explore what the techniques we have just described can reveal about the atmosphere as represented in the 2020 ECMWF analysis.

Chapter 3

Local mesoscale slopes in the atmosphere

In this Chapter, we present local mesoscale spectral slopes derived from ECMWF analysis. We show spatial inhomogeneity in these spectral slopes and how they vary under different conditions. Spectral slopes are one of three clues to the mesoscale dynamics that we investigate over the next three Chapters. The other two are spectral fluxes and the rotational-divergent decomposition. The discussion section of Chapter Five gives a summary of what our results imply about mesoscale dynamics.

3.1 Introduction

Motivation

As discussed in Chapter One, power spectra can provide clues to the underlying dynamics. In this chapter, we will find that different phenomena are associated with energising the mesoscales, leading to shallower spectral slopes. The next three chapters will mirror the order of discovery that occurred for the global atmospheric spectrum. We start by measuring spectral slopes in the mesoscales as was performed in [Nastrom et al. \(1984\)](#), then we quantify the direction of kinetic energy flux as in [Cho and Lindborg \(2001\)](#), and finally, we uncover the relative magnitudes of rotational and divergent motion as in [Callies et al. \(2014\)](#). Once we finish this journey, we will be able to confirm that the current global understanding of the spectrum holds, while presenting the possibility of turbulence theories being applicable to the mesoscales under certain conditions.

Apart from diagnosing mesoscale activity, this chapter will also provide useful insights for future model diagnostic methods. We can imagine that with a more complete picture of how the spectrum behaves under different conditions, we could build more precise tools for model validation. For example, [Bolgiani et al. \(2022\)](#) derive effective resolutions for the tropics, midlatitudes and poles from reanalysis data using Fourier techniques on periodic strips. By using truly local spectra, we could go beyond this, evaluating local effective resolution or effective resolution in the presence of particular conditions such as precipitation. This idea is discussed further in the future work section of Chapter Ten.

An exciting motivation for analysing local spectral slopes is their possible connection to predictability. It is important to be aware of some pitfalls to avoid drawing premature conclusions about predictability from the spectra and spectral slopes in this chapter. Lorenzian predictability theory is explained in detail in Chapter Six and conclusions one can draw about predictability from the results in this chapter are explored in Chapter Seven. For the moment, it is important to point out that shallower spectral slopes do not automatically imply less predictable flow. A dimensional model that grasps the essence of error growth in purely turbulent systems is presented in [Lilly \(1973\)](#). In that model, error starts at the smallest scales of the system and grows sequentially, in a cascade, to larger and larger scales. It fully infects a scale before growing to a larger one, like a Russian doll burning from the inside out. The time a scale takes to be saturated is related to its eddy turnover time, that is, the time it would take for a vortex at that scale to rotate once given the speed implied by the kinetic energy at that scale. Hence, by this argument's logic, it is the magnitude of the spectrum that relates to the eddy turnover time and so estimates of predictability, not the spectral slope.

The confusion arises because when dealing with an arbitrarily small error, the spectral slope determines if the flow is predictable arbitrarily far into the future or only until a particular time. However, this only holds if that slope extends to arbitrarily small scales and so is not relevant to the analysis we present here. See Chapter Six for a more complete explanation of this concept.

The coarse-graining methodology

In Chapter Two, we explained how a coarse-graining methodology can be used to derive local and conditioned spectra as well as local and conditioned spectral fluxes. In the rest of Part One, we will apply these techniques to ECMWF analysis for the year 2020 to see what insights and patterns we discover. To reproduce the analysis in the following chapters for a dataset of interest, one can use the open-source software FlowSieve ([Storer and Aluie, 2023](#)), consulting the documentation for that software found on GitHub.

Some details of our analysis

The ECMWF 2020 analysis consists of hourly means sampled four times per day at 00, 06, 12 and 18 UTC, we included all of these times in our analysis. This study was confined to the 200 and 600 hPa pressure levels. Precipitation and orography data were sourced from ERA5 for 2020 ([Hersbach et al., 2020](#)). ERA5 precipitation broadly agrees with observations. However, it has limitations in reproducing extreme events and is more accurate in the extratropics than the tropics ([Lavers et al., 2022](#)). It has the advantage of decomposing the precipitation into that arising from its convective and large-scale parametrisations, which we use later to investigate the spectra associated with these two conditions.

Although the ECMWF model changes frequently, in 2020, it only changed once in June. This upgrade improved stratospheric temperature biases by improving the data assimilation scheme and by improving vertical interpolation in the advection scheme, it also incorporated a more accurate surface albedo climatology. The impact of this change on our analysis is minimal with spectra broadly consistent before and after the upgrade.

We evaluated the spectrum, spectral slope and flux at 19 scales. In kilometres, these start at 10 km and increase by a factor of 1.3, with each increasing scale until 1000 km. This data was then coarsened from $1/12^\circ \times 1/12^\circ$ to $2^\circ \times 2^\circ$ resolution (this is done by an energy-weighted mean for spectral slopes, as outlined in Chapter Two). It is on these coarsened fields on which all further analysis is performed. At various points we average “over the mesoscales”. This means averaging over the scales 167 km, 215 km, 278 km and 359 km, which are the well-resolved scales within the mesoscales that we evaluated.

Bootstrapping statistics

We employ bootstrapping to obtain estimates of the standard deviation associated with our sampling of the spectrum, spectral slopes and fluxes. This is done by resampling the data with replacement, generating an array of equal size to the original data but potentially with some repeated entries. The procedure for finding the spectral slope (for instance) is then applied to each resampled dataset to give a particular value for the slope. Repeating this process sufficiently many times produces a distribution of results. As we generate more members, the distribution converges towards an estimate of the distribution of the slope we may have calculated over that year. After a parameter sweep, bootstrapping with 256 members was found to have converged sufficiently.

Outline of Chapter Three

We will present maps of local spectral slopes in the 2020 ECMWF analysis for 200 and 600 hPa. We will also present the average energy in the mesoscales at both pressure levels and the unfiltered average kinetic energy, as this gives the necessary context for interpreting local slopes. In addition, we will explore a range of conditioned spectra. We present the orography and ERA5 precipitation-conditioned spectra but also condition on whether the precipitation was classed as large-scale precipitation or convective precipitation in ERA5. Further, we condition on MSLP, producing results that align with those found in aircraft measurements by [Vonich and Hakim \(2018\)](#).

Contribution statement

Some of the work in this chapter has been published in [Kouhen et al. \(2024\)](#). That paper was a collaboration with the Flowsieve team. All results in this thesis are the result of my own postprocessing of coarse-grained fields I generated using the Flowsieve open source software.

3.2 Mesoscale spectral slopes across the globe

3.2.1 200 hPa, the upper troposphere, lower stratosphere

In Figure 3.1, we present the slopes averaged over the mesoscales at 200 hPa. This spans the upper troposphere and lower stratosphere (UTLS), and so is indicative of the pressure levels [Nastrom et al. \(1984\)](#) observed. In particular, in the tropics, 200 hPa is in the troposphere, in the extratropics (more than 30 North or South), it is within the stratosphere. In Figure 3.2 the ERA5 orographic and precipitation fields are provided for reference. Figure 3.3 and 3.4 show the average unfiltered kinetic energy at 200 hPa and the energy averaged over the mesoscales respectively.




Figure 3.1: A global map of 200 hPa mesoscale spectral slope in analysis data centered on a slope of -2.2. Shallow spectral slopes are red while steep ones are blue.

Shallow tropics is not a solely latitudinal effect but correlates with precipitation

The most striking feature of Figure 3.1 is the red slash painted across the equator. In Figure 3.2, the global

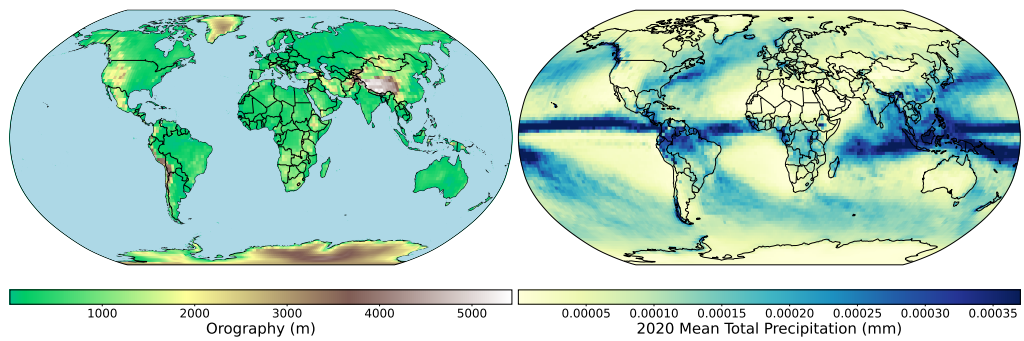


Figure 3.2: ERA5 average hourly precipitation for 2020 and an orography field for reference.




Figure 3.3: Global map of kinetic energy in analysis at 200 hPa, representing upper-tropospheric (within the tropics) and lower stratospheric levels (in the extratropics).

average precipitation traces a shape in the equator that is visually close to the red shape in Figure 3.1. It is interesting to note that although extratropical regions can have similar levels of rainfall to locations in the tropics, they do not correlate with similarly shallow slopes. This seems to be due to much larger scale features having more energy than in the tropics (Figure 3.3), causing a much steeper distribution of energy to form despite the energetic small scales shown in Figure 3.4.

The subtropical jet is associated with particularly steep mesoscale slopes

The subtropical jet stream, in particular, is associated with notably steep spectral slopes. This is evident over North Africa and Australia but less in regions of higher latitude. This is likely due to the combination of lower kinetic energy in the tropical regions at scales of 100 km and higher kinetic energy associated with the subtropical jets, leading to a steep distribution of energy. This contrast is not as strong in the extratropics, leading to the weaker jet signal.




Figure 3.4: A global map of 200 hPa mesoscale energy in analysis data on a logarithmic colour bar. Mesoscale energy is lowest in the polar regions and in the doldrums.

Local orographic shallowing of the UTLS spectrum is small but noticeable

In the extratropics, when one studies Figures 3.1 and 3.2, one notices a correlation between instances of shallowing and orography. This is particularly noticeable for the Andes, Rockies and Himalayas but also appears in other mountainous regions. This orographic shallowing is probably due to terrain-induced gravity waves propagating into the lower stratosphere, and its spatial inhomogeneity is likely due to the intermittent absorption of upwardly propagating gravity waves by critical layers in some regions before they reach the shown pressure level. If we were to analyse a greater range of pressure levels we could trace this vertical transfer of mesoscale energy and test this hypothesis.

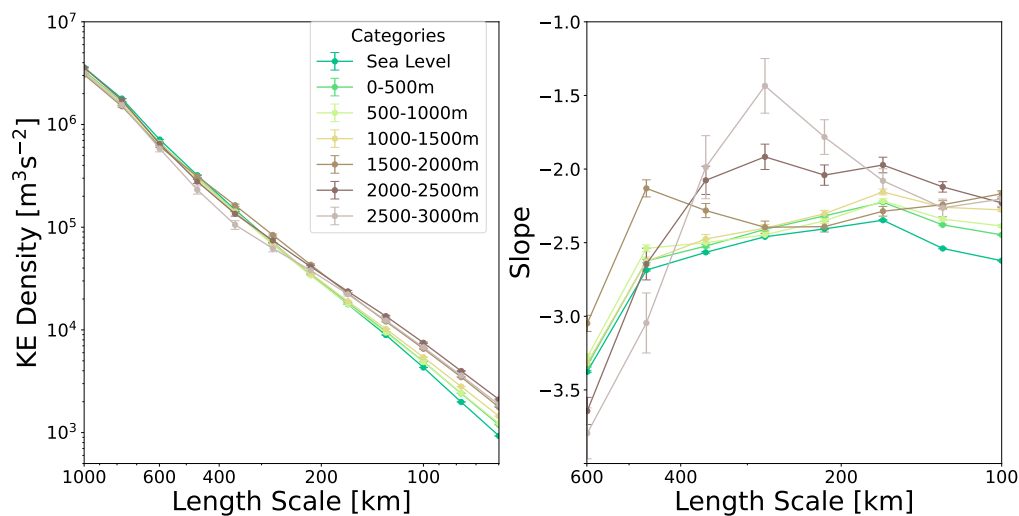


Figure 3.5: Spectra built only from points with an orographic height in the given range are shown in the left panel. The poles and precipitating regions were excluded from the analysis. In the right panel, the spectral slopes of the conditioned spectra are shown, quantifying the shallowing with increased orographic height shown in the left panel. Error bars here and in all future figures were obtained by bootstrapping and show the standard deviation (68% confidence interval) on the mean of the estimated variable at each scale.

In Figure 3.5, the spectra conditioned on the average height of orography within our $2^\circ \times 2^\circ$ grid box is calculated. The points are also chosen not to be precipitating to separate any shallowing effect of orography from those associated with precipitation. In addition, the polar regions, which have a high altitude but are qualitatively different to mountainous regions at lower latitudes, are also excluded. Due to smaller sample sizes leading to larger variances, orography categories above 3000 m were excluded as well. We observe in Figure 3.5 that greater orographic height tends to be associated with greater small-scale kinetic energy at 200 hPa. This leads to a mesoscale shallowing, which is seen in the slope plot on the right side of the figure. There is a systematic shallowing from around -2.6 to -2 and shallower for greater orographic height. The spectral kink is visible with the steepening of the spectrum past 400 km. Spectra were also conditioned upon the derivative of the orography map and the subgrid variability. Perhaps due to the somewhat arbitrary scale separation (the resolution of our orography map is 1/4 degree), these yielded noisier results than simply orographic height.

Local slopes at 200 hPa are noisier in the extratropics than the tropics

The general character of the extratropics is noisier than the slopes seen in the tropics. This is likely due to 200 hPa being a pressure level that straddles the troposphere and stratosphere, with the tropics here being within the troposphere and the extratropics being in the stratosphere. In the stratosphere, the greater stratification increases the role played by gravity waves, which are by their nature intermittent in space and time. This is shown in a map of the ratio of mesoscale rotational and divergent energy in Chapter Five.




Figure 3.6: Spectra built only from points with average precipitation rates in the given ranges are shown in the left panel. Only sea-level points were used in the analysis. In the right panel, the spectral slopes for these conditioned spectra are plotted, showing the shallowing associated with greater intensities of precipitation.

Precipitation is correlated with shallower spectral slopes

In Figure 3.6, we condition based on the average rain rate in our $2^\circ \times 2^\circ$ regions. We consider exclusively sea-level regions to try to mitigate the confounding effects of orographic shallowing. We see that there is a strong relationship between this rain rate and the mesoscale energy with a systematic increase in kinetic energy with greater rain intensities. This leads to a clear shallowing trend in the slope, with slopes ranging from -2.6 to around -5/3. Again, we exclude greater rain rates due to an insufficient sample size.

The net effect of conditioning to exclude orography and precipitation is a steepening of the global spectrum

by 0.2

We have demonstrated that, locally, orography and precipitation are correlated with shallower spectral slopes, but it remains to ask what the net influence of these two factors is on the global spectrum. It could be possible that these two effects entirely explain the spectral kink at 400 km. In Figure 3.7, we show this is not the case. Figure 3.7 shows the total spectrum compared with a spectrum calculated only using non-precipitating, sea-level points. Although the mesoscales are globally significantly spectrally steeper than they otherwise would be by around .1 to .2, the transition to an even steeper slope at larger scales still occurs.




Figure 3.7: The left panel shows the total filter spectrum in black and the spectrum built only from points with sea-level orography and little to no precipitation in green. The right panel quantifies the slight steepening of the global spectrum by 0.2 that this induces.

3.2.2 600 hPa, the mid-troposphere

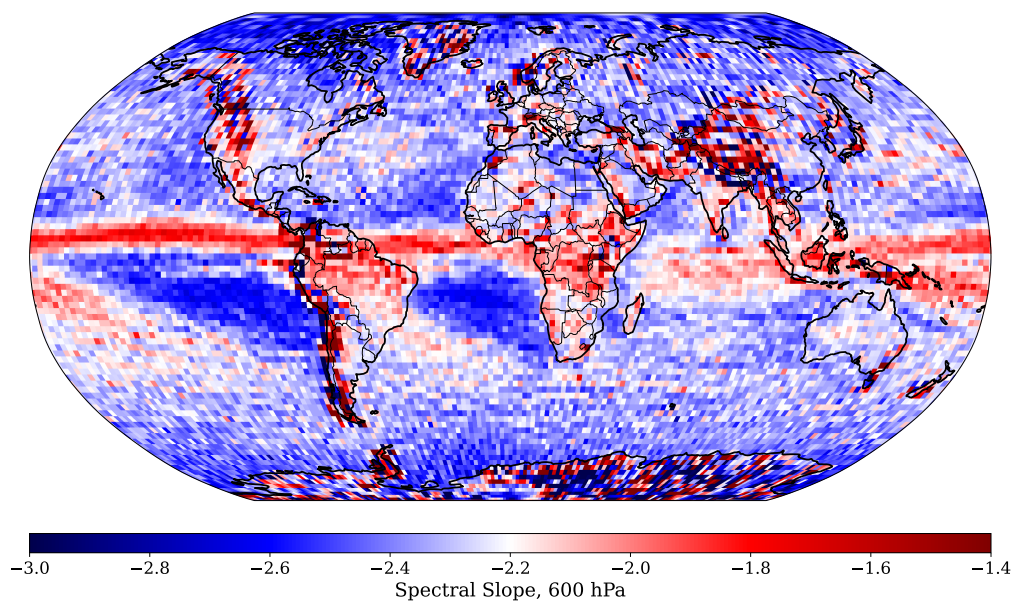


Figure 3.8: A global map of 600 hPa mesoscale spectral slope in analysis data centered on a slope of -2.2. Shallow slopes are displayed as red while steep slopes are blue.




Figure 3.9: ERA5 average hourly precipitation for 2020 and orography field for reference.




Figure 3.10: Global map of kinetic energy in 2020 ECMWF analysis at 600 hPa, showing the distribution at mid-tropospheric levels.

Now, we turn our attention to the mid-troposphere. Figure 3.8 shows the local spectral slopes at 600 hPa. The orography and precipitation reference maps have been reproduced for convenience in Figure 3.9. Figures 3.10 and 3.11 show the unfiltered total kinetic energy and the energy averaged over the mesoscales, respectively.

Shallow tropics correlates with precipitation

Again, we observe a red slash across the equator corresponding to the Intertropical Convergence Zone (ITCZ). This slash is more tightly correlated with the precipitation rates we see in Figure 3.2 than the red regions seen in Figure 3.1. This may be due to gravity waves radiating energy away from the convection at 200 hPa.

Doldrums exhibit the steepest local slopes in the mid-troposphere

The next feature that draws the eye is the large blue regions in the subtropical Atlantic and Pacific. These regions have extremely low rainfall, but the exact mechanism of their steepness is unclear. One may

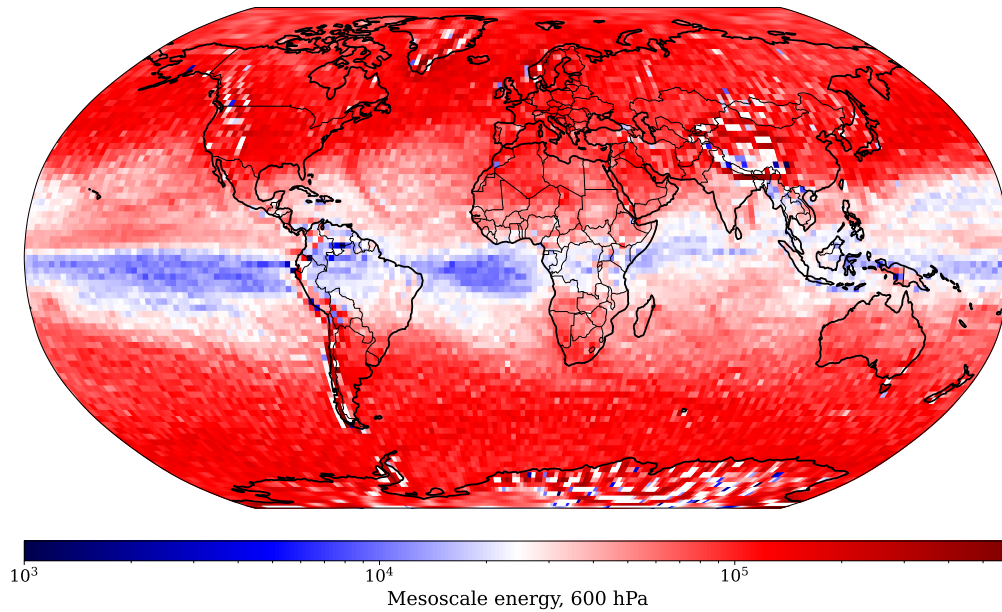


Figure 3.11: A global map of 600 hPa mesoscale energy in analysis data on a logarithmic colour bar. The tropics are the least energetic region of the world at this pressure level.

conjecture that these regions are experiencing an enstrophy direct cascade, which is the same mechanism that leads to the larger scale -3 slope. In order to test this hypothesis, we can evaluate the local energy flux. If the flux were inverse, then this would discount our theory, as then if two-dimensional turbulence dynamics were dominant a $-5/3$ slope should appear. We display the result of this analysis in Chapter Four. Like most other regions, we find the flux of energy in this region is direct, which does not discount this hypothesis. In our analysis in Chapter Five, we find that rotational motion does dominate divergent motion at these locations, which also does not discount our hypothesis since two-dimensional turbulent theories are horizontally divergence-free. In future investigations, we could look at local enstrophy fluxes to pursue this hypothesis further.

Latitudinal steepening seems to occur

We also notice a latitudinal effect with steeper slopes at higher latitudes, excluding the land effects seen over Greenland and Antarctica. In Chapter Five, we see this is associated with a dominance of rotational over divergent motion, so this steepening may also be due to a two-dimensional turbulence enstrophy cascade.

Orographic effects in the mid-troposphere are pronounced

The orographic effect at 600 hPa is also striking, with a clear shallowing associated with higher orography. This association is clearer than at 200 hPa. This may be due to differing mechanisms being responsible for the two cases' shallowing. It may be that at 600 hPa, the shallowing is more associated with turbulent wakes (ala [Lilly \(1983\)](#)) or with direct imprinting (the direct energising of a variety of lengthscales - that is, without a cascade - by external forcing, in this case, interactions with orography) than at 200 hPa, where it could be more of a wave phenomenon.

Figure 3.12 shows the orography-conditioned spectra at 600 hPa excluding the poles. This confirms that the shallowing associated with orography at 600 hPa is more extreme than at 200 hPa, with slopes shallowing by

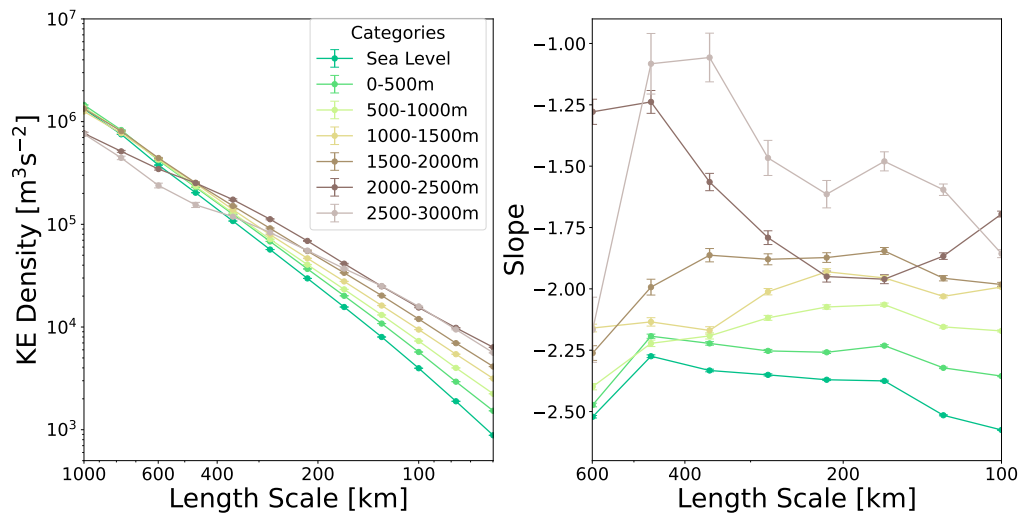


Figure 3.12: Spectra built only from points with an orographic height in the given range are shown in the left panel. The poles and precipitating regions were excluded from the analysis. On the right, spectral slopes are shown, quantifying the shallowing with increased orographic height in the left panel.

a unit and a clear increasing trend in mesoscale kinetic energy being associated with the higher orography.

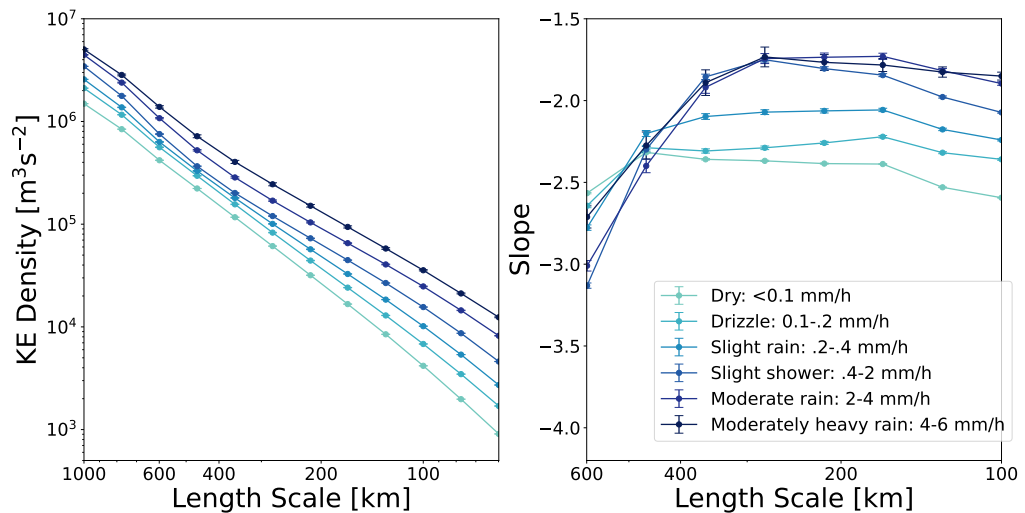


Figure 3.13: Spectra built only from points with average precipitation rates in the given ranges are shown in the left panel. Only sea-level points were used in the analysis. In the right panel the spectral slopes are shown, demonstrating the shallowing of the kinetic energy power spectrum at 600 hPa associated with greater intensities of precipitation.

Precipitation is correlated with shallow spectral slopes

Figure 3.13 shows the conditioned spectra for precipitation for points with sea-level orography. In contrast to previous conditioned spectra (c.f Figures 3.5, 3.6, 3.12), we see that at 600 hPa, greater precipitation is clearly associated with not only greater mesoscale kinetic energy but greater kinetic energy in the synoptic scales. It seems logical to conclude that this is due to larger-scale weather patterns associated with precipitation being more energetic than the average flow, for example, low-pressure systems like storms. A clear shallowing also occurs concurrently with the increase in energy. This shallowing goes from a value of around -2.5 for dry points and seems to tend towards a value closer to -5/3 (-1.67). This does not occur at 200 hPa, likely because the energising effect of the precipitation is felt through vertically propagating gravity waves at that pressure level, and the larger scale energetic motions do not project into the higher altitudes.

The net effect of conditioning to exclude orography and precipitation is a steepening of the global spectrum by 0.1

We now ask what the combined effect of orography and precipitation are on the spectrum. Figure 3.14 shows the effect of excluding precipitation and orography from the global calculation. There is a more marked decrease in energy across the mesoscales and beyond at 600 hPa compared with 200 hPa. There is a similar steepening of the spectrum in both cases, with the exclusion of orography and precipitation leading to average mesoscale slopes of around -2.4, while without this, the slopes in the analysis are closer to -2.2. Although it is subtler than at 200 hPa, a steepening of the spectrum at larger scales is still observed for both spectra.




Figure 3.14: The left panel shows the total filter spectrum in black and the spectrum built only from points with sea-level orography and little to no precipitation in green. The right panel quantifies the slight steepening of the global spectrum by a little over 0.2 that this induces.

3.3 Convective and large-scale precipitation

The precipitation data we use is from ERA5. This precipitation is model-derived, the two main sources being the large-scale and convective precipitation schemes. Given the potential differences in the spectrums associated with two kinds of precipitation, namely that convective precipitation is more likely to be intense and localised, whereas precipitation from the large-scale scheme is likely to come from large-scale fronts, it is interesting to see if these differences appear in our analysis.

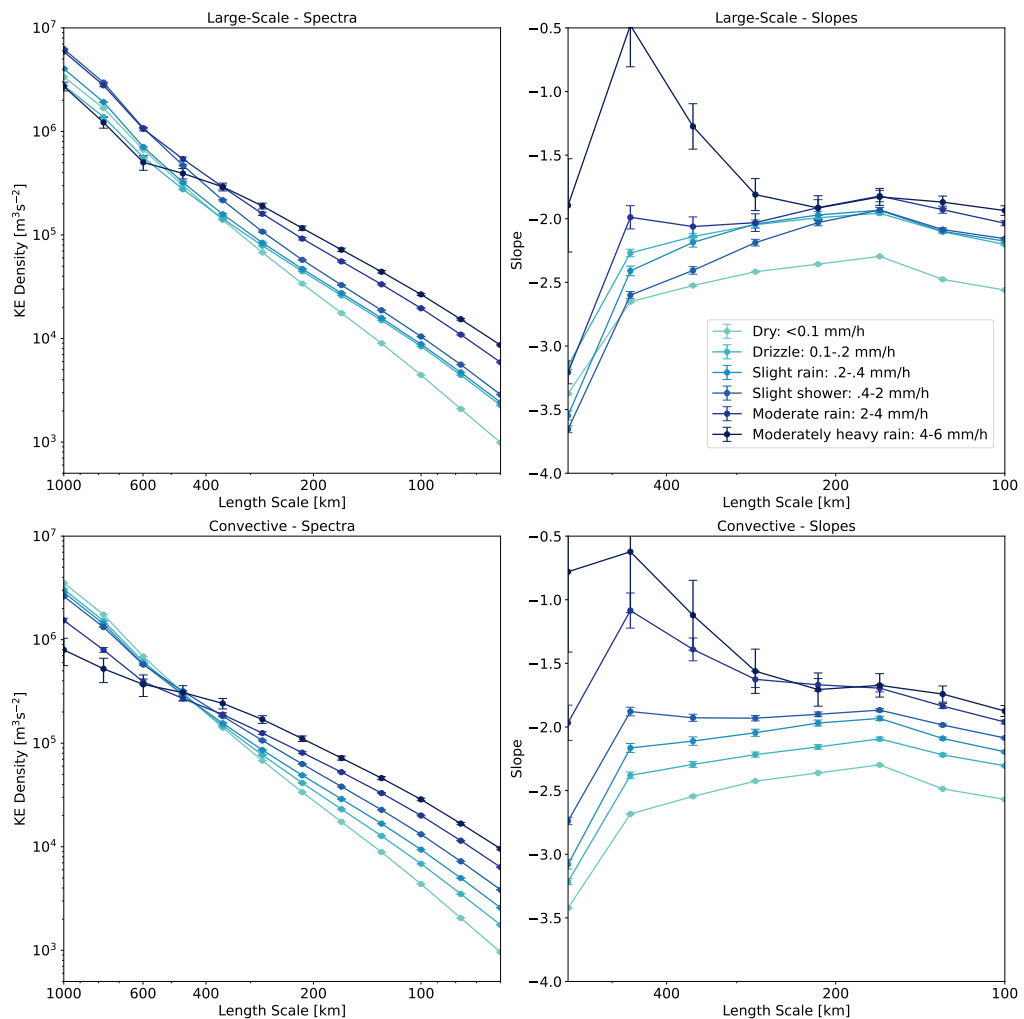


Figure 3.15: 200 hPa spectra are conditioned based on the intensities of rain from the convective or large-scale precipitation (LSP) parameterisations of ERA5. Spectra are aggregated only from sea-level points with precipitation in specified bands. In the top left and right, the spectra conditioned on large-scale precipitation are shown, with spectra on the left and spectral slopes on the right. There is no clear relationship between large-scale behaviour and the intensity of precipitation generated by this scheme. In the bottom two panels, the spectra and slopes of convective precipitation-conditioned spectra are shown. In the left panel, we see a systematic shallowing of spectral slope at the large scales with increased convective precipitation intensity, corresponding to less energy between 1000 and 600 km associated with more intense convective precipitation. This aligns with intense convection being a localised phenomenon, and so generating gravity waves that propagate vertically in a localised region, while rain generated by the LSP scheme is associated with larger-scale systems, which may be associated with gravity waves generated over a larger region.

Figure 3.15 shows these conditioned spectra for 200 hPa. As we may have expected, at larger scales between 1000 and 600 km, greater convective precipitation is associated with less large-scale energy than greater large-scale precipitation. Applying the same analysis at 600 hPa yields Figure 3.16. As we noted when

discussing Figure 3.13, large-scale energy associated with weather systems appears to be present with increased precipitation levels. This effect is enhanced when we restrict ourselves to precipitation from the large-scale precipitation parametrisation and almost disappears when considering only convective precipitation, essentially confirming our assumption.

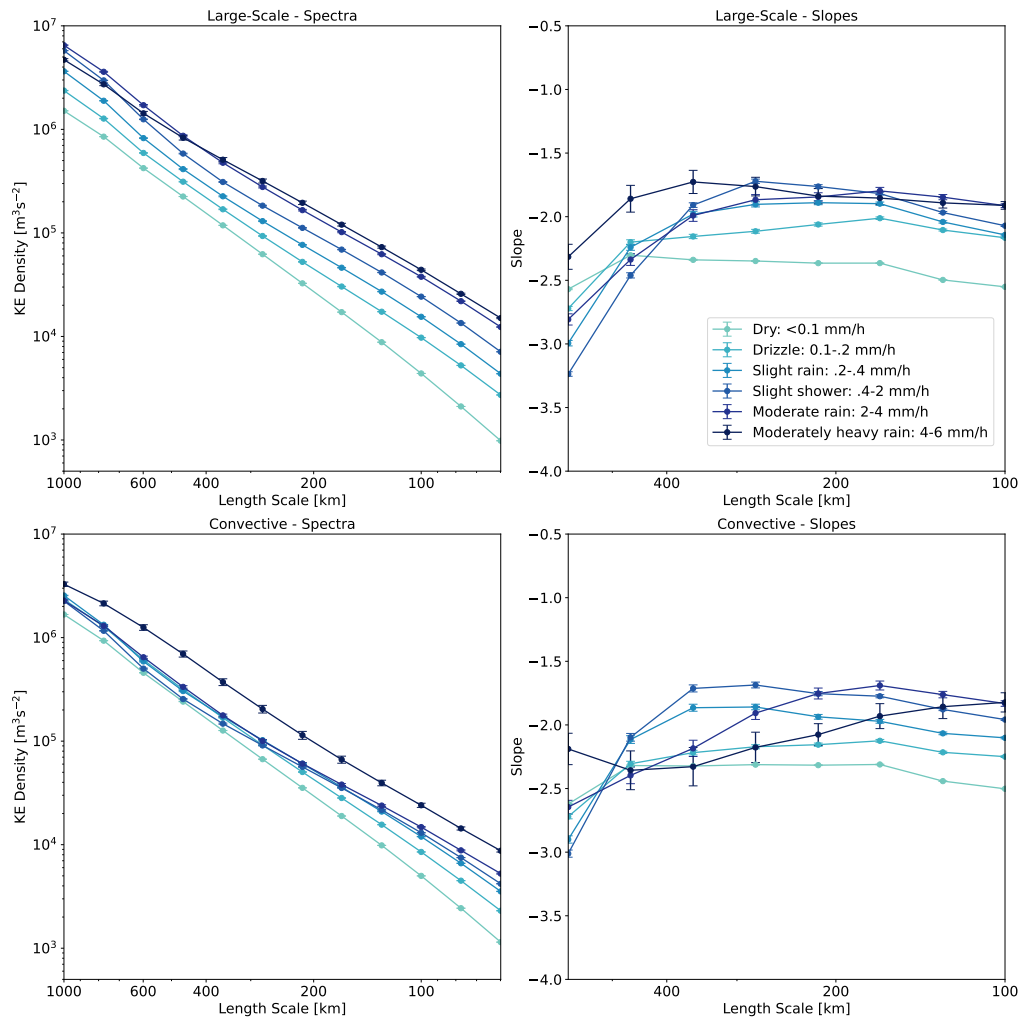


Figure 3.16: 200 hPa spectra are conditioned based on the intensities of rain from the convective or large-scale precipitation (LSP) parameterisations of ERA5. Spectra are aggregated only from sea-level points with precipitation in specified bands. The top panels show the LSP-conditioned spectra (left) and slopes (right). The bottom panels show the convective conditioned spectra (left) and slopes (right). In contrast to 200 hPa shown in Figure 3.15, both increased convective and LSP precipitation are associated with greater energy at the largest scales. However, while the increase of energy is quite uniform for the LSP scheme, for the convective precipitation, there is significantly more energy in the lower mesoscales around 100 km with increased intensity than at the larger scales, as one can see from the bottom left panel.

