Stride Search: A general algorithm for storm detection in high resolution climate data

P. A. Bosler[1], E. L. Roesler[2], M. A. Taylor[1], and M. R. Mundt[3]

[1]Center for Computing Research, Sandia National Laboratories, Albuquerque NM, USA
[2]Geophysics and Atmospheric Science, Sandia National Laboratories, Albuquerque NM, USA
[3]Department of Mathematics, University of California Los Angeles, Los Angeles CA, USA

Dear Editorial Staff,

We submit the attached revisions to our Stride Search paper. Based on the referees' comments, we have made the following revisions to [1]:

- The paper was revised to focus on the independence of the Stride Search algorithm from the layout and resolution of the data. We have removed our previous general assertions that grid point searches can fail and the emphasis on the definition of the angular distance function, which the reviewers objected to, and instead focused the writing on what Stride Search provides. The Introduction and Conclusions sections are therefore mostly new.

- Equations (3) through (6) that describe Stride Search were rewritten (in a mathematically equivalent manner) to emphasize this new focus. The code was also rewritten to match the paper.

- In rewriting the code, we found two bugs that did not affect the output of Stride Search, but did affect its run time. Figure 6 and Table 4 have been corrected to match the two bug fixes.

- We now use only one resolution of data (Section 3) to avoid confusion that one of the referees mentioned.

- We added Figure 5 and Table 3 to show the tropical cyclones detected during the entire dataset, as suggested by a referee.

- We updated Figure 9 to show spatial scale and revised the discussion about it to show its importance (one referee suggested removing it).

- Stride Search has completed the copyright assertion process and is now available online as open source software. The Code Availability section has been updated accordingly.

We have also made numerous minor changes throughout the manuscript to address the rest of the reviewers' comments. For a point-by-point discussion of each comment by each reviewer, please see our final response to the Interactive Discussion [2].

Sincerely,

Peter A. Bosler

References

[1] P. A. Bosler, E. L. Roesler, M. A. Taylor, and M. R. Mundt, Stride Search: A general algorithm for storm detection in high resolution climate data, *Geosci. Model Dev. Disc.*, 2015.

[2] P. A. Bosler, E. L. Roesler, M. A. Taylor, and M. R. Mundt, Interactive comment on "Stride Search: A general algorithm for storm detection in high resolution climate data," *Geosci. Model Dev. Disc.*, 8, C3548C3558, 2016.

Stride Search: A general algorithm for storm detection in high resolution climate data

Peter A. Bosler[1], Erika L. Roesler[2], Mark A. Taylor[1], and Miranda R. Mundt[3]

[1]Sandia National Laboratories, Center for Computing Research, PO Box 5800 , Albuquerque NM, 87185-1321, USA
[2]Sandia National Laboratories, Geophysics and Atmospheric Sciences, PO Box 5800 , Albuquerque NM, 87185-0750, USA
[3]University of California Los Angeles, Department of Mathematics, Box 951555, Los Angeles CA, 90095-1555, USA

Correspondence to: Peter Bosler (pabosle@sandia.gov)

Abstract. This article discusses the problem of identifying extreme climate events such as intense storms within large climate data sets. The basic storm detection algorithm is reviewed, which splits the problem into two parts: a spatial search followed by a temporal correlation problem. Two specific implementations of the spatial search algorithm are compared: the commonly used grid point search algorithm is reviewed, and a new algorithm called Stride Search is introduced. The Stride Search algorithm is defined independently of the spatial discretization associated with a particular data set. Results from the two algorithms are compared for the application of tropical cyclone detection, and shown to produce similar results for the same set of storm identification criteria. Differences between the two algorithms arise for some storms due to their different definition of search regions in physical space. The physical space associated with each Stride Search region is constant, regardless of data resolution or latitude, and Stride Search is therefore capable of searching all regions of the globe in the same manner. Stride Search's ability to search high latitudes is demonstrated for the case of polar low detection. Wall clock time required for Stride Search is shown to be smaller than a grid point search of the same data, and the relative speed up associated with Stride Search increases as resolution increases.

1 Introduction

The identification of extreme events in climate data sets is a fundamental objective of many climate scientists. Data sets may be a reanalysis product or a particular model's output, and an extreme event may be any event classified as an important deviation from a subjective normal state – loosely, a "storm." End-users of climate data and model developers alike frequently investigate prevalent

storm tracks, intensity, or formation areas within a given data set, e.g., Williamson (1981); Hodges (1994); Vitart et al. (1997). Annual means and statistical averages regarding the frequency of a particular type of storm in a particular region are another frequent subject of study, e.g., Sinclair (1994); Blender et al. (1997); Raible and Blender (2004); Bracegirdle and Gray (2008); Kleppek et al. (2008). Individual storms' structures are often investigated to evaluate how well a model captures realistic physical features, e.g., Nordeng and Rasmussen (1992); Walsh et al. (2007); Reed and Jablonowski (2011); Føre et al. (2012). Common to all such efforts is the need to search data sets for quantifiable, objective storm *identification criteria*.

Identification criteria are defined as a small set of variables that together give a basic characterization of storm's location, intensity, and size (Williamson, 1981; Hodges, 1994). Each variable is paired with a threshold value used to filter the data into a small number of categories. The particular variables and their appropriate threshold values vary greatly by application, and many studies have proposed and compared different sets of criteria; see Raible et al. (2008) and Neu et al. (2013) for a discussion of extratropical cyclone criteria. Walsh et al. (2007) and Horn et al. (2014) provide similar analyses for tropical cyclones.

