

An Integrated Geometric and Topological Approach for the Identification and Visual Analysis of Rossby Wave Packets

KARRAN PANDEY

Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India

JOY MERWIN MONTEIRO

Department of Earth and Climate Science, Indian Institute of Science Education and Research, Pune, India

VIJAY NATARAJAN

Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India

(Manuscript received 16 January 2020, in final form 20 May 2020)

ABSTRACT

A new method for identifying Rossby wave packets (RWPs) using 6-hourly data from the ERA-Interim is presented. The method operates entirely in the spatial domain and relies on the geometric and topological properties of the meridional wind field to identify RWPs. The method represents RWPs as nodes and edges of a dual graph instead of the more common envelope representation. This novel representation allows access to both RWP phase and amplitude information. Local maxima and minima of the meridional wind field are collected into groups. Each group, called a v -max cluster or v -min cluster of the meridional wind field, represents a potential wave component. Nodes of the dual graph represent a v -max cluster or v -min cluster. Alternating v -max clusters and v -min clusters are linked by edges of the dual graph, called the RWP association graph. Amplitude and discrete gradient-based filtering applied on the association graph helps identify RWPs of interest. The method is inherently robust against noise and does not require smoothing of the input data. The main parameters that control the performance of the method and their impact on the identified RWPs are discussed. All filtering and RWP identification operations are performed on the association graph as opposed to directly on the wind field, leading to computational efficiency. Advantages and limitations of the method are discussed and are compared against (transform-based) envelope methods in a series of experiments.

1. Introduction

Rossby wave packets (RWPs) are localized contiguous regions of significant meridional flow with alternating signs that have a maximum near the tropopause (Wirth et al. 2018). RWPs have a group velocity that is larger than the phase velocity of an individual wave component. The faster propagation of energy generates new wave components at the leading (eastern) edge of the wave packet, resulting in the phenomenon of “downstream development” (Simmons and Hoskins 1979;

Chang and Orlanski 1993; Hakim 2003). Eddy variance generated in localized baroclinically active regions (predominantly over the oceans) is transported over long distances in the form of RWPs, affecting weather and climate at planetary scales. Observational analyses show that RWPs preferentially propagate in the zonal direction. This preferential zonal propagation is attributed to the focusing of these RWPs by baroclinic waveguides, whose location correlates strongly with the seasonally varying location of the subtropical and polar jets (Wallace et al. 1988; Chang and Yu 1999; Chang 1999; Martius et al. 2010).

There has been an increasing interest in the dynamics of RWPs due to their role in the variability of midlatitude weather by the chaotic mixing of air in regions adjacent to the baroclinic waveguides (see, e.g., Swanson and Pierrehumbert 1997; Schneider et al. 2015).

Supplemental information related to this paper is available at the Journals Online website: <https://doi.org/10.1175/MWR-D-20-0014.s1>.

Corresponding author: Karran Pandey, karran13@gmail.com

DOI: 10.1175/MWR-D-20-0014.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

The anomalous advection of water vapor, potential vorticity, and temperature by winds associated with the RWPs can result in a dynamic and thermodynamic environment favorable to extreme weather events (Schubert et al. 2011; Parker et al. 2014; Dimri et al. 2015; Ratnam et al. 2016; Hunt et al. 2018; Fragkoulidis et al. 2018; Monteiro and Caballero 2019). Therefore, understanding the dynamics of RWPs and their predictability is essential to predict these extremes and how they might evolve in a changing climate.

a. Related work

A first step in such an analysis of RWPs is to identify and track them in gridded data. Due to the difficulty in providing a precise algorithmic description of RWPs, their identification and tracking has proved to be a challenging and interesting aspect of RWP research. Initial attempts used time–longitude maps that averaged the geopotential field over a certain latitude band (Hovmöller 1949) or used one-point correlation maps derived from the geopotential field at different levels (Wallace et al. 1988). While these approaches are ideal for case studies, attempts were made to find a more “objective” method to aid in automated identification of RWPs in large datasets. Toward this end, Fourier-transform-based methods such as complex demodulation (Lee and Held 1993), Hilbert transforms (Zimin et al. 2003, 2006) and filtered local finite-amplitude wave activity (LWA) (Ghinassi et al. 2018) have been developed. These algorithms use the meridional wind field as input and provide as output the envelope of the RWP, which is referred to as an RWP object. Combining this identification step with a tracking algorithm (Souders et al. 2014b) allows for an automated way to extract information about RWPs in large datasets.

The above algorithms provide a concise description of RWPs from noisy data (usually the upper-tropospheric meridional winds) at the expense of losing phase information. However, the phase information is often important for characterizing the local synoptic situation at a location. For example, we have observed that extreme wet-bulb temperature events in south-west Pakistan are associated with northerly winds at 300 hPa to the north of the Indus valley (Monteiro and Caballero 2019). An algorithm that is capable of objectively identifying RWPs while still providing access to the phase information is essential to automate the identification and prediction of such extreme events. Furthermore, there is evidence that transform-based methods that analyze the meridional wind struggle to capture the full nonlinear evolution of RWPs, and that the LWA field might be a better way to track RWPs through their entire life cycle (Ghinassi et al. 2018). However, the LWA

field is not readily available from either reanalysis data or model output.

b. Contributions

In this paper, we describe a method to identify RWPs from the meridional wind field v in the spatial domain without the use of Fourier transforms. We instead use the topology of the v field to obtain a concise description of RWPs in the form of a graph. A node of this graph represents an individual wave component, called a v -max cluster or v -min cluster. A v -max cluster (v -min cluster) is a cluster of maxima (minima) of the wind field. An edge of the graph represents the spatial adjacency between a v -max cluster and v -min cluster. The use of topological features (defined by local maxima and minima of v) helps avoid making assumptions about the wavelike behavior of RWPs.

The simplicity of the graph representation enables fast recomputation of RWPs for different values of parameter thresholds. Furthermore, the graph representation allows a global description of the RWP field (entire graph) without loss of phase information (individual nodes). Finally, tracking methods that work with graph-based representations (e.g., Valsangkar et al. 2019) could be used to follow the evolution of RWPs over time.

Table 1 lists various features and characteristics of the proposed method and compares the proposed method's relative strengths and weaknesses against previous approaches.

2. RWP identification

In this section, we describe our proposed approach to RWP identification, representation, and interactive visual analysis. The method operates on the meridional wind field defined on a 2D grid with a fixed resolution. For the current study, we use the 300 hPa meridional winds from ERA-Interim, which has a spatial resolution of 0.75° in latitude and longitude space and a temporal resolution of 6 h (Dee et al. 2011). The method does not make any assumptions on the resolution or the projection of the input data. Further, the method provides inherent support for controlled noise removal and does not require prior smoothing of the meridional wind field v in space or time (see, e.g., Souders et al. 2014b).

a. Method overview

We utilize the geometric and topological properties of the meridional wind field v to extract RWPs. We define an RWP as alternating clusters of maxima and minima of the meridional wind field. These high-intensity clusters of maxima and minima in the meridional wind field

TABLE 1. A summary of the features of the envelope method (Zimin et al. 2003, 2006), the LWA method (Ghinassi et al. 2018), and the proposed method. The terms u , v , and T in the table have their usual meaning.

Feature	Envelope method	LWA	Geometric + topological method (proposed)
Detection algorithm	Hilbert transform (along latitude or streamlines)	Filtered LWA	Identification and clustering of maxima and minima of the v field
Input data	v on one pressure level	u , v , T on one isentrope	v on one pressure level
Smoothing of input data	Yes	Yes (LWA is filtered)	No
Output data	Envelope field	Filtered LWA field	A graph representing the RWP and its phase information
Noise sensitivity	Yes, so the method requires filtering of input in a preprocessing step	Yes, so the method requires filtering of input in a preprocessing step	Yes, so the method filters low-amplitude extrema before clustering but does not filter input data
RWP life cycle	Works best in the linear and wavelike stages of the RWP life cycle	Works in all stages of RWP life cycle, including the finite-amplitude stage of the life cycle	Works in all stages of RWP life cycle, including the finite-amplitude stage of the life cycle
Dependence of output size (M) on input size (N)	$M = N$; envelope field is sampled on same grid as input v	$M = N$; filtered LWA field is sampled on same grid as input v	$M \approx \text{constant}$, $M \ll N$; the nodes in the output graph corresponds to the RWP phase; number of nodes is much smaller than the input v
Phase information	Not available; further nontrivial processing may be required to extract phase information	Not available (filtered out from LWA); further nontrivial processing may be required to extract phase information	Directly available in output graph
Ease of computation	Simple	Complicated (even when isentropic data are available)	Moderate
Separation of field into different RWPs	Simple (thresholding)	Simple (thresholding)	Involved (calculating importance of edges by estimating curvature vorticity)
Visualization of RWP field	Direct plotting of envelope field	Direct plotting of filtered LWA	Calculation of graph representation of RWP from association graph

are henceforth referred to as wave components. Figure 1 presents an overview of the different steps toward RWP computation. Figure 2 provides a visual representation of various terms and illustrates the algorithm by showing the output of different steps. The method first computes local maxima and minima of v and uses them to identify wave components of RWPs. The collection of local maxima and minima is partitioned into clusters. Each cluster represents a potential wave component and is called a v -max cluster or v -min cluster, respectively. Next, the method searches for coherent associations between the identified v -max clusters and v -min clusters. Spatial adjacencies between the clusters are stored as edges in an association graph. Nodes of this graph represent the individual v -max clusters and v -min clusters. We associate a cost or weight with each edge of the graph that depends on the value of v at the nodes that the edge connects and the distance between these nodes. The specific formulation of the edge weights is presented later in this section. These weights are used to identify and prune irrelevant edges, thereby reducing the graph to a collection of paths. A path is an ordered sequence of alternating nodes and edges, where each edge connects its predecessor and successor node. Finally, each connected component in this collection of

paths is processed to extract representative paths to display the identified RWPs.