3.4 Mean sea level pressure

Mean sea level pressure (MSLP) can indicate the presence of storms or fine days. It would be reasonable to expect a relationship between the MSLP and the spectral slope. Indeed, Figure 3.17 shows there is a relationship. At 200 hPa, high-pressure systems are associated with greater energy; this is the case even when conditioning the data only using points in the extratropics (not shown), so it is not solely due to a correlation between low pressures and tropical conditions. Why this relationship exists between energy and MSLP in the stratosphere requires further investigation. One possibility is that in high-pressure regions we are seeing the troposphere rise locally (Morgan and Nielsen-Gammon, 1998). Since tropospheric dynamics are more energetic than stratospheric, this may explain the greater magnitude of the spectrum, as well as its more tropospheric profile without a clear kink.

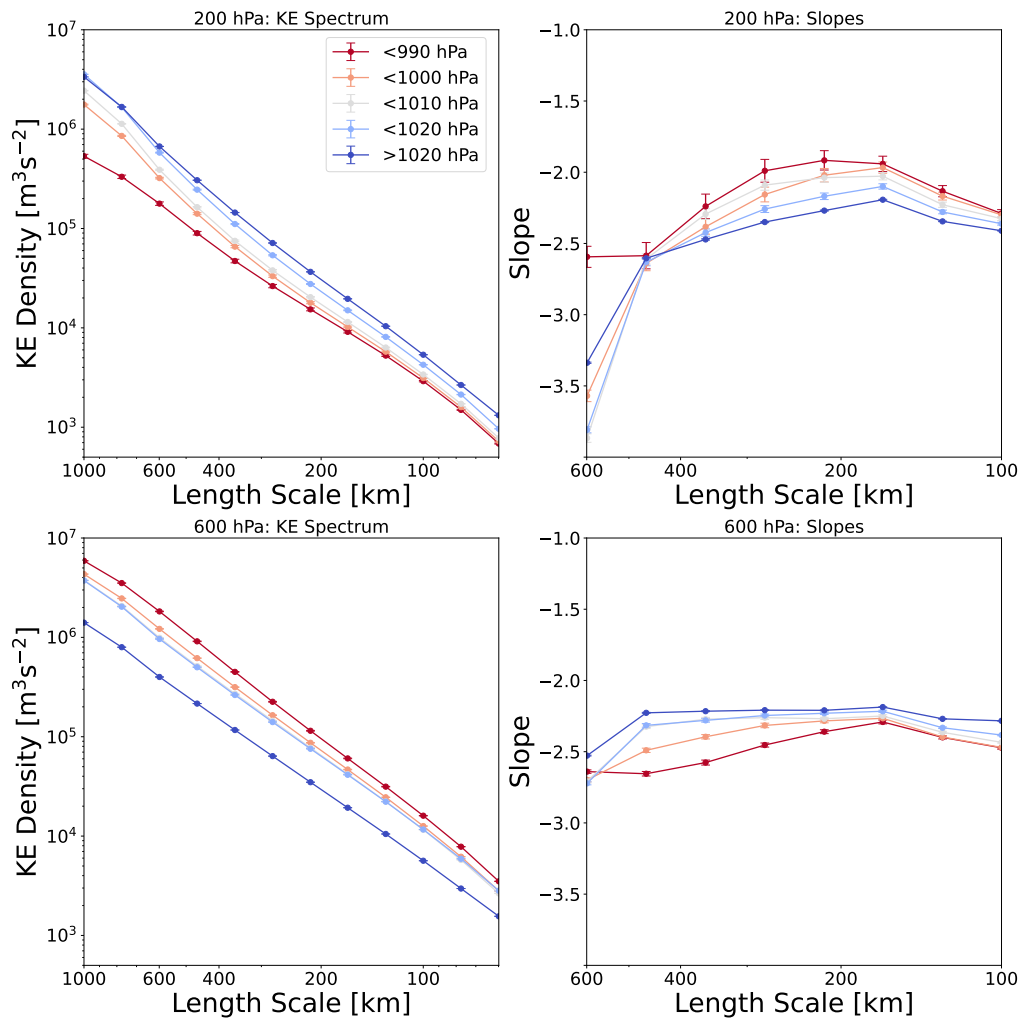


Figure 3.17: Mean sea level pressure (MSLP) conditioned spectra for 200 and 600 hPa. Spectra are aggregated only from points with MSLP in the specified bands. The top panels show 200 hPa with spectra (left) and spectral slopes (right). The bottom panels show 600 hPa with spectra (left) and spectral slopes (right). We see opposite relationships between MSLP and the spectrum at the different altitudes. At 200 hPa, low pressures are associated with low energy and shallower slopes than high pressures. At 600 hPa, low pressures are associated with high energy and steeper slopes than high pressures.

The situation at 600 hPa is reversed. This aligns with our general knowledge of high and low-pressure systems, with lower pressures being associated with energetic phenomena in the tropospheric mesoscales, such as storms.

3.5 Discussion and conclusions

In this chapter, we have shown the first consistent maps of local spectral slopes in the mesoscales and the first conditioned spectra. We now give an overview of our key observations.

The shallowing in the tropical mesoscales is associated with precipitation

A relatively shallower spectral slope in the tropics has been noted in simulations (Judt, 2018) and reanalysis Bolgiani et al. (2022). It is potentially present in direct observations as well, see Figure 10 of Frehlich and Sharman (2010). However, it has not been recognised that there is a strong correlation between this tropical slope and precipitation. A similar shallowing pattern was seen in Wang and Sardeshmukh (2021). However, their methodology has some serious drawbacks, which are discussed in Chapter Two.

In our analysis, we have observed a strong correlation between the shallowing in the equator and the rainfall patterns in the tropics at 200 and 600 hPa, with the pattern at 200 hPa slightly more diffuse than at 600 hPa. Within the troposphere, it is unclear how or why this shallowing occurs. One possibility we can rule out is an inverse cascade, as we shall discuss in Chapter Four when we analyse the local spectral flux direction. It is probably some combination of wave and turbulent dynamics or simply the imprint of forcing at various scales. In the tropics at 200 hPa, we are still within the troposphere. Repeating our analysis at a higher altitude would be interesting for observing the tropical stratosphere. Gravity wave propagation is likely responsible for the more spread out shallowing at 200 compared to 600 hPa.

It would be beneficial for this field if instead of the functional fits displayed in Frehlich and Sharman (2010), the data was reanalysed by transforming the structure functions into power spectra (for consistency with the majority of simulation, observation and analysis work) and then publishing slopes over octave ranges as in our Figure 1.2 in Chapter One.

Jets at 200 hPa are associated with strong local steepening

The fact that jets are associated with steep spectra is well known. This is due to their concentrations of energy at large scales. Measurements from observations of Jupiter and spectra of jet-dominated simulations record slopes that can become as steep as -5 or -7 (see, for example, Choi and Showman (2011) and references therein). This feature is not observed in Wang and Sardeshmukh (2021), possibly due to the nonlocality of their approach.

I am unaware of any observations of steepening associated with jets in the Earth's atmosphere. However, extracting this signal from currently existing aircraft observations may not be difficult. We can conjecture that analysing only data passing through the subtropical jet will produce steeper spectra in the mesoscales, likely a -3 although possibly as steep as -5.

Orography causes significant local shallowing

At both 200 and 600 hPa, significant shallowing was associated with orography. In the troposphere, this may be due to wake effects energising turbulent motions and wave activity. At higher altitudes, it is likely that vertically propagating orographic waves are responsible. This is supported by a clear resemblance between shallowing at 200 hPa and the spatial distribution of gravity wave activity (see the climatology of gravity wave activity at 100 hPa in Hocke et al. (2016)). We could further test our hypothesis by investigating how critical layers influence the spectrum. If gravity waves energise the spectrum, the shallow slope would be intermittent in time and space, depending on whether a critical layer absorbed the gravity wave before it

arrived at the particular pressure level. By analysing a range of vertical levels, we could trace this energy transfer and test this hypothesis. That orography significantly energises the mesoscales was observed in [Nastrom et al. \(1987\)](#).

Precipitation is associated with local shallowing

At both pressure levels, precipitation is associated with spectral shallowing. Once again, this has not been observed in the atmosphere. However, it is a well-known effect in simulations, for instance, [Selz et al. \(2019\)](#) and [Sun et al. \(2017\)](#). Our conditioned spectra add to this body of work and show the relationship between increasing precipitation levels and slope in the analysis. We also show how the source of precipitation can be associated with different shapes of the power spectrum. It would be interesting to repeat a similar investigation in a convection-resolving model to check if there are any qualitative differences in our conclusions.

Lower pressure is associated with steeper slopes in the troposphere and shallower in the UTLS

In agreement with the observations of [Vonich and Hakim \(2018\)](#), we have observed that steeper slopes in the troposphere are associated with low MSLP. However, we also observe that in the UTLS shallower slopes are observed. Why the latter should be the case is unclear but could be due to a higher tropopause in high pressure regions.

Chapter 4

Energy flux and spectral slope

In this Chapter, we analyse the local spectral flux of kinetic energy in the atmosphere. This can provide us with useful clues to the underlying dynamics and be of use in future investigations into parametrisation design.

Subscale flux

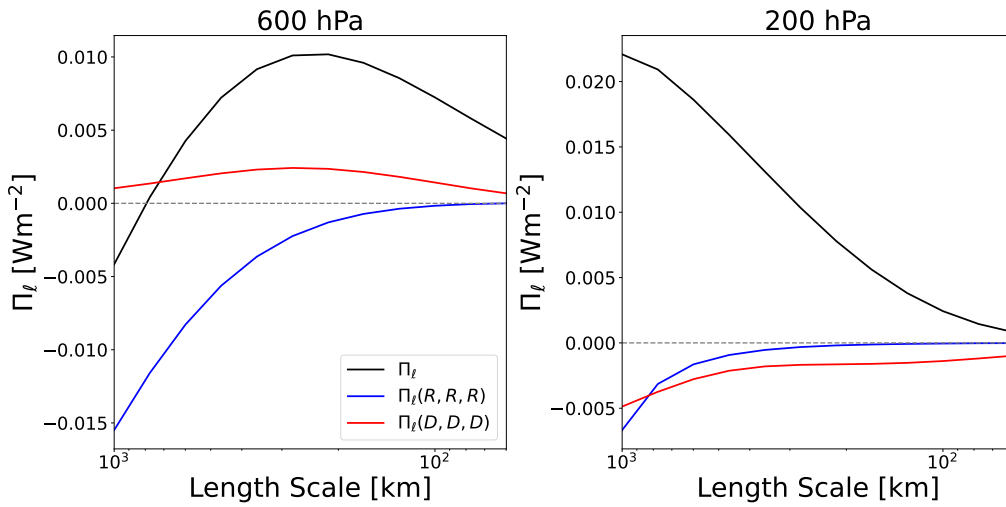


Figure 4.1: The average subscale flux of kinetic energy for 2020, Π_ℓ is shown in black. Negative values indicate a net transfer of kinetic energy to larger scales, while a positive value indicates a transfer to smaller scales. Two components (of a total of eight detailed in Chapter Two) are also shown; the pure rotational (blue) and pure divergent (red) components. $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$ in blue is particularly significant as it indicates the flux due to only rotational motion. The two-dimensional turbulent theories for the atmosphere are all nondivergent, and so are theories for this component of the flow. 600 hPa is displayed on the left pane, and 200 hPa on the right. We see that at both altitudes, on average the rotational component of the flow is acting to transfer kinetic energy to larger scales, in line with the inverse cascade of two-dimensional turbulence theory. The average energy, which includes all divergent-rotational interactions as well as the purely divergent flux in red, is on average, transferring kinetic energy from large to small scales.

As detailed in the last section of Chapter Two, the subscale flux Π_ℓ , is the kinetic energy transferred from scales larger than ℓ to those smaller at a given time and location. It can be broken down into various combinations arising from the interaction of rotational and divergent motions. In particular, $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$ is the transfer due to purely rotational motion. Since the purely turbulent theories for the spectral slope described in Chapter One are nondivergent, any energy transfers due to that explanation must be associated with $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$. So if we were to detect both a shallow slope and a negative value of $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$, we would have strong reason to suspect the inverse cascade of energy of classical 2D turbulence theory.

[Storer et al. \(2023\)](#) applied coarse-graining to the oceans to find local fluxes. They observed bands around the equator of energy moving to larger scales. By the breakdown described above, it was surmised that this inverse transfer was not due to an inverse cascade but divergent Ekman flow near the equator stretching vortices horizontally, thereby causing energy to be transferred to larger scales. This is called the “piston effect” in that work, terminology we will borrow here.

Figure 4.1 shows the subscale flux for 600 and 200 hPa. We see that at both altitudes, on average, energy is transferred towards smaller scales in the mesoscales. In both, the purely rotational interactions transfer energy towards the larger scales.

The $\Pi_\ell(\mathbf{D}, \mathbf{D}, \mathbf{D})$ term is more difficult to interpret. It may be thought of as a result of wave-wave interactions, but these are spread across several other terms as well. Interestingly, it acts to transfer energy downscale at 600 hPa and upscale at 200 hPa. From this point, when we are interested in the effects of divergent motion we will consider $\Pi_\ell - \Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$ rather than $\Pi_\ell(\mathbf{D}, \mathbf{D}, \mathbf{D})$ since all theories of gravity wave shallowing acknowledge some level of interplay between waves and the balanced flow.

Previous work

Similar maps of local spectral flux have been produced by a slightly different methodology presented in [Faranda et al. \(2018\)](#). In that work, it was found that in the troposphere (500 hPa), downscale transfers dominated in the upper mesoscales while in the stratosphere (100 hPa) upscale transfers dominated. With our higher resolution analysis, we report a slightly different picture. Our analysis at 600 hPa yields a similar pattern to that in [Faranda et al. \(2018\)](#), but we do not see the strong inverse transfers presented in that work at 200 hPa. The reason for this discrepancy requires further investigation.

4.1 Energy flux in the mesoscales

4.1.1 Flux in the stratosphere and upper troposphere: 200 hPa

Figure 4.2 shows the average transfer of energy through the mesoscales at 200 hPa. The reader is reminded that a positive red value indicates a transfer towards smaller scales while a negative blue value indicates a transfer towards larger scales. Note that energy fluxes are not explicitly weighted by the corresponding kinetic energy as the spectral slopes were. Implicitly, however, this is the case, since greater spectral energy corresponds to larger spectral fluxes, making the slope maps in the previous chapter and the maps shown here comparable.

Energy is transferred to larger scales at the equator

The inverse transfer across the equator is the most striking feature of Π_ℓ . Since this feature is not present in the purely rotational component (the bottom panel of Figure 4.2), a possible explanation for this is that it is due to a piston effect, where the general circulation introduces a divergence that stretches vortices meridionally to larger scales.

Alternatively, the strong correlation between the areas of inverse transfer and greatest precipitation may indicate we are seeing a more convectively driven phenomenon. Several suggestions are compatible with this theory. We may be seeing the effects of organised convection, or the spreading of energy away from localised heating through tropical waves (Kelvin and Rossby) in a Gill response or even the equilibration of temperature gradients through gravity waves.

Strong direct transfers are associated with orography

Another interesting feature of Figure 4.2 is the stronger direct transfers of energy that seem to be associated with orography. We can speculate that these are due to the breaking of orographically induced gravity waves. This interpretation finds some support in the absence of this intensification in the purely rotational component of the flux. We will show conditioned fluxes later in this chapter (Figure 4.4) that will confirm

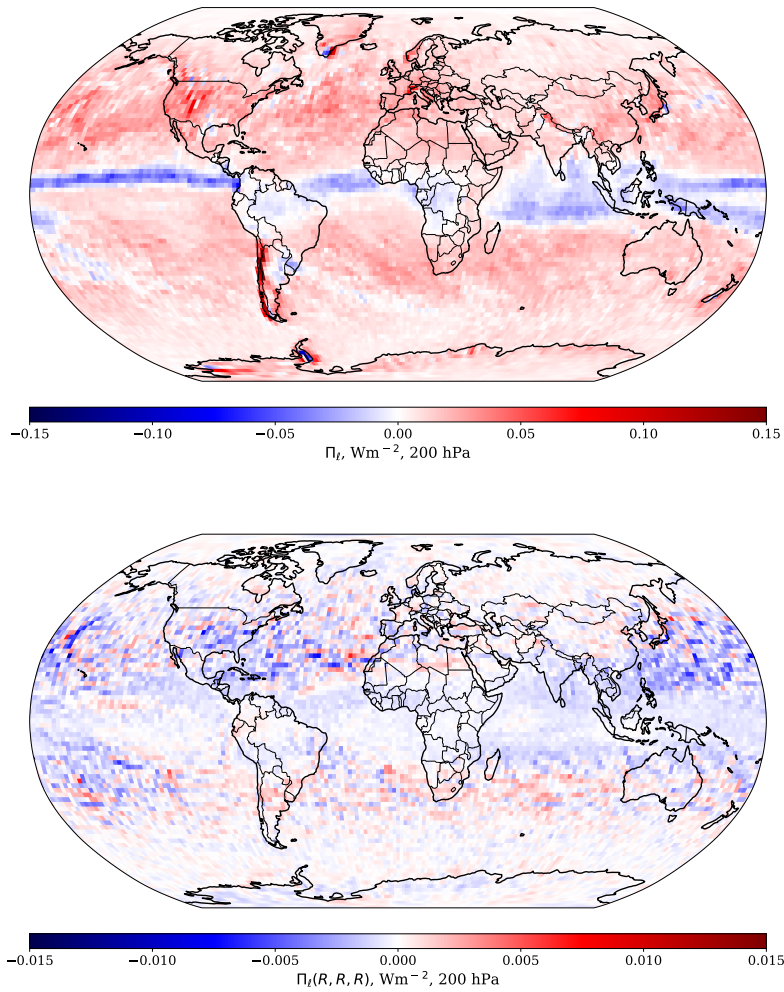


Figure 4.2: A global map of 200 hPa average mesoscale energy flux in analysis data. The top panel shows the total transfer of energy through the mesoscales, Π_ℓ . Negative blue values indicate a transfer of kinetic energy to larger scales while positive red values indicate a transfer of kinetic energy to smaller scales. The bottom panel shows $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$, the transfer due to the rotational component of the flow only. This must contain the contribution from two-dimensional turbulence as these theories are purely rotational. Note the order of magnitude difference in the colour bar between the upper and lower panel, indicating the dominance of divergent-rotational interaction over rotational motion in explaining the transfer of kinetic energy at the pressure level. Notice also that the lack of a similar feature in the lower panel demonstrates the strong band of negative transfer at the equator in the top panel is not due to a two-dimensional turbulence inverse cascade.

this observation.

Precipitation is associated with inverse transfer in the tropics and direct in the extratropics

We can also observe that although precipitation in the tropics correlates well with transfer of energy to larger scales, the opposite is true in the extratropics. This is explained well by the fact that the overall circulation in the tropics is divergent at this altitude. In the extratropical regions, we are viewing the stratosphere and might expect a different relationship with precipitation led by gravity wave activity.

A turbulent inverse or direct cascade remains a possibility in some regions

This is the average over a year, and so even if turbulent inverse cascades are common, they may not appear here. However, there are some regions that exhibit both an overall inverse transfer and a non-insignificant

rotational inverse transfer. Here, we see only Texas and a region in the North Pacific.

4.1.2 Flux in the troposphere: 600 hPa

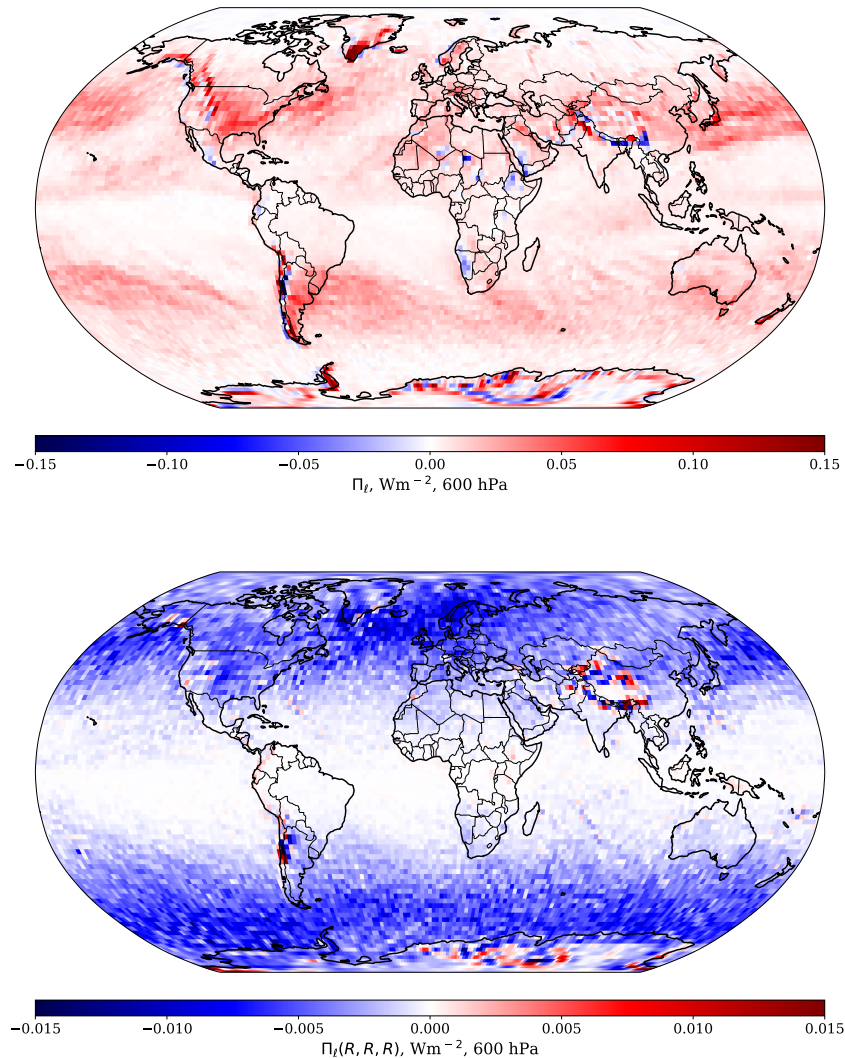


Figure 4.3: A global map of 600 hPa mesoscale energy flux in analysis data. The top panel shows the total transfer of energy through the mesoscales, Π_ℓ . Negative blue values indicate a transfer of kinetic energy to larger scales while positive red values indicate a transfer of kinetic energy to smaller scales. The bottom panel shows $\Pi_\ell(\mathbf{R}, \mathbf{R}, \mathbf{R})$, the transfer due to the rotational component of the flow only. As in Figure 4.2 note the order of magnitude difference in the colour bar, indicating that at this pressure level also, transfers are dominated by divergent-rotational interactions rather than purely rotational flow.

Figure 4.3 shows the equivalent results for 600 hPa. As at 200 hPa the dominant transfer is direct. There is no broad inverse band across the equator but the transfer is much weaker there than in the extratropics.

Storm tracks are associated with downscale energy transfer

The storm tracks appear as particularly powerful sources of downscale energy transfer. We see that this transfer is an essentially divergent process.

Orography is associated with both inverse and direct transfer

In contrast to 200 hPa, orographic effects cause direct and inverse transfers of energy. This may be speculated

to be a divergence effect where a mountain range disrupts the flow which takes on larger scale characteristics as it is diverted. That is, the anisotropic pattern of red on the windward side of the range and blue on the leeward may indicate compressive and divergent motion, respectively. Why and how orographic features impact the purely rotational flux requires further investigation but may be related to vortex stretching effects.

There are patches of inverse transfer over Africa

We also see an inverse transfer of energy in some locations over Africa. These correlate to some degree with regions of greater relative divergent energy, as we will show later in Chapter Five. They also correlate with orographic fields, except in Southern Africa. Here, a correlation appears in Namibia but breaks down in South Africa and the more Eastern countries directly above it. Since it is clear $\Pi_\ell(R, R, R)$ cannot be responsible from the lack of activity in this region in the bottom panel of Figure 4.3, this may be a case of divergent convective motion causing a stretching upscale energy transfer, with smaller divergent motions associated with convective cells acting to stretch the rotational motion.

There is a strong latitudinal relationship with the rotational inverse flux

It is interesting to note the strong inverse cascade in the extratropics for $\Pi_\ell(R, R, R)$. This is likely a 2D turbulent inverse cascade, in which case the quiescence towards the equator is in line with a growing Rossby radius, or put another way, the diminishing role of planetary rotation in the mesoscales at lower latitudes.

4.1.3 Orography, precipitation, pressure and flux

In the previous results, we saw imprints of orography and precipitation on the flux. Although the orographic relationship is clearly anisotropic at 600 hPa, being dependent on whether we are analysing the windward or lee side of the range, we can begin to investigate these effects through conditioned spectra.

Flux conditioned on orography

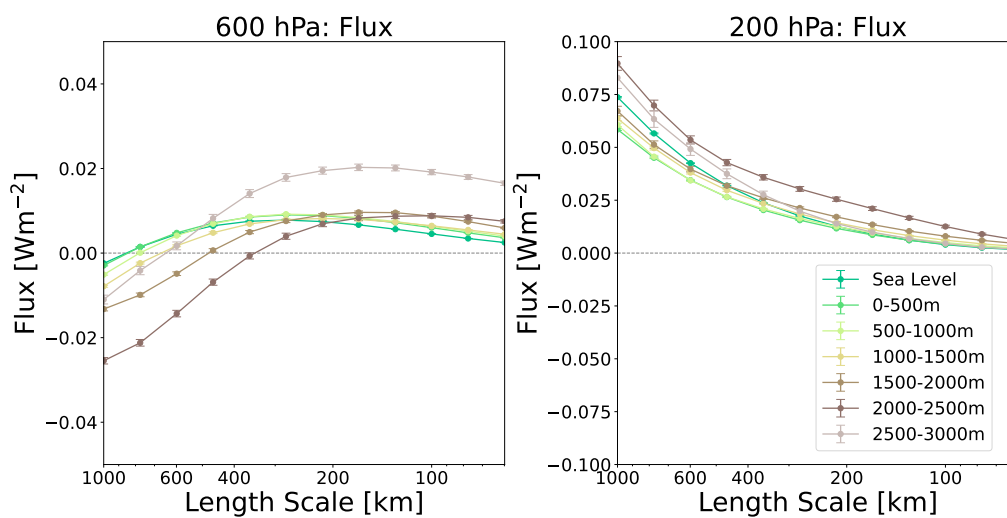


Figure 4.4: Flux aggregated only from points with orographic height in a particular range is shown. Polar points have been excluded. In the left panel, flux at 600 hPa is shown, demonstrating an unclear relationship between orographic height and flux that may result from the anisotropy of the situation, with upscale transfers on the windward and downscale on the lee sides of ranges. At 200 hPa in the right panel, an intensification of transfer to smaller scales occurs, which may be the result of orographically induced gravity waves.

Figure 4.4 shows the flux under different orography conditions. At 600 hPa, the relationship is anisotropic, with upscale transfers on the windward and downscale on the lee sides of ranges. Still, for the smallest scales present, it is clear that we generally have a greater flux of energy towards small scales with greater orography. At 200 hPa we have a general intensification of downscale energy transfer with height, which confirms the observation we made of Figure 4.2. Whether this is direct forcing from vertically propagating gravity waves or a gravity wave induced transfer within the 200 hPa pressure level requires further investigation.

Flux conditioned on precipitation

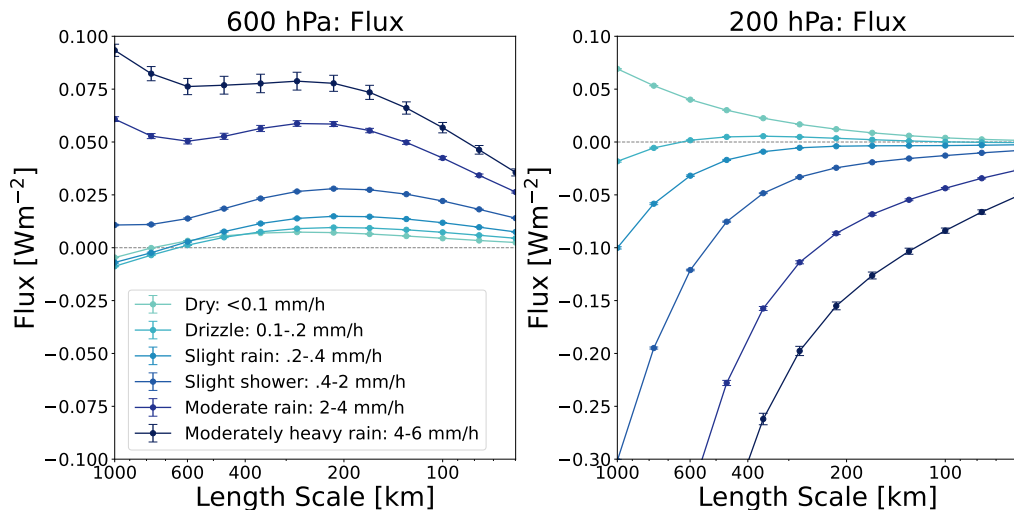


Figure 4.5: Flux aggregated only from points with precipitation in a particular range is shown. Points with orography greater than sea-level height were excluded. The left panel shows flux at 600 hPa. There is a clear increased transfer to smaller scales with intensification of precipitation at this pressure level. In the right panel, the conditioned flux for 200 hPa is shown. More intense precipitation is associated with increased flux to large scales. This is likely due to the concomitance of higher levels of precipitation along the inter-tropical convergence zone (ITCZ) with a broad zone of inverse energy transfer likely associated with the general circulation.

Figure 4.5 is precipitation-conditioned flux. At 600 hPa, in contrast to orography where more extreme orography was associated with a large-scale inverse transfer of energy, more intense precipitation at 600 hPa is associated with a stronger direct energy flux at small and large scales. In Figure 4.6, we investigate the troposphere further by conditioning the flux based on whether the precipitation was due to convective or large-scale ERA5 precipitation. We see particularly large direct transfers associated with convection. At 200 hPa there is a strong inverse relationship where greater precipitation rates are associated with greater transfers of energy to large scales. From Figure 4.2, this is likely due to the concomitance of convective motion associated with rain and the large-scale divergent flow in the tropics associated with the Hadley circulation.

The combined effect of precipitation and orography on spectral flux

We now investigate how important these orographic and precipitation trends are for the global energy flux in Figure 4.7. At 600 hPa they slightly enhance the transfer of energy to smaller scales. At 200 hPa, likely due to their concomitance with the strong divergent inverse energy flux at the equator, removing these points from the analysis significantly increases the direct transfer of energy at the larger scales considered.

Flux conditioned on MSLP

Figure 4.8 shows the fluxes conditioned on MSLP for 600 and 200 hPa. This relationship between flux and

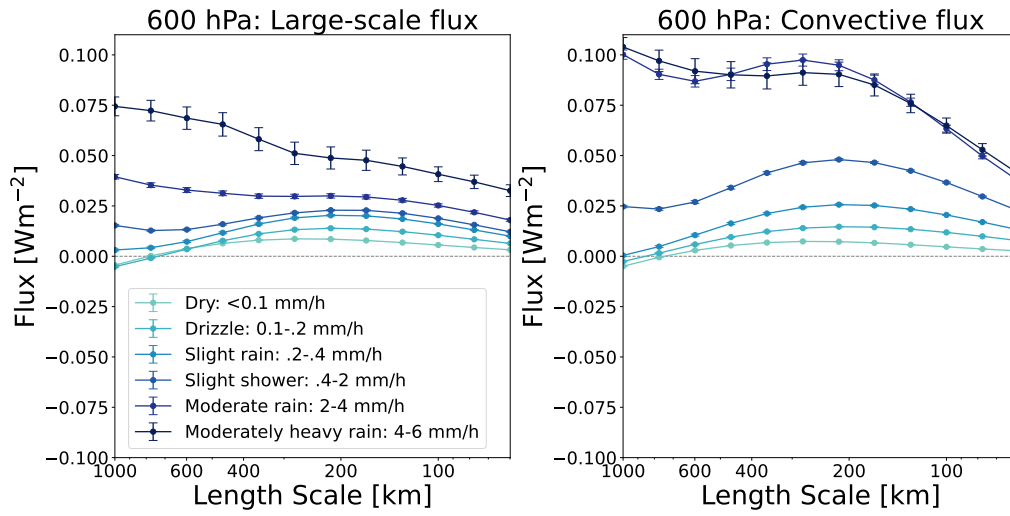


Figure 4.6: Flux conditioned on the ERA5 Large-scale precipitation (LSP) (left panel) and convective schemes (right panel) are displayed for 600 hPa. Flux has been aggregated only from sea-level points with precipitation in specified bands. More intense convective precipitation is associated with more intense downscale kinetic energy transfer in the mesoscales than the LSP scheme.




Figure 4.7: Flux aggregated only from dry points with sea-level orography (green) are compared against the total filter or coarse-graining flux (black). The left panel shows 600 hPa, where the exclusion of precipitation and orography results in a small decrease in transfer to small scales. At 200 hPa (right panel), likely due to the concomitance of precipitation and the general circulation-driven inverse transfer in the tropics (Figure 4.2), excluding precipitating and orographic points results in an increased positive flux of energy to small scales.

MSLP persists even when we restrict our analysis to the extratropics (not shown). At 200 hPa, it is plausible that the upwelling associated with low-pressure systems causes divergent motion in the stratosphere, leading to a more negative $\Pi_\ell(D, R, R)$ term. The situation at 600 hPa is also interesting. High-pressure systems are generally associated with weaker pressure gradients than low-pressure systems. By gradient wind balance, this implies a smaller radius of curvature is reached in low-pressure systems so that Coriolis can balance pressure. This stronger vortical motion at small scales may be associated with a turbulent inverse cascade. Indeed, we see a clear trend of an increasingly strong inverse transport associated with $\Pi_\ell(R, R, R)$ at 600 hPa as pressure decreases (not shown), lending some support to this interpretation.

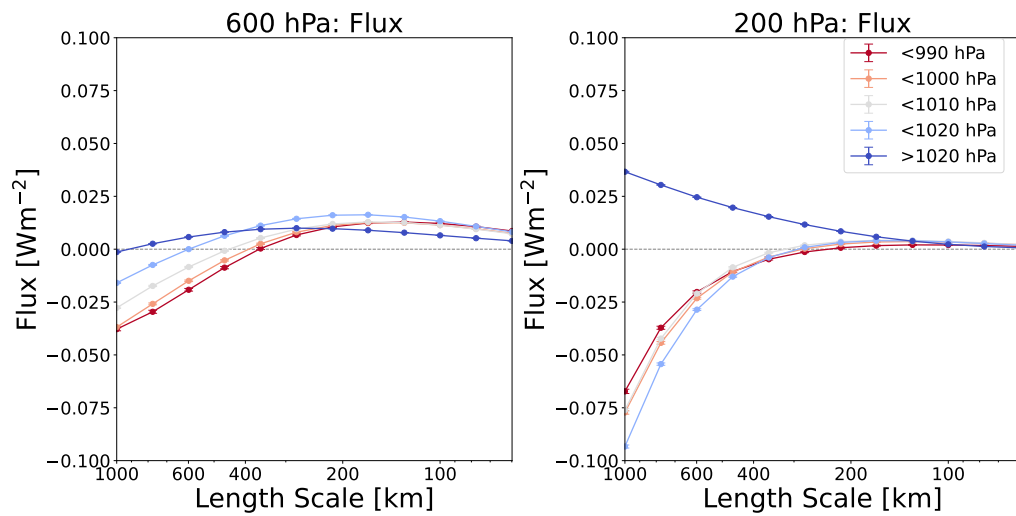


Figure 4.8: MSLP conditioned flux for 600 (left) and 200 hPa (right). Fluxes are aggregated only from points with MSLP in the specified bands. Lower sea level pressures were associated with more energy transfer to larger scales above the mesoscales at both pressure levels.

