Computers & Geosciences 93 (2016) 36–44

Contents lists available at ScienceDirect

Computers & Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Feature tracking in high-resolution regional climate data

Neil R. Massey

Environmental Change Institute, University of Oxford, Oxford University Centre for the Environment, South Parks Road, Oxford OX1 3QY, UK

ARTICLE INFO

Article history:
Received 16 September 2015
Received in revised form
15 April 2016
Accepted 28 April 2016
Available online 4 May 2016

Keywords:
Storm tracking
Hierarchical mesh
Regional climate modelling
Sparse grid

ABSTRACT

In this paper, a suite of algorithms are presented which facilitate the identification and tracking of storm-indicative features, such as mean sea-level pressure minima, in high resolution regional climate data. The methods employ a hierarchical triangular mesh, which is tailored to the regional climate data by only subdividing triangles, from an initial icosahedron, within the domain of the data. The regional data is then regridded to this triangular mesh at each level of the grid, producing a compact representation of the data at numerous resolutions. Storm indicative features are detected by first subtracting the background field, represented by a low resolution version of the data, which occurs at a lower level in the mesh. Anomalies from this background field are detected, as feature objects, at a mesh level which corresponds to the spatial scale of the feature being detected and then refined to the highest mesh level. These feature objects are expanded to an outer contour and overlapping objects are merged. The centre points of these objects are tracked across timesteps by applying an optimisation scheme which uses five hierarchical rules. Objects are added to tracks based on the highest rule in the scheme they pass and, if two objects pass the same rule, the cost of adding the object to the track. An object exchange scheme ensures that adding an object to a track is locally optimal. An additional track optimisation phase is performed which exchanges segments between tracks and merges tracks to obtain a globally optimal track set. To validate the suite of algorithms they are applied to the ERA-Interim reanalysis dataset and compared to other storm-indicative feature tracking algorithms.

1. Introduction

The automated tracking of extra tropical cyclones, by identifying and tracing the evolution of storm-indicative features in gridded meteorological data is an active and challenging research topic, with over 15 research groups having developed their own algorithms (Neu et al., 2013). Although these algorithms use different methods to identify and track the features, they all split the problem into two distinct phases. The first is the identification of the storm-indicative features in the gridded data. The second phase is the tracking of these features across timesteps, associating features in one timestep with features in subsequent timesteps so as to track the feature as it moves through time and across a domain, which is known as the correspondence problem (Post et al., 2003). The analysis of the resulting feature tracks, which include acception or rejection of some tracks based on inclusion criteria may be performed before calculating the track statistics.

A common problem facing all tracking algorithms stems from the representation of data on a regular latitude–longitude grid. This causes grid boxes to become smaller toward the poles, in effect increasing their resolution and decreasing the spatial scale that features occur at. In addition, the data exhibits a singularity at the poles, with the first and last row of data having the same value, which makes searching for localised features difficult. Tracking algorithms overcome these problems in a number of ways. To correct the spatial discrepancies, the data may be smoothed by a Cressman filter (Murray and Simmonds, 1991; Sinclair, 1994), transformed to spectral space and truncated at a wavenumber (Hodges, 1994; Benestad and Chen, 2006), convolved with a filter (Hewson and Titley, 2010; Rudeva and Gulev, 2007) or regridded (Serreze, 1995; Wang et al., 2006). In order to account for the pole singularities, the data may be reprojected to a stereographic polar projection (Murray and Simmonds, 1991; Hoskins and Hodges, 2002, 2005), at the expense of only being able to track storms in one hemisphere at a time.

Massey (2012) shows that both of these problems can be overcome by regridding the latitude–longitude data to a hierarchical triangular mesh obtained by repeatedly sub-dividing an icosahedron contained within the unit sphere. Constructing the icosahedron so that a centroid of a triangle corresponds to the locations of the poles ensures that there are no singularities.

This paper extends and improves Massey (2012) to enable the regridding of regional climate data to a sparse version of the hierarchical triangular mesh. Improvements to both the identification of storm-indicative features and their tracking are

---

E-mail address: neil.massey@ouce.ox.ac.uk

http://dx.doi.org/10.1016/j.cageo.2016.04.015

presented and incorporated into a new suite of regridding, feature identification and tracking software. To validate the methods, the software is applied to reanalysis data compared to other tracking algorithms, and is also applied to output from a high-resolution (50 km) regional climate model.

2. Regional hierarchical mesh generation

2.1. Generating the mesh

Generation of the mesh commences with the construction of an icosahedron bound by the unit sphere and rotated so the poles lie at the centroid of a triangle. The coordinates of this icosahedron are provided in Table 1 of Massey (2012). The vertices of the triangles are stored in a point cloud and the indices of the point cloud for the vertices of the triangles are stored in an array of 20 quad trees, as in Massey (2012). In this revised mesh generation scheme, the user is required to supply an example of the data they wish to track features within, in the form of a netCDF file with latitude and longitude dimensions. This can be defined for regions including hemispheres or the entire globe, either on a regular or rotated pole grid.