We now describe the four steps of the RWP identification pipeline. First, we introduce the data structures used to store and efficiently access the input and intermediate objects computed by the algorithm.

b. Data structure and representation

The v field is stored as a fixed resolution grid, samples are available at each grid point and we assume linear interpolation along each axis in all computations. Maxima and minima of the v field are stored as a list of n -tuples, each containing the point coordinates, value of the v field, and cluster ID tag for the point. The association graph constructed in the process of computing the RWPs is stored as a list of line segments (edges) with associated weights. We use the Python library NetworkX (Hagberg et al. 2008) for storing and processing the graph. The spatial region associated with a v -max cluster and v -min cluster is stored as a list of vertices and edges that bounds the region.

c. Extract critical points

Critical points of a scalar field together with the associated gradient field can be used to infer important

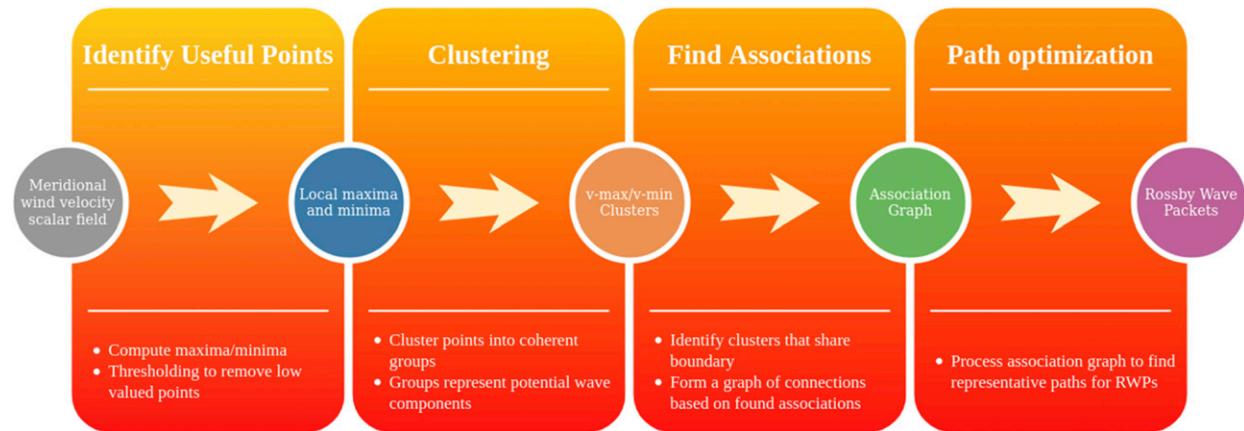


FIG. 1. RWP identification pipeline. The meridional wind field is processed to identify v -max clusters and v -min clusters. Spatial adjacencies between the v -max clusters and v -min clusters is captured by constructing an association graph. In the final step, this graph is processed to identify RWPs and to compute a representative path for each RWP.

structural information. Topological analysis based on critical points and their interrelationships have been successfully used for feature identification, analysis, and tracking (Heine et al. 2016), specifically for extratropical cyclone identification and tracking (Valsangkar et al. 2019), visualization of cloud system movement (Doraiswamy et al. 2013), and tracking pressure perturbations (Widanagamaachchi et al. 2017).

An important characteristic of these methods is that they do not require numerical computations of the gradient, instead the critical points are computed based on combinatorial characterizations. Here, we are interested in capturing v -max cluster and v -min cluster like behavior in the meridional wind velocity field. High-valued local maxima (minima) serve as starting points for capturing v -max clusters (v -min clusters).

A local maximum is located at a grid point whose scalar value is higher than neighboring grid points. Similarly, a local minimum is located at a grid point whose scalar value is lower than neighboring grid points. A large number of low-amplitude, structurally irrelevant local maxima and minima are reported by a method that is directly based on this definition. Since the focus is on identifying relevant starting points for significant v -max clusters and v -min clusters, we remove all local minima and maxima with a value of v smaller than 5 m s^{-1} .

Figure 2a shows the distribution of local maxima and minima within a small region in the Northern Hemisphere at 0600 UTC 6 January 2007. We can see how they can be potentially used as markers for the different v -max clusters and v -min clusters in the scalar field. While at least one local maximum is guaranteed per v -max cluster (one local minimum per v -min cluster), we observe a

total of 18 local maxima and 21 local minima spread over 5 v -max clusters and 4 v -min clusters, with at least 4 local maxima (minima) belonging to the same v -max cluster (v -min cluster).

d. Compute v -max clusters and v -min clusters

The collection of local maxima and minima identified in the previous step is partitioned into clusters. A v -max cluster consists of a collection of local maxima that are not separated by a region of negative v values. Additionally, we require two local maxima that are spatially distant to belong to distinct v -max clusters. Similarly, a v -min cluster consists of a collection of local minima that are not separated by a region of positive v values. The aim is to obtain a partition where each v -max cluster and v -min cluster contains a collection of proximal maxima and minima, respectively. This partition is computed by clustering the local maxima/minima using a suitable measure of similarity that is described later in this section.

We need a flexible clustering algorithm with a geometry- and topology-aware similarity measure to compute such a partition. We choose the affinity propagation clustering algorithm (Frey and Dueck 2007) because of its ability to automatically select the number of clusters and its support for different similarity measures.

The affinity propagation algorithm takes as input the similarity measures between points. It treats all input points as part of a network, exchanging real valued messages between the points iteratively. It aims to identify one exemplar per cluster. The similarity between a pair of points is used to compute how well one point represents the other. The exemplar is a point that is the

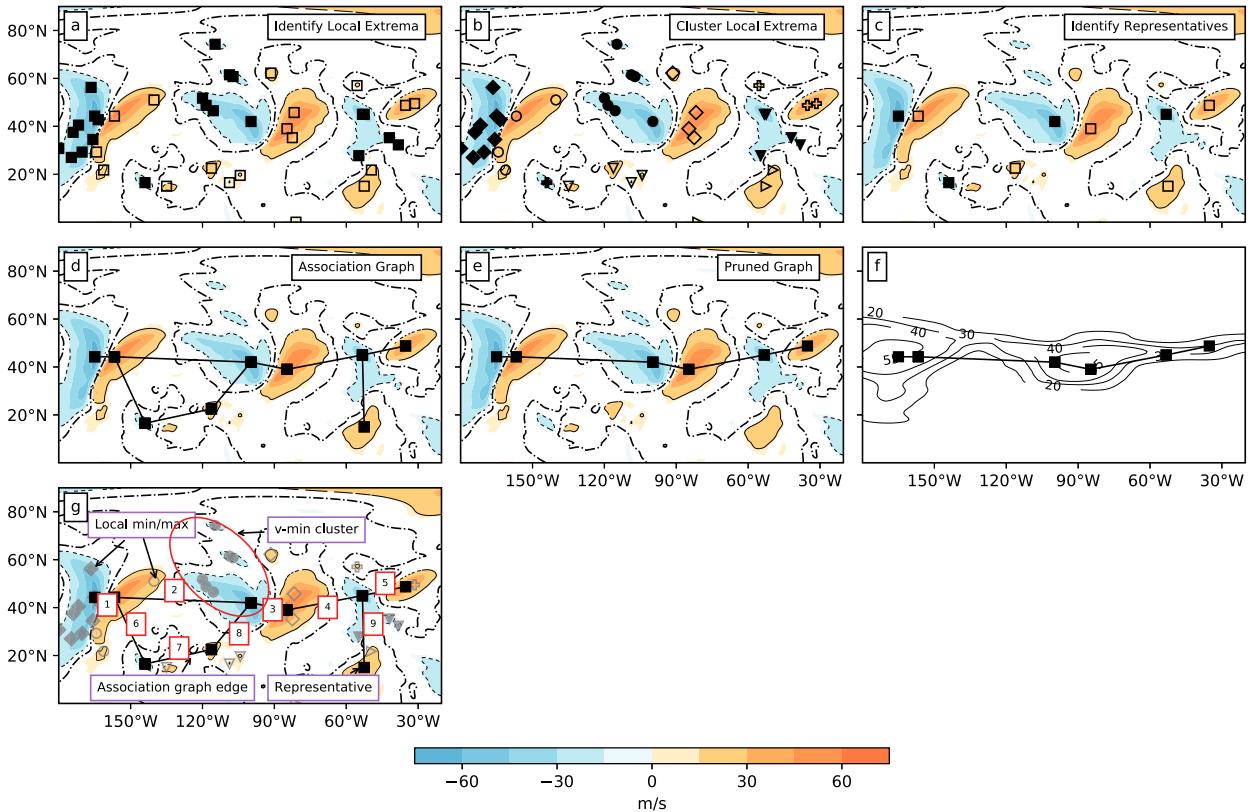


FIG. 2. The result of various steps of the algorithm applied to the meridional wind field at 0600 UTC 6 Jan 2007. (a) Identification of local maxima and minima. Local maxima are represented by unfilled squares and local minima by filled squares. (b) Clustering local maxima and minima into distinct v -max clusters and v -min clusters. Local maxima and minima belonging to different clusters are represented by different markers. (c) Choosing a representative local maximum and minimum per cluster. (d) Computing the association graph between different clusters. (e) Processing the association graph to obtain the RWP paths. (f) The RWP path overlaid on the RWP envelope representation (m s^{-1}) as a comparison. Using a threshold of 20 m s^{-1} to identify RWP objects, the envelope contours show the RWP object that would be identified using an algorithm such as Souders et al. (2014b). The envelope method highlights two centers of activity at 150° and 90°W . The RWP path highlights centers of activity where the graph nodes are clustered. The centers of activity identified by the RWP path are slightly displaced when compared to those identified by the envelope method, and this displacement depends on the shape of individual clusters. (g) A visual representation of various terms: the numbers label edges of the graph. An RWP path is a sequence of edges, say 1–2–3–4–5 or 1–6–7–8–3–4–9. The weight associated with each edge is called the edge weight, and the sum of weights of all edges in a path is used in the path optimization step for identifying individual RWPs. The dash-dotted line represents the zero-isocontour in all panels. The thin black contour represents the 30 m s^{-1} contour and the thin dashed contour represents the -30 m s^{-1} contour.

best representative of other points within the cluster. The messages sent in each iteration of the algorithm contain information about the suitability of a point to be an exemplar. These iterations continue similar to a voting process, until a set of stable exemplars and their corresponding clusters emerge.