4.2 Spectra conditioned on spectral flux

Flux categories

Now we condition spectral fluxes on the average spectral flux. That is, we find the typical spectral flux profile when the transfer of the energy averaged over the mesoscales falls in a particular band. It is not immediately obvious what this analysis will return. It may be that on average, a negative flux in the mesoscales is associated with a transition from an inverse to direct transfer at some scale within the mesoscales. Figure 4.9 shows that in the aggregate, no such transition occurs.




Figure 4.9: Flux aggregated only from points with an average mesoscale spectral flux in a particular range. Those ranges are fluxes greater than 0.1 W/m^2 , between 0.1 and 0.01 W/m^2 , between 0 and 0.01 W/m^2 (denoted by $<0.01 \text{ W/m}^2$), between 0 and -0.01 W/m^2 (denoted by $>-0.01 \text{ W/m}^2$), between -0.01 and -0.1 W/m^2 and less than -0.1 W/m^2 . The left panel shows 600 hPa and the right 200 hPa. This shows that when accumulating spectral flux profiles that, averaged over the mesoscales, are negative or positive, the profile is negative or positive over the whole range of scales. It is not obvious this should happen, many of the previous conditioned spectra we have examined cross the x-axis, particularly around 400 km.




Figure 4.10: 200 hPa spectra built only from points with a spectral flux in a particular range are shown (ranges are as in Figure 4.9). Spectra are on the left panel, and associated spectral slopes are on the right. As the magnitude of the transfer, both positive and negative, increases, the energy in the associated spectrum increases. At the lowest magnitude of transfer, a strong kink emerges with slopes that become shallower than $-5/3$ (-1.67) despite the biases towards steep slopes in our methodology and dataset.

We demonstrate that there is only a weak relationship between slope and flux in Figure 4.10. An interesting feature of the flux-conditioned spectra is that greater flux is concomitant with greater energy. Hence, although the character of the spectrum does not differ greatly depending on the direction of spectral flux, it does differ based on the energy in the spectrum and for that reason, on the magnitude of flux. Figure 4.11 shows a similar relationship between flux and spectra exists at 600 hPa to that at 200 hPa. It does seem there is a slightly shallower slope associated with a stronger inverse transfer, this is likely explained by the tendency of some convective events to be associated with an inverse transfer, as we noticed over Africa in Figure 4.3. This is an interesting result because it shows that there is space to revive some of the pre-[Cho and Lindborg \(2001\)](#) theories for the shallow mesoscale slope that invoke an inverse transfer of energy.




Figure 4.11: 600 hPa spectra built only from points with a spectral flux in a particular range are shown (ranges are as in Figure 4.9). Spectra are on the left panel, and associated spectral slopes are on the right. As the magnitude of the transfer, both positive and negative, increases, the energy in the associated spectrum increases. Although less pronounced than at 200 hPa (Figure 4.10), at the lowest magnitude of transfer, a distinct kink emerges which can be seen as slopes shallow from -3 to close to $-5/3$.

4.3 Discussion and conclusions

In this chapter, we have shown local spectral flux maps and conditioned fluxes for the atmosphere for the first time in high resolution analysis. We have reproduced the result that overall energy transfers from large to small scales in the mesoscales as reported in observations by [Cho and Lindborg \(2001\)](#). Our key conclusions follow.

There is a band of inverse transfer across the equator at 200 hPa

At 200 hPa, we have observed a band of inverse transfer in the atmosphere across the equator. This may be due to a piston effect as was proposed to explain a similar phenomenon seen in the ocean in [Storer et al. \(2023\)](#). Alternatively, it could be a fingerprint of convective organisation, the action of tropical waves induced in a Gill response, or be due to gravity wave equilibration of temperature gradients.

Orography has an anisotropic relationship with energy transfer at 600 hPa

We have also observed an altitude-dependent anisotropic relationship between orography and energy transfer at 600 hPa, with upscale transfers on the windward and downscale on the lee sides of ranges. This effect is present in both the rotational and total flow. At the higher altitude of 200 hPa, the relationship was solely one of increasing direct transfer. This may be due to vertically propagating gravity waves as in [Callies et al. \(2014\)](#) and [Malardel and Wedi \(2016\)](#) or some kind of wave driven horizontal transfers of energy may be occurring as in earlier theories of [VanZandt \(1982\)](#).

Turbulent explanations remain plausible locally in the UTLS

As mentioned in Chapter One, theories such as those of [Gage \(1979\)](#), which posit a two-dimensional inverse energy cascade is responsible for the shallow mesoscale slope, became unpopular when it was realised that the overall flux of energy was direct in the mesoscales ([Cho and Lindborg, 2001](#)). However, we have seen that in some specific locations, both shallow spectral slopes and a total and rotational inverse transfer of energy coexist. This opens the door to the possibility that under some specific circumstances, a 2D turbulent inverse cascade may be responsible for local shallowing in the atmosphere's power spectrum. We will discuss turbulence theories further in Chapter Five.

Chapter 5

A local divergent rotational decomposition

5.1 Introduction

Helmholtz decomposition roughly corresponds to decomposing a flow into balanced and unbalanced motion in the atmosphere, with unbalanced motion mostly corresponding to gravity waves and balanced motion including turbulence and Rossby waves. This decomposition has proved a useful perspective for investigating the mesoscale kink, with the divergent component of the flow following a roughly $-5/3$ slope and the rotational a roughly -3 , the crossover of these two spectra in the UTLS (that is, the scale the magnitude of the divergent spectrum begins to dominate the magnitude of the rotational one) is the scale at which the global spectrum begins to shallow, 400 km (Callies et al., 2014). In this Chapter, by investigating the magnitudes of the local spectra of each of these components, we assess how the relative importance of these two kinds of motion differs at different locations. In this way, we show that the dominance of divergent over rotational motion in the mesoscales is location-specific and, as we have hinted towards in certain results in Chapters Three and Four, this allows us to resurrect the possibility of turbulence theories being applicable to the atmospheric mesoscales.

The Helmholtz decomposition

Helmholtz decomposition decomposes a vector field into divergence-free $\nabla \wedge \Psi$ (the divergence of a curl is zero) and irrotational $\nabla \Phi$ (the curl of a gradient is zero) components. So $\mathbf{u} = \nabla \Phi + \nabla \wedge \Psi$. As explained in Chapter Two, the Helmholtz decomposition of the velocity is a necessary step in coarse-graining the velocity. This provides an opportunity for us to examine the local spectra of the divergent and rotational components of the flow. While a Helmholtz decomposition is not strictly a wave decomposition, the divergent component strongly aligns with geostrophically unbalanced motion, mostly consisting of gravity waves at our scales, and the rotational with balanced motions (Avalos, 2024; Waite, 2020). So, we can draw conclusions about the relative role of waves in the mesoscales at different locations using this decomposition.

Helmholtz decomposition has been used in models (Skamarock and Klemp, 2008) and in observations (Callies et al., 2014, 2016). In both cases, at 200 hPa the decomposition shows that globally, the transition from a steeper to shallower mesoscale slope coincides with divergent energy beginning to dominate (in the stratosphere) or becoming of the same magnitude as (in the upper troposphere) the rotational energy.

Given the success of this method, we conjecture that a similar relationship between the relative magnitudes of the divergent and rotational energy in the mesoscales and spectral slope holds locally. In this chapter, we investigate this conjecture and conclude that at 200 hPa a higher local ratio of divergent energy is associated with a greater crossover scale for steepening.

Reproducing the global result

Figure 5.1 shows the spectrum of the rotational (horizontally divergence-free) and divergent (irrotational) components of the global spectrum for 2020. This is obtained by performing a Helmholtz decomposition on the flow and then finding the filtering spectrum associated with the resulting divergence-free and curl-free

components. It is consistent with the spherical harmonic analysis of IFS shown in Li et al. (2023), implying that differences with direct observations are independent of our coarse-graining methodology. We note that the divergent component of the flow has a shallower spectrum than the rotational component. The rotational component is closer to -3, while the divergent is closer to -5/3. Notice that at 200 hPa, the flow's divergent component dominates the rotational component at small length scales, whereas this is not the case at 600 hPa. In observations such as Callies et al. (2016), the divergent component of the flow at 200 hPa tends to be more energetic, leading to a greater cross-over scale at 200 hPa. It is likely that this discrepancy is due to underactive gravity wave activity in IFS (Avalos, 2024). However, qualitatively, our results match observations.




Figure 5.1: The average coarse-graining spectrum for 2020 (black) decomposed into rotational (blue) and divergent (red) components for 600 hPa (left) and 200 hPa (right). The divergent component is less energetic than in observations, but the co-occurrence of mesoscale shallowing of the total spectrum and increased divergent energy relative to rotational kinetic energy remains.

5.2 Results

The local balance between rotational and divergent motion at 200 hPa

In Figure 5.2 we show the average mesoscale divergent energy divided by the rotational energy. In the stratospheric regions at 200 hPa (that is, approximately all regions outside the tropics), given the stratification of the stratosphere and the more prominent role waves play in this environment, we may expect the divergent motion to dominate rotational in the mesoscales. Indeed, in analysis, in the regions outside the tropics, we see divergent motion is of comparable magnitude to rotational or greater. In regions of high orography, the scales are tipped further towards divergent energy. One can assume this is due to orographically induced gravity waves. This is seen particularly in the stratosphere, where the strong stratification lends itself to wave activity. In the upper troposphere (tropics), we see either rotational dominance or divergence and rotational motion of comparable magnitude in regions of high precipitation. There are some regions in the extratropics, particularly on the boundary between the two regions, where rotational motion is quite dominant.

This has confirmed that in the UTLS, wave activity plays a key role. We can assume, due to the bias in IFS towards less wave activity, that in the real atmosphere, values are tipped more towards divergent energy. In the conclusion, we will give an overview of the implications for mesoscale dynamics at 200 hPa of this figure and those in previous chapters.

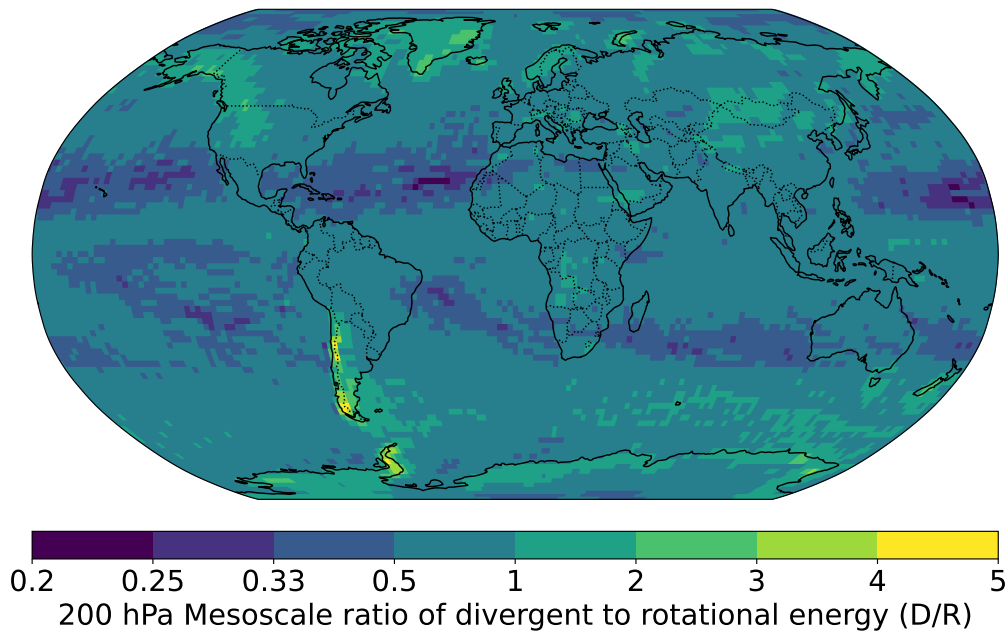


Figure 5.2: A global map of the ratio of the divergent to rotational kinetic energy at 200 hPa averaged over the mesoscales. A logarithmic colour bar is used to ensure equal colour changes correspond to proportional differences in values, i.e. so that the colour change between 1/2 and 1 is the same as that between 1 and 2. At 200 hPa at most locations, mesoscale rotational energy exceeds divergent energy slightly in the ECMWF analysis, likely due to the underrepresentation of gravity wave activity at this altitude.

The balance at 600 hPa

Figure 5.3 shows the equivalent results to Figure 5.2 for 600 hPa. This is a purely tropospheric case. In the convectively active tropics with the reduced influence of planetary rotation, we see the divergent motion is often comparable to or exceeds rotational energy. We also see strong relative divergence associated with orography and precipitation. As we move to higher latitudes, we see rotational energy begin to dominate. Dark blue tracks of strong rotational energy can be seen. These are the result of storms that pass through the tropics. For example, the track in the western Atlantic is Hurricane Teddy, and the strong rotational signal in the Bay of Bengal is Cyclone Amphan. That these storms have such a visible effect on a year's average highlights both the relative calmness of the tropics and the importance of intermittent events. The dominance of rotational energy in the extratropics in Figure 5.3 opens intriguing possibilities about the role of turbulence in the atmosphere. Again, we revisit this in the conclusions of this chapter. This figure also raises the question of how much the slopes we saw at 600 hPa correlate with the energy in divergent motion. We begin to answer this by looking at conditioned spectra.

Spectra conditioned on the ratio of divergent and rotational energy

In Figure 5.4 we investigate the relationship between relative divergent energy and spectra through conditioned spectra. At 200 hPa, we produce a more general version of the classic observation that the ratio of divergent and rotational energy determines the cross-over scale when shallowing occurs (e.g. [Skamarock and Klemp \(2008\)](#); [Callies et al. \(2014, 2016\)](#)). We see that when the ratio of divergent to rotational energy is very high (>1.5), the shallow portion of the spectrum extends to around 600 km, and as the ratio decreases to very low (<0.1), the shallow portion ends much earlier, around 300 km. In addition, the spectrum steepens at scales smaller than 200 km. However, even when rotational motion is dominant, the spectral slope is still shallow. This could imply surface quasi-geostrophic turbulent dynamics, which we revisit in the conclusions.

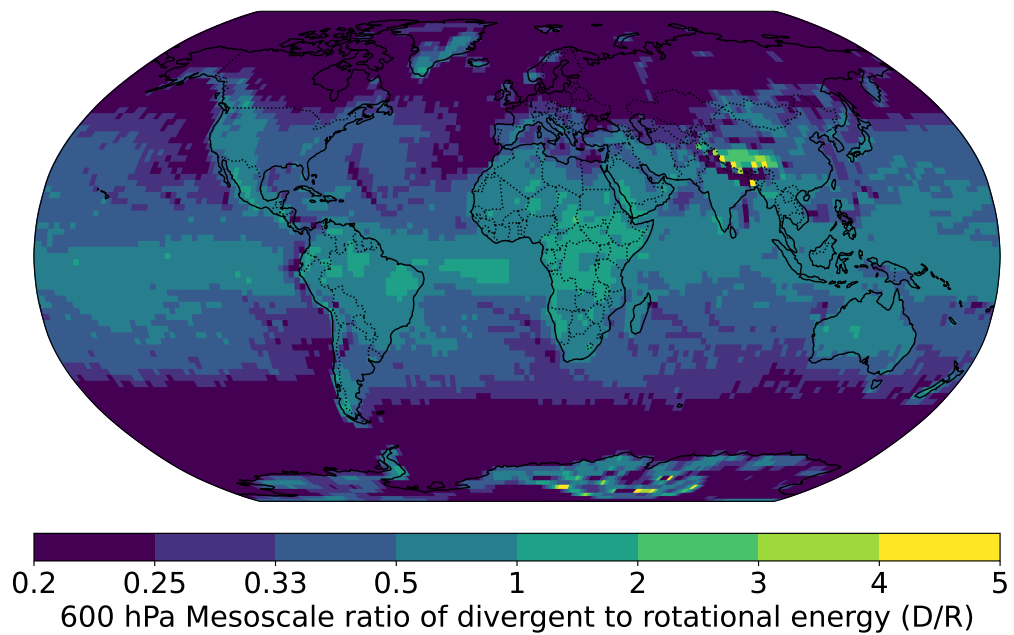


Figure 5.3: A global map of 600 hPa mesoscale divergent divided by rotational energy in analysis data on a logarithmic colour bar. In the extratropics, rotational energy makes up a much larger proportion of total kinetic energy than in the tropics.

At 600 hPa, higher levels of divergent energy are associated with a strong mesoscale shallowing. This is likely due to the association between divergent motion and convection powering small-scale motion having been generated either by moist processes or orography. Higher levels of rotational energy are associated with steeper mesoscale slopes but more energy overall. The increased relative magnitude is likely due to rotational motion dominating in the more energetic extratropics; however, the steep slow may be due to a classic two-dimensional turbulence enstrophy cascade.

Flux conditioned on divergent energy ratio

Figure 5.5 shows the flux conditioned on the relative energy in the mesoscales of divergent to rotational energy. At 600 hPa, we notice that greater rotational energy correlates with a less intense downscale transfer of energy and even an inverse transfer in part of the mesoscales. This may indicate that when rotational turbulent dynamics are dominant, energy injected in the mesoscales can grow to larger scales in an inverse cascade while energy moves to smaller scales in a direct cascade. This would align with the spectral slopes we see in Figure 5.4, which matches the unusual shape of the famous 2D turbulence $-5/3$, -3 spectrum supposing energy was injected around 300 km. When rotational motion is less dominant, we see a strong direct transfer of energy to small scales across the mesoscales. At 200 hPa by contrast, even when rotational motion is dominant, we see an overall direct flux of energy. This is in alignment with surface quasi-geostrophic dynamics, which exhibit a direct cascade of kinetic energy to small scales.

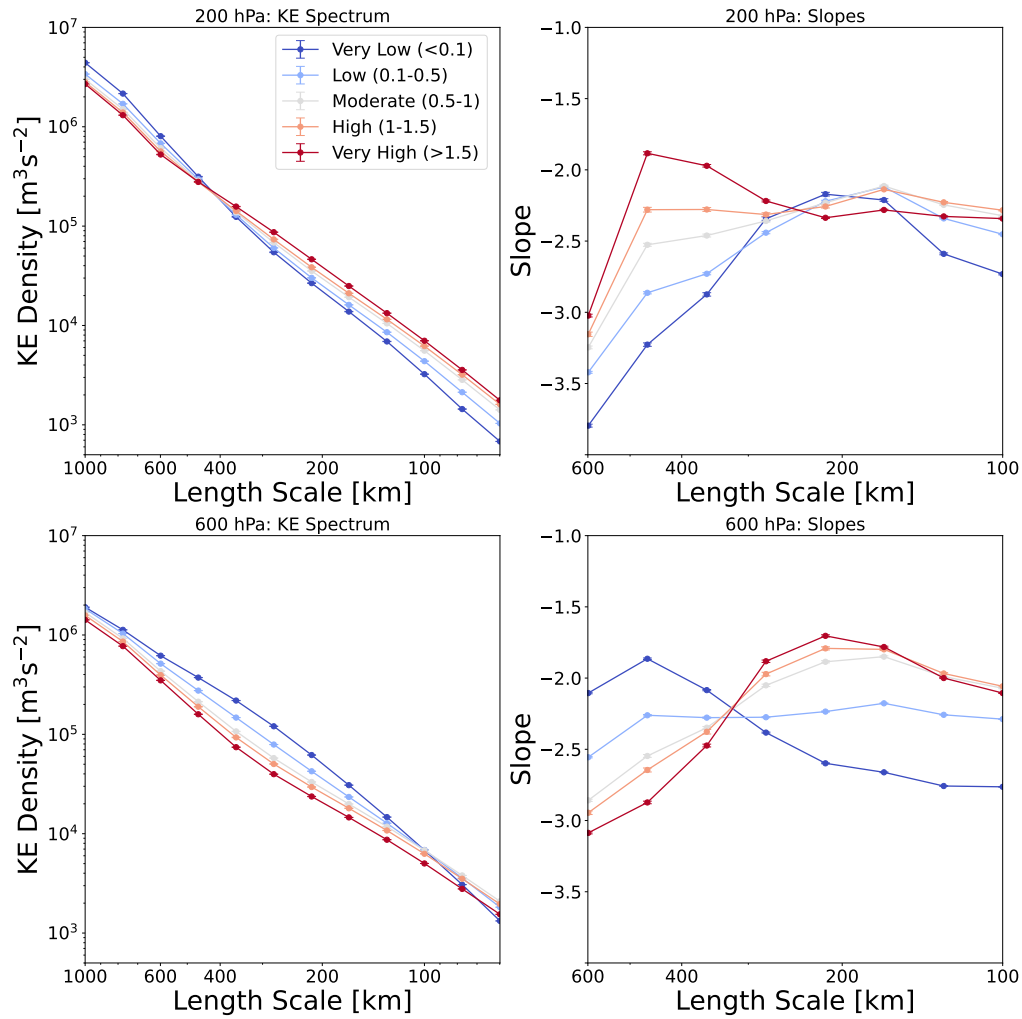


Figure 5.4: Spectra are conditioned on the relative energy in the mesoscales of divergent against rotational energy so that the red line indicates high relative energy in divergent motion and blue indicates high relative energy in the rotational motion. The top panels show 200 hPa with spectra on the left and spectral slopes on the right. Divergent spectra are more energetic and associated with shallower slopes at larger scales. The bottom panels show 600 hPa with spectra on the left and slopes on the right. Rotational spectra are more energetic and display the classic 2D turbulence spectral slopes of shallow large scales and steep small scales (Figure 1.4). By contrast, the divergent spectrum displays a Nastro and Gage-esque spectrum of -3 at large scales and -5/3 at smaller scales.




Figure 5.5: Flux conditioned on the relative energy in the mesoscales of divergent against rotational energy, 600 hPa is shown in the left panel and 200 hPa on the right. Positive values indicate a transfer of energy to smaller scales, negative values indicate a transfer of kinetic energy to larger scales. At 600 hPa we see that spectral fluxes line up with the explanation of the rotational spectral slopes in Figure 5.4 as being caused by two-dimensional turbulence as we see direct and inverse kinetic energy transfers with a crossover roughly corresponding to the steepening of the rotational spectrum at 600 hPa. At 200 hPa, the transfer is direct regardless of the ratio but significantly reduced when rotational is dominant.

5.3 Conclusions

In this chapter, we have extended the common Helmholtz analysis of spectra to local spectra. We now collate the evidence presented over the past three chapters to present our understanding of local mesoscale dynamics.

200 hPa

At 200 hPa we straddle the troposphere and stratosphere with the boundary typically being that of the tropics. However, local changes in mean sea level pressure can cause the tropopause to rise or fall. In particular, high MSLP in the extratropics is associated with a high tropopause, leading to more energetic tropospheric-like power spectra under these conditions (Figure 3.17).

In the stratosphere, divergent motion plays a leading role due to the high stratification lending itself to gravity wave dynamics. This is seen in Figure 5.2. We also notice the relative intensification of divergent motion in orographic regions in the extratropics, implying the shallowing associated with orography in Figure 3.5 is due to topographically induced gravity waves. The gravity wave activity theory is further supported by the global map of slopes we produced in Chapter Three (Figure 3.1). There is a clear resemblance between shallowing at 200 hPa and the spatial distribution of gravity wave activity (see the climatology of gravity wave activity at 100 hPa in [Hocke et al. \(2016\)](#)). While gravity waves can be generated by orography and convection ([Hocke et al., 2016](#)), they can be dampened by other features, such as the suppression of gravity waves over the Himalayas caused by the tropical jet. There appear to be steeper slopes at 200 hPa associated with the jet, particularly when comparing the shallowing seen near the Rockies to that over the Himalayas. This may be due to the suppression of the gravity waves produced by orography, although the interactions between critical layers and topographically induced gravity waves can be complex depending on the details of the atmospheric conditions present ([Clark and Peltier, 1984](#)). In addition, convective activity is a strong source of wave activity in the stratosphere ([Hocke et al., 2016](#); [Dhaka et al., 2006](#)), consistent with the strong correlation between precipitation and shallow spectral slopes at 200 hPa in Figure 3.6. Our results are, therefore, in alignment with studies that posit that upwardly propagating gravity waves primarily energise the mesoscales in the stratosphere ([Callies et al., 2014](#); [Avalos, 2024](#)).

In the tropics, which are within the troposphere, we see divergent motion associated with precipitation. This is likely due to gravity waves induced by convection or divergent motion directly caused by the convective motion. In the tropics and in some regions of the extratropics, there are also regions where rotational motion dominates. In these regions, although weaker than the total flux, the flux of energy due to rotational motion in Figure 4.2 has direct transfer patches, in contrast to the almost purely upscale fluxes in Figure 4.3. We see from Figures 5.4 and 5.5 that rotational motion at 200 hPa is associated with a direct flux of kinetic energy to small scales and a shallow spectral slope. These could match with surface quasi-geostrophic dynamics, proposed as an explanation for the mesoscale shallowing in the UTLS by [Tulloch and Smith \(2006\)](#). These results were discarded due to the observational analysis of [Callies et al. \(2014\)](#), but we now see that they could still explain local spectral behaviours.

600 hPa

600 hPa is entirely within the troposphere. However, Figure 5.3 shows that flow in the upper midlatitudes and poles is heavily rotationally dominated, while divergent motion plays a prominent role in the tropics and lower midlatitudes. The divergent regions could undergo any of the wave explanation theories for the spectral shallowing outlined in Chapter One, however, it is interesting to consider the dynamics in

the rotationally dominated extratropics. The dominance of rotational motion opens the possibility that turbulence theories could be a good explanation for the dynamics in these regions. Indeed, in Figures 5.4 and 5.5, we note that when rotational motion is dominant, we see a spectrum that resembles a two-dimensional canonical spectrum (as in Figure 1.4 in Chapter One). We see an inverse transfer shallow $-5/3$ at large scales and a direct transfer steep -3 at small scales. The mechanism for the associated injection of energy under these conditions at around 300 km remains obscure but could be associated with baroclinic instability. In Figure 4.3, we see that the purely rotational flux of energy is towards larger scales, while the average mesoscale flux is direct in the extratropics, implying the rotational flux of energy is dominated by flux due to rotational and divergent processes. This implies that even if two-dimensional turbulence is a good model for mesoscale dynamics in some regions of the troposphere, divergent processes are likely to be crucially important even in regions dominated by rotational motion.

Higher relative levels of divergent energy are associated with a locally more extensive shallow range

We also note that in this chapter we have shown that in regions where there is a mix of rotational and divergent motion, the ratio of divergent to rotational energy can indicate at what scale the mesoscale shallowing occurs locally. It would be interesting to aggregate direct observations based on their levels of divergent and rotational energy to see if this relationship holds in aircraft observations.

Part II

Predictability

Chapter 6

Spectral error growth theory in 2D turbulence

6.1 Introduction

In the first part of this thesis, we examined the power spectrum of the atmosphere. A commonly cited motivation for work on the spectrum is its link to predictability. This link was established in the work of [Lorenz \(1969\)](#) (L69). In this chapter, we will start by explaining how Lorenz derived the link between spectra and predictability. Our presentation follows that in L69 but is presented slightly more didactically. In particular, we derive a version of his result only using Fourier series, whereas in the original presentation there are two final steps to convert his Fourier series into Fourier transforms and then re-discretise by converting back into discrete wavenumber octaves. I feel these final steps only add unnecessary difficulty in grasping the derivation and they have been omitted. Also omitted is some of the careful mathematical framing Lorenz provides that I feel distract from the core derivation. We then engage in some original research and investigate the effect including a condensate in the power spectrum has on predictability in this model.

The assumed background kinetic energy power spectrum is centrally important to the predictions the L69 model makes for error growth. In 1969, the power spectrum of the atmosphere was practically unknown and Lorenz assumed a universal $-5/3$ slope. His method has been subsequently updated to reflect the kinked -3 to $-5/3$ spectrum observed in [Nastrom et al. \(1984\)](#) ([Rotunno and Snyder, 2008](#); [Durrán and Gingrich, 2014](#)). However, as discussed in Chapter One, this spectrum is quite specific to the UTLS, where, in reality, there is also a significant condensate (flattening of the spectrum) at the largest scales (see Figure 1.2).

We end this chapter by showing how including a condensate of realistic size can more than double the outer limit for predictability in these models while decreasing the predictability times of the mesoscales. This modification to L69 is performed using code shared by Tsz Yan Leung. I modified the spectrum used in that code to include a condensate. This is a result that may have interesting implications for seasonal forecasting. However, it must be acknowledged that turbulence models like L69 have severe limitations in explaining error growth in the atmosphere. We have already seen in Chapters Three, Four and Five that turbulent dynamics are not generally dominant in the atmospheric mesoscales and the beta effect dominates the dynamics at these largest scales, calling into question the relevance of our model.

6.1.1 Key concepts

Error and predictability

Error is the distance by some metric between a particular prediction of the atmosphere and the true evolution. In a thought experiment, a prediction may be generated by a perfect replica of the Earth system but with an imperfect initial condition. The deviation of this alternate universe atmosphere from our own would be called the error.

By predictability, we broadly refer to the time we expect to wait until error saturates, either completely - a situation where any plausible state of the system is equally likely - or partially, either regionally or up to a

particular length scale. For example, a particular state of the midlatitude weather is highly predictable if it takes a relatively long time for error to saturate the synoptic weather scales in that region when initialised from that state.

Error as the foundational object of study in weather prediction

As will be explained in this chapter, in [Lorenz \(1969\)](#), it was argued that certain turbulent flows possess a finite horizon of predictability. That is, no matter how accurately one measures the initial conditions, there is a time beyond which there will be no way to predict the motion.

If this is to be taken seriously, then the object of forecasting turbulent motion, including weather prediction, is not to obtain a high-resolution prediction far into the future but to fully and accurately represent the inevitable growth of error in the forecast. This perspective led to a shift away from deterministic and towards probabilistic forecasting in the early 90s ([Toth and Kalnay, 1993](#); [Palmer et al., 1993](#); [Buizza et al., 1993](#)). Understanding how and why error grows in the atmosphere, both in simulations and theory, remains an active area of research.

Theoretical investigations start from simple turbulence models, building on L69. Experiments in predictability are typically achieved by perturbing full physics simulations of the atmosphere and quantifying the growth of error compared to a control forecast. An advantage of the theoretical approach is the simplicity of the models, which essentially allows for arbitrarily high resolutions, (in this chapter, our smallest scale will be 38 meters). By contrast, high-resolution experiments in error growth are expensive, with contemporary studies limited to a few kilometres ([Judt, 2018, 2020](#)).

6.2 The Lorenz 1969 model

The progenitor of all theoretical turbulence error growth models is the L69 model. We will present the model, explain how it is derived, and explain Lorenz's conclusions.

6.2.1 The model

In the original L69 model and subsequent work on this topic, the scales in the flow are partitioned into n power of two octaves, as in Figures 1.1 and 1.2 of Chapter One. The magnitude of the error kinetic energy power spectrum within octave k is given by Z_k . For simplicity, our pedagogical derivation will arrive at an evolution equation for error at a particular wavenumber instead of octave and so we will refer to wavenumber for the remainder of this section.

The L69 model is to evolve an initial spectral distribution of error Z_k , through the equation

$$\partial_{tt} Z_k = \sum_{l=1}^n C_{kl} Z_l \quad (6.2.1.1)$$

and fix the error at a given scale once it reaches the amplitude of the background spectrum.

The fixing or saturation step must be done since Equation 6.2.1.1 is linear. It is nonlinearity in the turbulent system that causes error growth to slow down as its amplitude increases. If this did not occur, the error would grow arbitrarily large, and the perturbed flow would quickly become physically unrealistic.

In L69, error is defined as the distance of the perturbed run from an average value rather than the control run. This is why error saturates when it reaches the amplitude of the background spectrum. If the latter definition had been used, the error would saturate when it had reached twice the background spectrum, corresponding to the distance of the control from the average plus the distance of the perturbed from the average.

C_{kl} is a constant matrix (given a fixed kinetic energy power spectrum of the background flow) that describes how the error magnitude at wavenumber l contributes to the acceleration of error growth at wavenumber k .

The key to understanding this model and its later modifications is to understand how C_{kl} is derived. In essence, C_{kl} is the product of an interaction term that determines how strongly scales k and l are linked and the power spectrum of the background flow, which tells us how vigorous the motion is at a relevant scale.

6.2.2 The derivation of L69

The L69 model is based on the 2D vorticity equation. This is a poor model for the flow in the atmosphere; it is completely barotropic and only a single layer. However, it contains the nonlinear advection term, which is at the heart of turbulence and therefore it serves as a logical first model of study.

The vorticity equation can be written compactly simply as the conservation of vorticity $\frac{D\omega}{Dt} = 0$. However, it is usually expressed

$$\frac{\partial \Delta \psi}{\partial t} = -J(\psi, \Delta \psi), \quad (6.2.2.1)$$

where ψ is the stream function ($\partial_x \psi = v$, $\partial_y \psi = -u$), $\Delta \psi$ is the vorticity ($\partial_{xx} \psi + \partial_{yy} \psi = \partial_x v - \partial_y u = \omega$) and $J(A, B)$ is the determinant of the Jacobian matrix formed from derivatives of A and B , $J(A, B) = \partial_x A \partial_y B - \partial_y A \partial_x B$.

The assumptions of the L69 model

Three key assumptions must be made in the L69 model:

1. HIT: The averaged turbulent field is homogenous, isotropic and the background flow is stationary.
2. Small errors: The error is small compared to the background flow so our error growth equation can be linearised.
3. Weakened quasi-normal closure: In order to close the spectral error growth equation, Lorenz assumes that quadratic functions of error and quadratic functions of the streamfunction are independent.

The spectral error growth equation for small errors

Suppose we introduce a small perturbation ε to our initial stream function. We wish to learn how this small perturbation grows in time. Substituting $\psi + \varepsilon$ into Equation 6.2.2.1 we find

$$\frac{\partial \Delta \varepsilon}{\partial t} = -J(\psi, \Delta \varepsilon) - J(\varepsilon, \Delta \psi) - J(\varepsilon, \Delta \varepsilon) \quad (6.2.2.2)$$

where the terms involving only ψ cancel, as they are just the original governing equation. Since we assume ε is small we can linearize to remove the last term leaving us with the governing equation

$$\frac{\partial \Delta \varepsilon}{\partial t} = -J(\psi, \Delta \varepsilon) - J(\varepsilon, \Delta \psi). \quad (6.2.2.3)$$

This equation still dictates the complicated motion of a specific pair of realisations of a fluid. We wish to derive results about the general properties of error growth and how that error is communicated through different spatial scales of the fluid. Therefore, we should study the evolution of the error amplitude at each scale.

To this end we write ψ and ε as Fourier series, so we have

$$\psi(\mathbf{x}, t) = \sum_{\mathbf{k}} \hat{\psi}(\mathbf{k}, t) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (6.2.2.4)$$

$$\varepsilon(\mathbf{x}, t) = \sum_{\mathbf{k}} \hat{\varepsilon}(\mathbf{k}, t) e^{i\mathbf{k}\cdot\mathbf{x}}. \quad (6.2.2.5)$$

By substituting in these forms for ψ and ε , Equation 6.2.2.3 can be transformed into a tendency equation for the Fourier coefficient of ε at a wavevector \mathbf{k}

$$\partial_t \hat{\varepsilon}(\mathbf{k}) = \sum_{\mathbf{p}} \hat{J}_\varepsilon(\mathbf{p}, \mathbf{k}) \hat{\psi}(\mathbf{k} - \mathbf{p}, t) \hat{\varepsilon}(\mathbf{p}, t). \quad (6.2.2.6)$$

This is analogous to the derivation of the spectral tendency equation (e.g Vallis (2017)). Here we define $\hat{J}_\varepsilon(\mathbf{p}, \mathbf{k})$ to be $(p_x k_y - p_y k_x) \frac{|\mathbf{k} - \mathbf{p}|^2 - |\mathbf{p}|^2}{|\mathbf{k}|^2}$.

$\hat{J}_\varepsilon(\mathbf{p}, \mathbf{k})$ is equivalent to the interaction coefficient in the spectral tendency equation. It determines how much the error at wavevector \mathbf{p} contributes to the evolution of error at wavevector \mathbf{k} just as the interaction coefficient determines the strength of triad interactions.