For the purpose of the mesh generation and, later, the regridding of data to this mesh, grid points in the latitude–longitude grid are taken to be at the centre of the grid box, and the grid box is assumed to have the same value within the grid box. No interpolation to the triangular mesh is attempted and a triangle value is taken to be at the centroid of the triangle and throughout the triangle.

The grid of latitude–longitude coordinates are converted to 3D Cartesian coordinates and then assigned to one of the 20 triangles in the icosahedron by way of a point inclusion test (Eq. (5) in Massey (2012)), or a nearest point test if the point inclusion test fails. Each of these 20 triangles is split into 4 sub triangles if it contains a point, by projecting the midpoints of each triangle side onto the unit sphere and adding to the point cloud. The four new triangles are added to the quad tree as children of the initial triangle. Additionally, the point inclusion test is applied to the latitude–longitude coordinates that are within the original triangle to determine which child triangle they should be assigned to. This splitting process continues iteratively until all triangles contain one point or a maximum mesh level is reached. Fig. 1 shows this iterative process of subdividing the mesh on a Dymaxion projection (Gray, 1994, 1995) and Fig. 2 contains a pseudo-code algorithm of the splitting process. Each subdivision of a triangle is referred to as a level in the mesh, the 20 base triangles are at level 0, a triangle that has been split once is level 1, twice level 2, etc.

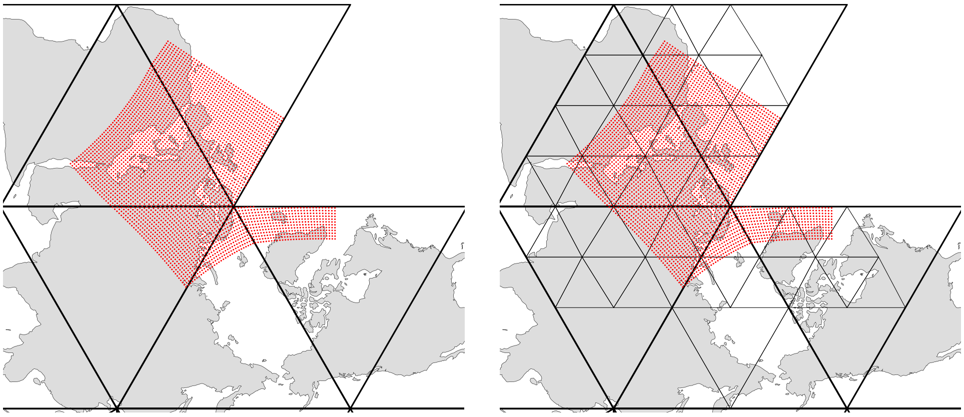

Maximum mesh level: 0 2

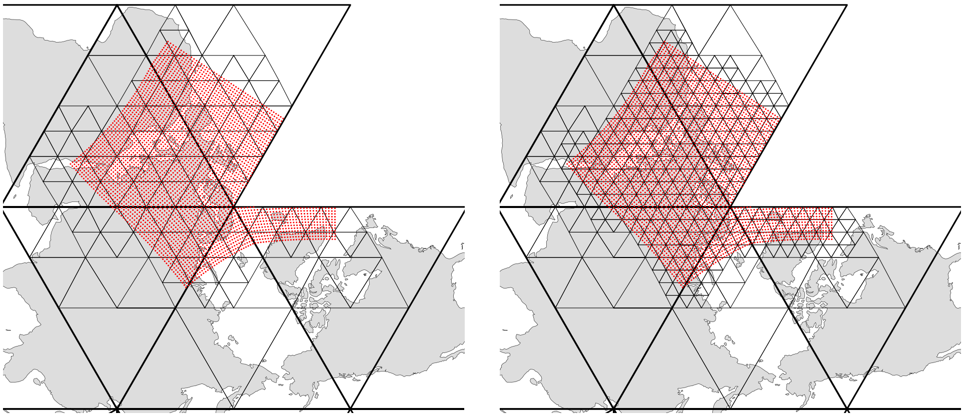

Maximum mesh level: 3 4

Fig. 1. The repeated subdivision of the triangular mesh based on the grid from the HadRM3P regional climate model over Europe. Red dots show the latitude–longitude points of the HadRM3P grid and each panel shows the triangles at the maximum mesh level after the subdivision. Only triangles with at least one latitude–longitude point are subdivided per iteration. The triangular mesh is shown on a Dymaxion projection (Gray, 1994, 1995), which represents the world map on an icosahedron and so is ideal for displaying the hierarchical triangular mesh. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)