We formulate a similarity measure that assumes high values for local maxima (minima) that belong to a common v -max cluster (v -min cluster) and low values if they belong to different v -max clusters (v -min clusters).

To understand the similarity measure, we first define a superlevel set and sublevel set for a given scalar value.

Given a scalar function f defined over a domain D , a superlevel set for a value c consists of all $x \in D$ for which $f(x) \geq c$. Similarly, a sublevel set would consist of all $x \in D$ for which $f(x) \leq c$. Intuitively, a superlevel set is obtained by clipping the scalar field using a scalar value and including all points in the domain above the clip. Similarly, a sublevel set is obtained by including all points below the clip. In the following discussion, we use a clipping parameter value of 0 to simplify the exposition. We discuss parameter selection in the following section.

The dissimilarity between two local maxima and minima x_i and x_j is defined as

$$\text{Dissimilarity}(x_i, x_j) = \begin{cases} D_{\text{supl}}(x_i, x_j) - v_{\text{min}}(x_i, x_j) & \text{for maxima} \\ D_{\text{subl}}(x_i, x_j) + v_{\text{max}}(x_i, x_j) & \text{for minima} \end{cases},$$

where $D_{\text{supl}}(x_i, x_j)$ is the length of the shortest path between x_i and x_j that lies within the superlevel set of 0, $D_{\text{subl}}(x_i, x_j)$ is the length of the shortest path between x_i and x_j that lies within the sublevel set of 0, $v_{\text{min}}(x_i, x_j)$ is the minimum of v within the straight line joining x_i and x_j , and $v_{\text{max}}(x_i, x_j)$ is the maximum of v within the straight line joining x_i and x_j .

The similarity between two local maxima and minima is defined as the negative of the above dissimilarity measure. Figure 3a shows the path connecting two pairs of local maxima, the first pair within a common v -max cluster and second pair from different v -max clusters. Such an approximate shortest path is computed by considering the shortest path in a graph whose nodes are the grid points of the superlevel (sublevel) set and edges connect all horizontally or vertically adjacent nodes. The superlevel set consists of multiple regions that are pairwise disjoint. If two local maxima belong to two different regions of the superlevel set of 0, there necessarily lies a region of negative v values that separates them. We therefore assume that they belong to disparate v -max clusters. Based on this assumption, we run affinity propagation clustering algorithm individually within each region of the superlevel set of 0 to identify v -max clusters. Similarly, we run the clustering algorithm individually within each region of the sublevel set of 0 to identify v -min clusters.

Figure 2b shows the typical output of the clustering step. Each v -max cluster (v -min cluster) has an associated region, namely, the interior of the zero-isocontour that bounds all its constituent maxima (minima). We observe that the clustering effectively segments the v field into v -max clusters and v -min clusters.

e. Compute association graph

The v -max clusters and v -min clusters identified in the previous step could together form a wave packet. In this step, the method identifies pairwise associations between v -max clusters and v -min clusters if there is evidence in the form of a shared boundary between them. The zero-isocontour of the v field is a natural boundary between v -max clusters and v -min clusters, since any path from a v -max cluster to a v -min cluster would necessarily pass through this zero-isocontour. We declare that a v -max cluster and v -min cluster share a common boundary when we find a point on the zero-isocontour whose closest local maximum and closest local minimum belong to the two clusters that represent

the v -max cluster and v -min cluster, respectively. Figure 3b shows how the representative maximum within a v -max cluster and minimum within a v -min cluster necessarily lie on either side of the shared boundary, which is a segment of the zero-isocontour. A path between them has to pass through the shared boundary, and therefore the zero-isocontour.

We compute a graph that stores all such associations by iterating through every point on the zero-isocontour and inserting an edge between the closest v -max cluster and v -min cluster. If a v -max cluster has more than one v -min cluster as a neighbor, then the corresponding node in the association graph has two edges associated with it. This situation is illustrated in Fig. 2d. To layout this graph within the spatial domain, we use the point with the highest magnitude of v within each cluster as the representative node.

Figure 2c illustrates the representative points for each v -max cluster and v -min cluster. Figure 2d shows the full association graph. The association graph acts as a succinct representation of the segmentation of the field and connectivity between segments.

f. Prune association graph

The association graph may contain many unwanted edges between v -max clusters and v -min clusters that are weakly associated and belong to separate RWPs. We therefore subject the association graph to a further pruning step to ensure the separation of the individual RWPs and validity of each identified connection.

To prune the graph, each edge of the association graph is assigned two weights. The scalar weight, that depends on the meridional wind within the v -max cluster and v -min cluster that it connects, and an estimated gradient, computed based on the maximum two-point gradient across all local maxima and minima pairs between a given v -max cluster and v -min cluster. More specifically, if e_{ij} is the edge between the i th and j th cluster,

$$\text{Scalar Weight}(e_{ij}) = \min(|v_i|, |v_j|),$$

where v_i and v_j are the highest absolute values of the meridional wind within the i th and j th clusters, respectively; and

$$\text{Estimated Gradient}(e_{ij}) = \max_{m_i \in M_i, m_j \in M_j} \frac{|v_{m_i}| + |v_{m_j}|}{\text{dist}(m_i, m_j)},$$

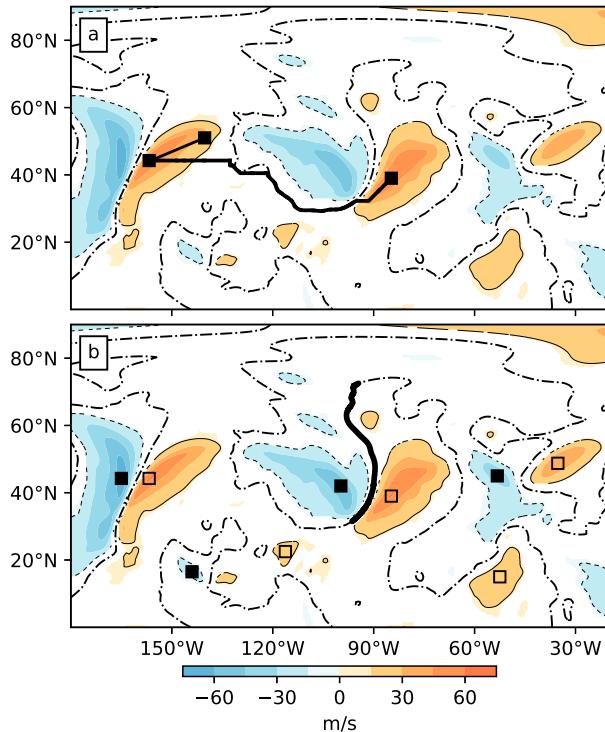


FIG. 3. (a) The shortest path between local maxima lying within the superlevel set of 0. One shortest path connects two local maxima that lie within the same v -max cluster. Another shortest path connects two local maxima that lie in two different v -max clusters. Note that neither path crosses the zero-isocontour (dot-dashed line); that is, they lie within the superlevel set of 0. (b) Shared boundary between two clusters (bold black). The shared boundary is a segment of the zero-isocontour (dot-dashed line). Points within this segment are such that their closest local maxima and minima belong to the respective clusters that define the v -max cluster and v -min cluster. All notations are as in Fig. 2.

where M_i and M_j are the set of all maxima or minima belonging to the i th and j th clusters, respectively; v_{m_i} and v_{m_j} are the meridional wind speeds for the maximum m_i and minimum m_j , respectively; and $\text{dist}(m_i, m_j)$ are the haversine distance between the points m_i and m_j .

Since RWPs are typically associated with high values of v , the clusters that belong to an RWP can be identified by choosing a high threshold for the scalar weights while pruning the graph. While the scalar-weight threshold only uses geometric information to prune the graph, the use of the gradient weight is physically motivated. The gradient of v provides an estimate of the curvature vorticity associated with a v -max cluster– v -min cluster pair. Pruning edges based on a threshold gradient weight works on the assumption that RWPs are separated by regions of low curvature vorticity.

The intensity of meridional wind along with the estimated curvature vorticity act as good pruning measures to separate the association graph into connected

components representing individual regions of RWP activity given the large number of associations found in the initial graph.

g. Extract and display RWPs

While the connected components of the pruned association graph represent regions of RWP activity, further processing is required to identify and display clear representative RWPs for each such region. We emphasize that choosing representative nodes for RWPs is purely for visualization, and the information about all maxima and minima in a cluster is retained for scientific analysis. The use of information contained in a cluster is illustrated in one of the subsequent case studies.

Nodes in a connected component represent clusters of local maxima and minima while edges represent their spatial connections. Therefore, in extracting an optimal representative RWP, we are not only faced with a choice over the set of clusters but also a choice of representative maxima or minima for these clusters. We thus model this decision as an optimization problem across all simple paths between the connected clusters constructed using all possible representative critical points for each cluster. First, we filter and retain only those paths where the longitudes form an increasing sequence. This prevents backward connections. Next, we compute path scores as the sum of edge weights of the path's constituent edges. The edge weights are calculated using the estimated gradient described previously. This weight penalizes meridional associations by scaling latitude coordinates by a constant factor of 2 for the distance computation. Finally, the path with the highest score within each connected component is extracted and displayed as the representative RWP. Figure 2e shows the extracted RWP and Fig. 2f presents a visual comparison with the envelope representation.