The magnitude of the error at wavevector \mathbf{k} is given by $|\hat{\varepsilon}(\mathbf{k}, t)|^2 = \hat{\varepsilon}(\mathbf{k}, t) \hat{\varepsilon}(\mathbf{k}, t)^* = \hat{\varepsilon}(\mathbf{k}, t) \hat{\varepsilon}(-\mathbf{k}, t)$. The last equality follows from the property that if ε is real, $\hat{\varepsilon}(\mathbf{k})$ and $\hat{\varepsilon}(-\mathbf{k})$ must be complex conjugates. We can therefore derive, by using the product rule and some algebra, that

$$\partial_t (|\hat{\varepsilon}(\mathbf{k}, t)|^2) = \sum_{\mathbf{p}} \hat{J}_\varepsilon(\mathbf{p}, \mathbf{k}) (\hat{\psi}(\mathbf{k} - \mathbf{p}, t) \hat{\varepsilon}(\mathbf{p}, t) \hat{\varepsilon}(-\mathbf{k}, t) + \hat{\psi}(\mathbf{p} - \mathbf{k}, t) \hat{\varepsilon}(-\mathbf{p}, t) \hat{\varepsilon}(\mathbf{k}, t)). \quad (6.2.2.7)$$

Closing the equations

In a simpler world, the evolution of the third-order terms that appear in Equation 6.2.2.7 would be easily expressible in terms of the error and background flow coefficient magnitudes. This would allow us to write

a self-contained equation for the evolution of $|\widehat{\varepsilon}(\mathbf{k}, t)|^2$ given a constant background flow (ie Equation 6.2.1.1).

Instead, we find that the evolution equation for the third-order terms involves fourth-order terms. In turn, those fourth-order terms evolve depending on fifth-order terms, and so on, ad infinitum. This is a manifestation of the closure problem in turbulence. In particular,

$$\begin{aligned} \partial_t(\widehat{\psi}(\mathbf{k} - \mathbf{p}, t) \widehat{\varepsilon}(\mathbf{p}, t) \widehat{\varepsilon}(-\mathbf{k}, t)) &= \sum_{\mathbf{m}} \widehat{J}(\mathbf{m}, \mathbf{k} - \mathbf{p}) \widehat{\psi}(\mathbf{m}, t) \widehat{\psi}(\mathbf{k} - \mathbf{p} - \mathbf{m}, t) \widehat{\varepsilon}(\mathbf{p}, t) \widehat{\varepsilon}(-\mathbf{k}, t) \\ &+ \sum_{\mathbf{m}} \widehat{J}_{\varepsilon}(\mathbf{m}, \mathbf{p}) \widehat{\psi}(\mathbf{p} - \mathbf{m}, t) \widehat{\psi}(\mathbf{k} - \mathbf{p}, t) \widehat{\varepsilon}(\mathbf{m}, t) \widehat{\varepsilon}(-\mathbf{k}, t) \\ &+ \sum_{\mathbf{m}} \widehat{J}_{\varepsilon}(\mathbf{m}, \mathbf{k}) \widehat{\psi}(\mathbf{m} - \mathbf{k}, t) \widehat{\psi}(\mathbf{k} - \mathbf{p}, t) \widehat{\varepsilon}(-\mathbf{m}, t) \widehat{\varepsilon}(\mathbf{p}, t). \end{aligned} \quad (6.2.2.8)$$

Here $\widehat{J}(\mathbf{p}, \mathbf{k})$ is the usual interaction coefficient, $\widehat{J}(\mathbf{p}, \mathbf{k}) := (p_x k_y - p_y k_x) \frac{|\mathbf{k} - \mathbf{p}|^2}{|\mathbf{k}|^2}$. The first term is just $\partial_t(\widehat{\psi}(\mathbf{k} - \mathbf{p}, t) \widehat{\varepsilon}(\mathbf{p}, t) \widehat{\varepsilon}(-\mathbf{k}, t))$.

The path we take to close these equations is to consider the statistics of the flow by taking an ensemble average, which we assume yields homogeneous, isotropic and, for the background flow, stationary results and then making a weaker kind of quasi-normal approximation - assuming we can split the fourth-order statistics into second-order statistics in exactly the way we want.

Our homogeneity and isotropy assumptions are useful because they imply that $\overline{\widehat{\psi}(\mathbf{m})\widehat{\psi}(\mathbf{p})} = \delta(\mathbf{m} + \mathbf{p})|\overline{\widehat{\psi}(\mathbf{m})}|^2$ (and the same for error). Here, the overline denotes an ensemble average. If the expression did not include a Dirac delta, the Fourier coefficients of the ensemble average at different length scales would be correlated, and some non-homogenous or isotropic spatial pattern would emerge.

To use this fact and the stationarity, which will allow us to remove the first term in Equation 6.2.2.8, we need to be able to separate out fourth-order statistics into second-order ones. So, we assume we can do this (this is our weaker quasi-normal approximation). This cannot be justified from the equations; it is just a working approximation. This yields us, for example, that

$$\overline{\widehat{\psi}(\mathbf{m})\widehat{\psi}(\mathbf{p})\widehat{\varepsilon}(\mathbf{k})\widehat{\varepsilon}(\mathbf{q})} = \overline{\widehat{\psi}(\mathbf{m})\widehat{\psi}(\mathbf{p})} \overline{\widehat{\varepsilon}(\mathbf{k})\widehat{\varepsilon}(\mathbf{q})}. \quad (6.2.2.9)$$

Applying these assumptions to Equation 6.2.2.8 we find that

$$\partial_t(\overline{\widehat{\psi}(\mathbf{k} - \mathbf{p}, t) \widehat{\varepsilon}(\mathbf{p}, t) \widehat{\varepsilon}(-\mathbf{k}, t)}) = \overline{\widehat{\psi}(\mathbf{k} - \mathbf{p})^2 (\widehat{J}_{\varepsilon}(\mathbf{k}, \mathbf{p}) \widehat{\varepsilon}(\mathbf{k}, t)^2 + \widehat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k}) \widehat{\varepsilon}(\mathbf{p}, t)^2)}. \quad (6.2.2.10)$$

Taking the time derivative of Equation 6.2.2.7 and substituting in our expression for the tendency of the fourth-order statistics we find

$$\partial_{tt}(\overline{\widehat{\varepsilon}(\mathbf{k}, t)^2}) = 2 \sum_{\mathbf{p}} \overline{\widehat{\psi}(\mathbf{k} - \mathbf{p})^2} (\widehat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k}) \widehat{J}_{\varepsilon}(\mathbf{k}, \mathbf{p}) \overline{\widehat{\varepsilon}(\mathbf{k}, t)^2} + \widehat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k})^2 \overline{\widehat{\varepsilon}(\mathbf{p}, t)^2}). \quad (6.2.2.11)$$

The coefficients in Equation 6.2.2.11 are independent of time. Observing that by Plancherel's theorem

$$E(t) = \sum_{\mathbf{k}} \mathbf{k}^2 |\hat{\psi}(\mathbf{k}, t)|^2, \quad (6.2.2.12)$$

$$Z(t) = \sum_{\mathbf{k}} \mathbf{k}^2 |\hat{\varepsilon}(\mathbf{k}, t)|^2, \quad (6.2.2.13)$$

where E is the total kinetic energy and Z is the error energy, it follows we can define the average kinetic energy and error spectrum as $E(\mathbf{k}) = \mathbf{k}^2 |\overline{\hat{\psi}(\mathbf{k})}|^2$ and $Z(\mathbf{k}, t) = \mathbf{k}^2 |\overline{\hat{\varepsilon}(\mathbf{k}, t)}|^2$ respectively. Thus we have

$$\partial_{tt} Z(\mathbf{k}, t) = 2 \sum_{\mathbf{p}} |\mathbf{k} - \mathbf{p}|^{-2} E(\mathbf{k} - \mathbf{p}) (\hat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k}) \hat{J}_{\varepsilon}(\mathbf{k}, \mathbf{p}) Z(\mathbf{k}, t) + \hat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k})^2 |\mathbf{k}|^2 |\mathbf{p}|^{-2} Z(\mathbf{p}, t)). \quad (6.2.2.14)$$

Here we diverge from L69. Lorenz then takes the limit to convert from Fourier series to Fourier transforms. Then, he re-discretizes the model by integrating over continuous wavenumber bands.

However, we can simply write

$$\partial_{tt} Z(\mathbf{k}, t) = \sum_{\mathbf{p}} C(\mathbf{k}, \mathbf{p}) Z(\mathbf{p}, t) \quad (6.2.2.15)$$

with

$$C(\mathbf{k}, \mathbf{p}) = 2 \frac{|\mathbf{k}|^2}{|\mathbf{k} - \mathbf{p}|^2 |\mathbf{p}|^2} \hat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k})^2 E(\mathbf{k} - \mathbf{p}) + \delta(\mathbf{k} - \mathbf{p}) \sum_{\mathbf{p}} 2 \hat{J}_{\varepsilon}(\mathbf{p}, \mathbf{k}) \hat{J}_{\varepsilon}(\mathbf{k}, \mathbf{p}) |\mathbf{k} - \mathbf{p}|^{-2} E(\mathbf{k} - \mathbf{p}). \quad (6.2.2.16)$$

If we then sum across wavenumber bands, we will arrive at an equation of the same qualitative form as Equation 6.2.1.1.

$C(\mathbf{k}, \mathbf{p})$ may appear complicated, but it is independent of time and depends only on the background kinetic energy power spectrum. We have therefore derived a highly simplified model for error growth in turbulence.

6.2.3 Interpreting the L69 model

We have reduced the complicated problem of turbulent error growth into a simple system of linear equations. The acceleration of error growth at wavenumber k is given by the weighted sum of error at all other scales. To understand how our model behaves, we can inspect the weighting coefficients $C(\mathbf{k}, \mathbf{p})$.

We shall stay with the discrete model we just derived since our conclusions will be completely analogous to Lorenz's conclusions. Summing across wavenumber bands and setting our energy spectrum to follow a $-5/3$ or -3 power law (normalised so the total energy is 1 in either case) we find that $C(k, p)$ takes the values indicated in Figure 6.1.

In Figure 6.1 the k wavenumber band is the scale of the error whose evolution we are modelling due to the action of the error at scale p . We have a few observations. In this model, as in Lorenz's, the coefficients along the diagonal are generally negative and large while those on the off-diagonal are large and positive. This indicates that the error will spread out to nearby scales. Hence, if the error starts at a small scale, it

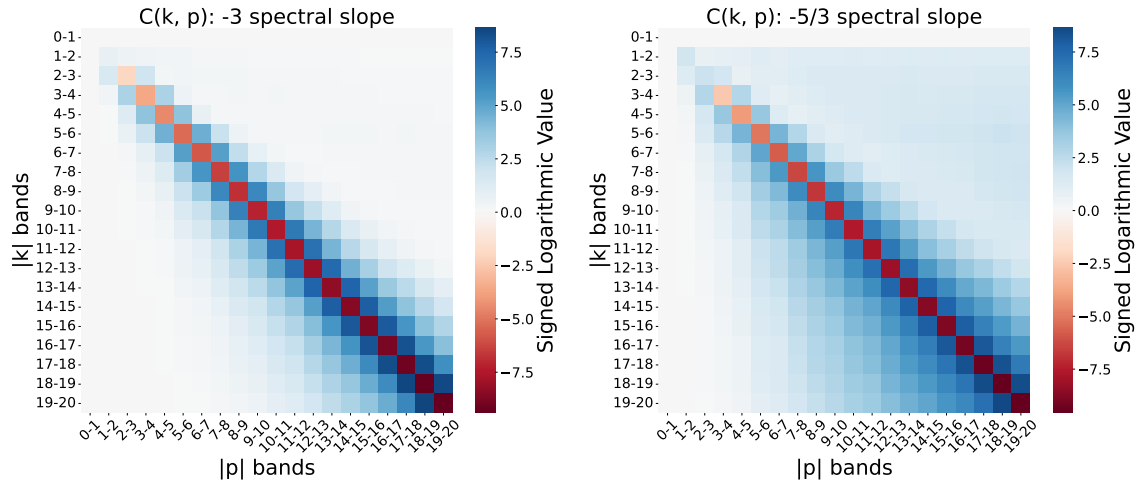


Figure 6.1: $C(k, p)$ is the weighting function that determines how much error at wavenumber p contributes to the acceleration of error growth at wavenumber k . Plotted is the signed logarithmic value of $C(k, p)$ (diagonal values tend to be heavily negative). The y-axis k values show the wavenumber band of the error whose temporal evolution the model is yielding. This means that reading vertically up a column tells us how error at scale p is coupled to the evolution of error at other scales k . The left plot corresponds to a spectral slope of -3 , while the right plot corresponds to a slope of $-\frac{5}{3}$.

will strongly affect error growth in the scale immediately larger than it. This is where the justification for a “cascade” of error in turbulence originates. The -3 spectral slope, with the more significant differences in energy between neighbouring scales, is more local in error growth coupling than the $-5/3$ regime. This nonlocality manifests asymmetrically, with the $-5/3$ case exhibiting a more substantial downscale nonlocal transfer of error than a nonlocal upscale transfer. The consequences of this strong downscale coupling were explored in [Durrán and Gingrich \(2014\)](#), where it was noted that small errors in the large-scale flow can swamp out large errors at the small scales. This led them to argue that “butterflies are not of practical importance”.

The next step in exploring this model is to solve it and see how quickly it predicts error will grow. When Lorenz did this, he observed that the time it took for each scale to saturate with error went like

$$\tau(k) \sim k^{-\frac{3+\beta}{2}} \quad (6.2.3.1)$$

where β is the spectral slope.

Lorenz concluded that the total predictability time could be given by summing up these saturation times and therefore that, since this is a geometric series, if the base in the series is less than one, there must be a finite horizon of predictability.

He concluded that flows with a spectral slope less than -3 are predictable arbitrarily far into the future, while flows with a slope of $-5/3$ have a finite point beyond which they cannot be predicted. We will return to this in more detail in the next section.

6.2.4 Limitations of the L69 model

Updates to the model

Since Lorenz's time, his model has only undergone minor modifications. First, in [Rotunno and Snyder \(2008\)](#), the model was extended to the surface-quasi geostrophic equations. These are a simple 2D model for surface heating-driven flow in the atmosphere that simplify to the conservation of buoyancy.

The reason for the modification was a criticism of the L69 model that it assumed a $-5/3$ spectral slope for the 2D vorticity equations, which develop a -3 slope under large-scale forcing and so, unless forced at an arbitrarily small scale, will exhibit a -3 slope for small enough scales and hence are not an example of equations with a finite predictability limit. However, we should still examine the 2DV equation with a $-5/3$ slope in the L69 model. Until the work of [Cho and Lindborg \(2001\)](#), this was a competitive explanation for the $-5/3$ observed in the atmosphere and may still be relevant in certain locations or under certain conditions (Part One of this thesis).

The SQG equations do not suffer from the same criticism since they exhibit a $-5/3$ slope under large-scale forcing, so they may have a finite horizon of predictability. This feature also puts them in line with the post-2000 atmospheric observations. Unfortunately, they assume geostrophic balance and so were ruled out as an explanation for the observed global mesoscale slope in [Callies et al. \(2014\)](#). In either case, qualitative differences between SQG and 2DV error growth were found to be small in [Rotunno and Snyder \(2008\)](#), with the spectral slope playing the dominant role rather than the dynamics.

The second modification was in [Durrán and Gingrich \(2014\)](#) (the ssLRS model). This was to introduce a gradual saturation of error rather than a sudden imposition of nonlinearity through a cutoff. Again, the qualitative effect of this change was minimal.

A further version of the model was presented in [Sun and Zhang \(2020\)](#). This was to consider the kinked spectrum as observed in the upper troposphere and lower stratosphere. Using the functional fit to this spectrum presented in [Lindborg \(1999\)](#), the -3 portion is used to evolve error under 2DV dynamics while the $-5/3$ portion evolves error under SQG dynamics. This approach makes little difference compared to simply using the L69 model with the [Lindborg \(1999\)](#) spectrum, as is implied in the work of [Leung et al. \(2020\)](#).

In [Leung et al. \(2020\)](#), the original L69 model, which, as we have seen, is just as valid for the atmosphere as the subsequent modifications, was compared with 2DV simulations with forcing at both large and small scales to form a kinked spectrum. The results were favourable, with the error growth well modelled by L69 when the simulation's power spectrum was provided as input for the model.

We have kindly been allowed to use Leung's implementation of the L69 model in the next part of this chapter.

Limitations of Lorenz models

The most fundamental limitation of Lorenz models is that they are based only on simple 2D turbulence models (2DV or SQG). These do not capture the three error propagation processes observed by [Zhang et al. \(2007\)](#). There is no representation of moist physics, gravity waves or baroclinic instability. The flows we study are divergence-free and evolve on a single layer.

Limiting our discussion to 2D turbulence, the next major issue is the homogeneity assumption. Even

if 2DV was a good model of the atmosphere, we should still like to account for features such as mean flows, mountain ranges and inhomogeneous forcing from radiation, for example. In the atmosphere, these features change the local energy distribution (Chapter Three), which may impact the predictability at different locations. In the next chapter, we will discuss improving this aspect of the model.

The next limitation is the assumption of small error. Ideally, one would repeat the present derivation but include the nonlinear self-interaction of error. However, dealing with this nonlinearity is nontrivial. This issue is partially ameliorated in the ssLRS model through the tapering off of error growth as it approaches the scale of the background flow. However, an issue that persists is that if we were to consider a background flow with no energy below a particular scale (a simple model of a jet, for instance), our current model would suggest that perturbing this model at scales smaller than the cutoff would result in little to no error growth. This is because a perturbation in our model cannot interact with itself and grow up to the large scales as it can in a 2DV simulation.

Finally, there is the issue of the turbulence closure used. Work using other closures ([Leith and Kraichnan, 1972](#); [Métais and Lesieur, 1986](#)) has come to the same qualitative conclusions, indicating at least some robustness to the closure employed.

6.2.5 A dimensional model for error growth

We will now arrive at the same conclusions as Lorenz by a more straightforward method ([Lilly, 1973](#); [Palmer et al., 2014](#)). If we link the growth rate of error at a particular scale to the eddy turnover time at that scale, we will recover the observation of Lorenz that the time it takes for error to saturate goes like

$$\tau(k) \sim k^{-\frac{3+\beta}{2}} \quad (6.2.5.1)$$

where β is the spectral slope.

A simple dimensional model for error growth

The eddy turnover time is the timescale for an eddy of lengthscale $\sim 1/k$ (wavelength is $2\pi/k$ wavenumber) to rotate once. Since the kinetic energy at wavenumber k is given by the spectrum $E(k) \sim U^2/k$, this implies that since $t = l/U \sim 1/Uk$,

$$\tau(k) = A k^{-3/2} E(k)^{-1/2}. \quad (6.2.5.2)$$

Supposing that error is initially confined to the smallest scale of the system $2^N k_0$, and that it progresses from scale to scale through a scale doubling cascade, saturating the previous scale before the one twice as large begins to saturate, we may estimate the total predictability time T as

$$T = \sum_{n=0}^N \tau(2^n k_0). \quad (6.2.5.3)$$

Seagulls, butterflies and grasshoppers

The key insight of [Lorenz \(1969\)](#) was that there might be a finite time beyond which there is no way to predict the flow, no matter how small the initial error.

This has been put poetically (but as Lorenz himself points out misleadingly) as the flap/flight of a seagull/butterfly/grasshopper changing the global weather in about two weeks. A simple argument to see that the model cannot be interpreted so straightforwardly is that it also predicts that after a few hours, every storm on Earth would have been altered by the critter. There is no mechanism for such a signal to propagate so quickly.

Regardless, if we consider a spatially uniform error at a small spatial scale, like noise in the details of the flow, then, supposing the kinetic energy spectrum follows a power law $E(k) \sim k^\beta$,

$$T \sim \sum_{n=0}^N \left(2^{-\frac{3+\beta}{2}}\right)^n. \quad (6.2.5.4)$$

This is a geometric series, so it will converge if the base is less than one. Hence, we will have a finite predictability time. Hence we can recover Lorenz's conclusion that a slope of -3 is the boundary between finite and unlimited predictability.

6.3 The effect of a condensate on L69

We will now consider how including a realistic condensate affects the predictions L69 makes for the total predictability time of the atmosphere. For details of the coded implementation, see [Leung et al. \(2020\)](#).

For comparison, we first provide the [Lindborg \(1999\)](#) UTLS observational spectrum fit to the L69 model. We normalised and then adjusted the magnitude of the spectrum so that the dimensionalised saturation time in the mesoscales matched with full twin experiments such as [Judt \(2018\)](#) and yielded saturation in between five and six days for the 400 km mesoscale boundary. This is shown in Figure 6.2.

Several lengthscales are annotated with a black dashed line. One kilometre represents the current cutting edge of weather prediction resolution. Nine kilometers is around the current standard for forecasting center operational resolution for global models. We see that our turbulence model predicts that if a weather model could perfectly represent features up to its resolution, a 1 km model would be able to predict around 7 hours further ahead than a 9 km model. However, if we assume that the effective resolution heuristic of [Skamarock \(2004\)](#) holds, then a 1 km model represents features up to 7 km while a 9 km model up to 63 km. In that more realistic case, we would expect a 1 km resolution model to predict on the order of 1 day further ahead than a 9 km model.

We also show 6000 km, this scale is the weather scale. Once it has been saturated we no longer have any information that could reasonably be called weather and are totally concerned with more general properties, for example, one may imagine retaining skill in predicting average wind speed over the equator. This model gives 12 days for this to occur. Once the outer scale is saturated, there is no predictability left in the system. In this model total loss of predictability occurs after just over 19 days.

In Figure 6.3 we show the effect of including a condensate in the Lindborg fit to observations. The Lindborg fit is based on aircraft measurements like those in [Nastrom et al. \(1984\)](#). The start of a condensate is seen in these results but the tracks are generally not long enough to properly resolve the condensate and it is clearly not considered in the empirical fit. The condensate is well resolved in analysis. In Figures 1.1 and 1.2 of Chapter One, we show the spectrum in analysis and note the condensate. We have set the first three octaves (up to wavenumber 8), to equal the value of the spectrum at octave four. If we considered lower

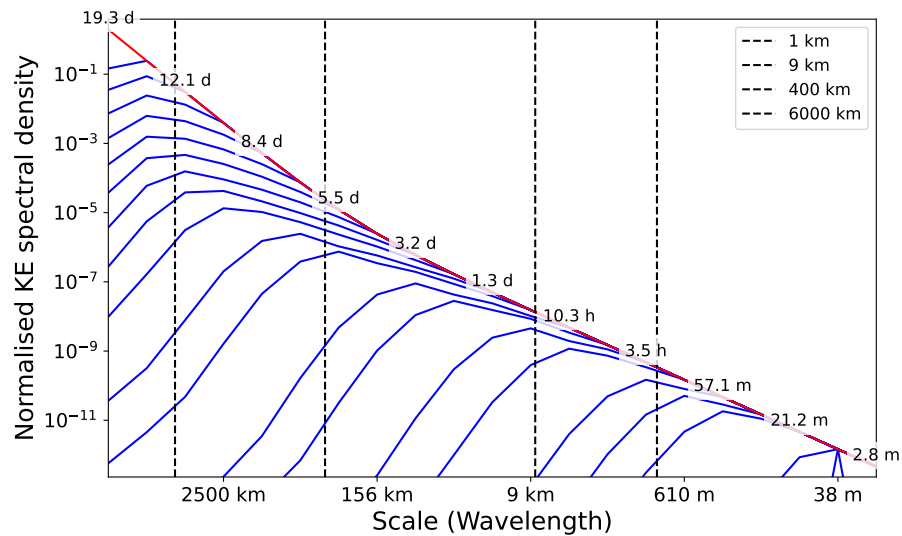


Figure 6.2: L69 error growth predictions from Lindborg fit. Magnitude adjusted to match mesoscale saturation times observed in twin experiments of Judd (2018). The background kinetic energy power spectrum is shown in red. In blue are the profiles of error at the moment when a scale is saturated. Next to every other of these blue profiles is the dimensionalised time it took for the error to reach the given scale. For instance, the error saturates 6000 km after 12.1 days and is significant at larger scales as seen from the height of the blue line originating at the spectrum at 6000 km. Dashed black lines have been added at 6000 and 400 km, which are the synoptic and mesoscale boundaries, and 9 and 1 km, which indicate the current standard weather model resolution and a next-generation resolution. By design, the mesoscales saturate in just over 5 days, the weather scales are then implied to saturate in under two weeks at 12 days. All skill is predicted to be lost after 19 days.

altitudes, an earlier flattening would have been appropriate, but we would also have to shallow the spectral slope in the synoptic scale portion (see Chapter One).

The effect is reduced total kinetic energy in comparison with Figure 6.2. From a dimensional analysis perspective, it stands to reason that lower kinetic energy leads to longer eddy turnover times and so greater predictability. Indeed, this is exactly what we see at the outer and synoptic scales, with the predictability of weather improving by 3 days when a condensate is included and the outer limit of predictability more than doubling to over a month at 43 days.

It is interesting to note that the predictability at the small scales is apparently worsened by the inclusion of a condensate, given that this inclusion reduces the total kinetic energy. The saturation time of the mesoscales decreases by almost a day. In the previous case, we estimated that a 1 km model could predict 7 hours further ahead than a 9 km model, but when effective resolution was accounted for, this increased to around one day. When a condensate is included, this estimate decreases to 4 hours and around half a day respectively. We can explain this phenomenon by inspecting the error growth profiles more closely.

If one looks closely at the total error, one can see it is significantly greater when a scale is saturated in Figure 6.2 compared to Figure 6.3. Take, for example, the scale that is saturated after 21.1 hours in Figure 6.3. In Figure 6.2 this is saturated after 1.3 days but the associated blue line, which shows the total error at all scales when that scale was saturated, is peaked around 400 km whereas, in the condensate figure, it is peaked much closer to the saturated scale.

This is due to the nature of the matrix of coefficients $C(\mathbf{k}, \mathbf{p})$. $C(\mathbf{k}, \mathbf{p})$ is characterised by strong negative values along the diagonal buttressed by strong positive values, as in Figure 6.1. The magnitude of both the positive and negative values are determined in part by the magnitude of the background spectrum. When




Figure 6.3: L69 error growth predictions from Lindborg fit with a condensate added. As in Figure 6.2, blue lines indicate error growth profiles and the background spectrum is shown in red. Black dashed lines indicate 6000 and 400 km, which are the synoptic and mesoscale boundaries, and 9 and 1 km, indicate the current standard weather model resolution and a next-generation resolution. The time for mesoscale saturation has decreased despite the overall decrease in kinetic energy in the system compared to the spectrum without a condensate. Saturation time decreased to 4.7 days from 5.5 days in Figure 6.2. However, note that the error profile associated with the saturated mesoscales has significantly less error at larger scales than in Figure 6.2. The error at 6000 km saturated at 15 and a half days, an increase in predictability of 3 days compared to when the condensate was excluded. The outer scale of 40000 km saturated after 43 days, more than double the outer limit of predictability without the condensate.

the total spectrum decreases, in this case, not only has the magnitude of the offdiagonal values decreased, but so has the magnitude of the negative values. This decreases the diffusion of error out of a scale to the neighbouring ones, and so error at a particular scale can reach saturation more quickly. In this way, even though error saturates more quickly in the mesoscales when a condensate is included, the overall error is less when that scale saturates than in the case without the condensate. What this corresponds to in physical space is unclear. It may represent weakened large-scale shears that allow vortices to grow with less external interruption.

6.4 Discussion and conclusions

In this chapter, we have shown how Lorenz derived the L69 model, how that model has been modified since its conception and its limitations. We have then compared the predictions the model makes when provided with a more realistic spectrum that includes a condensate with those based on the Lindborg spectrum. Using the Lindborg spectrum is contemporary and has been considered recently in [Sun and Zhang \(2020\)](#). We conclude that including a condensate decreases the predictability of the mesoscales by around a day, but increases the predictability of the weather scales by 3 days. Most impressively, it increases the outer limit of predictability by more than double to 43 days.

It seems to the author that a common intuition amongst those working in seasonal predictability is that any signal further ahead than two weeks must be due to some phenomena with a longer scale predictability than turbulent weather. Whether this be sea surface temperatures, atmospheric phenomena with longer time scales such as the QBO or some other source. However, the model we have just used is purely turbulent. It has no complicated mechanisms, such as gravity wave breaking or coupled ocean dynamics. Yet in this model, skill originates from the initial conditions for more than 40 days, simply by adjusting the magnitude

of the power spectrum to more accurately match that of the atmosphere.

If this result is found to apply in more complicated settings (specifically in a version of Lorenz's model formulated on a beta plane), it may mean that an emphasis should be placed on accurately reproducing the spectrum of the condensate in seasonal models, to best exploit the predictability that can be gained from the initial conditions of the system. Support for the role of the condensate extending predictability may be present in [Judt \(2018\)](#). In that study, spectra saturated the largest scales most slowly when the spectrum had the most prominent condensate. Further, many spectra did not fully saturate over the 20-day period studied. Two-dimensional turbulence is associated with non-divergent rotational motion. It is this motion that dominates the condensate and this component of the error that grew slowest at the large scales in [Judt \(2018\)](#).

Chapter 7

Towards a local nonlinear model for error growth

7.1 Introduction

In the previous chapter, we introduced the turbulent error growth model of [Lorenz \(1969\)](#). This showed how error growth can be related to the power spectrum of the background flow. In Part One of this thesis, we examined the local properties of the kinetic energy power spectrum. We discovered large differences in the spectrum at different locations and under different conditions. We now start investigating the relationship between the local spectrum and error growth.

This is a long and important project that should involve the eventual local analysis of error growth in full physics simulations of the atmosphere using coarse-graining techniques and improvements to the work we present in this chapter. We discuss future directions in Chapter Ten. However, we shall begin by investigating a toy model for local error growth.

First, we give a broader background on predictability in the atmosphere, so that we have the proper context to understand the limitations of our toy model.

A holistic understanding of error growth in the atmosphere

The following is a conceptual model summarising the understanding of error growth in numerical weather prediction gained from numerical simulation investigations (for instance, perturbed twin experiments in full physics models of the atmosphere) and case studies. Error is observed to grow in the following stages:

1. Convective instability: Initially, errors are associated strongly with convective instability and moist processes.
2. Geostrophic adjustment: These errors propagate, likely through gravity waves, and begin to reach the synoptic scale.
3. Barotropic instability: Barotropic or nearly barotropic motion dominates the final and slowest error growth phase.

This three-stage model was first presented in [Zhang et al. \(2007\)](#), although at that time it was assumed that the synoptic scale error growth was dominated by baroclinic rather than barotropic motion. It was the work of [Baumgart et al. \(2018\)](#); [Baumgart and Riemer \(2019\)](#) that showed the final stage was actually dominated by quasi-barotropic motion. Aspects of this perspective have also been reinforced in the work of [Selz and Craig \(2015\)](#) and [Bierdel \(2017\)](#), which examined the problem spectrally and via Helmholtz decomposition. Further support for this pattern of error growth is found in [Judt \(2018\)](#) and [Selz et al. \(2022\)](#).

These three processes are, with the possible exception of the final stage, not spectrally local; they are spatially local. This challenges the utility of the spectrum-based perspective for predicting error growth in the atmosphere. However, our toy model marries both spatial and spectral information, and so has a chance of reproducing these local characteristics, unlike L69.

Our dimensional model

In the previous chapter, we showed a dimensional model for error growth based on L69 (Lilly, 1973; Palmer et al., 2014). There are two tuneable parameters in this dimensional model, the first is the step up the error is assumed to take after it saturates a scale. In Lilly (1973) and Palmer et al. (2014), the model begins to saturate the scale twice as large once it has saturated the current scale. Our dimensional model operates similarly. The second tuneable parameter is the coefficient A , the proportionality between the turnover time and the time for error to saturate a scale (Equation 6.2.5.2).

By comparing our time estimates with both the L69 model, as tabulated in Leung et al. (2020), and the full physics simulations of Judt (2018), we choose $A = 4$. So our model for the predictability time of scale $2^N k_0$ is

$$T_N = 4 \sum_{n=0}^N (2^n k_0)^{-3/2} E(2^n k_0)^{-1/2}. \quad (7.1.0.1)$$

We can use this along with local kinetic energy spectra to estimate the time error takes to saturate a particular scale, or, the scale error saturates by a particular time. By integrating the area of the local spectrum less than this scale, we then have a prediction for the total error growth after a given period of time.

We do not use L69 directly since there is no reason to suspect this should be a better model for local error growth than our dimensional model (L69 was not derived in the context of local spectral slopes), yet it is significantly more computationally expensive and complicated to implement. If we could derive a similar model to L69 through coarse-graining, then this would be a worthwhile upgrade to our system. However, for the time being, our objective is to assess the patterns of error growth that emerge when local eddy turnover time is of central importance, and this approach makes the most sense for that goal.

A weakness of the model is that it does not consider advection. Error may grow in a region with short eddy turnover times, for example, over the Andes, and then be advected by the flow to a position over Argentina. This would cause greater levels of error than one would expect solely from the local power spectrum.

In this chapter, we are interested in comparing patterns of error growth when global errors are within the mesoscales. To this end, we investigate forecast errors after a single day (mesoscales are saturated in ECMWF IFS forecasts after around two days) and errors in our local dimensional model after five days (mesoscales are saturated in L69 models after around five days).

Choosing the parameter A

Judt (2018) investigated error spectra in a convection-permitting, 4km resolution, global simulation. The error was seeded as tiny Gaussian noise in the temperature field (standard deviation 0.01 K). The mesoscale error was observed to be mostly saturated in around five days. However, the largest scales had not yet saturated by 20 days, which was the longest time considered. Employing Lorenz-type models using the Lindborg (1999) fit to the atmospheric spectrum, Sun and Zhang (2020) predicted mesoscale saturation in around five days and large-scale saturation in 20 days.

In Figure 7.1, we demonstrate that the dimensional model with $A = 4$ matches these predictions. To obtain an estimate for the predictability of the smallest scale, the spectrum beyond that point was assumed to

follow a pure $-5/3$ law whose magnitude was determined so there was no discontinuity in the spectrum. The value of the spectrum at doublings of the smallest scale, which was 16 km - roughly the resolution of a $1/8$ th degree model - was then used to estimate predictability time. We see that the picture is completely consistent with the results of the Lorenz-type model. Mesoscale saturation is complete in 5 to 6 days, and the largest scales saturate at around 20.




Figure 7.1: Dimensional error growth predictions using the Lindborg fit to observations for the power spectrum. The background spectrum is shown as a blue line, dashed red lines show the assumed error growth profile. Black arrows indicate the saturation time for each scale. By design, the mesoscales saturate in a similar amount of time to Figure 6.2. Our dimensional model then implies similar results at 6000 km and the outer scale to Figure 6.2, indicating that it produces roughly comparable results while still being a dimensional model, and so is somewhat more general and flexible for our purposes.

The unreasonable effectiveness of dimensional arguments

Since this dimensional argument is untethered from the arguments made in L69, beyond receiving its justification for considering a cascade of error from that work, it is, in some sense, more general. This argument will be valid, so long as error grows through a cascade and the relevant timescale is determined by the kinetic energy at the relevant lengthscale.

The same kind of generality is seen when using a dimensional argument to derive the $-5/3$ power law. If one assumes a constant flux of kinetic energy from scale to scale, and timescales given by the eddy turnover times, one must conclude that the power spectrum obeys a $-5/3$ power law (see Chapter Eight). But this argument is equally valid whether the cascade mechanism is turbulence or a turbulence-catalysed gravity wave-wave interaction (Lindborg, 1999).

Of course, there is no reason to assume that a local spectral cascade, whose timescale depends on the eddy turnover time, is the mechanism by which error grows in the real atmosphere. As we have touched on, the eddy turnover time is not the only relevant timescale observed in error growth in full physics simulations. For instance, the time for gravity waves to undergo geostrophic adjustment, which may well be a relevant timescale for error growth in the mesoscales (Selz and Craig, 2015), is

$$T = \frac{\pi}{f}. \quad (7.1.0.2)$$

However, being mindful of the generality of our model prevents us from concluding that where it behaves well, a turbulent cascade must be the mechanism for error growth. Although we are able to conclude that where it performs poorly, a different mechanism is likely responsible.

Comparison with other methods

Local predictions for error growth are not new. A well-established approach is to linearise the equations of motion and analyse perturbation growth using singular vectors or Lyapunov exponents. These approaches are fundamentally different from the toy model we propose in that they are linear. The saturation of scales in our approach makes our model nonlinear. If this model can be further developed, for example, by accounting for advection by the background flow (as is done in linear stability methods), it would be interesting to compare its predictions with these methods, particularly in a more straightforward setting such as the 2DV equations.