Our final representation for an identified RWP is therefore in the form of a graph. However, it is important to note that each node in the graph represents the cluster of local maxima or minima belonging to the corresponding v -max cluster or v -min cluster. This collection of critical points facilitates access to the geometric properties of the v -max clusters or v -min clusters, like the spatial extent (computed based on the spatial distribution of local maxima and minima), amplitude (computed as the mean or median value of the amplitude of the local maxima and minima), and orientation (computed as a weighted least squares fit to the local maxima and minima). The geometric properties support further analysis of the identified RWPs. We plan to elaborate on the visualization and interaction aspects of our framework in a companion

paper, which documents the open-source software package that we have developed.

3. Sensitivity to tunable parameters

Our RWP identification algorithm has three tunable parameters:

- The scalar-weight threshold (ST), which determines which of the identified local maxima or minima are included in the RWP computation.
- The gradient-weight threshold (GT), which is used to prune edges from the association graph to identify the RWPs.
- The clipping parameter, which controls the spatial extent of each v -max cluster and v -min cluster.

A clustering parameter is also available. However, our analysis suggests that varying the clipping parameter results in physically more intuitive results. Hence the clustering parameter is fixed to a constant value for all our sensitivity tests.

To find appropriate values for ST and GT, we used 6-hourly meridional winds at 300 hPa during winter (December–February) for four arbitrarily chosen years (1990, 1995, 2000, and 2005) and calculated the statistics of the identified RWPs for a range of values of ST and GT. Each time step was considered independently and RWPs were not tracked across time steps. While this method implies multiple counting of RWPs, this is not directly relevant to the question of choice of thresholds since we are more interested in the spatial structure of the RWPs than their lifetime characteristics.

The summary statistics obtained by varying GT and ST is presented in Fig. 4. The number of time steps in Figs. 4d and 4h containing RWPs displays a weak dependence on GT and ST until a certain value (up to 0.04 s^{-1} for GT and up to 40 m s^{-1} for ST), then reduces rapidly. Thus, there appears to be an upper bound on the gradient and amplitude of the meridional wind in the chosen months. We observed similar behavior during the other seasons. The rapid reduction in number of time steps containing RWPs begins at lower values of GT (0.03 s^{-1}) and ST (35 m s^{-1}) in the summer (see Figs. S1–S3 in the online supplemental material). This difference suggests that there exist RWPs of relatively lower amplitude in the summer as compared with the other seasons. Furthermore, at a given value of the scalar threshold, there exist RWPs with a lower value of the estimated curvature vorticity in the summer as compared with other seasons. This observation is consistent with lower RWP activity and less intense RWPs observed in the summer in Souders et al. (2014a, their Figs. 4 and 5). It would be interesting

to see if such behavior is exhibited in longer-term datasets.

Spatially, low values of GT tend to merge multiple RWPs (observed visually) whereas high values tend to identify only intense dipole structures—see changes in identified RWPs between 140° and 240°E in Fig. S4. The merging of distant RWPs is also evidenced by the fact that the median edge length is shorter than the mean edge length in Fig. 4b for values of GT lower than 0.03. For values of GT higher 0.05 this behavior is observed again, but is likely due to the very small edge lengths associated with the intense dipole structures. The shorter median length implies the presence of longer-than-average edges in the pruned association graph. Figure 4f shows that the difference between mean and median values of edge lengths is also seen as ST decreases, suggesting that decreasing ST (keeping GT fixed) connects distant RWPs. A high value of GT can be used if the requirement is to identify regions of intense RWP activity.

Figures 4a and 4e shows that increasing GT and ST decreases the number of edges in each pruned association graph almost monotonically. This is because increasing values of ST results in the rejection of lower-amplitude maxima and minima, and increasing values of GT prunes edges associated with a lower gradient. Since the mean extent of the RWPs is simply the sum of the length of the edges in the pruned association graph, this metric decreases monotonically with increasing GT and ST as well, as seen in Figs. 4c and 4g.

The following analysis and case studies focus on the range of ST and GT where the number of identified RWPs is relatively constant: 0.1–0.4 for GT and 30–40 for ST. Specifically, we use a GT of 0.3 and ST of 30. For these values of the thresholds, the algorithm produces RWPs that have two edges (three v -max clusters or v -min clusters) on average and have an average wavelength (equal to twice the edge length, Fig. 4) of around 4000 km, a value supported by other methods as well (Chang 1999).

We performed a sensitivity analysis to determine the effect of the clipping parameter. The RWP statistics are not very sensitive to different clip values except for low-amplitude RWPs. For low-amplitude RWPs, increasing the clipping parameter results in smaller edge lengths since we no longer connect clusters that are far apart. This is seen in Fig. S5. Furthermore, low-amplitude clusters are correctly separated from the high-amplitude ones, whereas high-amplitude clusters are largely insensitive to the choice of the clip value (see Fig. S6). Based on this analysis, the framework uses a constant 2 m s^{-1} as the clip value to ensure fidelity with the

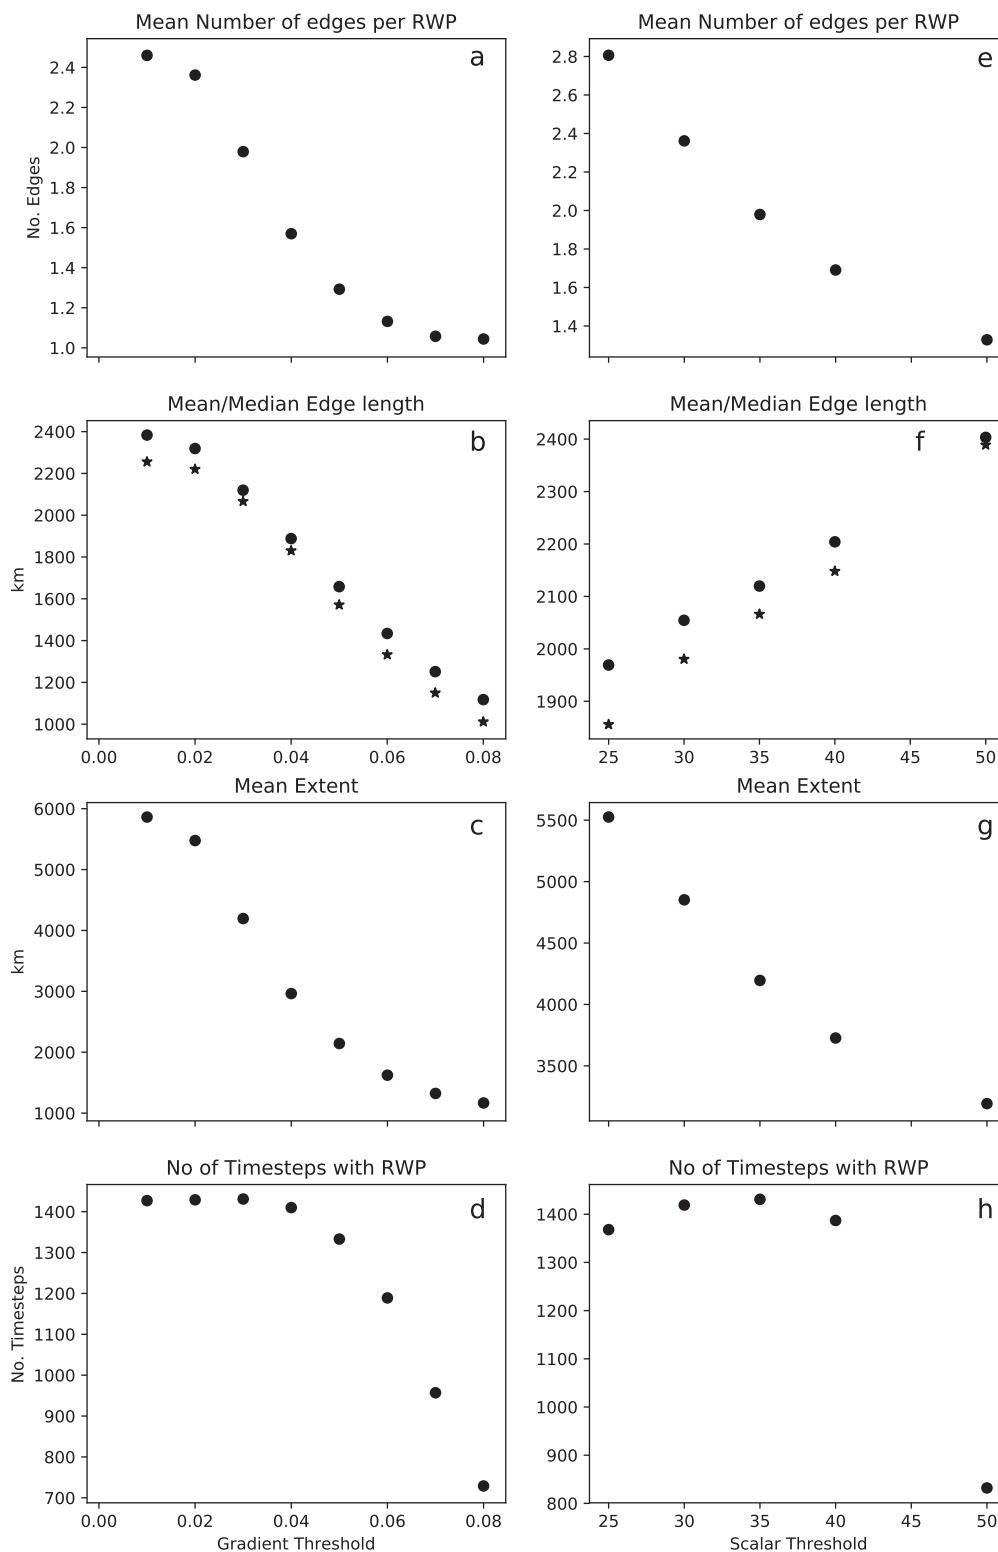


FIG. 4. The spatial statistics of the identified RWPs when (a)–(d) the gradient weight threshold GT is varied and (e)–(h) the scalar-weight threshold ST is varied. The edge length is calculated as the haversine distance between the end-point nodes. The mean extent of the RWP is calculated by adding the lengths of all edges in the pruned association graph. In (b) and (f) the median value is represented by stars and the mean value by circles. The total number of time steps in the data is 1440.