7.2 Applying the dimensional model to the atmosphere

We will now apply our dimensional model to the atmosphere. We will start by applying it to the spherical harmonic spectra shown in Chapter One. We will then apply it to the local coarse-grained spectra explored in Chapter Three.

Considering the spectral condensate

In the previous chapter, we saw that including the condensate significantly increased the total predictability of the atmosphere. We now investigate what the dimensional model predicts when provided with the background spectra from IFS analysis for 2020 at two pressure levels. The background spectra, along with the values fed to the model shown as red dots, are shown in Figure 7.2.




Figure 7.2: Spherical harmonic power spectra at 200 (left) and 600 (right) hPa. Values at wavenumber octaves, which are given to the dimensional model, are shown as red crosses. Red crosses do not extend further since the method will be compared with local power spectra, whose quality begins to be compromised at the next highest wavenumber octave.

The result of our dimensional model applied to these background spectra is shown in Figure 7.3. Since the model takes no information from larger scales while calculating error growth in small scales, the decrease

in mesoscale predictability we found in the previous chapter cannot be replicated. However, in agreement with the previous results, a significant increase in the outer limits of predictability is observed.

In the previous chapter, we applied the L69 model to the Lindborg empirical fit to observations, with the addition of a condensate based on the background spectra extracted from ECMWF analysis for 2020 (as in Figure 7.2). This spectrum best corresponds to that found in the UTLS, and so is most comparable to the 200 hPa spectrum from ECMWF analysis. When applying our dimensional model to the 200 hPa spectrum (Figure 7.3, left panel), we find that the model predicts synoptic scale weather to become saturated in around 20 days, in contrast with the prediction of L69 of 14 days. The largest discrepancy is that the outer limit of predictability is 96 days compared to 43 in the L69 model. This bias towards longer predictability times is almost certainly the result of assuming error grows successively through scales. In our L69 integration (Figure 6.3), the blue error lines show that the error had already grown significantly at 40000 km once it saturated 20000 km.

The right panel of Figure 7.3 shows the predicted error growth at 600 hPa. As is shown in Figure 7.2, the magnitude of the 600 hPa spectrum is equal to or greater than the 200 hPa spectrum in the mesoscales, while it is smaller than at 200 hPa at the large scales. This leads to longer eddy turnover times at 600 hPa at the larger scales. We see that in Figure 7.3, the predicted time for the mesoscales to saturate is consistent at the two altitudes. As the synoptic scale is reached, the predictions begin to diverge, with a prediction of 32 days for the saturation of the weather scale of 6000 km as compared to 20 days in the dimensional model at 200 hPa and 10 to 14 days according to simulations (Judt, 2018) and the L69 models run in the previous chapter. The outer limit of predictability is starkly increased, with 166 days being returned for the outer scale in the dimensional model. This value is more than 1.7 times larger than the outer limit predicted at 200 hPa.

If the average 600 hPa spectrum was provided to L69, according to the results of the previous chapter, we would likely see more cascade-like error growth due to a decrease in the divergence of error out of its current wavenumber. This may decrease the predictability time for the mesoscales compared to the 200 hPa. In addition, due to the lower kinetic energy in the condensate, the outer limit of predictability obtained from the 600 hPa integration will likely be longer than the 43 days obtained in the previous chapter. If an unrealistically long saturation time is predicted, it may be an indication that 2DV is inadequate for studying error growth at the largest scales or that vertical propagation of error is playing an important role at large scales.




Figure 7.3: The dimensional model applied to the ECMWF analysis for 200 (left) and 600 hPa (right) is shown. The model predicts weather scale saturation (around 6000 km) in around 2 weeks and a half for 200 hPa and around 3 weeks for 600 hPa. The outer scales are significantly longer with the outer limit of predictability being 96 days at 200 hPa and 166 days at 600 hPa.

RMSE error as comparison for spatial distributions of error

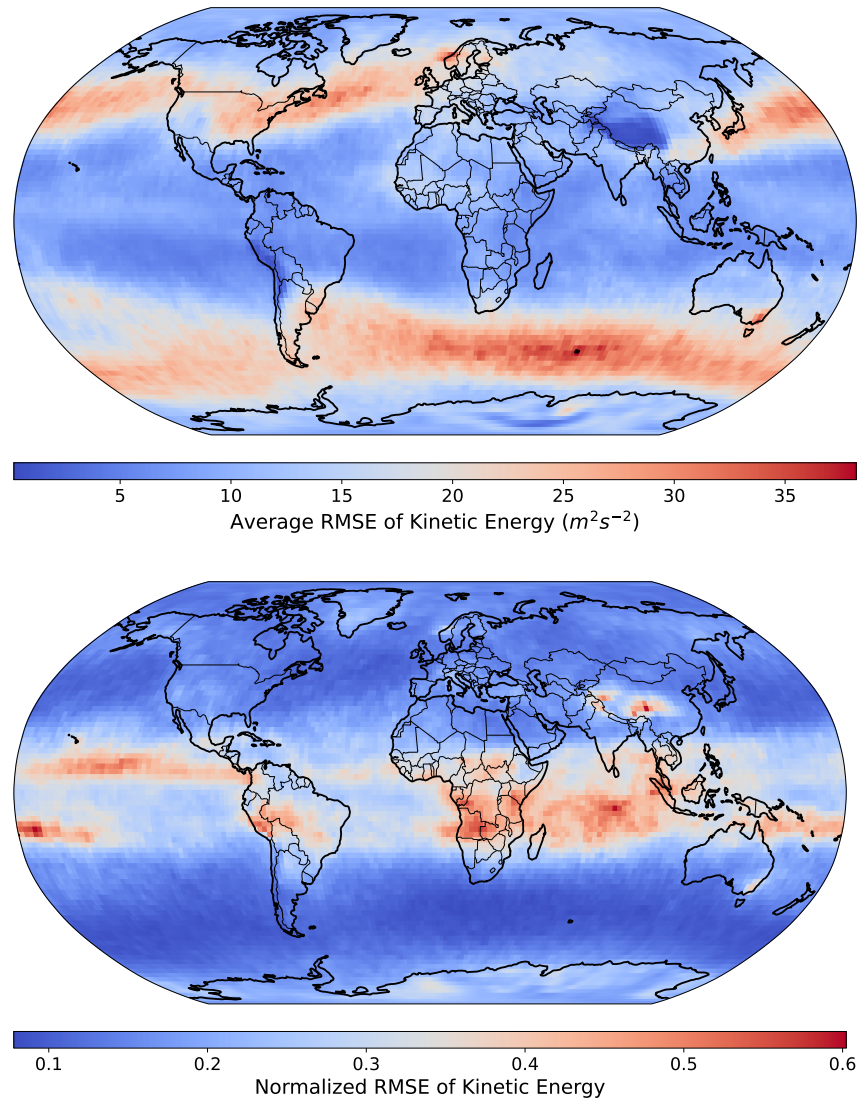


Figure 7.4: Root mean square error (RMSE) of kinetic energy calculated from the ECMWF integrated forecasting system (IFS) and averaged over 2020 for 600 hPa is shown in the top panel. Error is calculated between the 24 hour forecast and the corresponding analysis (estimate of the true atmospheric state). The bottom panel shows that RMSE divided by the averaged analysis kinetic energy to show the error relative to the total kinetic energy at a given location. Red regions indicate high RMSE and blue low.

In order to assess the success of our model in capturing the spatial characteristics of error growth in the atmosphere, we need to understand what these patterns might be. We have chosen to do this by presenting the root mean square error (RMSE) of the IFS forecasting system.

Figure 7.4 (upper panel), shows the average one-day RMSE for the ECMWF IFS forecast over 2020. In order to calculate this, the forecast of U and V at step 24 and the initial conditions were retrieved at two times daily; the forecast initialised at 00 UTC and 12 UTC. The kinetic energy of each of these fields at 600 hPa was then calculated. Then the pointwise difference between the kinetic energy of the forecast field and that of the initial condition corresponding to that forecast was calculated. This difference was averaged over the full year, and then square-rooted to give the RMSE value displayed.

In the upper panel of Figure 7.4 we see several clear features. We see strong errors associated with the

storm tracks in the midlatitudes in both the northern and southern hemispheres. We see a weak signal associated with the ITCZ in the tropics, and we see low levels of RMSE in places of extreme orography. This is due to the low surface pressure in these regions, meaning that 600 hPa is close to the surface. Close to the surface, wind speeds are generally slower. This means that the capacity for error in kinetic energy to grow is diminished. In the lower panel, we see that error is much higher in these high orographic regions relative to total kinetic energy than elsewhere.

The lower panel shows the result of dividing the RMSE by the average total kinetic energy at that location. This gives the relative size of the error growth in different regions after a forecast period of one day, compared to the maximum it could attain. As we have just stated, we see that over the Andes and Himalayas, even though the absolute RMSE is small, the relative error is large. This is an important consideration since if the error reaches the magnitude of the background kinetic energy, the flow is no longer predictable.

We also note that in the tropics, relative error is larger than in the extratropics. This corresponds to the experience of weather forecasters that forecasts in the tropics are less reliable than in the extratropics. Even though the absolute size of the error is smaller, since the motion is generally less energetic in the tropics, this leads to larger relative errors. A convective signal also appears slightly more strongly, with the Pacific branch of the ITCZ particularly visible.

Spatial patterns of error growth predicted by the model

In Part One we discussed the theory of extracting local kinetic energy power spectra from ECMWF analysis and showed that significant differences in the spectrum exist for different locations and under different conditions in the atmosphere. Applying the same dimensional model we just applied to the global spectrum to each of these local spectra yields a predictability time for each scale at each location. By interpolating the output predictability times for each scale, we can estimate the spatial scale reached by a given time. If we then integrate the background spectrum between the highest wavenumber and the scale we have just estimated, we arrive at a prediction for the total error growth.

As we have noted, a significant drawback of this approach, beyond it being most applicable to 2DV rather than the full physics of the atmosphere, is that error is not advected. For this reason, we only seek to investigate mesoscale error saturation patterns, assuming that errors will remain quite spatially localised over the corresponding timescale (in this case, a single day) in the forecast. For our dimensional model, we use 5 days, the typical time for the global mesoscales to saturate.

The upper panel of Figure 7.5 shows the prediction for the total error growth after 5 days at different locations across the globe. The most striking point is the magnitude of the error is significantly smaller in Figure 7.5 than Figure 7.4 proportional to the total kinetic energy. The RMSE normalised by KE is shown in the lower panel of Figure 7.4 while our equivalent ratio is shown in the lower panel of Figure 7.5. While RMSE comfortably sits at close to 50% of the total kinetic energy in the tropics, in most locations the error we predict is significantly smaller in magnitude. Given that we only account for error at scales smaller than the one currently saturated and that the background spectrum follows a power law, we can expect our predicted error to be smaller than that observed. It is also possible that in the real atmosphere, we are seeing the effects of large-scale initial errors propagating down into smaller scales, in the way suggested by [Durran and Gingrich \(2014\)](#).

Now we address the spatial patterns of error growth. In the upper panel of Figure 7.5, we see greater error growth in the extratropics than in the tropics and the poles. This pattern is similar in the average RMSE. In

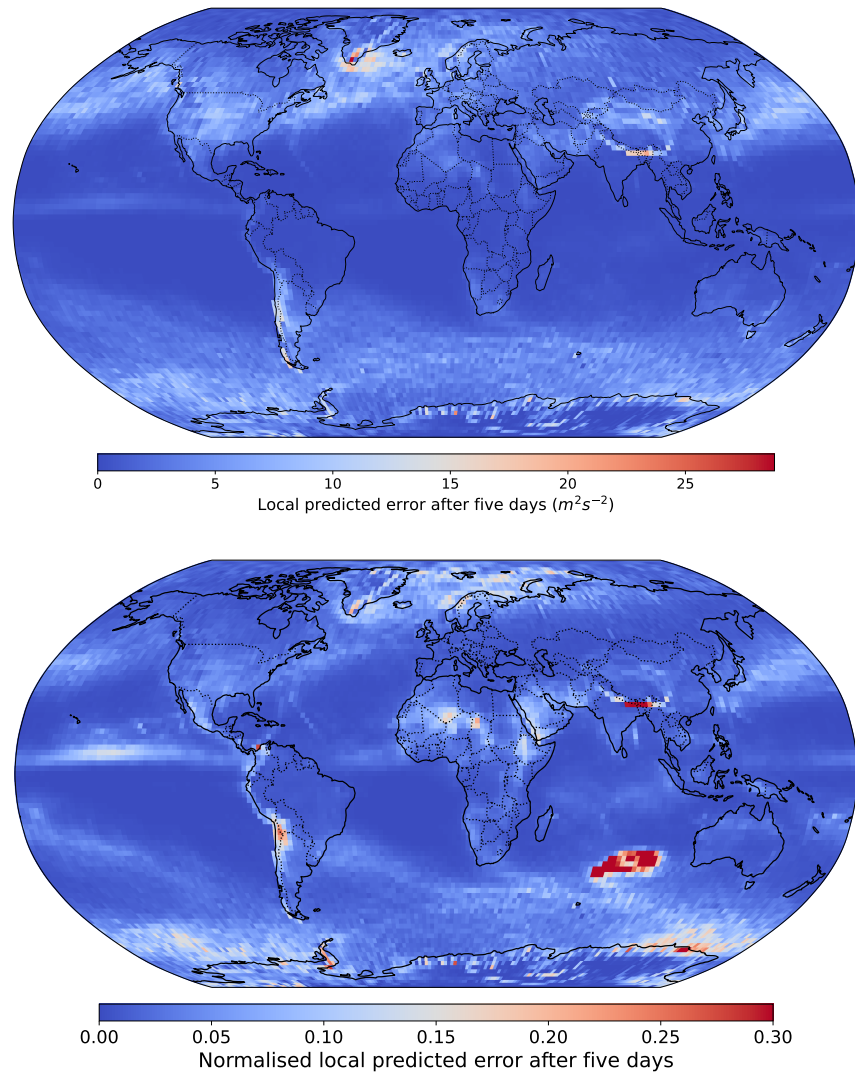


Figure 7.5: Upper panel shows a spatial map of local error growth estimated using the dimensional model. Coarse-grained local power spectra from 2020 were averaged for each location, then the dimensional model was applied to estimate the magnitude of error after five days. In the lower panel, the estimate above is normalised by dividing by the total kinetic energy in the analysis.

both cases, greater kinetic energy associated with storm tracks seems to lead to greater error growth. In addition, the smaller signal that appears to correspond to the ITCZ appears in both the dimensional model and the RMSE.

It appears that the two models diverge in the way orography and moist processes influence error. The relative error over orography is much larger in the forecast RMSE than in our model. In addition, the relative error across the tropics is much less pronounced than in the forecast RMSE. This is likely due to error growth from convective instability that is not accounted for in our model.

7.3 Summary and conclusions

In this chapter we have created a toy dimensional model for error growth that takes the coarse-graining local spectrum for the atmosphere as input to produce an estimate of local error growth rates. We have found that this model reproduces some of the general characteristics of error growth, a tendency towards

greater absolute error in the extratropics and in regions of convection, and greater relative error in regions of high orography. However, it has failed to reproduce the magnitude of these errors or the much larger relative errors in the tropics.

It has been noted ([Zhang et al., 2007](#)) that early error growth in full physics simulations of the atmosphere is associated with convection and moist processes. This is especially relevant in the tropics. Our approach is purely concerned with two-dimensional turbulence and is not capturing these 3D processes; this may be the root cause for the discrepancies. In the extratropics our error growth is also not as large as in the forecasts. This could be due to a lack of downscale error propagation in our model, the role of which was highlighted in [Durrán and Gingrich \(2014\)](#).

In order to uncover if there is any diagnostic value to the local background spectrum for local error growth, it would be desirable to apply coarse-graining to a high-resolution twin experiment of the atmosphere. This could reveal if a relationship between the local eddy turnover time and error growth exists and if so, under what conditions.

Part III

Turbulence theory

Chapter 8

Two-dimensional turbulence theory

In Part One of this thesis we examined local power spectra in the atmosphere. We found that divergent motion was of central importance in the mesoscales. However, we also found regions where two-dimensional turbulence fit with our local spectral profiles. In this Chapter, we explain how one can derive these spectral characteristics of two-dimensional turbulence models. In particular, we explain the characteristics of both the classic two-dimensional turbulent system, the 2D vorticity equations (2DV) and the surface quasi-geostrophic equations (SQG), which were proposed as explanations for the mesoscale shallowing in [Gage \(1979\)](#) and [Tulloch and Smith \(2006\)](#) respectively. We will also outline how we can deduce a steep -3 slope holds for the quasi-geostrophic equations downscale from an injection of energy, as is the accepted explanation for the synoptic steepening of the power spectrum ([Charney, 1971](#)).

We will do this by working from a general framework, we introduce the family of equations called the alpha turbulence models ([Pierrehumbert et al., 1994](#)). These include several physically realisable systems as special cases. We will show how to derive the spectral slopes these models produce and show that kinetic energy moves to larger scales in 2DV but smaller scales in SQG. In Chapter Nine, we investigate how well our theoretical spectra are reproduced in simulations of alpha models and what these models can teach us about error growth in turbulent systems.

This chapter provides theoretical background relevant to Chapter Nine and the rest of this thesis. The general arguments for finding slopes and the direction of energy transfer originated in the work of Kolmogorov, Kraichnan, Leith and Batchelor ([Kolmogorov, 1941](#); [Kraichnan, 1967](#); [Leith, 1968](#); [Batchelor, 1969](#)). I found the presentations in [Vallis \(2017\)](#) and [Burgess \(2014\)](#) useful for understanding the topic.

8.1 Alpha models

The framework we will be using in this part of the thesis is the alpha turbulence model ([Pierrehumbert et al., 1994](#)). These models have the advantage of containing several physically realisable turbulent systems that can be moved between by changing the value of a single parameter, α . In particular, the alpha models include the SQG equations ($\alpha = 1$), the 2D vorticity equations ($\alpha = 2$) and the rotating shallow flow (RSF) equations under internal heating ($\alpha = 3$) ([Tran, 2004](#)). Negative values of alpha can also correspond to physical systems ([Bhushan Jha et al., 2023](#)).

Two alpha models of particular interest to us are the classic 2D turbulence system (we will call this 2DV since it is described fully by the two-dimensional vorticity equation) and the surface quasi-geostrophic (SQG) equations (see Appendix B for a derivation of the SQG equations and a proof that they are an alpha model). Both of these equations have been historically relevant to mesoscale -5/3 research ([Gage, 1979](#); [Tulloch and Smith, 2006](#)) and are still contemporary in atmospheric predictability investigations. [Rotunno and Snyder \(2008\)](#) generalised the L69 model to the SQG equations and in so doing to all alpha models, [Sun and Zhang \(2020\)](#) introduced a mixed model for error growth where small scale error grew according to the L69 model generalised to the SQG equations, while large scale error grew according to the original 2DV version of L69. As noted in [Rotunno and Snyder \(2008\)](#), the spectrum itself was the dominant determiner

of predictability rather than whether the dynamics were SQG or 2DV.

Why many models can be fit into a general framework

The motion in 2DV is completely determined by the vorticity. The vorticity is advected by and advects the flow. It is an active tracer, a variable that is transported by the fluid and influences its motion. Similarly, in the SQG system, the flow is completely determined by surface buoyancy.

In any alpha model, all the dynamics are contained in the relationship between the active tracer and the stream function ψ . The vorticity ω is minus the laplacian of the stream function ($\omega = -\Delta\psi$). In the SQG equations, surface buoyancy Θ is related to the first derivative of the stream function $\Theta = \frac{\partial\psi}{\partial z}|_{z=0}$ (see Appendix B for how this equation leads to the simple relationship we state below).

After Fourier transforming the vorticity formula (and a more involved derivation for the SQG equations), we find that $\hat{\omega} = k^2\hat{\psi}$ and $\hat{\Theta} = k\hat{\psi}$. We can therefore generalise. We define the generalised vorticity ζ by its relationship to the stream function in spectral space, $\hat{\zeta} = k^\alpha\hat{\psi}$. The full alpha turbulence equations are therefore

$$\frac{D\zeta}{Dt} = 0, \quad (8.1.0.1)$$

$$\text{where } \hat{\zeta} \equiv k^\alpha\hat{\psi}. \quad (8.1.0.2)$$

The behaviour of α models

A general form for the Green's function of alpha models (their response to a point source of generalised vorticity) was found by [Iwayama and Watanabe \(2010\)](#). This demonstrated that interactions become longer range as one increases the parameter alpha. In particular, the velocity field generated by a point source of generalised vorticity decays like r^{-2} for SQG, r^{-1} for 2D Euler and is constant with distance for $\alpha = 3$. Interactions are more spatially local for smaller alpha, allowing local structures to form more easily without being destroyed by external shears.

This gives systems with a smaller value for alpha a more "turbulent" flavour, which can be quantified by the shallower spectral slope associated with these systems. We will derive this relationship via a dimensional argument in this chapter. In addition, we will see that errors in higher alpha systems evolve more nonlocally in spectral space in Chapter Nine. This can also be seen to be due to the spatial nonlocality of interactions in higher-alpha models, which causes an error localised to a given location to affect flow at physically distant locations. We will discuss this further in Chapter Nine.

8.2 Deriving the spectrum

We will now derive the spectral power laws and energy flux directions for the alpha models. The key to deriving spectral power laws in this context is exploiting invariants in the system. In alpha models, two quantities, generalised energy and generalised enstrophy, are invariant in time. By leveraging this fact, we will show that generalised energy fluxes to large scales and the generalised enstrophy to small scales. If we then assume that the flux of generalised energy and generalised enstrophy through scales is constant, we can make a dimensional argument to show that this implies a particular kinetic energy spectrum with a spectral slope that is a function of alpha.

This cascade-based framework performs well for $0 < \alpha < 2$, where spatial interactions remain relatively local. However, as we will explore through simulations in Chapter Nine, the accuracy of the predictions diminishes as α increases and interactions become more nonlocal (Pierrehumbert et al., 1994).

8.2.1 Proving invariance of generalised enstrophy and generalised energy

The first step of our derivation is to show that generalised enstrophy, which we will also simply call enstrophy where confusion is unlikely and generalised energy, which we also call energy, are both quadratic invariants of the system. This will allow us to deduce which direction they move in spectral space. Note that generalised energy and kinetic energy and generalised enstrophy and enstrophy are only the same quantities in the case of the 2DV equations. Our ultimate goal is to derive the kinetic energy spectrum in any alpha model, not the generalised energy spectrum.

Enstrophy invariance

Generalised enstrophy is defined to be

$$Z_g = \frac{1}{2} \int_A |\zeta|^2 dA \quad (8.2.1.1)$$

where ζ is generalised vorticity and A is the domain. Any function of generalised vorticity is conserved over the domain, not just the enstrophy. The proof of this is the same as that for 2D Euler. Following Vallis (2017),

$$\frac{d}{dt} \int f(\zeta) dA = \int f'(\zeta) \frac{d\zeta}{dt} dA, \quad \text{chain rule} \quad (8.2.1.2)$$

$$= - \int f'(\zeta) \nabla \zeta \cdot u dA, \quad \text{generalised vorticity conservation (8.1.0.1)} \quad (8.2.1.3)$$

$$= - \int \nabla f(\zeta) \cdot u dA, \quad \text{chain rule} \quad (8.2.1.4)$$

$$= - \int \nabla \cdot (f(\zeta)u) dA, \quad \text{velocity divergence-free} \quad (8.2.1.5)$$

$$= - \int f(\zeta)u \cdot n ds, \quad \text{Gauss's theorem} \quad (8.2.1.6)$$

$$= 0. \quad \text{periodic or no inflow BC} \quad (8.2.1.7)$$

Therefore, in particular, enstrophy is a quadratic invariant

$$\boxed{\frac{d}{dt} \frac{1}{2} \int |\zeta|^2 dA = 0.} \quad (8.2.1.8)$$

Energy invariance

I wrote the following proof of the invariance of generalised energy after a correspondence with Raymond Pierrehumbert. Generalised energy is defined to be

$$\mathcal{E}_g = \frac{1}{2} \int_A \psi \zeta \, dA \quad (8.2.1.9)$$

(Burgess, 2014). We can see that for $\alpha = 2$, this generalised energy is the kinetic energy by integration by parts:

$$\frac{1}{2} \int \psi \zeta \, dA = \frac{1}{2} \int \psi (-\nabla^2 \psi) \, dA, \quad \text{when } \alpha = 2, \zeta = \omega, \omega = -\nabla^2 \psi \quad (8.2.1.10)$$

$$= \frac{1}{2} \int |\nabla \psi|^2 \, dA, \quad \text{integration by parts} \quad (8.2.1.11)$$

$$= \mathcal{E}. \quad \text{definition of kinetic energy } (\mathcal{E}) \quad (8.2.1.12)$$

We now show generalised energy is conserved:

$$0 = \int \psi \frac{D\zeta}{Dt} \, dA, \quad \text{conservation of generalised vorticity (8.1.0.1)} \quad (8.2.1.13)$$

$$= \int \psi \frac{\partial \zeta}{\partial t} + \psi \nabla \zeta \cdot u \, dA, \quad \text{expanding material derivative} \quad (8.2.1.14)$$

$$= \int \psi \frac{\partial \zeta}{\partial t} + \psi \nabla \cdot (\zeta u) \, dA, \quad \text{divergence-free condition} \quad (8.2.1.15)$$

$$= \int \psi \frac{\partial \zeta}{\partial t} + \nabla \cdot (\psi \zeta u) - \nabla \psi \cdot \zeta u \, dA, \quad \text{rewriting using product rule} \quad (8.2.1.16)$$

$$= \int \psi \frac{\partial \zeta}{\partial t} \, dA, \quad \text{since } \nabla \psi \cdot u = 0 \text{ and Gauss's theorem eliminates the divergence term} \quad (8.2.1.17)$$

$$= \int \hat{\psi} \frac{\partial}{\partial t} \bar{\zeta} \, dK, \quad \text{applying Parseval's theorem} \quad (8.2.1.18)$$

$$= \int \hat{\psi} \frac{\partial}{\partial t} k^\alpha \bar{\psi} \, dK, \quad \text{using inversion relation (8.1.0.2)} \quad (8.2.1.19)$$

$$= \frac{1}{2} \int k^\alpha \frac{\partial}{\partial t} \hat{\psi}^2 \, dK, \quad \text{symmetry in Fourier terms} \quad (8.2.1.20)$$

$$= \frac{1}{2} \int \frac{\partial}{\partial t} (\hat{\psi} \bar{\zeta}) \, dK, \quad \text{rearranging terms} \quad (8.2.1.21)$$

$$= \frac{1}{2} \frac{\partial}{\partial t} \int \psi \zeta \, dA. \quad \text{inverse Parseval's theorem} \quad (8.2.1.22)$$

In line 8.2.1.16, $\nabla \psi \cdot u = 0$ since $\nabla \psi = (v, -u)$ so $\nabla \psi \cdot u = vu - uv = 0$. We have therefore shown that generalised energy is conserved,

$$\boxed{\frac{\partial}{\partial t} \frac{1}{2} \int \psi \zeta \, dA = 0.} \quad (8.2.1.23)$$

8.2.2 Argument for upscale energy flux and downscale enstrophy flux

We will now deduce that energy moves to larger scales while enstrophy moves to smaller scales in spectral space. We will do this by considering how a distribution of energy in spectral space corresponds with

a distribution of enstrophy in spectral space. We will find the direction of flux follows from the global conservation of both quantities.

The generalised energy, kinetic energy and generalised enstrophy spectra

We now derive a formula that relates the distribution of energy in spectral space with the distribution of enstrophy. When we later apply a dimensional argument in this context, we will need the local eddy turnover time, so we will also derive the relationship between the energy distribution and the kinetic energy distribution in spectral space.

We recall that wavenumber is the magnitude of the wavevector, so $k^2 = k_x^2 + k_y^2$ in 2D and $k^2 = k_x^2 + k_y^2 + k_z^2$ in 3D. The generalised energy density is the value of the generalised energy spectrum, $\mathcal{E}_g(k)$. Assuming this is only a function of wavenumber is equivalent to assuming homogeneity and isotropy. The central property of the generalised energy spectrum is that integrating over all wavenumbers will return the total generalised energy. Similarly to the generalised energy spectrum, there is a generalised enstrophy spectrum, $\mathcal{Z}_g(k)$. The two can be related to one another via Parseval's theorem, which we now show:

$$\int \mathcal{E}_g(k) dk = \frac{1}{2} \int \psi \zeta dA, \quad \text{definition of generalised energy } (\mathcal{E}_g) \quad (8.2.2.1)$$

$$= \frac{1}{2} \int \widehat{\psi} \widehat{\zeta} dk, \quad \text{Parseval's theorem} \quad (8.2.2.2)$$

$$= \frac{1}{2} \int k^{-\alpha} \widehat{\zeta} \widehat{\zeta} dk, \quad \text{substitute } \widehat{\psi} = k^{-\alpha} \widehat{\zeta} \quad (8.2.2.3)$$

$$= \int k^{-\alpha} \mathcal{Z}_g(k) dk. \quad \text{definition of generalised enstrophy spectrum } (\mathcal{Z}_g(k)) \quad (8.2.2.4)$$

This implies, as in Burgess (2014), that

$$\mathcal{Z}_g(k) = k^\alpha \mathcal{E}_g(k). \quad (8.2.2.5)$$

We are also interested in the kinetic energy power spectrum $\mathcal{E}(k)$. Since $\nabla \wedge \psi = \mathbf{u}$, the kinetic energy can be written in terms of the stream function as $\frac{1}{2} \mathbf{u}^2 = \frac{1}{2} \left(\left(\frac{\partial \psi}{\partial x} \right)^2 + \left(\frac{\partial \psi}{\partial y} \right)^2 \right)$.

Using Parseval,

$$\int \mathcal{E}(k) dk = \frac{1}{2} \int \psi \Delta \psi dA, \quad \text{definition of kinetic energy } (\mathcal{E}(k)) \text{ in terms of } \psi \text{ and } \Delta \psi \quad (8.2.2.6)$$

$$= \frac{1}{2} \int \widehat{\psi} \widehat{\Delta \psi} dk, \quad \text{Parseval's theorem} \quad (8.2.2.7)$$

$$= \frac{1}{2} \int k^2 \widehat{\psi}^2 dk. \quad \widehat{\Delta \psi} = -k^2 \widehat{\psi} \quad (8.2.2.8)$$

$$= \frac{1}{2} \int k^{2-2\alpha} \widehat{\zeta}^2 dk. \quad \text{substitute } \widehat{\psi} = k^{-\alpha} \widehat{\zeta} \quad (8.2.2.9)$$

$$= \int k^{2-2\alpha} \mathcal{Z}_g(k) dk. \quad \text{definition of generalised enstrophy spectrum } (\mathcal{Z}_g(k)) \quad (8.2.2.10)$$

Hence

$$\mathcal{Z}_g(k) = k^\alpha \mathcal{E}_g(k) = k^{2\alpha-2} \mathcal{E}(k). \quad (8.2.2.11)$$

As is presented in [Pierrehumbert et al. \(1994\)](#).

Argument that energy moves to larger scales

The following graphical motivation for the direction of cascades is inspired by the conservation argument in [Vallis \(2017\)](#). The argument has been simplified for didactic purposes by only appealing to the properties of the graph of the spectrum. The argument is also generalised since it is applied to alpha turbulence rather than only 2DV.

Suppose energy is concentrated at some small range of wavenumbers. The generalised energy spectrum will look like a bump as in Figure 8.1. This energy will evolve to spread out in spectral space, as it is transferred to smaller and larger scales by triad interactions.

Since the generalised energy must be conserved, the total area under the curve must remain the same. We may naively hypothesise that the spectrum will evolve to a symmetric state, like that illustrated in Figure 8.2.




Figure 8.1: We suppose that kinetic energy is initially injected at some small range of wavenumbers. We may stir a stationary fluid at a particular length-scale.

The symmetric spread in Figure 8.2 fails to account for enstrophy. From the relationship between the generalised energy and generalised enstrophy spectrum (8.2.2.11) and given $\alpha > 0$, high wavenumbers contribute more to the enstrophy than low ones. As energy spreads in spectral space, it must increase more in low wavenumbers than high ones to conserve the enstrophy. A sketch of a distribution that achieves this is shown in Figure 8.3.

The effect is generalised energy moves to larger scales on average as time progresses. The same graphical argument plotting the enstrophy reveals that enstrophy will, on average, move to smaller scales.

When $\alpha = 2$, the generalised energy is the kinetic energy, and so we have shown that in 2DV, kinetic energy moves to larger scales (inverse cascade). When $\alpha = 1$, the generalised enstrophy corresponds to kinetic




Figure 8.2: The kinetic energy initially distributed as in Figure 8.1 evolves in time in a way that preserves kinetic energy (area under the curve remains constant). By turbulent interactions, it spreads to other lengthscales. This figure illustrates the spread if it were completely symmetric with the same amount of energy going to large and small scales (this is not what happens, in reality kinetic energy primarily goes to larger scales).




Figure 8.3: The correct spectral spread. Energy is conserved since the area under the graph is the same as in the initial state and enstrophy is conserved by the proportion of energy in the low wavenumbers being higher than in the high wavenumbers.

energy, and so we have also shown that kinetic energy moves to smaller scales (direct cascade) in the SQG equations.

Generalised enstrophy is kinetic energy in the SQG system since

$$\mathcal{E}_g = \frac{1}{2} \int \zeta^2 \, dA, \quad \text{definition of generalized enstrophy} \quad (8.2.2.12)$$

$$= \frac{1}{2} \int |\hat{\zeta}|^2 \, dk, \quad \text{Parseval's theorem} \quad (8.2.2.13)$$

$$= \frac{1}{2} \int k^2 |\hat{\psi}|^2 \, dk, \quad \alpha = 1 \text{ so substitute } \hat{\zeta} = k\hat{\psi} \quad (8.2.2.14)$$

$$= \frac{1}{2} \int \psi \omega \, dk, \quad \text{Parseval's theorem} \quad (8.2.2.15)$$

$$= \mathcal{E}. \quad \text{equivalent definition of KE (8.2.1.10)} \quad (8.2.2.16)$$

8.2.3 Finding the spectral slopes

We now know the direction of flux of two quadratic invariants of the system, as well as how their magnitude at a given wavenumber is related to the kinetic energy at that wavenumber. To derive spectral power laws from this information, we will now make Kolmogorov-like assumptions.

We assume that for wavenumbers above and below an injection of generalised energy, the flux of generalised energy is independent of scale and time. That is, we are assuming that at these scales the motion is independent of the nature of the forcing or dissipation. We are assuming the formation of inertial ranges. This assumption corresponds to Kolmogorov's first similarity law (Kolmogorov, 1941).

The only timescale at wavenumber k is the eddy turnover time τ_k . By definition, the flux of energy ϵ must go like

$$\epsilon \sim \frac{\psi_k \zeta_k}{\tau_k}. \quad (8.2.3.1)$$

A key property of the kinetic energy spectrum $\mathcal{E}(k)$ is that its integral is the total kinetic energy

$$\int \mathcal{E}(k) \, dk = \int \frac{u^2}{2} \, dA. \quad (8.2.3.2)$$

By scale analysis this gives

$$\mathcal{E}(k)k \sim u_k^2. \quad (8.2.3.3)$$

Since $\nabla \wedge \psi = \mathbf{u}$ we have $\mathcal{E}(k) \sim \psi_k^2 k$. Using (8.2.3.3) and $k \sim 1/L$ to construct a quantity with the dimensions of time, we arrive at the eddy turnover time

$$\tau_k \sim k^{-3/2} \mathcal{E}(k)^{-1/2}. \quad (8.2.3.4)$$

Using (8.1.0.2), $\zeta_k \sim k^\alpha \psi_k$, so the numerator of (8.2.3.1) is

$$\psi_k \zeta_k \sim k^\alpha \psi_k^2 \sim k^{\alpha-1} \mathcal{E}(k). \quad (8.2.3.5)$$

Together this gives

$$\epsilon \sim \frac{\psi_k \zeta_k}{\tau_k}, \quad \text{definition of energy flux} \quad (8.2.3.6)$$

$$\sim \frac{k^{\alpha-1} \mathcal{E}(k)}{k^{-3/2} \mathcal{E}(k)^{-1/2}}, \quad \text{substitute } \tau_k \text{ and } \psi_k \zeta_k \quad (8.2.3.7)$$

$$\sim k^{\alpha+(1/2)} \mathcal{E}(k)^{3/2}. \quad \text{simplify} \quad (8.2.3.8)$$

Rearranging, the kinetic energy spectrum upscale of injected energy follows the power law

$$\boxed{\mathcal{E}(k) \sim \epsilon^{2/3} k^{-(2\alpha+1)/3}}. \quad (8.2.3.9)$$

Exactly analogously, we can consider enstrophy flux η and find

$$\eta \sim \frac{\zeta_k^2}{\tau_k}, \quad \text{definition of enstrophy flux} \quad (8.2.3.10)$$

$$\sim \frac{k^{2\alpha-1} \mathcal{E}(k)}{k^{-3/2} \mathcal{E}(k)^{-1/2}}, \quad \text{substitute } \tau_k \text{ and } \zeta_k^2 \quad (8.2.3.11)$$

$$\sim k^{2\alpha+(1/2)} \mathcal{E}(k)^{3/2}. \quad \text{simplify} \quad (8.2.3.12)$$

When rearranged, this gives the power law obeyed by the kinetic energy spectrum downscale of injected energy,

$$\boxed{\mathcal{E}(k) \sim \eta^{2/3} k^{-(4\alpha+1)/3}}. \quad (8.2.3.13)$$

As alpha is increased, the spectral slopes become steeper. Note also that the gradient in the enstrophy range is always steeper than that in the energy range, with the difference in gradient being $2\alpha/3$. The picture of the spectrum we have just derived is illustrated in Figure 8.4.

The 2DV equations

When $\alpha = 2$ we have the 2DV equations. We have shown that kinetic energy moves to larger scales while enstrophy moves to small scales in this system. In addition, we have shown the existence of a -5/3 range above an injection of energy and a -3 range below it.

The SQG equations

We have shown that kinetic energy move to small scales in this system while its generalised energy moves to large scales. We see a -5/3 range below an injection of energy and a -1 range above it.




Figure 8.4: The kinetic energy spectrum of a generalised 2D fluid in equilibrium plotted in log-log space. Energy is injected in some spectral neighbourhood. At smaller wavenumbers, there is a shallower slope, which ends once large-scale drag starts to dominate or the outer scale is reached. At higher wavenumbers, the spectrum has a steeper gradient. This range ends when the small-scale dissipation range begins.

The -5/3 slope in 3D turbulence

The argument for the -5/3 downscale of forcing in 3D turbulence is analogous to this 2D case. In 3D, there is no conservation of enstrophy, only kinetic energy. Assuming that the energy fluxes to smaller scales and that flux is constant, we find ourselves with the same dimensional argument explored above. Notice that any spectrum where something with the dimensions of kinetic energy is being fluxed in a cascade at a constant rate must have a -5/3 slope by the above dimensional argument. This is why the slope of the 2DV inverse cascade and of the SQG direct cascade are both -5/3.

The -3 slope of the atmosphere

In Chapter One, it was mentioned that the justification of the -3 slope observed in the atmosphere was given when [Charney \(1971\)](#) adapted the arguments for the spectral slopes of the 2DV equations to the quasi-geostrophic equations. We are now in a position to outline how he did this.

The quasi-geostrophic equations are the result of including the next lowest order terms in geostrophic balance, allowing the flow to evolve in time. It is a good model for synoptic scale weather. It can be written as the material conservation of the potential vorticity

$$\frac{Dq}{Dt} = 0, \quad (8.2.3.14)$$

$$\text{where } q \equiv \Delta_h \psi + f + \frac{f^2}{N^2} \frac{\partial^2}{\partial z^2} \psi. \quad (8.2.3.15)$$

By rescaling the z coordinate to $\bar{z} = \frac{N}{f}z$, and noting that since only derivatives of q appear in Equation 8.2.3.14, the value of the constant f is irrelevant and may be set to 0, we can write

$$q = \Delta \psi \quad (8.2.3.16)$$

in the rescaled coordinates. Since the equivalents of both enstrophy and energy are conserved, we once again have a flux of kinetic energy to large scales. Since q has the dimensions of vorticity, dimensional arguments give us the same -3 slope at smaller scales than an injection of energy as in the 2DV case. Hence, given that the sun injects energy into the atmosphere at very large scales through differential heating (the scale of kinetic energy injection being determined by the scale of the baroclinic instability the temperature gradient generates), we can deduce a -3 spectral range forms in the synoptic scales.

8.3 Summary

In this chapter, we have given the theoretical background to predict classic turbulence spectral slopes. This is the theory underpinning the spectral characteristics we referenced in Part One of the thesis. We have outlined why a -5/3 slope can arise due to 2DV, SQG or 3D dynamics and explained how a synoptic -3 slope can arise in the atmosphere. We have presented this in generality by introducing the alpha model framework for two-dimensional turbulence. In the next chapter, we will show that these predicted spectra break down for $\alpha > 2$, where interactions become significantly non-local. We will also show how the character of error growth changes in these models as alpha is increased.

Chapter 9

Spectral slopes and error growth in simulated alpha models

Alpha models are a generalised setting to study two-dimensional turbulence. In the previous chapter, we showed how their spectral slopes can be predicted. We also mentioned that as alpha increases, the locality of interactions in physical space decreases ([Iwayama and Watanabe, 2010](#)); that is, motion at a particular location can quickly affect the evolution at distant points in the flow. In this chapter, we explore how our predictions for the spectral slope lose accuracy at higher alpha values and how error growth changes with larger alpha.

The results in this chapter are less tethered to the context of the atmosphere than those in the rest of the thesis, instead focusing on dynamics present in systems of equations, which, for some alpha values, are unlikely to have any physical realisability. However, due to the ability to tune the locality of interactions in these models, alpha models may be the perfect testbed for improving the local model of error growth presented in Part Two. We discuss this further later in this chapter. In addition, the SQG equations ($\alpha = 1$), can be used as a tropopause model due to the stratosphere's strong stratification making the tropopause a PV boundary ([Juckes, 1994](#)).

9.1 Simulating alpha models psuedo-spectrally

The simulations used in this chapter were produced by adapting a Python package called PYQG written by Malte Jansen, Ryan Abernathey, Cesar Rocha and Francis Pouline. We will give an outline of how the code functions and some details of our simulations.

Contributions to the code

PYQG contains a pseudo-spectral solver of the SQG equations. I contributed to the GitHub project by correcting an error and contributing to the documentation. In my forked version of the code, I also added an alpha turbulence solver, a way to add forcing at particular wavenumbers, the ability to choose between that forcing being stochastic or fixed, a large-scale drag which is necessary to reach a cascade in equilibrium and I changed the way the code used a C fast Fourier transform (FFT) package since before this change, the simulations were not bit reproducible. Note that without bit reproducibility, error growth cannot be studied through twin-experiments using PYQG.

Principles of a pseudo-spectral alpha model solver

At a high level, the solver we employ functions as follows. The alpha turbulence system in Jacobian form is given by

$$\frac{\partial \zeta}{\partial t} = -J(\psi, \zeta), \quad (9.1.0.1)$$

$$\text{where } \hat{\zeta} \equiv k^\alpha \hat{\psi}. \quad (9.1.0.2)$$

J is the Jacobian $J(A, B) = \partial_x A \partial_y B - \partial_x B \partial_y A$. We seek to solve the equation in spectral space. This is convenient since, in spectral space, derivatives are transformed into multiplications by powers of wavenumber. After applying a Fourier transform, Equation 9.1.0.1 becomes an ordinary differential equation (ODE) in time. This ODE is then stepped forward in a standard way (for example, Adams-Bashforth).

The method we employ is pseudo-spectral because some parts of this ODE are calculated in real space. This is done because although derivatives become convenient in spectral space, products become inconvenient since the Fourier transform of a product is the convolution of the Fourier transforms of the terms (Appendix A).

The algorithm

We now provide a more detailed breakdown of the algorithm.

As an initial condition, or at the start of each time step, we have the generalised vorticity field. The first step is to apply a fast Fourier transform (FFT) to this generalised vorticity ζ . Using $\hat{\zeta} = k^\alpha \hat{\psi}$ gives the Fourier transform of the stream function, $\hat{\psi}$. The x derivative of ψ , $\frac{\partial \psi}{\partial x} = v$, is then given by multiplying $\hat{\psi}$ by ik and applying an inverse fast Fourier transform (IFFT). Likewise we can obtain $\frac{\partial \psi}{\partial y} = -u$. Fourier transforming (9.1.0.1), denoted by \mathcal{F} here, gives

$$\frac{\partial \mathcal{F}(\zeta)}{\partial t} = -\mathcal{F}(J(\psi, \zeta)). \quad (9.1.0.3)$$

Since $J(\psi, \zeta) = \frac{\partial \psi}{\partial x} \frac{\partial \zeta}{\partial y} - \frac{\partial \zeta}{\partial x} \frac{\partial \psi}{\partial y} = \frac{\partial}{\partial y} \left(\frac{\partial \psi}{\partial x} \zeta \right) - \frac{\partial}{\partial x} \left(\frac{\partial \psi}{\partial y} \zeta \right)$ where the terms with double stream function derivatives cancelled,

$$\frac{\partial \mathcal{F}(\zeta)}{\partial t} = i l \mathcal{F} \left(\frac{\partial \psi}{\partial x} \zeta \right) + i k \mathcal{F} \left(\frac{\partial \psi}{\partial y} \zeta \right). \quad (9.1.0.4)$$

Therefore, since we have the derivatives of ψ and ζ , we can take the Fourier transform of their product and calculate the tendency of the Fourier transform of ζ . The value of $\mathcal{F}(\zeta) \equiv \hat{\zeta}$ is then stepped forward using Adams-Bashforth 3 (AB3), a common numerical method for approximating the solution to an ODE. AB3 requires two previous time steps, so for the first two time steps, Euler's method and then AB2 are used.

To parameterise viscosity, a spectral filter is applied to $\hat{\zeta}$. As in [LaCasce \(1996\)](#), it is given by multiplying the output of AB3, $\hat{\zeta}_{n+1}(\mathbf{k})$, by

$$\text{filter} = e^{-18.4(|\mathbf{k}| - c_k)^4} \quad (9.1.0.5)$$

when $|\mathbf{k}| > c_k$. c_k was chosen to be $0.65K_{max}$, the exponent four was chosen as the cutoff slope and -18.4 ($\ln 10^{-8}$) was selected so that for the largest present wavenumber, the filter is 0 to single precision machine

epsilon ($\approx 10^{-8}$). To remove energy from the largest scales, a similar filter is then employed. This has a cutoff wavenumber of three, larger scales are damped analogously to (9.1.0.5).

To sustain turbulence, energy is injected into the system via forcing in a predefined spectral range $[k_1, k_2]$. The forcing term in spectral space is $\hat{\zeta}_{\text{force}}(k) = A_{\text{force}} \cdot \text{Filter}(k) \cdot e^{i\phi(k)}$, where the filter ensures only wavenumbers within the spectral range specified are affected and ϕ is a random phase evolving according to a Brownian motion step $\phi(k) \rightarrow \phi(k) + \frac{\Delta t}{T_f} \Delta\phi$. T_f is the tuneable timescale of the forcing. The forcing is then added to the generalised vorticity in spectral space, $\hat{\zeta} \leftarrow \hat{\zeta} + \hat{\zeta}_{\text{force}}$. Finally, an IFFT is applied to the modified $\hat{\zeta}$ to return to real space.

Dimensionalisation

For toy atmospheric models, physical time scales are often found by comparing the error doubling time in the model to that of the atmosphere (for instance, [Arnold et al. \(2013\)](#)). Since changes in predictability are a variable of interest, this approach does not work here.

Instead, we perform the following trick. We assume the outer scale in our simulation is 1000 km. We then find the factor needed to match the kinetic energy power spectrum for the SQG equations ($\alpha = 1$) to the -5/3 component of the observed kinetic energy power spectrum in the atmosphere. Since this gives us the factor needed to convert the nondimensional kinetic energy to dimensional kinetic energy and we have defined the length scales present in the simulation, we can deduce how the timescale must be dimensionalised. We now explain this more completely.

A fit to observations we have used multiple times previously in this thesis is given in [Lindborg \(1999\)](#), equation 71: $\mathcal{E}(k) = d_1 k^{-5/3} + d_2 k^{-3}$ where $d_1 \approx 9.1 \times 10^{-4} \text{ m}^{4/3} \text{ s}^{-2}$ and $d_2 \approx 3.0 \times 10^{-10} \text{ s}^{-2}$. As is visible in Figure 9.1, the kinetic energy at the largest scales is comparable between alpha models, so we will assume that provided with the same large-scale source of energy, these models will arrive at the displayed power laws, with higher alpha models having less kinetic energy in their small scales than lower alpha models. Since the outer scale of our simulations is defined to be 1000 km, with our resolution, a grid box has a side length of about 4 km. The factor that the power spectrum must be multiplied by to match the -5/3 component of the observations was found. By dividing the expression for the kinetic energy power spectrum in terms of seconds by the expression in terms of model time units (MTU), we see that the multiplicative factor required for the two spectra to match in magnitude is equal to the ratio $(\text{MTU}/\text{s})^2$.

In our simulations, this gives an MTU to be close to 251 seconds or around 4 minutes; since there are 200-time steps in a single time unit, a time step corresponds to a little over a second. Therefore, 300 MTU approximately corresponds to the eddy turnover time of the synoptic scales, 18 hours, 50 MTU, used later in the twin experiments, is around 3 hours and the spectral slope measurement runs lasted 2000 MTU or about 5 and a half days.

Forcing details

The alpha turbulence models are forced randomly on large scales in the simulation (either $k = 3$ or 5, corresponding to 333 or 200 km respectively) from no flow initial conditions with either a fast or slow timescale associated with the random forcing. The simulations are on a 256 by 256 doubly periodic grid. The slow timescale runs had a length of 2.3 days (800 MTU), the faster timescale runs ran for 5.5 days (2000 MTU).




Figure 9.1: The instantaneous KE power spectrum for fully developed alpha turbulence. The slope is calculated over a range of wavenumbers 15 (67 km) to 70 (14 km) to represent the inertial range, this is denoted by the red dashed lines in the figure. In this chapter, forcing is applied at either wavenumbers 3 (333 km) or 5 (200 km). In this figure, forcing is applied at 200 km. A filter removes energy from wavenumbers greater than 83 (12 km) and less than 3 (333 km). In black is the kinetic energy power spectrum of the atmosphere, a fit to observations from Lindborg (1999). The fit is the sum of a $-5/3$ and -3 component which are shown as the blue dashed lines (the -3 is the steeper line since the power law it refers to becomes a gradient on a log-log plot).

The amplitude of the forcing was selected using the following scaling such that if there is no flow, the KE added is constant across different values of alpha or forcing wavenumber k . This amplitude was found by writing that for the kinetic energy power spectrum, $\mathcal{E}(k) \sim k^{-1}\hat{u}^2 \sim k\hat{\psi}^2$ which implies $\mathcal{E}(k) = k^{2(1-\alpha)}\hat{q}^2$ and so the amplitude should be $\hat{q} = \sqrt{\frac{\mathcal{E}(k)}{k^{1-2\alpha}}}$ where k here is the wavenumber of the forcing. The KE amplitude used in my experiments was 0.01.

The phase starts at some random value $e^{2\pi i\phi}$ where ϕ is a uniformly distributed random variable between 0 and 2π . Each time step the phase is updated by multiplying by $e^{2\pi i\phi}$ where ϕ is uniformly randomly distributed between $-2\pi/(T_f/dt)$ and $2\pi/(T_f/dt)$ where T_f is the forcing timescale in MTU: $T_f = 100$ (7 hours) was used for the slow case and 0.1 (25 seconds) for the fast case.

Before measuring the spectrum, the spectral slope was allowed to reach a steady state. This is not the same as the energy reaching equilibrium and in all cases the slope became steady before energy equilibrium was reached. For the slowly varying forcing case, energy never reached equilibrium, but it did for the fast cases for most alpha. Figure 9.1 shows the kinetic energy power spectrum for a range of alpha at a particular instant in fully developed turbulence. The inertial range was defined to lie between $k = 15$ (67 km) and 70 (14 km); a best fit line was applied there to determine the spectral slope. The figure demonstrates that even instantaneously, the power spectrum is well approximated with a power law, and our choice for the inertial range is appropriate.

Figure 9.2 shows the variation in spectral slope for various alpha for fast forcing at $k = 5$ (200 km), results

are qualitatively similar for the other forcing timescales and wavenumbers. The slope is initially steep since energy has been injected at large scales. Variation is small once the slope equilibrates.




Figure 9.2: The evolution in time of the best fit spectral slope of the kinetic energy power spectrum. The slope is initially steep since energy is injected at wavenumber 5. The slope reaches a steady state before the energy.

9.2 Spectral slopes in alpha models

Figure 9.3 shows the mean slope for the three forcing cases simulated (slow forcing at 333 km, fast forcing at 333 km and fast forcing at 200 km) and compares them to the results found in [Pierrehumbert et al. \(1994\)](#) (converted from enstrophy to KE using Equation 8.2.2.11).

Forcing at a larger scale or longer timescales led to a steeper slope. This could be due to the formation of coherent vortices, which trap energy and enstrophy at larger scales. Alternatively, larger forcing length or time scales could be associated with longer timescales of flow evolution, and the simulations may have eventually converged on the same spectral slope regardless of the details of the large-scale forcing.

As we have highlighted and was first pointed out by [Pierrehumbert et al. \(1994\)](#), the assumption of spectral locality breaks down for alpha greater than two. As alpha increases, the stream function is an increasingly smoothed version of the active scalar field. This decouples the details of the field from the dynamics, so the active tracer starts to resemble a passive tracer. Thus a deviation from the Kolmogorov prediction is expected. However, we show a new relationship between alpha and slope appears rather than the tracer acting exactly like a passive scalar as previously assumed.

Figure 9.4 shows the standard deviation of the measured instantaneous slope. Variation in the slope was approximately uniform between values of alpha, but changing the nature of the forcing did have a systematic effect. Slower forcing and forcing at larger scales had more variability associated with the spectral slope. This is interesting since variability in slope over time is rarely addressed. If we had seen some systematic relationship between variability in the slope and the value of alpha, it could have revealed that some models have slopes that are more driven by intermittent events than others.




Figure 9.3: The simulated and theoretical spectral gradient against the value of alpha for the enstrophy range. For slower forcing, or forcing at a greater scale, the spectrum was systematically steeper. The Kolmogorov prediction is followed best when forcing is fast and at the wavenumber 5. The prediction begins to fail for $\alpha > 1.5$. The “limit prediction” of the spectral slope refers to a prediction from a nonlocal dimensional argument (see text). Red points are spectral slopes measured in [Pierrehumbert et al. \(1994\)](#) (P94 in legend).




Figure 9.4: The standard deviation of the spectral slope time series. Slower forcing at larger scales was associated with a slope that varied more in time.

9.2.1 Speculation on extending dimensional arguments to Casimirs and nonlocal interactions

In Figure 9.3, we showed how the spectral slopes measured in alpha models with $\alpha > 2$ deviate from Kolmogorov predictions. We show a “limit prediction” in that figure which seems to match the observed spectral slopes in those higher alpha models better. This section is more speculative than the rest of this chapter, and I include it because I find the discussion interesting.

In Chapter Eight, we showed that not only enstrophy, but any function of vorticity (so-called Casimirs) are invariant. It is interesting to see what happens if we apply basic dimensional arguments (which we must emphasise are flawed in this case) to these invariant functions. We will then attempt to extend the spectrally local dimensional argument we used in Chapter Eight to a spectrally non-local argument more relevant to higher alpha models. A combination of these two arguments results in the “limit prediction” shown in Figure 9.3, although whether there is some more suitable argument for arriving at this prediction is unknown to me.

Kolmogorov arguments applied to Casimirs

In the previous chapter, we showed that not only generalised enstrophy, but any function of generalised vorticity is conserved. We call these Casimirs. In particular, any power of generalised enstrophy, $\zeta^{2\beta}$ is conserved. The problem with applying a dimensional argument, as we now do, is there is no guarantee that the dimension of the spectral flux of the Casimir is $\zeta_k^{2\beta}$. That this follows for quadratic invariants is due to Parseval's law. However, we continue

$$\eta_\beta \sim \frac{\zeta_k^{2\beta}}{\tau_k}, \quad \text{assumed Casimir flux} \quad (9.2.1.1)$$

$$\sim \frac{k^{(2\alpha-1)\beta} \mathcal{E}(k)^\beta}{k^{-3/2} \mathcal{E}(k)^{-1/2}}, \quad \text{substitute } \tau_k \text{ and } \zeta_k^2 \quad (9.2.1.2)$$

$$\sim k^{(2\alpha-1)\beta+(3/2)} \mathcal{E}(k)^{\beta+(1/2)}. \quad \text{simplify} \quad (9.2.1.3)$$

When rearranged, this gives a prediction for the power law obeyed by the kinetic energy spectrum downscale of injected energy,

$$\mathcal{E}(k) \sim \eta_\beta^{\frac{2}{2\beta+1}} k^{-\frac{4\alpha\beta-2\beta+3}{2\beta+1}}. \quad (9.2.1.4)$$

The first thing to note is that except for the 2D vorticity equations, the constant flux of Casimirs of different powers with the method we have just used, results in different predictions for the spectral slope of the kinetic energy. This may be because our assumed form for the Casimir flux is incorrect, or that Casimirs do not all flux in the same way as enstrophy. Figure 9.5 shows how these predictions compare with the Kolmogorov predictions. Within the range $0 < \alpha < 1$, we saw in Figure 9.3 that the Kolmogorov predictions (which, of course, match those predicted here when $\beta = 1$) work well. In the limit of large β , the predictions for the slope is (taking the limit as $\beta \rightarrow \infty$)

$$\mathcal{E}(k) \sim k^{-2\alpha+1}. \quad (9.2.1.5)$$




Figure 9.5: Dimensional Kolmogorov arguments applied naively to Casimirs. The dashed green line indicates the slopes predicted for alpha models according to the Kolmogorov predictions outlined in the previous chapter. Darker blue lines indicate a dimensional argument applied to a cascade of a Casimir corresponding to enstrophy to the given power (hence β of 1 and the Kolmogorov prediction coincide). Simulated spectral slopes shown in Figure 9.3 are reproduced as red crosses.

This prediction matches those for large alpha up to the highest alpha investigated (α of 5, 10 and 13 were investigated).

However, in Figure 9.3, the steepening of slopes compared to Kolmogorov begins to occur before $\alpha = 2$. In addition, our assumption about the dimension of the Casimir flux is not well grounded. However, these issues can be partially remedied if we assume a nonlocal transfer, as in the following argument.

A nonlocal dimensional argument

We now extend the dimensional argument of Kolmogorov to a setting where fluxes of energy or enstrophy are nonlocal in spectral space. In equilibrium, all injected energy must be transferred to the smallest scales in some manner, and so I speculate that instead of the flux through each wavenumber being constant, the flux from all scales greater than a given wavenumber must be constant.

We will remain with Casimirs in the following. The issue of assuming that the dimension of their flux at a particular wavenumber is $\zeta_k^{2\beta}$ is slightly improved by this more nonlocal approach. Since their integral over all wavenumbers is their total value, their dimension over some large range of wavenumbers is more likely to match their dimension in real space $\zeta^{2\beta}$, making our dimensional argument more applicable.

So, assuming the transfer of Casimirs from all scales larger than our scale to all scales smaller is constant and occurs on the timescale of the eddy turnover time, we have:

$$\eta_\beta \sim \frac{\int_0^k \zeta_k^{2\beta} dk}{\tau_k}, \quad \text{our assumed nonlocal Casimir flux} \quad (9.2.1.6)$$

$$\sim \frac{\int_0^k k^{(2\alpha-1)\beta} \mathcal{E}(k)^\beta dk}{k^{-3/2} \mathcal{E}(k)^{-1/2}}, \quad \text{substitute } \tau_k \text{ and } \zeta_k^2 \text{ as in previous chapter} \quad (9.2.1.7)$$

$$\sim \frac{k^{(2\alpha-1)\beta+1} \mathcal{E}(k)^\beta}{k^{-3/2} \mathcal{E}(k)^{-1/2}}, \quad \text{integral scales as } k \quad (9.2.1.8)$$

$$\sim k^{(2\alpha-1)\beta+(5/2)} \mathcal{E}(k)^{\beta+(1/2)}. \quad \text{simplify} \quad (9.2.1.9)$$

When rearranged, this gives a prediction for the kinetic energy spectrum

$$\mathcal{E}(k) \sim \eta_\beta^{\frac{2}{2\beta+1}} k^{-\frac{4\alpha\beta-2\beta+5}{2\beta+1}}. \quad (9.2.1.10)$$




Figure 9.6: Assuming flux from all greater scales, rather than from only a particular scale steepens the predicted spectral slopes. Predictions for different values of β are again plotted. Simulated spectral slopes shown in Figure 9.3 are reproduced as red crosses.

Figure 9.6 shows the slopes for the spectrum this argument predicts. The predicted slopes no longer align for the 2DV equations ($\alpha = 2$), instead aligning for the RSF equations with constant internal heating ($\alpha = 3$). In the limit of large β , the prediction for the slope remains the same, this is the “limit prediction” shown in Figure 9.3.

Since the different Casimirs give different predictions for the slope in this argument, it is unclear whether it has any relevance to the problem. In order to progress, it would be helpful to investigate if Casimirs are truly conserved in my simulations and if their spectral flux is constant.

9.3 Error growth

We now begin an investigation of error growth in alpha models. Error growth in this context has not received attention and we believe this is the first investigation. We are going to find that increased nonlocality leads to error growth of a more upamplitude than upscale nature. These two regimes of error growth are a contemporary topic in error growth research and we now give a background on the subject.

9.3.1 Upscale vs upamplitude error growth

It has been noted that in simple models, such as the L69 model discussed in Chapter Six, error tends to grow roughly upwards in scale. That is, it is initially concentrated at small scales, saturating these first so that there is no predictability left in these scales, and then gradually infecting larger and larger scales. By contrast, in twin-experiment simulations of the atmosphere such as those presented by [Judt \(2018\)](#), we see that error has grown so significantly at larger scales before it saturates the smallest scales that we can no longer really call this mode of error growth upscale. Instead, it is called upamplitude. As stereotypes, upscale growth could be seen as an error completely saturating a scale and then working its way up through larger and larger scales in a cascade. Upamplitude growth could be seen as error quickly infecting every aspect of the system by equal amounts so that the error spectrum looks like a miniature version of the background kinetic energy spectrum, and then growing uniformly larger until it matches the magnitude of the background spectrum.

However, in practice it can be difficult to distinguish between upscale and upamplitude growth when defined in the above way. Even in pure turbulence models, the initial spectral distribution of error can severely affect the profile of error as it evolves ([Durran and Gingrich, 2014](#)). In operational settings, initial condition errors are certainly not isolated to the smallest scales present.

There are also more technical reasons distinguishing between upscale and upamplitude errors can be difficult. If one investigates the error spectrum using spectra calculated along lines rather than using a 2-D Fourier transform, although the total kinetic energy spectrum evaluated in this way will be qualitatively equivalent, the error spectrum, which has a portion increasing with wavenumber, will not be ([Rotunno et al., 2023](#)). This caused qualitative differences between the error spectra presented in [Judt \(2018\)](#) and [Judt \(2020\)](#), with those in the later study exhibiting flat spectra past the saturation scale, which may hastily be interpreted as a variety of upamplitude growth.

To avoid these issues, we give the following definitions, which focus on the error growth mechanism rather than the error spectrum's visual properties, to distinguish between the two flavours of error growth:

Upscale growth

This error primarily grows through a cascade of errors through scales. Growth of error at a particular scale is primarily fueled by an error in scales immediately smaller than it.

Upamplitude growth

The error does not cascade. It is transferred upscale non-locally, perhaps through instability.

We now investigate how these two kinds of error growth manifest in alpha models.

9.3.2 Twin-experiments in alpha models

Initial conditions were collected from the runs described in the previous section and used as branch-off points for twin experiments. The initial conditions selected were separated from each other by 3.5 hours (50 MTU). To perturb the twins, small perturbations were added to the generalised vorticity initial conditions in the form of a sine wave of random phase and small amplitude with a wavenumber of 80 (12.5 km). The amplitude of the sine wave perturbation was scaled in the same way as the forcing so that the kinetic energy added by the perturbation was comparable between different alpha and forcing wavenumbers. The KE amplitude used here was 10^{-15} .

The metric used to measure error growth is the difference kinetic energy. The DKE field is found by calculating the kinetic energy of the difference of the velocity fields, the DKE is then the sum over all grid points of the difference between the two velocity vector fields, divided by the total kinetic energy of one of the twins

$$\text{DKE} = \frac{\sum |\mathbf{u}_1 - \mathbf{u}_2|^2}{\sum |\mathbf{u}_1|^2}. \quad (9.3.2.1)$$

DKE fields for randomly chosen twins for each alpha at 3.5 hours (50 MTU) are shown in Figure 9.7.




Figure 9.7: DKE fields for randomly chosen twins at 3.5 hours (50 MTU), note different colour bars. For smaller alpha, error remains spectrally local but as alpha increases, error reaches larger spatial scales more quickly (cf Figure 9.8).

There is a large difference in the magnitude of error grown after 3.5 hours (50 MTU). The largest amplitude errors occurred for $\alpha = 4$, whereas the smallest in this case were seen in 2D Euler ($\alpha = 2$). The qualitative nature of the error is also different. Error spreads to all scales much faster for large alpha, whereas it remains concentrated at small scales for smaller alpha.

Spatial differences in the distribution of error

The differences in the shape of error growth seen in Figure 9.7 can be quantified by looking at the Fourier power spectrum of the DKE. Figure 9.8 shows the DKE power spectrum at equal time intervals of 42 minutes (10 MTU) in red and the background spectrum in black.




Figure 9.8: Power spectrum of the DKE for various alpha, each subsequent red line is the error power spectrum separated by 42 minutes (10 MTU). The initial red line is the error after 42 minutes, the initial error was localised to wavenumber 80 (12.5 km). The black line is the kinetic energy power spectrum of one of the twins. These have been averaged over 8 runs each. For small alpha error grows up through wavenumbers in a cascade whereas for greater alpha the error spectrum quickly resembles that of the background spectrum (black) and then grows up magnitude instead of up scale.

For small alpha, error grows fastest at smaller scales. This aligns with the theory of an error cascade where error saturates the smallest scales first and then saturates subsequent scales in an error cascade. In the atmospheric context, this can be called upscale error growth. For larger alpha, error quickly resembles a miniature version of the background spectrum. It then grows in magnitude across all scales. In the atmosphere, this style of error growth can be called upamplitude error growth.

Differences in total error growth

Figure 9.9 shows the logarithm of the ratio of the DKE after 3.5 hours (50 MTU) to the initial DKE, plotted against alpha. This shows an interesting valley shape with a minimum around $\alpha = 2$. For small alpha, as alpha gets smaller, error growth is on average more intense. For alpha larger than 2, increasing alpha corresponds with increasing error growth.

We hypothesise that the minimum of predictability is due to the competing effects of non-locality and a local cascade of error. When the fluid is especially non-local (high alpha), error in one part of the fluid strongly influences all other locations, enhancing the growth of error. For small alpha, error primarily cascades up through wavenumbers, this cascade is proportional in speed to the eddy turnover time and so occurs more quickly for a shallower spectral slope. Hence, error grows more quickly for smaller alpha.




Figure 9.9: The logarithm of the ratio of DKE at 3 hours (50 MTU) to the initial DKE against alpha. Symbols indicate different twin-experiments. Blue cross indicates slow forcing injected at wavenumber three, orange plus indicates fast forcing at wavenumber three and green dot are twin experiments with fast forcing at wavenumber five. Alpha 2 sits at a maximum of predictability likely caused by the competing effects of non-locality and spectral slope shallowness.

9.4 Summary and conclusions

In this chapter, we have presented simulations and twin experiments of alpha models. We have speculated on how the slopes of the more nonlocal alpha models can be predicted with dimensional arguments similar to those of Kolmogorov. We have then shown how their increased nonlocality affects the flavour of error growth in these models. If a cascade of error were responsible for error growth in the higher alpha models, higher alpha values would correspond to slower error growth. Instead, higher-alpha models exhibit the fastest error growth due to nonlocal error transfers to the largest scales.

Upscale vs upamplitude error growth is a topic of current interest in error research. We have proposed clear definitions for the two flavours. That upscale error growth should be defined as an error that primarily grows through a cascade with error at a larger scale primarily depending on an error at a scale just smaller than it. Upamplitude error should be defined as error growing nonlocally in spectral space, the opposite of upscale growth. Through our more general perspective, we can see that even purely turbulent systems such as 2DV lie on a continuum between these two error growth forces. Perhaps by diagnosing the flux of error to different scales, which of these kinds of error growth exists in simulations of the atmosphere can be answered in the future.

Directions for future investigation

Alpha models provide an excellent testing ground for the ideas presented in Part Two. This is because that model was based entirely on an upscale understanding of error growth and so should be most applicable to low alpha models. When attempting to extend the model, a logical direction could be to attempt to

add a nonlocal error growth aspect. This can be tested by increasing the value of alpha and seeing how predictions for error growth are altered.

In addition, it would be of interest to better understand the dynamics of Casimirs in the system and how they move in spectral space. This can bring clarity to our speculations on applying nonlocal and local Kolmogorov arguments to these Casimirs.

In Part One, we discussed that in different locations, SQG or 2DV dynamics may be relevant. This work suggests that error growth under these different conditions may differ in how upamplitude or upscale it is, with SQG dynamics leading to more upscale growth while the 2DV dynamics would be associated with more upamplitude growth. By applying coarse-graining to global twin experiments, we may be able to probe this possibility. This could complement future investigations into the magnitude of error flux in simulations of the atmosphere.

We now give our overall conclusions from the thesis.

Chapter 10

Conclusions

In this chapter, we present a summary and the key conclusions of this thesis. We then give an outline for future work we wish to see in this area.

10.1 Summary of key results

This thesis has been split into three parts: local spectra in the atmosphere, predictability, and turbulence theory. Each of these parts is, at its core, about trying to leverage a spectral perspective to gain insights into the nature of turbulent systems. In Part One, this was done by applying coarse-graining to ECMWF analysis, revealing local differences in the mesoscales to a level of detail that has not previously been seen. In Part Two, we used the spatially local work of Part One to extend the ideas of [Lorenz \(1969\)](#) and [Lilly \(1973\)](#) to a local model. The success of this toy model was limited but laid the groundwork for future work we propose in the sections below. In Part Three, we analysed error growth in alpha models and found that high alpha models demonstrate a more non-spectrally-local upamplitude flavour of error growth. This family of turbulent models could be a testbed for developing the model in Part Two further and investigating the dynamics of coarse-grained error growth.

The main conclusions of this thesis are as follows:

10.1.1 Part One

Spectra are spatially inhomogeneous

In Chapter Three, we showed the first consistent spatial maps of mesoscale spectral slopes and average mesoscale energy. We clearly show the spatial inhomogeneity of the spectrum, and it is visually apparent that the shallower slopes in the tropics measured in previous studies are correlated with patterns of average precipitation. Showing the spatial differences in the spectrum provides a more nuanced understanding of the spectral properties of the atmosphere. Unlike a simple two-dimensional turbulence toy model, whose spectrum is independent of location, in the atmosphere, different mesoscale distributions of energy prevail in different areas. This could have applications in model diagnostic procedures, where, by comparing local spectra across the globe, one can grasp differences between models. It also gives us clues to the underlying dynamics in the mesoscale, with orography and moist processes clearly playing a role in energising mesoscale motion.

Atmospheric conditions can be correlated with systematic changes in the mesoscale spectral slope

In Chapter Three, we showed that increased precipitation and orographic height were both associated with shallower mesoscale spectral slopes at both 200 and 600 hPa. Through creating the concept of “conditioned spectra”, we were able to quantify these relationships. These results could form a starting point for comparisons with other models for use as a diagnostic to assess how a model represents topography and moist processes. Similarly, we showed that at 200 hPa, higher MSLP was associated with steeper slopes, while at 600 hPa, greater MSLP was associated with shallower slopes. In addition, we showed that the

direction of spectral kinetic energy flux was not strongly associated with a change in mesoscale slope. Although the connection between local spectral slopes and local predictability remains an open question, these results could have future relevance for assessing the impact of each of these features on the growth of error in weather forecasts. In addition, partitioning spectra based on the local flux of error growth has not been performed before. We have shown that even when local flux is on average to larger scales, we still see the characteristic $-5/3$ slope. One cannot assume that this $-5/3$ is due to a turbulent inverse cascade of kinetic energy, and indeed, in Chapter Four, we found that strong inverse transfers of energy could be associated with divergent processes.