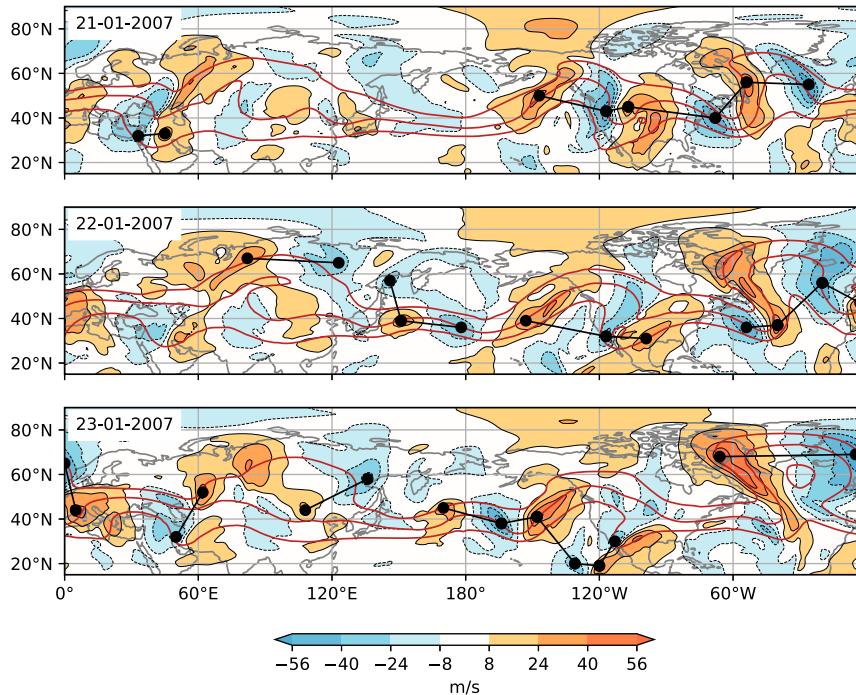


FIG. 5. The development of a locally forced RWP captured using the v field. The filled contours represent v , with blue-filled contours representing negative values. The brown contours represent the 300-hPa geopotential height (870, 900, and 930 dam). The pruned association graph representing the RWP is shown using black nodes and edges. The panels show v at 0600 UTC.

geometric intuition behind the clustering analysis, while obtaining physically appropriate clusters.

4. Case studies

We present a series of case studies that demonstrate the use of the proposed graph-based representation of RWPs and the visualization framework for interactive exploration and visual analysis. We also present comparisons with previous approaches to identify and represent RWPs and discuss strengths and shortcomings of the integrated geometric and topological approach.

a. RWP genesis

We use the same cases used by Souders et al. (2014b, their Figs. 6 and 7) to illustrate the genesis of RWPs, with and without previous RWP activity. The first case study tracks the evolution of the RWP field between 21 and 23 January 2007. The second case study focuses on the time period 6–10 January 2007. In the first case, an RWP is forced locally to the east of Japan by a deepening cyclone. As seen in Fig. 5, there is no RWP activity visible (as indicated by the lack of nodes/edges) over the western part of the Pacific basin (120°–180°E) at 0600 UTC 21 January 2007. The following day, the RWP

is visible over the eastern Pacific and amplifies downstream on 23 January. The sequence of development is captured satisfactorily when compared with the description in Souders et al. (2014b). Our method also suggests an equatorward extension of the merging RWPs over the Pacific (between 170°E and 120°W), which is not captured by the envelope approach (Souders et al. 2014b, their Fig. 6c).

In the second genesis case study, the development of the RWP seems to be shifted in time compared to the description in Souders et al. (2014b).¹ There is the genesis of an RWP associated with a cutoff low in Fig. 6 on 6 January at 120°E. This weak RWP amplifies downstream on 8 January and merges with an existing RWP over the Pacific near 180°. This merged, amplifying RWP then merges with another RWP present over North America on 10 January, forming a contiguous RWP field that extends from the Pacific to Europe. The pruned association graph again represents the development satisfactorily, although a v -max cluster and

¹ Analyzing the NCEP data used by Souders et al. (2014b), their sequence of figures appears to start around 4 January.

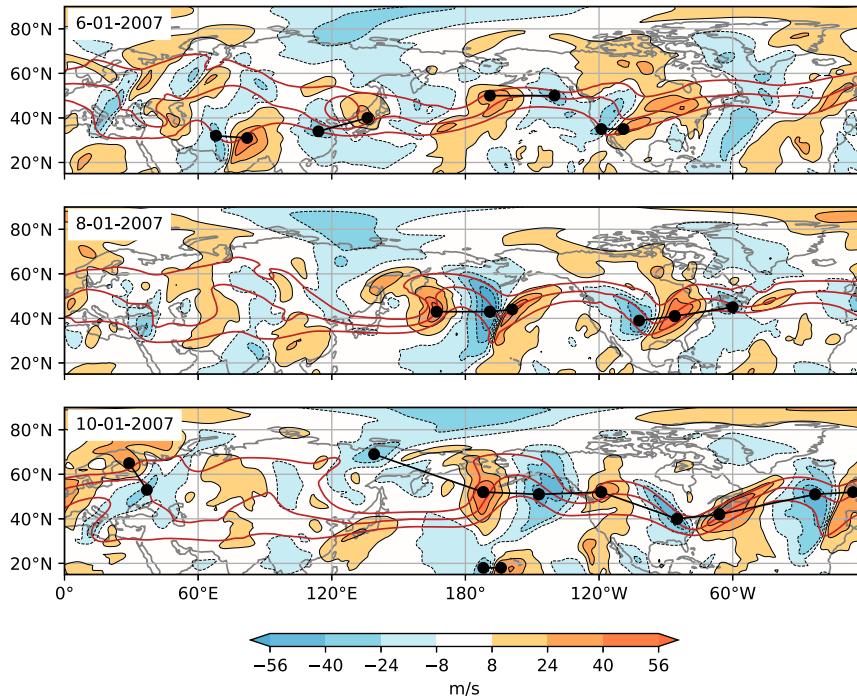


FIG. 6. The development of an RWP forced by a decaying RWP. Filled contours represent v , with blue corresponding to negative, and red/orange corresponding to positive values. The dark brown contours represent the 300-hPa geopotential height (870, 900, and 930 dam), and the graph representation of the RWP is shown using black nodes and edges. The panels show v at 0000 UTC.

v -min cluster connected around 120°E on 10 January do not seem to be part of the same RWP.

Overall, the representation of RWP development using our algorithm seems to be comparable to the envelope-based method in these case studies. Since our method is not constrained to envelopes along streamlines, it is able to capture the meridional extension of RWPs into the tropics as seen in the first case study, which may be important for studies of tropical–extratropical interactions.

b. Identifying nonwavelike RWPs

Identifying RWPs throughout their life cycle is desirable, especially during the finite amplitude evolution when the RWP no longer resembles a wave. Ghinassi et al. (2018) show that the Hilbert transform-based method fails to capture some parts of the evolution of the RWP involved in the “forecast bust” in April 2011.² This forecast bust was associated with a significant drop in forecast skill over Europe for both ECMWF and Met Office forecasts (see Ghinassi et al. (2018)

and references therein). We analyze the same event using our approach to investigate the reasons why the latter method fails in tracking the complete evolution.

Figure 7 traces the evolution of the v field and the corresponding RWP field between 12 and 17 April 2011. The algorithm initially identifies four RWPs, one over Europe (0°–60°E), one over Russia (120°E), a weaker one in the Pacific (180°–210°E), and one in the Atlantic Ocean (30°–100°W). This is in contrast to both the envelope method and the method in Ghinassi et al. (2018, their Figs. 4a and 4g), which show three centers of RWP activity. The first three RWPs propagate eastward on 13 April, whereas the RWP in the Atlantic Ocean evolves into a non-wavelike configuration as suggested by the zig-zag node-edge configuration of the pruned association graph (Ghinassi et al. 2018, their Figs. 4b and 4h). As pointed out in Ghinassi et al. (2018), even though there is a strong wave activity flux signal over the Atlantic Ocean, the Hilbert transform-based method does not capture the RWP probably due to the complicated spatial structure.

The RWP over Russia seems to grow in situ over the rest of the period before merging with the RWP over Europe on 15 April. The amplitude of this RWP reduces dramatically between 16 and 17 April, suggesting that it

²It is unclear if this picture changes if the Hilbert transform is defined along streamlines as described in Zimin et al. (2006).