Spectral fluxes are both spatially inhomogeneous and differ in correlation with atmospheric conditions

In Chapter Four, we showed the first maps of average mesoscale spectral transfer in high resolution analysis data, observing that this transfer was not spatially homogenous. This map provides a valuable starting point for future parametrisation design. We also showed only the rotational component of these fluxes, which is the component due to the two-dimensional turbulent motion. Although we see that in some regions it is possible turbulent dynamics play a leading role, in most regions, we confirmed the current understanding that divergent motion is centrally important for the flux of energy through the mesoscales. For instance, even though there was a band of inverse energy transfer found along the equator at 200 hPa, by assessing the rotational component of this flux, we concluded it was not due to an inverse turbulent cascade but rather was likely due to large-scale divergence. Using conditioned fluxes, we showed that the magnitude and direction of transfer could change depending on prevailing atmospheric conditions. In particular, we showed how MSLP, orography and precipitation systematically change the profile of spectral flux. This could again serve as a useful starting point for parametrisation design, as a model diagnostic, or to evaluate existing parametrisations.

Classic two-dimensional phenomenology likely occurs in some strongly rotational atmospheric flow

In Chapter Five, we decomposed the local spectra into rotational and divergent components and compared their magnitudes, creating the first map of the relative magnitude of divergent and rotational energy in the mesoscales. We saw that rotational motion was more dominant in some regions, opening the door for the return of two-dimensional turbulent theories to explain local properties of the spectrum in some locations.

We investigated further by conditioning spectra and fluxes based on the local ratio of divergent to rotational energy. At 200 hPa we saw fluxes and spectra that were consistent with surface quasi-geostrophic motion. At 600 hPa we saw characteristics consistent with the two-dimensional vorticity equations, with a strong inverse transfer for scales larger than 300 km and a shallow spectral slope close to $-5/3$, and a weak direct transfer with a steep spectral slope of -3 for scales smaller than 300 km. This result encourages further investigation to uncover atmospheric conditions when a two-dimensional turbulence phenomenology may dominate in the mesoscales. Under such conditions, arguments for local error growth through a cascade to larger error would be most applicable, making these conditions interesting targets for future local error growth studies.

10.1.2 Part Two

Purely turbulent atmospheric dynamics may retain predictability more than a month ahead rather than only two weeks

In Chapter Six, we explained the derivation of the [Lorenz \(1969\)](#) (L69) turbulent error growth model. We then

altered the spectrum used by the L69 model to include a more accurate representation of the condensate, the shallowing of the spectrum at large scales. The result was a decrease in mesoscale predictability by half a day, an increase in predictability of the weather scales by three and a half days and an increase of more than double, from 19.3 to 43.1 days for the outer limit of predictability, when all skill from the initial conditions is lost. However, it remains to be seen if this result holds in more realistic settings, such as in a formulation of the Lorenz error growth model on a beta plane.

The above result suggests that a purely turbulent model can retain skill from its initial conditions more than a month into the future. The current perspective in seasonal forecasting is that skill comes from components of the Earth system with longer timescales of predictability than the atmosphere, such as sea surface temperature or atmospheric phenomena that rely on more complicated mechanisms than two-dimensional turbulence, such as the quasi-biennial oscillation. However, our result suggests that some component of the skill may come from the turbulent evolution of the initial conditions, and so perspectives such as spectral error growth cascades may have relevance beyond weather prediction and into seasonal forecasting.

Local spectra offer a new direction for error prediction

In Chapter Seven, we created a toy model for local error growth based on the dimensional argument of Lilly (1973) and our results from Chapter Three. We compared estimates of error growth from our dimensional model with the RMSE of the ECMWF forecast. We found that some of the features of error growth observed in forecast RMSE were present in our model, such as increased absolute error in the storm tracks, and increased error over orography, but other features, such as an anomalously high relative error in a region of the Indian Ocean were less desirable. The largest discrepancies were the smaller absolute growth of error in the storm track regions of the extratropics, which were perhaps due to our dimensional model not accounting for downscale transfers of error and the much smaller relative errors in the tropics. The smaller relative errors in the tropics are likely a limitation of Lorenz's original model since error growth in that region is mostly driven by moist processes. Despite the shallowing of the spectrum associated with greater precipitation, it seems that the full extent of moist error growth is not so straightforwardly associated with the locally decreased eddy turnover time.

10.1.3 Part Three

Spectral slopes in the enstrophy inertial range are steeper than Kolmogorov arguments predict for high alpha models

In Chapter Nine, we showed the results of simulations of alpha models. We observed that spectral slopes were steeper than our predictions made using Kolmogorov arguments in Chapter Eight. It is known that due to the increased nonlocality of the higher alpha models, the dimensional arguments should begin to fail, but it had not been recognised that the slope may empirically be predicted by the formula $1 - 2\alpha$.

Error grows in a more upamplitude than upscale manner for higher alpha

In Chapter Nine, we show the result of twin experiments in alpha models. By observing the initial evolution of error profiles in spectral space, we noted that at high alpha, error profiles quickly appear to become similar to the background spectrum but at a smaller magnitude. This error profile, which resembles a miniature background spectrum, then grows in time, up in amplitude. Over the same period in lower alpha models, the error remains relatively peaked around its initial small scale as it evolves. Interestingly, the

partition point between these spectrally local upscale models and upamplitude models is the 2D vorticity equations, which are the classic two-dimensional turbulent system. Since more energetic spectra are associated with lower α , upscale error grew fastest in low α models, whereas errors grew faster when models were increasingly nonlocal. As a result of these two effects, the 2D vorticity experienced the least error growth of the models we investigated. In Part One, we noted that there might be conditions in the atmosphere where SQG ($\alpha = 1$) phenomenology and 2DV ($\alpha = 2$) phenomenology are dominant. If this is the case, our results imply that in future local error growth studies of the atmosphere, we will see differences in how upscale or upamplitude error growth is under these conditions.

10.2 Future work

In Chapter One, we outlined some of the reasons kinetic energy power spectra are academically interesting. These reasons were understanding mesoscale weather, the potential for model validation and parametrisation design, and improving understanding of predictability. We discuss future work towards each of these goals.

10.2.1 Understanding mesoscale dynamics

We give three suggestions for extending the analysis found in Part One in order to gain further insights into mesoscale dynamics.

Extending the investigation of the ECMWF analysis

In Chapters Three, Four and Five, we uncovered much about the inhomogeneity of spectra and atmospheric fluxes. However, there are several limitations to the work. The most important may be the restriction to two pressure levels. If the analysis could be repeated on all available pressure levels, it would be possible to look at the vertical profile of spectra in a region or under a condition of interest. This would be especially interesting above orography. The spectrum could be evaluated above points with sufficiently high orography and compared with critical layer heights to test the hypothesis that vertically propagating gravity waves energise the shallow spectra in the stratosphere. In particular, figures like Figure 1.2, which show spectral slopes at each pressure level and length scale, could be produced from conditioned spectra. This would show how conditions affect spectra throughout the depth of the atmosphere. Another limitation is only a single year was analysed. Extending the analysis to multiple years would allow for an investigation into the climatology of the data. Seasonal spectral maps could be produced. Figure 10.1 shows 2020 average mesoscale energy decomposed into the four seasons present. With the addition of more data, this could be transformed into a climatology of spectra.

Improved breakdown of spectral fluxes

In [Augier and Lindborg \(2013\)](#), a more comprehensive breakdown of spectral energy flux is presented. They break down the flux into total energy, potential energy, vertical fluxes and divergent and rotational kinetic energy components. The potential energy and vertical transfers are the key improvements compared to our divergent and rotational breakdown. The most interesting aspect of this breakdown for mesoscale dynamics may be the inclusion of the vertical flux term. If this approach can be extended to a coarse-graining context, then maps of flux due to vertical transfers could be created. Since the leading explanation for the shallowing of the mesoscale spectrum in the upper troposphere and lower stratosphere is upwardly propagating gravity waves, such a decomposition could neatly answer where on the globe and under which conditions this

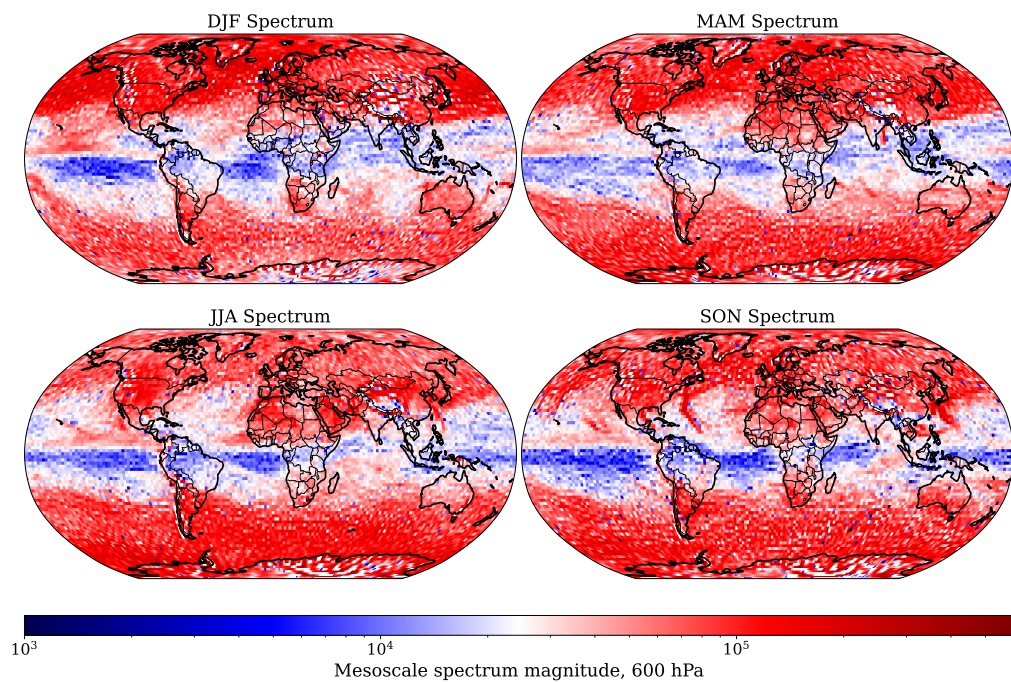


Figure 10.1: A global map of 600 hPa seasonal average mesoscale energy in 2020 ECMWF analysis data on a logarithmic colour bar.

explanation holds more or less weight.

Application to other datasets

As we showed in Chapter One, ECMWF analysis has a clear bias towards steeper mesoscale spectra than observations, at least in the upper troposphere and lower stratosphere where observations are abundant. Other systems, such as ICON (a numerical weather prediction model used by several weather services), produce shallower mesoscale slopes (Avalos, 2024). Reproducing our analysis on a different dataset could help us understand the differences between these systems and give us more confidence in the generality of our results. Direct observations are not readily available. However, we may gain additional insights by analysing the spectra of satellite cloud observation products such as the International Satellite Cloud Climatology Project (ISCCP).

10.2.2 Parametrisation improvement and model evaluation

The use of coarse-graining for model evaluation can help with the development of parametrisations. This is because a more comprehensive understanding of how the model differs from observation or expectation can inform refinements to parametrisation schemes. We start by outlining how such diagnostics may be carried out and then some particular cases that would be interesting first targets for analysis.

Coarse-grain diagnostics

By evaluating local and conditioned spectra and spectral fluxes before and after a change, one can identify the effect of a parametrisation update, resolution change or the difference between how two distinct models perform under a given condition. In practice, one would wish to perform many test simulations to aggregate comparable spectra for the two cases. However, consider the June 2020 update to the IFS (ECMWF, 2020) to illustrate how the diagnostic may look. This update improved stratospheric forecasts, likely due to

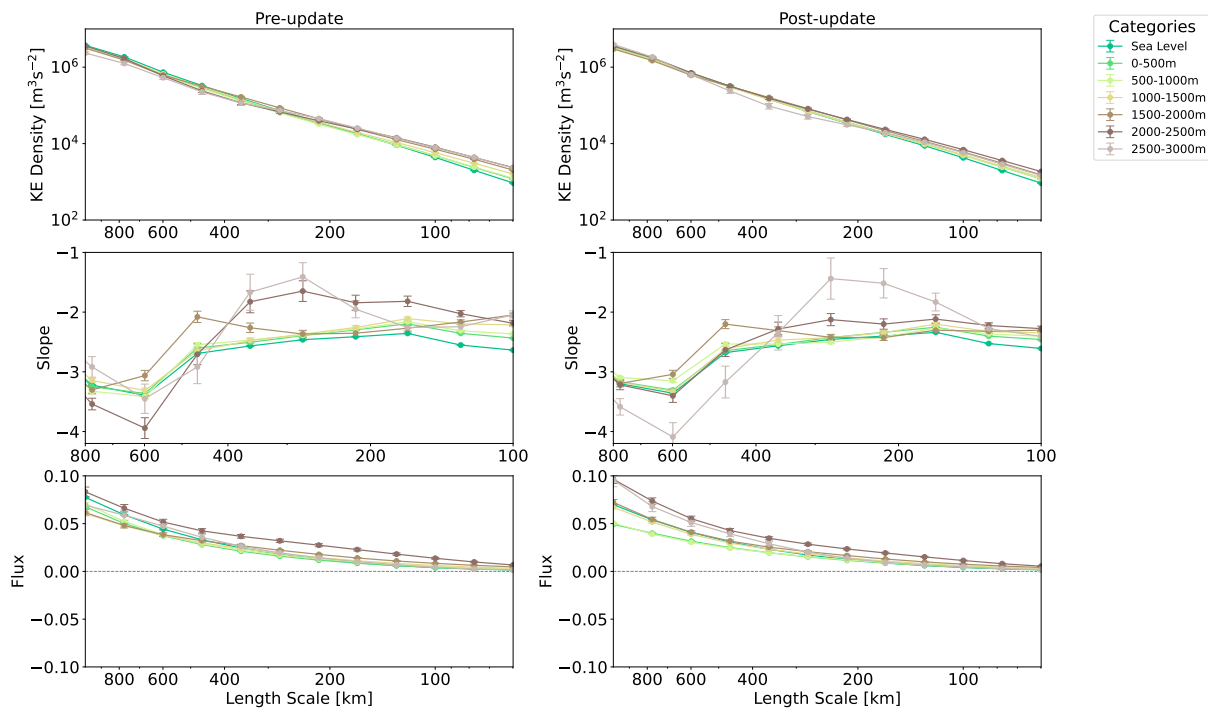


Figure 10.2: Conditioned spectra on orographic height (excluding precipitating points) between January and June on the left (pre-update) and July to December (post-update) from ECMWF analysis for 2020 at 200 hPa. The top panels show the conditioned spectra, the middle panels show spectral slopes, and the bottom panels show the spectral energy transfer. Positive values indicate a transfer of energy to smaller scales.

an improved representation of vertically propagating waves by improving the vertical interpolation used. The convective scheme was also updated to correct energy conservation and momentum diffusion issues. However, the updates to the convection scheme were reported as having a neutral impact on performance. We may conjecture that an improved representation of vertically propagating waves will lead to a more energetic spectrum in the stratosphere over orography. To test this, we show the spectra conditioned on orography before and after the update in Figure 10.2. Given the asymmetry between Boreal winter and summer, differences appear minimal before and after the update. However, suppose we applied this analysis to the test data produced to compare the pre and post-update prediction system. We could evaluate if the scheme systematically affected the spectra over orography.

Using local spectral fluxes to evaluate parametrisations

Parametrisations are used to represent the effects of unresolved or partially resolved scales or processes in weather models. For instance, [Shutts \(2005\)](#) presents a scheme for representing the transfers of energy from subgrid scales to larger scales. By producing maps showing where such schemes inject energy and comparing this to the local coarse-grained energy dynamics in a higher resolution “truth” run, the effectiveness of these and similar parametrisations could be quantified and thereby improved.

Local effective resolution

Effective resolution is an estimate of the scale at which features in a simulation are resolved. A successful heuristic for estimating this scale has been $7\Delta x$, where Δx is the grid spacing of the simulation. To evaluate the effective resolution in GCMs, [Skamarock \(2004\)](#) introduced the method of comparing the global spherical harmonic spectrum to the observational fit presented in [Lindborg \(1999\)](#), that is, evaluating




Figure 10.3: A schematic illustrating a proposal for local or conditioned effective resolution. The black solid line indicates the unknown true power spectrum of the atmosphere. It is assumed the models (in this case ERA5 and the ECMWF analysis (denoted by IFS)) match this spectrum until they cease to properly model small scales and their spectra steepen. Red dashes denote the scale at which this steepening occurs. This scale is the effective resolution. Numerical values are qualitative rather than quantitative. Instead of evaluating against the theoretical spectrum, which is only well constrained by observations at the cruising altitude of airliners, local spectra from a lower resolution simulation (ERA5) can be evaluated against local spectra from a higher resolution analysis (denoted by IFS). This allows the proposed method to extend to locations or conditions for which the observed spectrum is unknown.

when the mesoscale spectral slope begins to deviate from $-5/3$. A limitation of this methodology is that the observations used are only valid in the upper troposphere and lower stratosphere global average. As discussed in Chapter One, spectra steepen significantly at small scales in the analysis in the stratosphere. This may indicate a lower effective resolution in the stratosphere or, more likely in the author's view, it could be a physical consequence of less vertically propagating gravity waves making it to higher altitudes. This issue is sidestepped if we use a higher resolution simulation coarse-grain spectrum climatology as the comparison point. Figure 10.3 illustrates the idea. The effective resolution for a particular condition or location is defined as the scale when the average spectrum deviates from the climatology of the spectra produced by a higher-resolution simulation. This assumes the higher-resolution simulation more accurately represents the mesoscales than the lower-resolution one.

Work has been done in this direction, in [Boljani et al. \(2022\)](#), effective resolution was defined regionally in the tropics, midlatitudes and poles by finding the scale that spectra in ERA5 deviated from a $-5/3$ slope. However, for the reasons stated, this could be improved through coarse-graining and by reducing our certainty in the expected shape of the mesoscale spectrum by using a higher resolution simulation as our comparison benchmark, rather than assuming a $-5/3$ slope. If we carried out this analysis, we may be able to draw conclusions such as the effective resolution of ERA5 under clear conditions over the ocean is around 150 km, whereas in convectively active regions, it is closer to 300 km. Such statements could bring more clarity to when ERA5 is an appropriate dataset for the investigation of a given scientific question.

Evaluating AI forecasts

Artificial intelligence (AI) forecasting has rapidly begun to revolutionise weather prediction. A persistent issue with AI forecasts is they are blurry. This is a consequence of them being built to optimise root mean square error or similar metrics. In applications such as downscaling, where the purpose of the AI system is to add more detailed features, this is the central problem that must be solved; for work on this topic,

see [Xie et al. \(2020\)](#) and [Ravuri et al. \(2021\)](#). However, a similar drive to reduce the blur of forecasts is present in some work on weather prediction AI, for instance, [Liu and Lee \(2020\)](#) and [Hu et al. \(2021\)](#). It is not obvious that reducing blur is a necessary or desirable goal for an AI prediction system. It may be that blur accounts for the intrinsic unpredictability of smaller scales, if it does, then it may be a feature of the method and not a bug. To test if this is the case, we could compare the coarse-grained RMSE to the coarse-grained AI forecast spectra. If scales where error is saturated are simply not represented by the AI forecast, then this is debatably a better form of solution than the output of a traditional fluid simulation. We discuss this further in the next section on predictability.

10.2.3 Predictability

In the view of the author, the central task of weather prediction is to quantify the growth of uncertainty about the state of the atmosphere. The perfect weather forecast starts from the probability distribution of possible current states and tracks the evolution of this distribution until it matches that of climatology and all the predictive power of the initial measurements have evaporated. This is not currently plausible¹. However, suppose that we could approximate at a given location and time, either through theoretical insight or an empirical approach such as AI, how error in an infinite ensemble should grow at that point. Then, with a single simulation, we could approximate the spatial growth of error up to the point of error saturation, without running an ensemble. This would be an attempt to solve the full initial condition problem, that is, to extract all possible information about the system we possibly can. As mentioned under the heading of evaluating AI forecasts, it is possible we are already moving in this direction. In Chapter Seven, we presented a step along a more theoretical avenue towards the same goal. We now outline how this project may be carried forward.

Developing a local model of error growth in alpha models

In Part Three, we presented numerical work on alpha models. Our results suggested that error grows in more of a cascade manner for small alpha and in a less spectrally local way for larger alpha. This makes alpha models an attractive framework for developing the local error growth model presented in Chapter Seven. First, we could start by evaluating the local characteristics of error growth in alpha models. For the SQG equations ($\alpha = 1$) that seem to demonstrate a more cascade style of error growth, we can ask if error locally grows in a cascade or if this only appears to be the case in aggregate. In this more simple setting, it may be possible to find a connection between error growth and local spectral slopes. Increasing alpha to two and studying the classic two-dimensional turbulence system introduces more spectral nonlocality to the error growth and we can examine how this changes any relationship we find at the lower value of alpha.

Analysing error growth in full physics simulations or forecasts through coarse-graining

The complement of analysing error in the simple alpha-model context is analysing it in full physics simulations, such as the twin experiments presented in [Judt \(2018\)](#), or in an ensemble forecasting system. This could be an important step towards understanding the growth of uncertainty in our forecasting systems. One may imagine a similar analysis to that presented in Chapters Three, Four and Five, but with kinetic energy replaced by difference kinetic energy profiles at fixed intervals of time. We could create maps showing the time into a forecast or simulation when a particular length-scale saturates at each spatial location. We could

¹One can write the equations that govern the evolution of this distribution, they are the Liouville equations, or with the addition of noise, the Fokker-Plank equations. However, the huge dimensionality of these equations for the atmosphere makes this method intractable.

create conditioned error growth spectra showing how error tends to grow in the presence of orographic features, precipitation, differing directions of energy flux, MSLP or other variables of interest, such as total kinetic energy. It is a more speculative idea, but it may also be possible to formulate an expression for the flux of error through scales, we could then directly quantify how much different flow features contribute to our loss of predictability. A key target of the analysis would be to assess under which, if any, conditions, error grows in a scale-to-scale rather than upamplitude manner.

Appendix A

Fourier transform and structure function refresher

A.1 Fourier transform refresher

A.1.1 Introduction

The Fourier transform allows us to write a function $f(x)$ as¹

$$f(x) = \int_{-\infty}^{\infty} e^{2\pi i k x} \hat{f}(k) dk. \quad (\text{A.1.1.1})$$

The term $e^{2\pi i k x} = \cos(2\pi k x) + i \sin(2\pi k x)$ is called a Fourier mode. Here, k is the wavenumber. If we move a distance $\lambda = 1/k$ then the periodicity will mean the value of the Fourier mode remains the same, so this tells us how wavenumber and wavelength are related.

The Fourier coefficient is a complex number $\hat{f}(k) = A e^{i\theta}$ so the integrand is $A(\cos(2\pi k x + \theta) + i \sin(2\pi k x + \theta))$.

Thinking of integration as an infinite sum, the formula can be interpreted as a way to construct a function² f by summing together an infinite number of sine and cosine waves with cleverly chosen amplitudes and phases.

A.1.2 The orthogonality of Fourier modes

The key property that allows functions to be decomposed into Fourier modes is that those modes are orthogonal. That is, if $m \neq k$,

$$\int_{-\infty}^{\infty} e^{2\pi i k x} e^{-2\pi i m x} dx = 0. \quad (\text{A.1.2.1})$$

We now show that $\int_{-\infty}^{\infty} e^{2\pi i(k-m)x} dx = \delta(k-m)$.

Evaluating the Fourier mode integral

We want to evaluate the integral

¹There are a few different conventions for the definition. If there is no 2π in the Fourier mode $e^{2\pi i k x}$, a multiplicative factor of $1/2\pi$ will appear somewhere in our definition. Whether the factor of $1/2\pi$ is put in front of the Fourier transform, the inverse Fourier transform or split between the two leads to distinct definitions. These different conventions are a source of confusion. Depending on which one is being followed, how k relates to length-scale changes, and factors can appear or disappear in various results that use the Fourier transform. For example, Plancherel's theorem, discussed later, will have no factor of $1/2\pi$ due to our choice of definition, but with an alternative definition, it will.

²When the Fourier transform converges and in what sense is the subject of a branch of maths called Fourier analysis. In a pointwise sense, continuous and integrable functions will have convergent transforms.

$$\int_{-\infty}^{\infty} e^{2\pi i k x} dx. \quad (\text{A.1.2.2})$$

We can do this using the following neat trick. In the limit as $\epsilon \rightarrow 0$, $e^{-\epsilon x^2} \rightarrow 1$. So our integral can be manipulated as follows,

$$\int_{-\infty}^{\infty} e^{2\pi i k x} dx = \lim_{\epsilon \rightarrow 0} \int_{-\infty}^{\infty} e^{2\pi i k x} e^{-\epsilon x^2} dx \quad \lim_{\epsilon \rightarrow 0} e^{-\epsilon x^2} = 1 \quad (\text{A.1.2.3})$$

$$= \lim_{\epsilon \rightarrow 0} e^{-\frac{\pi^2 k^2}{\epsilon}} \int_{-\infty}^{\infty} e^{-\epsilon(x - \frac{i\pi k}{\epsilon})^2} dx \quad \text{complete the square} \quad (\text{A.1.2.4})$$

$$= \lim_{\epsilon \rightarrow 0} e^{-\frac{\pi^2 k^2}{\epsilon}} \sqrt{\frac{\pi}{\epsilon}}. \quad \text{standard Gaussian integral} \quad (\text{A.1.2.5})$$

As ϵ approaches zero, the Gaussian $e^{-\frac{\pi^2 k^2}{\epsilon}}$ becomes increasingly peaked at k . However, its integral remains independent of ϵ as (again evaluating a standard Gaussian integral) we have

$$\sqrt{\frac{\pi}{\epsilon}} \int_{-\infty}^{\infty} e^{-\frac{\pi^2 k^2}{\epsilon}} dk = \sqrt{\frac{\pi}{\epsilon}} \sqrt{\frac{\epsilon}{\pi}} \quad (\text{A.1.2.6})$$

$$= 1. \quad (\text{A.1.2.7})$$

Hence, in the sense of distributions, our limit converges to

$$\lim_{\epsilon \rightarrow 0} e^{-\frac{\pi^2 k^2}{\epsilon}} \sqrt{\frac{\pi}{\epsilon}} = \delta(x). \quad (\text{A.1.2.8})$$

So,

$$\int_{-\infty}^{\infty} e^{2\pi i k x} dx = \delta(k). \quad (\text{A.1.2.9})$$

Leveraging orthogonality

Suppose the Fourier decomposition is possible, that is, there is some function $\hat{f}(k)$ such that

$$f(x) = \int_{-\infty}^{\infty} e^{2\pi i k x} \hat{f}(k) dk. \quad (\text{A.1.2.10})$$

If we multiply by a particular Fourier mode and integrate, orthogonality will ensure the only term that survives is the corresponding Fourier coefficient for that mode. That is

$$\int_{-\infty}^{\infty} e^{-2\pi i k x} f(x) dx = \hat{f}(k). \quad (\text{A.1.2.11})$$

A.1.3 Definitions

The following definitions, therefore, make sense;

$$\hat{f}(k) = \int_{-\infty}^{\infty} e^{-2\pi i k x} f(x) dx, \quad (\text{A.1.3.1})$$

$$\check{f}(x) = \int_{-\infty}^{\infty} e^{2\pi i k x} f(k) dk. \quad (\text{A.1.3.2})$$

Where \check{f} is the inverse Fourier transform and $\check{\check{f}} = f$.

Extension to \mathbb{R}^n

The same arguments apply if we extend to $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Our definitions become

$$\hat{f}(\mathbf{k}) = \int_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{x}} f(\mathbf{x}) d\mathbf{x}, \quad (\text{A.1.3.3})$$

$$\check{f}(\mathbf{x}) = \int_{\mathbb{R}^n} e^{2\pi i \mathbf{k} \cdot \mathbf{x}} f(\mathbf{k}) d\mathbf{k}. \quad (\text{A.1.3.4})$$

\mathbf{k} is called the wavevector. The magnitude of this vector is what we now call the wavenumber $\kappa = |\mathbf{k}|$. We can extend the definition to $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, simply by applying the previous definition to each component

$$\mathbf{f}(\mathbf{k}) = \int_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{x}} \mathbf{f}(\mathbf{x}) d\mathbf{x}, \quad (\text{A.1.3.5})$$

$$\mathbf{f}(\mathbf{x}) = \int_{\mathbb{R}^n} e^{2\pi i \mathbf{k} \cdot \mathbf{x}} \mathbf{f}(\mathbf{k}) d\mathbf{k}. \quad (\text{A.1.3.6})$$

A.1.4 Plancherel's theorem

We will make use of a few properties of the Fourier transform, one fantastically useful property is that

$$\int_{\mathbb{R}^n} \mathbf{f}(\mathbf{x}) \overline{\mathbf{g}(\mathbf{x})} d\mathbf{x} = \int_{\mathbb{R}^n} \mathbf{f}(\mathbf{k}) \overline{\mathbf{g}(\mathbf{k})} d\mathbf{k}. \quad (\text{A.1.4.1})$$

This is so useful in our context because of the special case³ with $\mathbf{f} = \mathbf{g}$. In that case, changing \mathbf{f} to \mathbf{u} to make clear the kinetic energy connection,

$$\int_{\mathbb{R}^n} |\mathbf{u}(\mathbf{x})|^2 d\mathbf{x} = \int_{\mathbb{R}^n} |\mathbf{u}(\mathbf{k})|^2 d\mathbf{k}. \quad (\text{A.1.4.2})$$

That is, the Fourier transform conserves the kinetic energy.

³Sometimes called Parseval's theorem.

Proof

The proof for higher dimensions is exactly the same as that for $f : \mathbb{R} \rightarrow \mathbb{R}$, so we show the 1D case for simplicity.

$$\int_{-\infty}^{\infty} f(x)\overline{g(x)} dx = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{2\pi i k x} \hat{f}(k) dk \int_{-\infty}^{\infty} e^{-2\pi i l x} \overline{\hat{g}(l)} dl dx \quad \text{properties of complex conjugate} \quad (\text{A.1.4.3})$$

$$= \iint_{-\infty}^{\infty} \hat{f}(k)\overline{\hat{g}(l)} \int_{-\infty}^{\infty} e^{2\pi i(k-l)x} dx dl dk \quad \text{rearrangement (Fubini)} \quad (\text{A.1.4.4})$$

$$= \iint_{-\infty}^{\infty} \hat{f}(k)\overline{\hat{g}(l)}\delta(k-l) dl dk \quad \text{orthogonality of Fourier modes} \quad (\text{A.1.4.5})$$

$$= \int_{-\infty}^{\infty} \hat{f}(k)\overline{\hat{g}(k)} dk. \quad \text{definition of the delta function} \quad (\text{A.1.4.6})$$

A.1.5 The power spectrum

For a particular wavevector, the power spectrum $E(\mathbf{k}) := |\mathbf{u}(\mathbf{k})|^2$. The power spectrum for a given wavenumber κ is then

$$E(\kappa) := \int_{|\mathbf{k}|=\kappa} |\mathbf{u}(\mathbf{k})|^2 d\mathbf{k}. \quad (\text{A.1.5.1})$$

By Plancherel (Equation A.1.4.2) we know that

$$\int_{\mathbb{R}^n} |\mathbf{u}(\mathbf{x})|^2 d\mathbf{x} = \int_0^{\infty} E(\kappa) d\kappa. \quad (\text{A.1.5.2})$$

A.1.6 The convolution theorem

Another useful Fourier result is that convolution is a multiplication under a Fourier transform. That is

$$\widehat{f * g}(k) = \hat{f}(k)\hat{g}(k). \quad (\text{A.1.6.1})$$

This result generalises to \mathbb{R}^n easily.

The convolution operation can be thought of as locally averaging using some weighting function $g(x)$ (by convention, it actually weights using the mirror image of g , $g(-x)$). The definition of the convolution is

$$f * g(x) = \int_{-\infty}^{\infty} f(y)g(x-y) dy. \quad (\text{A.1.6.2})$$

Proof

The proof is again just definitions and some simple rearrangement,

$$\begin{aligned}
\widehat{f * g}(k) &= \int_{-\infty}^{\infty} f * g(x) e^{-2\pi i k x} dx && \text{definition of Fourier transform} && \text{(A.1.6.3)} \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(y) g(x - y) dy e^{-2\pi i k x} dx && \text{definition of convolution} && \text{(A.1.6.4)} \\
&= \int_{-\infty}^{\infty} f(y) \int_{-\infty}^{\infty} g(x - y) e^{-2\pi i k x} dx dy && \text{rearrangement} && \text{(A.1.6.5)} \\
&= \int_{-\infty}^{\infty} f(y) \int_{-\infty}^{\infty} g(u) e^{-2\pi i k(u+y)} du dy && \text{substitution } u = x - y && \text{(A.1.6.6)} \\
&= \int_{-\infty}^{\infty} f(y) e^{-2\pi i k y} \int_{-\infty}^{\infty} g(u) e^{-2\pi i k u} du dy && \text{rearrangement} && \text{(A.1.6.7)} \\
&= \widehat{g}(k) \widehat{f}(k). && \text{definition of Fourier transform} && \text{(A.1.6.8)}
\end{aligned}$$

A.1.7 Isotropy and the Hankel transform

A common assumption in turbulence is that the field we are dealing with is isotropic. If this is the case, then the Fourier coefficient will only depend on the wavevector rather than the wavenumber.

This is because the wavevector gives the orientation of the corresponding Fourier mode; if two Fourier modes with the same wavenumber (but with different wavevectors) had different Fourier coefficients, a preferred direction to the flow would be introduced.

The Hankel transform

When we apply a Fourier transform to a 2D isotropic function, we can simplify the formula a little. This simplified form of the Fourier transform is called a Hankel transform.

This simplification is really just using the following integral from Abramowitz and Stegun [Abramowitz and Stegun \(1972\)](#)

$$J_0(r) = \frac{1}{2\pi} \int_0^{2\pi} e^{ir \cos \theta} d\theta. \quad \text{(A.1.7.1)}$$

Here J_0 is the 0th order Bessel function of the 1st kind. One can think of it as a ripple spreading out radially where the amplitude is somewhat like the positive x portion of a sinc function.

Applying this to a two-dimensional Fourier transform, we find

$$f(\mathbf{x}) = \int_{\mathbb{R}^2} e^{2\pi i \mathbf{k} \cdot \mathbf{x}} \widehat{f}(\mathbf{k}) d\mathbf{k} \quad \text{definition of Fourier transform} \quad \text{(A.1.7.2)}$$

$$= \int_0^{\infty} \int_0^{2\pi} e^{2\pi i \kappa r \cos \theta} \widehat{f}(\kappa) \kappa d\theta d\kappa \quad \text{isotropy and cylindrical coordinates} \quad \text{(A.1.7.3)}$$

$$= 2\pi \int_0^{\infty} J_0(2\pi \kappa r) \widehat{f}(\kappa) \kappa d\kappa. \quad \text{applying Formula A.1.7.1} \quad \text{(A.1.7.4)}$$

We will use this later to translate results from structure functions to power spectra.

A.2 Structure functions

A.2.1 Introduction

The following section draws on [LaCasce \(2016\)](#) and [Frisch \(1995\)](#).

Structure functions can be thought of as an alternative to power spectra for learning about how a function varies with scale. They tell us how much variation there is between values separated by a particular spatial (or temporal) distance.

In particular, the n -th-order structure function is given by the formula

$$S_n(\mathbf{r}, \mathbf{x}) = \langle |\mathbf{u}(\mathbf{x} + \mathbf{r}) - \mathbf{u}(\mathbf{x})|^n \rangle \quad (\text{A.2.1.1})$$

where $\langle \cdot \rangle$ is an ensemble average. If we assume that we are dealing with the velocity field of a homogeneous and isotropic flow, then if $r = |\mathbf{r}|$, this simplifies to

$$S_n(r) = \langle |\mathbf{u}(\mathbf{x} + \mathbf{r}) - \mathbf{u}(\mathbf{x})|^n \rangle. \quad (\text{A.2.1.2})$$

A.2.2 Relation to power spectra

We will show that if $E(\kappa) \propto \kappa^{-\alpha}$ with $1 < \alpha < 3$, then

$$S_2(r) \propto r^{\alpha-1}. \quad (\text{A.2.2.1})$$

So, the second-order structure function can reproduce the power spectrum⁴ for a certain range of power laws.

This result applies in 1, 2 or 3 dimensions. We focus on two dimensions in the following.

Integral relation between $E(\kappa)$ and $S_2(r)$

Starting from the second-order structure function we have

⁴In the following we will work with the energy spectrum of some variable \mathbf{u} , rather than the kinetic energy power spectrum. We are only concerned with scalings, so this choice allows us to simplify the presentation by removing $\rho/2$ coefficients.

$$\begin{aligned}
\langle |\mathbf{u}(\mathbf{x} + \mathbf{r}) - \mathbf{u}(\mathbf{x})|^2 \rangle &= \langle \mathbf{u}(\mathbf{x} + \mathbf{r})^2 - 2\mathbf{u}(\mathbf{x})\mathbf{u}(\mathbf{x} + \mathbf{r}) + \mathbf{u}(\mathbf{x})^2 \rangle && \text{expand the square} && \text{(A.2.2.2)} \\
&= 2(\langle \mathbf{u}(\mathbf{x})^2 \rangle - \langle \mathbf{u}(\mathbf{x})\mathbf{u}(\mathbf{x} + \mathbf{r}) \rangle) && \text{homogeneity of statistics} && \text{(A.2.2.3)} \\
&= 2 \left(\int_{\mathbb{R}^2} E(\mathbf{k}) \, d\mathbf{k} - R(\mathbf{r}) \right) && \text{definition of energy spectrum } E(\mathbf{k}) && \text{(A.2.2.4)} \\
& && \text{and autocorrelation } R(\mathbf{r}) && \\
&= 2 \left(\int_{\mathbb{R}^2} E(\mathbf{k}) \, d\mathbf{k} - \int_{\mathbb{R}^2} e^{2\pi i \mathbf{k} \cdot \mathbf{r}} E(\mathbf{k}) \, d\mathbf{k} \right) && \text{Wiener-Khinchin theorem, see below} && \text{(A.2.2.5)} \\
&= 2 \left(\int_0^\infty E(\kappa) \, d\kappa - \int_0^\infty J_0(2\pi\kappa r) E(\kappa) \, d\kappa \right) && \text{isotropy, Hankel transform (A.1.7.4)} && \text{(A.2.2.6)} \\
& && \text{noting } E(\kappa) = 2\pi\kappa E(\mathbf{k}) && \\
&= \int_{-\infty}^\infty (1 - J_0(2\pi\kappa r)) E(\kappa) \, d\kappa. && \text{defining } E(\kappa) = E(-\kappa) && \text{(A.2.2.7)}
\end{aligned}$$

The Wiener-Khinchin theorem

In deriving the relation above we used the fact that the autocorrelation and power spectrum are Fourier pairs.

This result is called the Wiener-Khinchin theorem and it can be proved as follows.

$$\begin{aligned}
E(\mathbf{k}) &= |\mathbf{u}(\mathbf{k})|^2 && \text{definition of } E(\mathbf{k}) && \text{(A.2.2.8)} \\
&= \mathbf{u}(\mathbf{k})\mathbf{u}(\mathbf{k}) && \text{basic property of complex conjugate} && \text{(A.2.2.9)} \\
&= \int_{\mathbb{R}^n} e^{2\pi i \mathbf{k} \cdot \mathbf{x}} \mathbf{u}(\mathbf{x}) \, d\mathbf{x} \int_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{y}} \mathbf{u}(\mathbf{y}) \, d\mathbf{y} && \text{definition of Fourier transform} && \text{(A.2.2.10)} \\
&= \iint_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot (\mathbf{y} - \mathbf{x})} \mathbf{u}(\mathbf{x})\mathbf{u}(\mathbf{y}) \, d\mathbf{x} \, d\mathbf{y} && \text{rearrangement} && \text{(A.2.2.11)} \\
&= \iint_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{r}} \mathbf{u}\left(\mathbf{x}' + \frac{\mathbf{r}}{2}\right) \mathbf{u}\left(\mathbf{x}' - \frac{\mathbf{r}}{2}\right) \, d\mathbf{x}' \, d\mathbf{r} && \text{change of variables, } \mathbf{r} = \mathbf{y} - \mathbf{x}, \mathbf{x}' = (\mathbf{y} + \mathbf{x})/2 && \text{(A.2.2.12)} \\
&= \int_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{r}} \int_{\mathbb{R}^n} \mathbf{u}\left(\mathbf{x}' + \frac{\mathbf{r}}{2}\right) \mathbf{u}\left(\mathbf{x}' - \frac{\mathbf{r}}{2}\right) \, d\mathbf{x}' \, d\mathbf{r} && \text{linearity of integrals} && \text{(A.2.2.13)} \\
&= \int_{\mathbb{R}^n} e^{-2\pi i \mathbf{k} \cdot \mathbf{r}} R(\mathbf{r}) \, d\mathbf{r}. && \text{definition of autocorrelation} && \text{(A.2.2.14)}
\end{aligned}$$

A.2.3 Structure function blowup

We have the integral form of the structure function. Now we suppose the spectrum follows a power law, so $E(\kappa) = A\kappa^{-\alpha}$. Then we have

$$S_2(r) = A \int_{-\infty}^{\infty} (1 - J_0(2\pi\kappa r)) \kappa^{-\alpha} \mathbf{d}\kappa. \quad (\text{A.2.3.1})$$

Ultraviolet blowup

For large κ , the exponential term is highly oscillatory and we can neglect its contribution. The integrand is simply $\kappa^{-\alpha}$, so the integral's tail will diverge unless $\alpha > 1$.

Infrared blowup

For small κ , we can Taylor expand the exponential term. The Taylor expansion of our Bessel function is

$$J_0(x) := \sum_{l=0}^{\infty} \frac{(-1)^l}{2^{2l} (l!)^2} x^{2l} \quad (\text{A.2.3.2})$$

$$= 1 - \frac{1}{2^4} x^2 + O(x^4). \quad (\text{A.2.3.3})$$

When applying this in our integrand to $J_0(2\pi\kappa r)$, we see that the one cancels. Hence, our integrand to leading order goes like $\kappa^{-\alpha+2}$. This will diverge for small κ unless $-\alpha + 2 > -1$.

Hence, the second-order structure function diverges unless $1 < \alpha < 3$.

A.2.4 Structure function scaling

If we make the substitution $x = 2\pi\kappa r$. Our integral becomes

$$S_2(r) = A(2\pi r)^{\alpha-1} \int_{-\infty}^{\infty} (1 - J_0(x)) x^{-\alpha} \mathbf{d}x. \quad (\text{A.2.4.1})$$

Hence if $E(\kappa) \propto \kappa^{-\alpha}$ with $1 < \alpha < 3$, then $S_2(r) \propto r^{\alpha-1}$.

Appendix B

SQG and alpha models derivation

The SQG equations describe the motion of a rotating, stably stratified fluid whose motion is driven by surface buoyancy. It has been used to describe motion at scales of 10s of kilometres in the ocean, has attracted attention for studies of geostrophic turbulence due to it being a 2D model that accounts for 3D effects ([Held et al. \(1995\)](#)), and has mathematical interest given its similarities to the Navier-Stokes equations and the possibility that the SQG equations are not globally regular. See [Lapeyre \(2017\)](#) for a review of the applications of the SQG equations.

Our interest in these equations stems from their use in research on the atmospheric -5/3 and predictability ([Rotunno and Snyder, 2008](#); [Tulloch and Smith, 2009](#)).

In this appendix, the origin of these equations is briefly explained, and we show how they can fit into the more general alpha model framework by deriving the simple relationship between the stream function and active tracer in the model.

B.1 The surface quasi-geostrophic (SQG) equation derivation

The chapters 1, 2 and 5 of [Vallis \(2017\)](#) provide an introduction to quasi-geostrophic dynamics. I will describe the hierarchy of assumptions and simplifications that lead to the quasi-geostrophic equation (QG) and then explain how to reach SQG from QG.

The Boussinesq equations

We start with an f-plane approximation. This means our theory will apply better away from the equator. It is also assumed gravity is uniform and the motion is adiabatic - that is, there is no external heating.

With only these assumptions, the set of governing equations is

$$\frac{D\mathbf{u}}{Dt} + \mathbf{f} \wedge \mathbf{u} = -\frac{1}{\rho} \nabla p - g\mathbf{k}, \quad (\text{B.1.0.1})$$

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{u}, \quad (\text{B.1.0.2})$$

$$\frac{D\rho}{Dt} - \frac{1}{c_s^2} \frac{Dp}{Dt} = 0. \quad (\text{B.1.0.3})$$

Each symbol above has its usual meaning, \mathbf{u} is the 3D velocity vector and c_s is the speed of sound.

Now, the Boussinesq approximation is applied. It is assumed density remains close to some mean state with only small spatial and temporal fluctuations so that $\rho = \rho_0 + \rho'(x, y, z, t)$ with ρ' small compared to ρ_0 . This assumption of almost constant density works extremely well in the ocean where surface density is about 1025 kg/m^3 while density in the depths is on the order of 1050 kg/m^3 but less well in the atmosphere where density varies greatly in the vertical.

In order to simplify the conservation of mass equation, it is assumed that within the material derivative, the time derivative of the density scales the same way as the advective derivative. This removes quickly varying wave solutions such as sound and gravity waves. This makes the QG equations computationally attractive and is why the first successful numerical weather predictions were performed using the even simpler barotropic vorticity conservation equation by Jule Charney, Agnar Fjørtoff and John von Neumann in 1950 (Charney et al. (1990), Lynch (2006), Harper et al. (2007)).

To simplify the thermodynamic equation it is assumed that density is not altered by the changing mass of fluid above a parcel when motion has a vertical component.

These 3 assumptions result in a closed system called the simple Boussinesq equations:

$$\frac{D\mathbf{u}}{Dt} + \mathbf{f} \wedge \mathbf{u} = -\nabla_z \frac{p'}{\rho_0}, \quad (\text{B.1.0.4})$$

$$\frac{\partial p'}{\partial z \rho_0} = b, \quad (\text{B.1.0.5})$$

$$\nabla \cdot \mathbf{u} = 0, \quad (\text{B.1.0.6})$$

$$\frac{Db}{Dt} = 0. \quad (\text{B.1.0.7})$$

∇_z refers to the horizontal gradient and the buoyancy is $b = -\frac{g\rho'}{\rho_0}$.

The quasi-geostrophic equations

To reach the QG equations, the velocity is decomposed into a geostrophic and ageostrophic component. The geostrophic component is in geostrophic balance, meaning it is travelling orthogonally to the gradient in pressure at a velocity such that the Coriolis and pressure gradient forces are in balance. It is assumed that the geostrophic component is dominant: in terms of nondimensional numbers, this means the Rossby number, $Ro = \frac{U}{fL}$, is small. This yields the quasi-geostrophic equations,

$$\frac{D_g}{Dt} \left(f_0 \frac{\partial \psi}{\partial z} \right) = -N^2 w_a, \quad (\text{B.1.0.8})$$

$$\frac{D_g \Delta_z \psi}{Dt} = f_0 \frac{\partial w_a}{\partial z}. \quad (\text{B.1.0.9})$$

The subscripts a and g indicate velocities are ageostrophic or geostrophic, respectively, so that D_g/Dt refers to the material derivative with respect to the geostrophic component of the velocity field. ψ is the stream function for the horizontal geostrophic velocity. Δ_z is the horizontal Laplacian, note $\Delta_z \psi$ is the vorticity of the geostrophic flow in the particular x - y plane. N is the Brunt-Väisälä frequency $N^2 = \frac{\partial b_0}{\partial z}$.

The QG equations are the first-order correction to geostrophic balance. Equation (B.1.0.8) shows that changes in the stratification of the fluid occur if vertical velocity is accounted for. Equation (B.1.0.9) is a vortex stretching term where stretching only occurs along the vertical axis.

Eliminating the vertical velocity component between the two equations and assuming N is constant gives the conservation of the quasi-geostrophic potential vorticity (PV),

$$\frac{D_g q}{Dt} = 0, \quad (\text{B.1.0.10})$$

$$\text{where } q = \Delta_z \psi + \frac{\partial}{\partial z} \left(\frac{f_0^2}{N^2} \frac{\partial \psi}{\partial z} \right). \quad (\text{B.1.0.11})$$

The surface quasi-geostrophic equations

Now assume a semi-infinite domain with $z > 0$, no influx at the flat lower boundary and 0 interior PV. The assumption of constant PV in the interior of the domain can be justified since turbulent motion tends to mix and so homogenise fluid variables. z is re-scaled to $\hat{z} = \frac{N}{f_0} z$ such that $\frac{\partial}{\partial \hat{z}} = \frac{f_0}{N} \frac{\partial}{\partial z}$. The hat is dropped in the following. This simplified the PV to $q = \Delta \psi$ (note this is the 3D, not 2D Laplacian). The problem has thus reduced to solving Laplace's equation, $\Delta \psi = 0$, subject to some boundary conditions on the lower boundary.

From (B.1.0.8), given the no influx condition, the flow must satisfy $\frac{D_g}{Dt} \frac{\partial \psi}{\partial z} = 0$.

Therefore, the SQG equations which govern this situation are

$$\Delta \psi = 0, \quad (\text{B.1.0.12})$$

$$\frac{D_g \Theta}{Dt} = 0 \quad \text{at } z = 0, \quad (\text{B.1.0.13})$$

$$\text{where } \Theta \equiv \frac{\partial \psi}{\partial z} \quad (\text{B.1.0.14})$$

$$\text{and } \psi \rightarrow 0 \quad \text{as } z \rightarrow \infty. \quad (\text{B.1.0.15})$$

The vertical derivative of the geostrophic stream function is equivalent to the buoyancy, Θ .

B.2 Showing SQG is an alpha model

Motion in the SQG system is completely determined by surface buoyancy Θ . Surface buoyancy is an active tracer, it is conserved following the geostrophic component of the flow whilst also determining that flow. Away from the surface, Laplace's equation smooths the stream function with height. All the dynamics are contained in the relationship between the active tracer and the stream function. Fourier transforming (B.1.0.12) in x and y gives

$$k^2 \hat{\psi}(\mathbf{k}, z) = \frac{\partial^2 \hat{\psi}}{\partial z^2}. \quad (\text{B.2.0.1})$$

Where $\mathbf{k} = (k_x, k_y)$ denotes the wavenumbers in x and y . The magnitude $|\mathbf{k}|$ is referred to as the scalar wavenumber k . The hat is the Fourier transform. This ODE has solution

$$\hat{\psi}(\mathbf{k}, z) = A e^{kz} + B e^{-kz}. \quad (\text{B.2.0.2})$$

$\Theta = \frac{\partial \psi}{\partial z} |_{z=0}$ so since $\frac{\partial \hat{\psi}}{\partial z} = A k e^{kz} - B k e^{-kz}$ and the exponentially growing component must be discarded due to (B.1.0.15), $A = 0$ and $B = -\frac{\hat{\Theta}}{k}$. The solution to (B.2.0.1) is therefore

$$\hat{\psi}(\mathbf{k}, z) = -\frac{e^{-kz}}{k} \hat{\Theta}(\mathbf{k}). \quad (\text{B.2.0.3})$$

As remarked (B.2.0.3) shows the solution damps and smooths with height. At the surface

$$\hat{\psi}(\mathbf{k}) = -\frac{\hat{\Theta}(\mathbf{k})}{k}. \quad (\text{B.2.0.4})$$

Any constant times the buoyancy will still be materially conserved by the flow, so redefined the buoyancy by multiplying by a constant alters nothing. The important part of this relationship is the division by k .

Following this observation, the SQG equations can be generalised to the alpha turbulence family of equations ([Pierrehumbert et al. \(1994\)](#))

$$\frac{D\zeta}{Dt} = 0, \quad (\text{B.2.0.5})$$

$$\text{where } \hat{\zeta} \equiv k^\alpha \hat{\psi}. \quad (\text{B.2.0.6})$$

Where ζ is the generalised vorticity. If $\alpha = 2$, then this is the 2D vorticity equation; if $\alpha = 1$, it is the SQG equation. If $\alpha = 0$ then $\frac{\partial \zeta}{\partial t} = 0$ so the flow is stationary.

Bibliography

- Abramowitz, M. and Stegun, I. A. (1972). Handbook of mathematical functions with formulas, graphs, and mathematical tables. national bureau of standards applied mathematics series 55. tenth printing.
- Aluie, H. (2019). Convolutions on the sphere: commutation with differential operators. *GEM - International Journal on Geomathematics*, 10(1):9.
- Aluie, H., Hecht, M., and Vallis, G. K. (2018). Mapping the Energy Cascade in the North Atlantic Ocean: The Coarse-graining Approach. *Journal of Physical Oceanography*, 48:225–244.
- Arnold, H., Moroz, I., and Palmer, T. (2013). Stochastic parametrizations and model uncertainty in the Lorenz'96 system. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1991):20110479.
- Augier, P. and Lindborg, E. (2013). A new formulation of the spectral energy budget of the atmosphere, with application to two high-resolution general circulation models. *Journal of the Atmospheric Sciences*, 70(7):2293–2308.
- Avalos, Y. (2024). *The role of inertia-gravity waves in the atmospheric energy spectrum: insights from global storm-resolving simulations*. PhD thesis, PhD Thesis.
- Bartello, P. (1995). Geostrophic adjustment and inverse cascades in rotating stratified turbulence. *Journal of the atmospheric sciences*, 52(24).
- Bartello, P. and Tobias, S. (2013). Sensitivity of stratified turbulence to the buoyancy Reynolds number. *Journal of Fluid Mechanics*, 725:1–22.
- Batchelor, G. K. (1969). Computation of the energy spectrum in homogeneous two-dimensional turbulence. *The Physics of Fluids*, 12(12):11–233.
- Baumgart, M. and Riemer, M. (2019). Processes governing the amplification of ensemble spread in a medium-range forecast with large forecast uncertainty. *Quarterly Journal of the Royal Meteorological Society*, 145(724):3252–3270.
- Baumgart, M., Riemer, M., Wirth, V., Teubler, F., and Lang, S. T. (2018). Potential vorticity dynamics of forecast errors: A quantitative case study. *Monthly Weather Review*, 146(5):1405–1425.
- Bhushan Jha, V., Seshasayanan, K., and Dallas, V. (2023). Cascades transition in generalised two-dimensional turbulence. *arXiv e-prints*, pages arXiv–2312.
- Bierdel, L. B. (2017). *On the relevance of rotational and divergent modes of motion to mesoscale dynamics and upscale error growth*. PhD thesis, Imu.
- Bolgiani, P., Calvo-Sancho, C., Díaz-Fernández, J., Quiñán-Hernández, L., Sastre, M., Santos-Muñoz, D., Farrán, J., González-Alemán, J. J., Valero, F., and Martín, M. (2022). Wind kinetic energy climatology and effective resolution for the ERA5 reanalysis. *Climate Dynamics*, 59(3-4):737–752.
- Buizza, R., Tribbia, J., Molteni, F., and Palmer, T. (1993). Computation of optimal unstable structures for a numerical weather prediction model. *Tellus A*, 45(5):388–407.
- Burgess, B. H. (2014). The application of Kraichnan-Leith-Batchelor similarity theory to inverse cascades in generalized two-dimensional turbulence. *Ph. D. Thesis*.
- Burgess, B. H., Erler, A. R., and Shepherd, T. G. (2013). The troposphere-to-stratosphere transition in kinetic energy spectra and nonlinear spectral fluxes as seen in ECMWF analyses. *Journal of the Atmospheric Sciences*, 70(2):669–687.
- Buzzicotti, M., Storer, B. A., Khatri, H., Griffies, S. M., and Aluie, H. (2023). Spatio-temporal coarse-graining decomposition of the global ocean geostrophic kinetic energy. *Journal of Advances in Modeling Earth Systems*, 15(6).

- Callies, J., Bühler, O., and Ferrari, R. (2016). The dynamics of mesoscale winds in the upper troposphere and lower stratosphere. *Journal of the Atmospheric Sciences*, 73(12):4853–4872.
- Callies, J., Ferrari, R., and Bühler, O. (2014). Transition from geostrophic turbulence to inertia-gravity waves in the atmospheric energy spectrum. *Proceedings of the National Academy of Sciences*, 111(48):17033–17038.
- Capet, X., McWilliams, J. C., Molemaker, M. J., and Shchepetkin, A. F. (2008). Mesoscale to submesoscale transition in the california current system. part iii: Energy balance and flux. *Journal of Physical Oceanography*, 38(10):2256–2269.
- Charney, J. G. (1971). Geostrophic turbulence. *Journal of the Atmospheric Sciences*, 28(6):1087–1095.
- Charney, J. G., Fjörtoft, R., and Neumann, J. v. (1990). Numerical integration of the barotropic vorticity equation. In *The Atmosphere—A Challenge*, pages 267–284. Springer.
- Cho, J. Y. and Lindborg, E. (2001). Horizontal velocity structure functions in the upper troposphere and lower stratosphere: 1. observations. *Journal of Geophysical Research: Atmospheres*, 106(D10):10223–10232.
- Choi, D. S. and Showman, A. P. (2011). Power spectral analysis of jupiter's clouds and kinetic energy from cassini. *Icarus*, 216(2):597–609.
- Clark, T. and Peltier, W. (1984). Critical level reflection and the resonant growth of nonlinear mountain waves. *Journal of Atmospheric Sciences*, 41(21):3122–3134.
- Cox, M. R., Kafiabad, H. A., and Vanneste, J. (2023). Inertia-gravity-wave diffusion by geostrophic turbulence: the impact of flow time dependence. *Journal of Fluid Mechanics*, 958:A21.
- Dewan, E. M. (1979). Stratospheric wave spectra resembling turbulence. *Science*, 204(4395):832–835.
- Dhaka, S., Yamamoto, M., Shibagaki, Y., Hashiguchi, H., Fukao, S., and Chun, H. (2006). Equatorial atmosphere radar observations of short vertical wavelength gravity waves in the upper troposphere and lower stratosphere region induced by localized convection. *Geophysical Research Letters*, 33.
- Dueben, P. D., Wedi, N., Saarinen, S., and Zeman, C. (2020). Global simulations of the atmosphere at 1.45 km grid-spacing with the integrated forecasting system. *Journal of the Meteorological Society of Japan. Ser. II*, 98(3):551–572.
- Durrán, D. R. and Gingrich, M. (2014). Atmospheric predictability: Why butterflies are not of practical importance. *Journal of the Atmospheric Sciences*, 71(7):2476–2488.
- ECMWF (2020). Implementation of ifs cycle 47r1. <https://confluence.ecmwf.int/display/FCST/Implementation+of+IFS+Cycle+47r1>. Accessed: 2024-11-19.
- Eyink, G. L. (2006). Turbulent cascade of circulations. *Comptes Rendus Physique*, 7(3-4):449–455.
- Faranda, D., Lembo, V., Iyer, M., Kuzzay, D., Chibbaro, S., Daviaud, F., and Dubrulle, B. (2018). Computation and characterization of local subfilter-scale energy transfers in atmospheric flows. *Journal of the Atmospheric Sciences*, 75(7):2175–2186.
- Frehlich, R. and Sharman, R. (2010). Climatology of velocity and temperature turbulence statistics determined from rawinsonde and acars/amdar data. *Journal of Applied Meteorology and Climatology*, 49(6):1149–1169.
- Frisch, U. (1995). *Turbulence: the legacy of AN Kolmogorov*. Cambridge university press.
- Gage, K. (1979). Evidence far a $k^{-5/3}$ law inertial range in mesoscale two-dimensional turbulence. *Journal of Atmospheric Sciences*, 36(10):1950–1954.
- Harper, K., Uccellini, L. W., Kalnay, E., Carey, K., and Morone, L. (2007). 50th anniversary of operational numerical weather prediction. *Bulletin of the American Meteorological Society*, 88(5):639–650.
- Held, I. M., Pierrehumbert, R. T., Garner, S. T., and Swanson, K. L. (1995). Surface quasi-geostrophic dynamics. *Journal of Fluid Mechanics*, 282:1–20.

- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., et al. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.
- Hocke, K., Lainer, M., Moreira, L., Hagen, J., Fernandez Vidal, S., and Schranz, F. (2016). Atmospheric inertia-gravity waves retrieved from level-2 data of the satellite microwave limb sounder aura/mls. In *Annales geophysicae*, volume 34, pages 781–788. Copernicus Publications Göttingen, Germany.
- Hu, Y., Chen, L., Wang, Z., Pan, X., and Li, H. (2021). Towards a more realistic and detailed deep-learning-based radar echo extrapolation method. *Remote Sensing*, 14(1):24.
- Iwayama, T. and Watanabe, T. (2010). Green's function for a generalized two-dimensional fluid. *Physical Review E*, 82(3):036307.
- Juckes, M. (1994). Quasigeostrophic dynamics of the tropopause. *Journal of Atmospheric Sciences*, 51(19):2756–2768.
- Judt, F. (2018). Insights into atmospheric predictability through global convection-permitting model simulations. *Journal of the Atmospheric Sciences*, 75(5):1477–1497.
- Judt, F. (2020). Atmospheric predictability of the tropics, middle latitudes, and polar regions explored through global storm-resolving simulations. *Journal of the Atmospheric Sciences*, 77(1):257–276.
- Kafiabad, H. A., Savva, M. A., and Vanneste, J. (2019). Diffusion of inertia-gravity waves by geostrophic turbulence. *Journal of Fluid Mechanics*, 869.
- Kolmogorov, A. N. (1941). The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Cr Acad. Sci. URSS*, 30:301–305.
- Kouhen, S., Storer, B. A., Aluie, H., Marshall, D. P., and Christensen, H. M. (2024). Convective and orographic origins of the mesoscale kinetic energy spectrum. *Geophysical Research Letters*, 51(21):e2024GL110804.
- Kraichnan, R. H. (1967). Inertial ranges in two-dimensional turbulence. *The Physics of Fluids*, 10(7):1417–1423.
- LaCasce, J. H. (1996). *Baroclinic vortices over a sloping bottom*. PhD thesis, Massachusetts Institute of Technology.
- LaCasce, J. H. (2016). Calculating energy spectra from drifters.
- Lapeyre, G. (2017). Surface quasi-geostrophy. *Fluids*, 2(1):7.
- Lavers, D. A., Simmons, A., Vamborg, F., and Rodwell, M. J. (2022). An evaluation of era5 precipitation for climate monitoring. *Quarterly Journal of the Royal Meteorological Society*, 148(748):3152–3165.
- Leith, C. and Kraichnan, R. (1972). Predictability of turbulent flows. *Journal of Atmospheric Sciences*, 29(6):1041–1058.
- Leith, C. E. (1968). Diffusion approximation for two-dimensional turbulence. *The Physics of Fluids*, 11(3):671–672.
- Lelong, M.-P. and Riley, J. J. (1991). Internal wave—vortical mode interactions in strongly stratified flows. *Journal of Fluid Mechanics*, 232:1–19.
- Leung, T. Y., Leutbecher, M., Reich, S., and Shepherd, T. G. (2020). Impact of the mesoscale range on error growth and the limits to atmospheric predictability. *Journal of the Atmospheric Sciences*, 77(11):3769–3779.
- Li, Z., Wei, J., Bao, X., and Sun, Y. Q. (2023). Intercomparison of tropospheric and stratospheric mesoscale kinetic energy resolved by the high-resolution global reanalysis datasets. *Quarterly Journal of the Royal Meteorological Society*, 149(757):3738–3764.

- Lilly, D. (1973). Lectures in sub-synoptic scales of motion and two-dimensional turbulence. In *Dynamic Meteorology: Lectures Delivered at the Summer School of Space Physics of the Centre National D'Etudes Spatiales, Held at Lannion, France, 7 August–12 September 1970*, pages 353–418. Springer.
- Lilly, D. K. (1983). Stratified turbulence and the mesoscale variability of the atmosphere. *Journal of Atmospheric Sciences*, 40(3):749–761.
- Lilly, D. K. (1989). Two-dimensional turbulence generated by energy sources at two scales. *Journal of Atmospheric Sciences*, 46(13):2026–2030.
- Lindborg, E. (1999). Can the atmospheric kinetic energy spectrum be explained by two-dimensional turbulence? *Journal of Fluid Mechanics*, 388:259–288.
- Lindborg, E. (2006). The energy cascade in a strongly stratified fluid. *Journal of Fluid Mechanics*, 550:207–242.
- Liu, H.-B. and Lee, I. (2020). Mpl-gan: Toward realistic meteorological predictive learning using conditional gan. *IEEE Access*, 8:93179–93186.
- Lorenz, E. N. (1969). The predictability of a flow which possesses many scales of motion. *Tellus*, 21(3):289–307.
- Lovejoy, S., Tuck, A., Schertzer, D., and Hovde, S. (2009). Reinterpreting aircraft measurements in anisotropic scaling turbulence. *Atmospheric Chemistry and Physics*, 9(14):5007–5025.
- Lynch, P. (2006). *The emergence of numerical weather prediction: Richardson's dream*. Cambridge University Press.
- Malardel, S. and Wedi, N. P. (2016). How does subgrid-scale parametrization influence nonlinear spectral energy fluxes in global nwp models? *Journal of Geophysical Research: Atmospheres*, 121(10):5395–5410.
- Maltrud, M. and Vallis, G. (1991). Energy spectra and coherent structures in forced two-dimensional and beta-plane turbulence. *Journal of Fluid Mechanics*, 228:321–342.
- Métais, O. and Lesieur, M. (1986). Statistical predictability of decaying turbulence. *Journal of the atmospheric sciences*, 43(9):857–870.
- Morgan, M. C. and Nielsen-Gammon, J. W. (1998). Using tropopause maps to diagnose midlatitude weather systems. *Monthly weather review*, 126(10):2555–2579.
- Nastrom, G., Fritts, D., and Gage, K. (1987). An investigation of terrain effects on the mesoscale spectrum of atmospheric motions. *Journal of Atmospheric Sciences*, 44(20):3087–3096.
- Nastrom, G., Gage, K., and Jasperson, W. (1984). Kinetic energy spectrum of large-and mesoscale atmospheric processes. *Nature*, 310(5972):36–38.
- O’Gorman, P. A. and Schneider, T. (2007). Recovery of atmospheric flow statistics in a general circulation model without nonlinear eddy-eddy interactions. *Geophysical Research Letters*, 34(22).
- Palmer, T., Döring, A., and Seregin, G. (2014). The real butterfly effect. *Nonlinearity*, 27(9):R123.
- Palmer, T., Molteni, F., Mureau, R., Buizza, R., Chapelet, P., and Tribbia, J. (1993). Ensemble prediction. In *Proc. ECMWF Seminar on Validation of models over Europe*, volume 1, pages 21–66.
- Pierrehumbert, R. T., Held, I. M., and Swanson, K. L. (1994). Spectra of local and nonlocal two-dimensional turbulence. *Chaos, Solitons & Fractals*, 4(6):1111–1116.
- Preusse, P., Ern, M., Bechtold, P., Eckermann, S., Kalisch, S., Trinh, Q., and Riese, M. (2014). Characteristics of gravity waves resolved by ecmwf. *Atmospheric Chemistry and Physics*, 14(19):10483–10508.
- Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Kashem, S., Madge, S., et al. (2021). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878):672–677.
- Rhines, P. B. (1975). Waves and turbulence on a beta-plane. *Journal of Fluid Mechanics*, 69(3):417–443.

- Rotunno, R. and Snyder, C. (2008). A generalization of Lorenz's model for the predictability of flows with many scales of motion. *Journal of the Atmospheric Sciences*, 65(3):1063–1076.
- Rotunno, R., Snyder, C., and Judt, F. (2023). Upscale versus “up-amplitude” growth of forecast-error spectra. *Journal of the Atmospheric Sciences*, 80(1):63–72.
- Sadek, M. and Aluie, H. (2018). Extracting the spectrum of a flow by spatial filtering. *Physical Review Fluids*, 3(12):124610.
- Selz, T., Bierdel, L., and Craig, G. C. (2019). Estimation of the variability of mesoscale energy spectra with three years of cosmo-de analyses. *Journal of the Atmospheric Sciences*, 76(2):627–637.
- Selz, T. and Craig, G. C. (2015). Upscale error growth in a high-resolution simulation of a summertime weather event over Europe. *Monthly Weather Review*, 143(3):813–827.
- Selz, T., Riemer, M., and Craig, G. C. (2022). The transition from practical to intrinsic predictability of midlatitude weather. *Journal of the Atmospheric Sciences*, 79(8):2013–2030.
- Shutts, G. (2005). A kinetic energy backscatter algorithm for use in ensemble prediction systems. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 131(612):3079–3102.
- Skamarock, W. C. (2004). Evaluating mesoscale NWP models using kinetic energy spectra. *Monthly Weather Review*, 132(12):3019–3032.
- Skamarock, W. C. and Klemp, J. B. (2008). A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *Journal of Computational Physics*, 227(7):3465–3485.
- Skamarock, W. C., Park, S.-H., Klemp, J. B., and Snyder, C. (2014). Atmospheric kinetic energy spectra from global high-resolution nonhydrostatic simulations. *Journal of the Atmospheric Sciences*, 71(11):4369–4381.
- Stephan, C. C., Duras, J., Harris, L., Klocke, D., Putman, W. M., Taylor, M., Wedi, N., Žagar, N., and Ziemann, F. (2022). Atmospheric energy spectra in global kilometre-scale models. *Tellus A: Dynamic Meteorology and Oceanography*, 74(1).
- Storer, B. A. and Aluie, H. (2023). Flowsieve: A coarse-graining utility for geophysical flows on the sphere. *Journal of Open Source Software*, 8(84):4277.
- Storer, B. A., Buzzicotti, M., Khatri, H., Griffies, S. M., and Aluie, H. (2022). Global energy spectrum of the general oceanic circulation. *Nature Communications*, 13(1):5314.
- Storer, B. A., Buzzicotti, M., Khatri, H., Griffies, S. M., and Aluie, H. (2023). Global cascade of kinetic energy in the ocean and the atmospheric imprint. *Science Advances*, 9(51).
- Sun, Y. Q., Rotunno, R., and Zhang, F. (2017). Contributions of moist convection and internal gravity waves to building the atmospheric-5/3 kinetic energy spectra. *Journal of the Atmospheric Sciences*, 74(1):185–201.
- Sun, Y. Q. and Zhang, F. (2020). A new theoretical framework for understanding multiscale atmospheric predictability. *Journal of the Atmospheric Sciences*, 77(7):2297–2309.
- Toth, Z. and Kalnay, E. (1993). Ensemble forecasting at NMC: The generation of perturbations. *Bulletin of the American Meteorological Society*, 74(12):2317–2330.
- Tran, C. V. (2004). Nonlinear transfer and spectral distribution of energy in α turbulence. *Physica D: Nonlinear Phenomena*, 191(1-2):137–155.
- Tulloch, R. and Smith, K. (2006). A theory for the atmospheric energy spectrum: Depth-limited temperature anomalies at the tropopause. *Proceedings of the National Academy of Sciences*, 103(40):14690–14694.
- Tulloch, R. and Smith, K. S. (2009). Quasigeostrophic turbulence with explicit surface dynamics: Application to the atmospheric energy spectrum. *Journal of the Atmospheric Sciences*, 66(2):450–467.
- Vallis, G. K. (2017). *Atmospheric and oceanic fluid dynamics*. Cambridge University Press.

- VanZandt, T. (1982). A universal spectrum of buoyancy waves in the atmosphere. *Geophysical Research Letters*, 9(5):575–578.
- Vonich, P. T. and Hakim, G. J. (2018). Hurricane kinetic energy spectra from in situ aircraft observations. *Journal of the Atmospheric Sciences*, 75(8):2523–2532.
- Waite, M. L. (2020). Untangling waves and vortices in the atmospheric kinetic energy spectra. *Journal of Fluid Mechanics*, 888:F1.
- Wang, J.-W. and Sardeshmukh, P. D. (2021). Inconsistent Global Kinetic Energy Spectra in Reanalyses and Models. *Journal of the Atmospheric Sciences*, 78(8):2589–2603.
- Wedi, N., Benard, P., Yessad, K., Untch, A., Malardel, S., Hamrud, M., Mozdzyński, G., Fisher, M., and Smolarkiewicz, P. (2010). Non-hydrostatic modeling with ifs: current status. In *ECMWF NON-HYDROSTATIC WORKSHOP*.
- Xie, P., Li, X., Ji, X., Chen, X., Chen, Y., Liu, J., and Ye, Y. (2020). An energy-based generative adversarial forecaster for radar echo map extrapolation. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5.
- Yessad, K. (2019). Digital filtering initialisation in the cycle 46t1r1 of arpege/ifs.
- Zhang, F., Bei, N., Rotunno, R., Snyder, C., and Epifanio, C. C. (2007). Mesoscale predictability of moist baroclinic waves: Convection-permitting experiments and multistage error growth dynamics. *Journal of the Atmospheric Sciences*, 64(10):3579–3594.