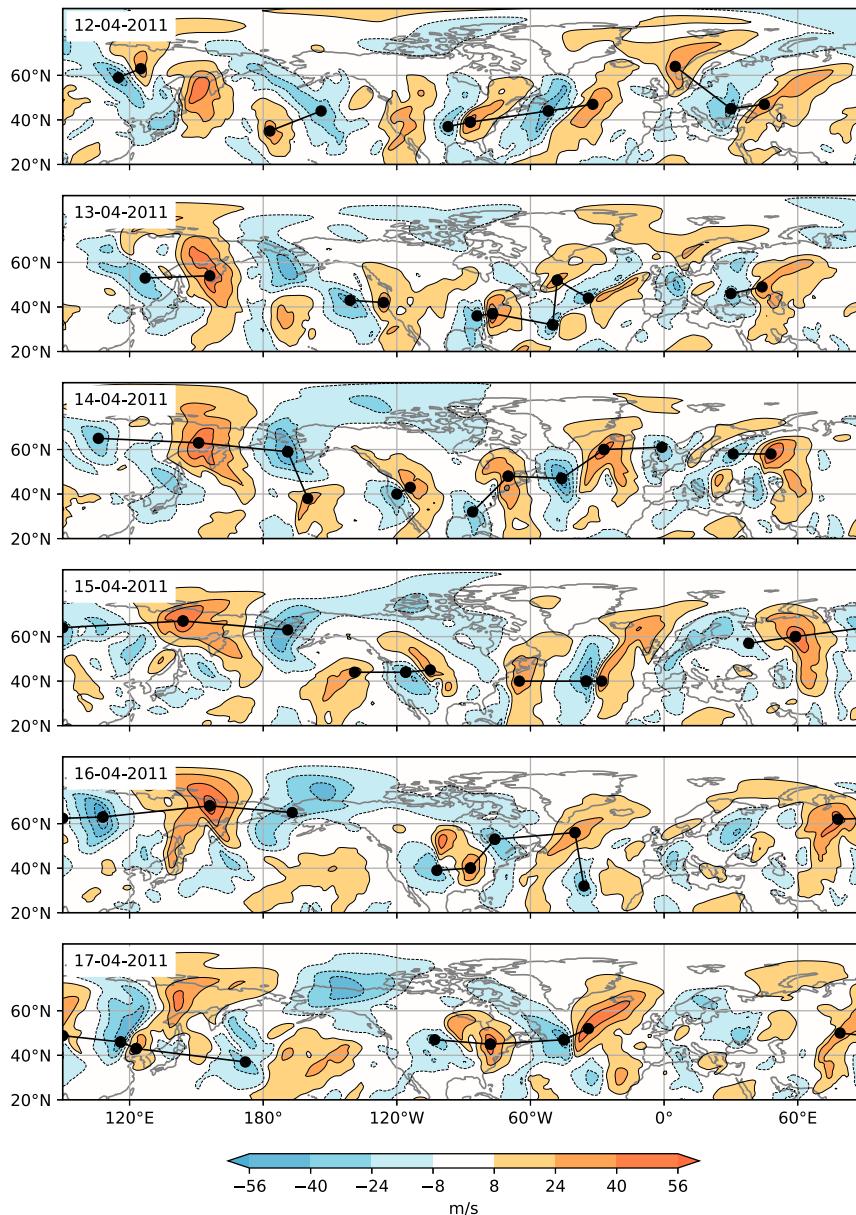


FIG. 7. The evolution of the meridional wind field between 0000 UTC 12 Apr and 0000 UTC 17 Apr 2011. The graph representation of RWPs identified by our method are shown using black nodes and edges. Negative values of v are denoted by blue-filled contours and positive values by red-/orange-filled contours.

is dissipating around 17 April. The RWP in the Atlantic seems to amplify and “reorganize” into a more wavelike pattern on 14 April. The RWP over the Pacific propagates toward and merges with the RWP over the Atlantic on 16 April, and leads to a reenergizing of RWP activity over the Atlantic on 17 April. Interestingly, the structure of this reenergized Atlantic RWP on 17 April is captured better by the envelope method rather than the LWA field (Ghinassi et al. 2018, compare Figs. 4f and 4l).

Our method seems to capture both the wavelike and non-wavelike parts of the evolution of the RWP field satisfactorily. However, the algorithm seems to miss connecting the Atlantic and Eurasian RWPs (via the v -min cluster over Scandinavia) on 15 April.

An important property of our algorithm is that the phase information is preserved and the amplitude and location of the RWP is represented directly in the spatial domain without any undesirable shifts that might occur

due to filtering. This property allows us to capture complicated RWP configurations while maintaining a direct correspondence to the raw data.

c. RWP driven wet-bulb temperature extremes in southwest Pakistan

The pruned association graph representing the RWPs can also be used for more statistically oriented studies. The structure of RWPs over multiple events is often studied by first averaging the meridional wind over these events, resulting in a composite field. The graph representation of the RWP is unusual but is capable of providing a more nuanced picture as the following case study illustrates.

South Asia contains regions with some of the highest wet-bulb temperatures observed globally (Im et al. 2017). The Sindh region of Pakistan contains a localized hotspot which is observed to have high wet-bulb temperatures in May and June, see Fig. S7 and Fig. 1 in Monteiro and Caballero (2019). Extreme wet-bulb temperature events in this region during May and June were identified by Monteiro and Caballero (2019) using the HadISD station data (Dunn et al. 2016). They define extreme events as time periods of 3 days or more when the mean wet-bulb temperature in three stations located in the hotspot is above the 90th percentile, and at least one station has a wet-bulb temperature that is above the 97th percentile. We refer the reader to their paper for further details. The circulation patterns associated with extreme wet-bulb temperatures suggest that the propagation of a RWP through this region is responsible for advection of moist oceanic air into the region, causing elevated wet-bulb temperatures. Southerly winds advecting oceanic air is mostly confined to the boundary layer and is in the opposite direction to the upper-tropospheric winds associated with the RWP, which is northerly. Here, we compare the signature of the RWPs using a composite field (averaging over all events) and using the pruned association graphs generated by our method.

The composite meridional winds for all extreme events (38 in total over the period 1979–2016) is shown in Fig. 8. We calculate the statistical significance of the meridional wind composites using a bootstrap test with 5000 samples, where each sample is the mean of 38 randomly selected meridional wind fields in May and June during the period 1980–2016. Regions with an amplitude that is significant at the $p < 0.05$ level are considered RWP-related wind anomalies. There is a weakening of the jet over central Asia associated with a weak northerly wind anomaly. The wave packet then develops over the days leading to the event, displaying a maximum on the day of the event and dissipates thereafter. The jet seems to weaken early on, which is puzzling given the weak amplitude of the

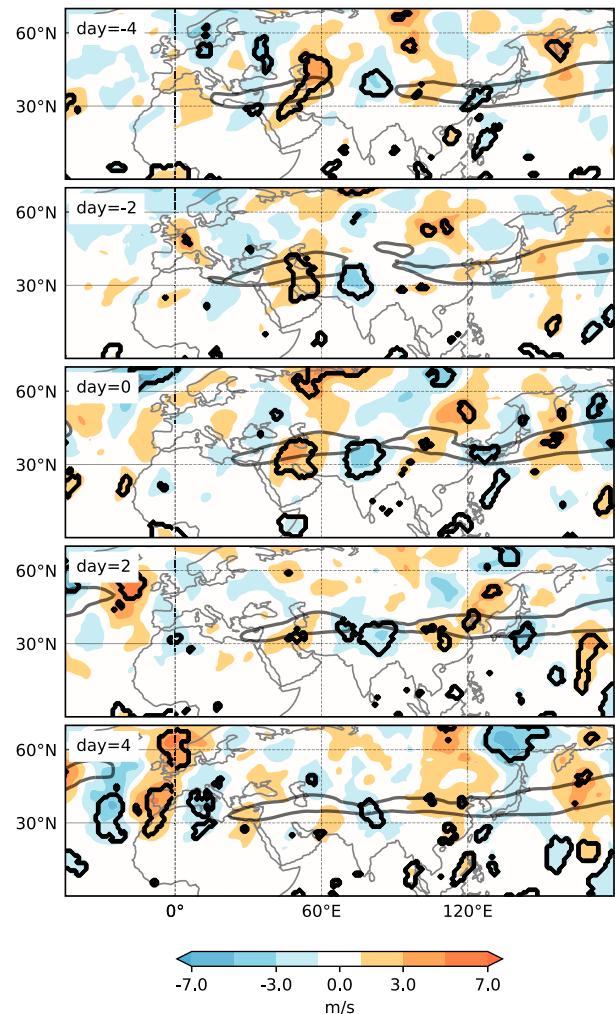


FIG. 8. RWP composites during extreme wet-bulb temperature events in South Asia. The filled contours represent the composite anomalous v associated with the event at 300 hPa. The contours are spaced 2 m s^{-1} apart starting from 1 and -1 m s^{-1} for the positive and negative values, respectively. The thick gray contour represents the region where the zonal winds exceed 20 m s^{-1} during the events. The thick black contours enclose regions with statistically significant meridional winds. The meridional wind anomalies are calculated using a smoothed daily climatology.

northerly winds in the region 4 days prior to the event. The jet remains weak in the region as the wave packet strengthens, and recovers a few days after the event. The amplitude associated with the wave packets appears weak in this composite analysis, with the meridional winds reaching a maximum of around 7 m s^{-1} close to the center of the weakened jet, $\sim 30^\circ\text{N}$, 80°E . The composite calculated here uses v anomalies calculated as the departure from a smoothed daily climatology. Using the anomalies was necessary in this case because the composite of the absolute meridional winds (not shown) is noisier and the v -max clusters and v -min clusters are not as clearly visible.

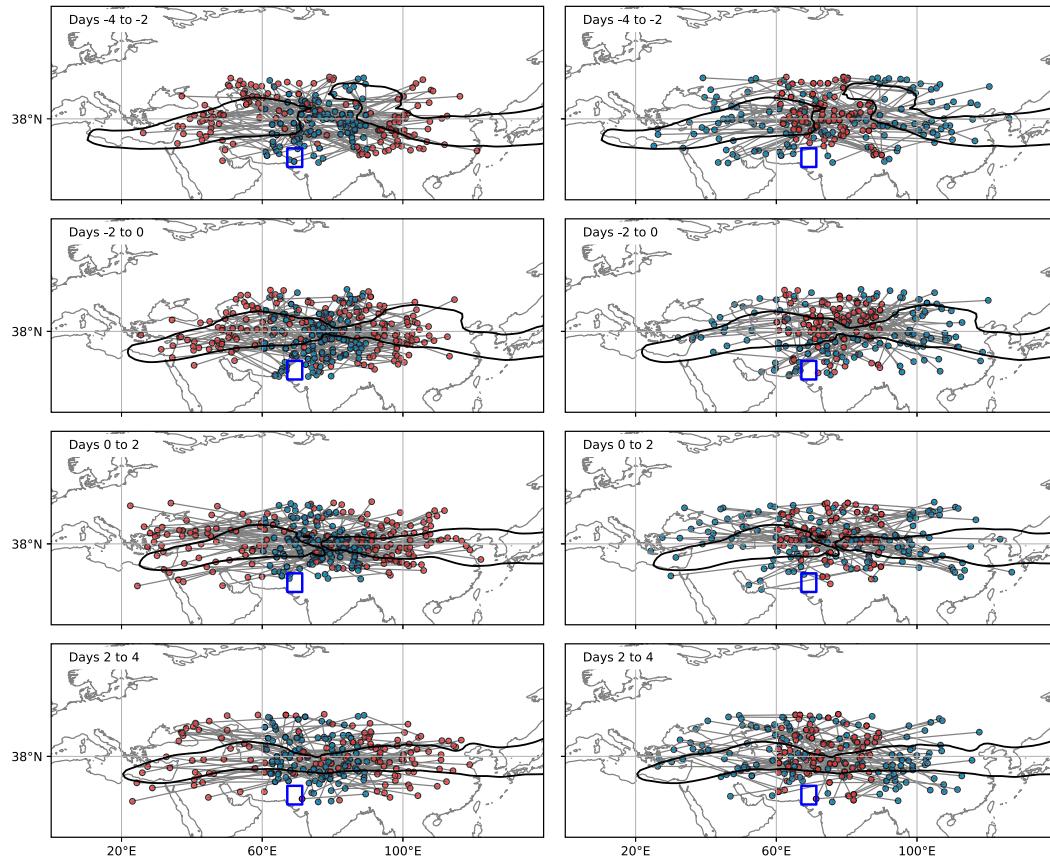


FIG. 9. Extreme wet-bulb temperature events in South Asia. Those RWP that have (left) a v -min cluster and (right) a v -max cluster between 60° and 90° E are shown. Blue dots represent minima or winds from the north and red dots represent maxima or winds from the south. The 20 m s^{-1} contour of the zonal winds associated with the events is plotted for reference. Each panel contains all RWP in the 2-day period indicated in the label. The region of interest is represented using a blue box. Out of a total of 38 events, 16 events exhibit predominantly northerly winds, 8 events exhibit predominantly southerly winds, and 14 events exhibit an equal frequency of northerly and southerly winds.

The association graph-based analysis provides a more nuanced picture of the event. In this case, using absolute or anomalous meridional winds does not significantly affect the results. We therefore present results using the absolute meridional winds.

The pruned association graph (or RWP paths) were filtered to retain only those v -max clusters and v -min clusters lying between 60° and 90° E. The immediate neighbors of these v -max clusters and v -min clusters were also retained. All paths were grouped in two groups: One group of paths whose v -max cluster lies between 60° and 90° E, and another whose v -min cluster lies in the same region. A small number of RWP paths are oriented meridionally, and have both a v -max cluster and a v -min cluster between 60° and 90° E. These RWP paths are included in both groups. The resulting pruned association graphs are shown in Fig. 9. The number of RWP paths plotted in each of the left panels of

Fig. 9 (top to bottom) are 64, 61, 73, and 53, respectively. The number of RWP paths plotted in each of the right panels of Fig. 9 (top to bottom) are 47, 48, 55, and 40, respectively. The presence of v -max clusters in the region of interest as seen in Fig. 9 implies that some events are actually associated with southerly winds, as opposed to what the composite picture may suggest. In contrast to the previous case studies, the nodes of the graph are placed at the “center of mass” of each cluster, which is defined as the weighted mean of the location of the local maxima or minima. The weights used are the value of the meridional wind at the location of the local maximum or minimum. It is observed that the RWP responsible for the event are much weaker in amplitude as compared to the RWP typically observed over the Pacific and Atlantic basins. To preferentially capture these RWP, we only identify RWP whose amplitude lies between 10 and 30 m s^{-1} . Correspondingly,

we use a gradient threshold of 0.01 to ensure we retain edges linking these weak v -max clusters and v -min clusters.

Four days prior to the event, it is seen that the region over central Asia with the weakened jet has multiple maxima and minima in the vicinity of the jet, as seen in both panels in the first row of Fig. 9. These mutually cancel each other in the composite analysis to suggest a weak minima. Thus, the composite analysis does not show us the actual structure of the circulation in this time step. Since envelope-based methods do not convey phase information, it is unclear whether they can capture such information either. However, since envelope-based methods accurately capture the amplitude of RWPs, they will be able to convey the fact that the amplitude of individual RWPs is much larger than what the composite analysis suggests. To understand the apparent equatorward amplification of the RWP signature in Fig. 8, we partition the maxima and minima based on their location relative to the jet axis. We choose a nominal latitude of 38°N as the location of the jet axis. The fraction of minima equatorward of this latitude for the four time periods (left panels in Fig. 9) are 0.53, 0.67, 0.61, and 0.41, respectively. Similarly, the fraction of maxima equatorward of this latitude (right-hand panels) for the four time periods is 0.27, 0.35, 0.3, and 0.25, respectively. These fractional distributions suggest that in the 2 days leading up to the event (left panel of second row of Fig. 9), the minima aggregate equatorward of the jet. This aggregation corresponds to the equatorward amplification of negative values of v around 70°E as seen in the top two panels of Fig. 8. This meridionally separated clustering continues to a lesser extent in the 2 days following the event. Eventually, the RWP signature dissipates as evidenced by lower density of v -max clusters and v -min clusters 2 days after the event.

During the entire period of the analysis, the individual v -max clusters and v -min clusters have a much higher amplitude of around 20 m s^{-1} (not shown), which is 3 times the value suggested by the composite analysis. This is because of the phase differences between the RWPs in each event, which leads to a substantially weaker signal on averaging.

As Fig. 9 suggests, some extreme events are associated with northerly winds at 300 hPa between 60° and 90°E , whereas others are associated with a southerly winds. The graph-based RWP representation makes it easy to group events based on whether each event is associated with predominantly northerly or southerly winds (see Figs. S8 and S10). We define an event to be associated predominantly with southerly (northerly) winds if the number of v -min clusters (v -max clusters) identified between 60° and 90°E is 0.6 times the number of v -max

clusters (v -min clusters) over the 8 days of the analysis. If neither criterion is met, then the event is defined to be associated with an equal frequency of both northerly and southerly winds. Based on these criteria, 16 events exhibit predominantly northerly winds, 8 exhibit predominantly southerly winds and 14 exhibit an equal frequency of both northerly and southerly winds from a total of 38 events. Though all these events are associated with a positive wet-bulb temperature anomaly in the region of interest, the evolution of the surface fields are quite different between these kinds of events (see Figs. S9 and S11 for the difference between events associated with predominantly northerly and southerly winds, respectively). Furthermore, events which exhibit northerly winds at 300 hPa are associated with larger wet-bulb temperature anomalies resulting from a more coherent advection of oceanic air into the region of interest (see Figs. S9 and S11). While a more detailed analysis is beyond of the scope of this paper, we note that the selection of extreme events with different life histories provides a more nuanced picture of the evolution of these extreme events and the differences between individual events.

This case study highlights another advantage of the current method, which is the capability to selectively identify RWPs within a particular amplitude range, which is challenging to achieve using other methods. Furthermore, the graph-based representation provides a more nuanced picture of the composite RWP structure since the representation is not affected by phase differences of RWPs between events.

5. Conclusions

The extraction of meteorologically interesting information from noisy observational data has always been a challenge. As the volume of data available to climate scientists grows ever larger, there is a further imperative to formulate algorithms that work with minimal but effective human supervision. Our paper is a contribution toward this end.

In the context of RWPs, the question of automated extraction is tricky since the definition of what constitutes an RWP is ambiguous: the spatial extent, wavenumber and amplitude all vary during the life cycle of an RWP, which makes it hard to formulate concise algorithms for their extraction. Furthermore, what RWP is “interesting” is also a question that can be answered only on a case-by-case basis. For instance, in our final case study the RWPs associated with high wet-bulb temperatures had amplitudes of around 20 m s^{-1} , which would be considered very low for the first and second case studies.

Given these challenges, we propose the use of topological methods as a novel alternative to the transform-based methods used for identification of RWPs. The topological description of input data is inherently robust against noise. This robustness eliminates the need to smooth the input data which can otherwise result in undesirable phase shifting of the RWP field.

A concise comparison of the relative strengths and weaknesses of the current method and other methods is presented in Table 1. Our algorithm is able to generate graph-based RWP descriptions that are concise, informative, and robust in the presence of noisy data.

The values of the tuning parameters in our algorithm were chosen based on a statistical analysis. The results for individual case studies based on these chosen values are satisfactory, which is heartening given that the statistical analysis was performed on an independent time period. This independent validation gives us confidence that these tuning values are valid for the general case.

A weakness of our method currently is that some v -max clusters or v -min clusters may not be connected by an edge after pruning the association graph. We have noticed that such edges are retained in future time steps, as the v -max cluster or v -min cluster amplitude and configuration change. Thus, including the time dimension will allow for more robust identification of RWPs. Keeping this in mind, a natural extension of our work is to follow the identified RWPs over time using a tracking algorithm. Algorithms to track graphs over time do exist (see, e.g., Doraiswamy et al. 2013; Valsangkar et al. 2019), but it remains to be seen which algorithm is both efficient and accurate in this context. In particular, tracking algorithms that work well during RWP splits and merges would be desirable. We leave this extension to future work. On the other hand, our method is able to capture meridional excursions of the RWP quite naturally since we are not constrained to work along streamlines.

The development of our new identification scheme along with a tracking algorithm would also provide an independent validation of RWP statistics calculated previously (Souders et al. 2014a; Hunt et al. 2018) and would pave the way toward a better understanding of RWPs and their role in shaping the weather and climate.

Acknowledgments. This work is partially supported by a Swarnajayanti Fellowship from the Department of Science and Technology, India (DST/SJF/ETA-02/2015-16), a Mindtree Chair research grant, an IoE research grant from Indian Institute of Science, and the Robert Bosch Centre for Cyber Physical Systems, IISc. JMM was partially funded by the Swedish Research Council (Vetenskapsrådet) Grant E0531901 and by IISER Pune.

Part of this work was carried out when the first author was visiting IISc Bangalore as an intern from BITS-Pilani, Hyderabad Campus. Part of this work was carried out when the second author was in Stockholm University and the Divecha Centre for Climate Change at the Indian Institute of Science.

Data availability statement: All data used in this paper are available from the ECMWF website: <https://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=pl/>. The source code for the RWP identification method proposed in this paper is available from https://bitbucket.org/vgl_iisc/rossby-wave-packet-identification.

REFERENCES

- Chang, E. K. M., 1999: Characteristics of wave packets in the upper troposphere. Part II: Seasonal and hemispheric variations. *J. Atmos. Sci.*, **56**, 1729–1747, [https://doi.org/10.1175/1520-0469\(1999\)056<1729:COWPIT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1999)056<1729:COWPIT>2.0.CO;2).
- , and I. Orlanski, 1993: On the dynamics of a storm track. *J. Atmos. Sci.*, **50**, 999–1015, [https://doi.org/10.1175/1520-0469\(1993\)050<0999:OTDOAS>2.0.CO;2](https://doi.org/10.1175/1520-0469(1993)050<0999:OTDOAS>2.0.CO;2).
- , and D. B. Yu, 1999: Characteristics of wave packets in the upper troposphere. Part I: Northern Hemisphere winter. *J. Atmos. Sci.*, **56**, 1708–1728, [https://doi.org/10.1175/1520-0469\(1999\)056<1708:COWPIT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1999)056<1708:COWPIT>2.0.CO;2).
- Dee, D. P., and Coauthors, 2011: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, **137**, 553–597, <https://doi.org/10.1002/qj.828>.
- Dimri, A. P., D. Niyogi, A. P. Barros, J. Ridley, U. C. Mohanty, T. Yasunari, and D. R. Sikka, 2015: Western disturbances: A review. *Rev. Geophys.*, **53**, 225–246, <https://doi.org/10.1002/2014rg000460>.
- Doraiswamy, H., V. Natarajan, and R. S. Nanjundiah, 2013: An exploration framework to identify and track movement of cloud systems. *IEEE Trans. Vis. Comput. Graph.*, **19**, 2896–2905, <https://doi.org/10.1109/TVCG.2013.131>.
- Dunn, R. J. H., K. M. Willett, D. E. Parker, and L. Mitchell, 2016: Expanding HadISD: Quality-controlled, sub-daily station data from 1931. *Geosci. Instrum. Methods Data Syst.*, **5**, 473–491, <https://doi.org/10.5194/gi-5-473-2016>.
- Fragkoulidis, G., V. Wirth, P. Bossmann, and A. H. Fink, 2018: Linking Northern Hemisphere temperature extremes to Rossby wave packets. *Quart. J. Roy. Meteor. Soc.*, **144**, 553–566, <https://doi.org/10.1002/qj.3228>.
- Frey, B. J., and D. Dueck, 2007: Clustering by passing messages between data points. *Science*, **315**, 972–976, <https://doi.org/10.1126/science.1136800>.
- Ghinassi, P., G. Fragkoulidis, and V. Wirth, 2018: Local finite-amplitude wave activity as a diagnostic for Rossby wave packets. *Mon. Wea. Rev.*, **146**, 4099–4114, <https://doi.org/10.1175/MWR-D-18-0068.1>.
- Hagberg, A. A., D. A. Schult, and P. J. Swart, 2008: Exploring network structure, dynamics, and function using networkx. *Proc. Seventh Python in Science Conf.*, Pasadena, CA, Entthought, [scipy.org](https://doi.org/10.1109/SciPy.2008.4454882), 11–15.
- Hakim, G. J., 2003: Developing wave packets in the North Pacific storm track. *Mon. Wea. Rev.*, **131**, 2824–2837, [https://doi.org/10.1175/1520-0493\(2003\)131<2824:DWPTITN>2.0.CO;2](https://doi.org/10.1175/1520-0493(2003)131<2824:DWPTITN>2.0.CO;2).
- Heine, C., H. Leitte, M. Hlawitschka, F. Iuricich, L. De Florian, G. Scheuermann, H. Hagen, and C. Garth, 2016: A survey of

- topology-based methods in visualization. *Comput. Graph. Forum*, **35**, 643–667, <https://doi.org/10.1111/cgf.12933>.
- Hovmöller, E., 1949: The trough-and-ridge diagram. *Tellus*, **1**, 62–66, <https://doi.org/10.3402/tellusa.v1i2.8498>.
- Hunt, K. M. R., A. G. Turner, and L. C. Shaffrey, 2018: The evolution, seasonality and impacts of western disturbances. *Quart. J. Roy. Meteor. Soc.*, **144**, 278–290, <https://doi.org/10.1002/qj.3200>.
- Im, E.-S., J. S. Pal, and E. A. B. Eltahir, 2017: Deadly heat waves projected in the densely populated agricultural regions of South Asia. *Sci. Adv.*, **3**, e1603322, <https://doi.org/10.1126/sciadv.1603322>.
- Lee, S., and I. M. Held, 1993: Baroclinic wave packets in models and observations. *J. Atmos. Sci.*, **50**, 1413–1428, [https://doi.org/10.1175/1520-0469\(1993\)050<1413:BWPIMA>2.0.CO;2](https://doi.org/10.1175/1520-0469(1993)050<1413:BWPIMA>2.0.CO;2).
- Martius, O., C. Schwierz, and H. C. Davies, 2010: Tropopause-level waveguides. *J. Atmos. Sci.*, **67**, 866–879, <https://doi.org/10.1175/2009JAS2995.1>.
- Monteiro, J. M., and R. Caballero, 2019: Characterization of extreme wet-bulb temperature events in southern Pakistan. *Geophys. Res. Lett.*, **46**, 10 659–10 668, <https://doi.org/10.1029/2019GL084711>.
- Parker, T. J., G. J. Berry, M. J. Reeder, and N. Nicholls, 2014: Modes of climate variability and heat waves in Victoria, southeastern Australia. *Geophys. Res. Lett.*, **41**, 6926–6934, <https://doi.org/10.1002/2014GL061736>.
- Ratnam, J. V., S. K. Behera, S. B. Ratna, M. Rajeevan, and T. Yamagata, 2016: Anatomy of Indian heatwaves. *Sci. Rep.*, **6**, 24395, <https://doi.org/10.1038/srep24395>.
- Schneider, T., T. Bischoff, and H. Plotka, 2015: Physics of changes in synoptic midlatitude temperature variability. *J. Climate*, **28**, 2312–2331, <https://doi.org/10.1175/JCLI-D-14-00632.1>.
- Schubert, S., H. Wang, and M. Suarez, 2011: Warm season sub-seasonal variability and climate extremes in the Northern Hemisphere: The role of stationary Rossby waves. *J. Climate*, **24**, 4773–4792, <https://doi.org/10.1175/JCLI-D-10-05035.1>.
- Simmons, A. J., and B. J. Hoskins, 1979: The downstream and upstream development of unstable baroclinic waves. *J. Atmos. Sci.*, **36**, 1239–1254, [https://doi.org/10.1175/1520-0469\(1979\)036<1239:TDAUDO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1979)036<1239:TDAUDO>2.0.CO;2).
- Souders, M. B., B. A. Colle, and E. K. M. Chang, 2014a: The climatology and characteristics of Rossby wave packets using a feature-based tracking technique. *Mon. Wea. Rev.*, **142**, 3528–3548, <https://doi.org/10.1175/MWR-D-13-00371.1>.
- , —, and —, 2014b: A description and evaluation of an automated approach for feature-based tracking of Rossby wave packets. *Mon. Wea. Rev.*, **142**, 3505–3527, <https://doi.org/10.1175/MWR-D-13-00317.1>.
- Swanson, K. L., and R. T. Pierrehumbert, 1997: Lower-tropospheric heat transport in the Pacific storm track. *J. Atmos. Sci.*, **54**, 1533–1543, [https://doi.org/10.1175/1520-0469\(1997\)054<1533:LTHTTT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1997)054<1533:LTHTTT>2.0.CO;2).
- Valsangkar, A. A., J. M. Monteiro, V. Narayanan, I. Hotz, and V. Natarajan, 2019: An exploratory framework for cyclone identification and tracking. *IEEE Trans. Vis. Comput. Graph.*, **25**, 1460–1473, <https://doi.org/10.1109/TVCG.2018.2810068>.
- Wallace, J. M., G.-H. Lim, and M. L. Blackmon, 1988: Relationship between cyclone tracks, anticyclone tracks and baroclinic waveguides. *J. Atmos. Sci.*, **45**, 439–462, [https://doi.org/10.1175/1520-0469\(1988\)045<0439:RBCTAT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1988)045<0439:RBCTAT>2.0.CO;2).
- Widanagamaachchi, W., A. Jacques, B. Wang, E. Crosman, P.-T. Bremer, V. Pascucci, and J. Horel, 2017: Exploring the evolution of pressure-perturbations to understand atmospheric phenomena. *2017 IEEE Pacific Visualization Symp. (PacificVis)*, Seoul, South Korea, IEEE, 101–110, <https://doi.org/10.1109/PACIFICVIS.2017.8031584>.
- Wirth, V., M. Riemer, E. K. M. Chang, and O. Martius, 2018: Rossby wave packets on the midlatitude waveguide—A review. *Mon. Wea. Rev.*, **146**, 1965–2001, <https://doi.org/10.1175/mwr-d-16-0483.1>.
- Zimin, A. V., I. Szunyogh, D. J. Patil, B. R. Hunt, and E. Ott, 2003: Extracting envelopes of Rossby wave packets. *Mon. Wea. Rev.*, **131**, 1011–1017, [https://doi.org/10.1175/1520-0493\(2003\)131<1011:EEORWP>2.0.CO;2](https://doi.org/10.1175/1520-0493(2003)131<1011:EEORWP>2.0.CO;2).
- , —, B. R. Hunt, and E. Ott, 2006: Extracting envelopes of nonzonally propagating Rossby wave packets. *Mon. Wea. Rev.*, **134**, 1329–1333, <https://doi.org/10.1175/MWR3122.1>.