The basic storm detection algorithm consists of two stages (Hodges, 1994). First, a spatial search loops over all time steps in a data set and collects detection points where the spatially defined identification criteria are met or exceeded. Second, a temporal correlation procedure correlates detection points across adjacent time steps to construct storm tracks and apply temporally defined identification criteria.

In comparison to the number of studies concerned with identification criteria, literature regarding the analysis of spatial search algorithms is relatively sparse in the climate community. Such a discussion is of heightened importance due the growth of climate data sets in both size and number. As models and reanalysis products increase spatial and temporal resolution, and as ensembles are more commonly used forecasting tools, the need to efficiently and accurately search climate data sets is also increasing. Of equal concern to an algorithm's performance is its ability to produce repeatable, objective analysis of data (Hodges, 1994), regardless of the data layout and resolution. Contemporary models can incorporate advanced features such as variable resolution using unstructured grids (e.g. Zarzycki and Jablonowski (2014)) and frequently employ different representations of the sphere than a traditional latitude-longitude grid (e.g. Putnam and Lin (2007); Neale et al. (2012); Skamarock et al. (2012)). An ideal search algorithm would be agnostic to such details.

In such an ideal world, the choice of search algorithm would not affect the statistics associated with a particular data set. In practice, however, we find that just as a change to the identification criteria of a particular storm type can change the statistics found in a particular data set (Raible et al., 2008; Horn et al., 2014), the way that data set is divided and searched – independently of the identification criteria – can also affect the statistics.

2

One contributing factor is that some of the variables used to identify storms, such as vorticity, have a dependence on the scale of the data (Sinclair, 1997; Walsh et al., 2007). This dependence may also vary with location depending on the layout of the data set. On a uniform latitude-longitude grid, for example, the spatial scale of adjacent grid points varies with latitude. A search algorithm that does not account for this variation may inadvertently allow non-physical artifacts related to data resolution and grid type to influence its output. Researchers may interpolate a dataset to a different type of grid (Sinclair, 1994; Bracegirdle and Gray, 2008), employ a spatial smoothing procedure (Sinclair, 1997), or adjust threshold values as data resolution changes (Walsh et al., 2007) to alleviate some of these problems. We propose an alternative approach that separates the definition of an extreme event from its discrete representation in a data set.

The goal of this work is to provide an algorithm that allows identification criteria to be defined independently of the spatial resolution and layout of the data. The "Stride Search" algorithm facilitates searching data given on general unstructured grids as well as uniform latitude-longitude grids without alteration, and provides improved performance over the commonly used grid point search algorithm. Additionally, Stride Search treats all regions of the globe in the same manner, which allows users to search all latitudes including the poles efficiently. By decoupling the choice of identification criteria from the resolution and layout of the data, we aim to provide a robust objective search algorithm.

2 Storm detection algorithms

Basic descriptions of the two-stage storm detection algorithm are given as Algorithm 1, the spatial search, and Algorithm 2, the temporal correlation procedure. The majority of this work focuses on the spatial search strategy used to define Algorithm 1, which requires the most computational effort (Prabhat et al., 2012). Its input is a search domain and a set of per time step storm identification criteria, as well as the data. A key step in the algorithm is the division of the search domain into a set of search sectors (line 4); we will discuss this process in more detail in the following sections. For each file and each time step, the algorithm compares the data within each sector to the storm identification criteria. If the criteria are met, a storm is recorded to the list L at the current time step.

Identification criteria, particularly those used with a spatial search, are highly application dependent. Ideally storm detection software should be flexible enough to allow users to easily define identification criteria relevant to their area of study and should not be limited to any specific geographic region. In other words, users should be able to easily modify the implementation of Algorithm 1, line 7, in code.

The output of the spatial search (Algorithm 1) is a list of candidate storms. This list may contain false positives due to noisy data, topographic effects, or ambiguity within the identification criteria. The second step of storm detection and tracking, the temporal correlation problem, handles these

3

Algorithm 1 Spatial search algorithm

Input: search domain R, storm identification criteria (spatial), data

Output: List of potential storms, L, organized by increasing time step

  1: set L = empty list

  2: **for all** files in data set **do**

  3:      n_T = number of time steps per file

  4:      Divide search domain into search sectors

  5:      **for** $i = 1$ **to** n_T **do**

  6:         **for all** search sectors at time step i **do**

  7:            **if** sector meets or exceeds identification criteria **then**

  8:               set l_i = storm data at time step i

  9:               add l_i to L

10:            **end if**

11:         **end for**

12:      **end for**

13: **end for**

---

issues. The temporal correlation algorithm's task is to identify the same storm across adjacent time steps. It does so by building storm tracks and is outlined by Algorithm 2. Users define a maximum travel speed $U_{\max}$ appropriate to the type of storm under investigation. The algorithm uses that speed to define $D_{\max} = U_{\max} \cdot \Delta t$, the maximum possible distance a storm may travel per time step. Beginning with a storm entry in the spatial search output list L at time step i, the algorithm searches all storms detected at time step $i+1$. Any storms at time step $i+1$ separated by a distance less than $D_{\max}$ from the storm at time step i are marked as candidate successors. If zero candidates are found, the track is ended at time step i. If one candidate is found, that candidate is linked to its predecessor at time step i and the algorithm continues to build the track by looking for candidates at time step $i+2$. If several candidates are found at time step $i+1$, the algorithm chooses the closest candidate to the entry at time step i and disregards the others. Track building proceeds until either zero candidates are found at the next time step or the data are exhausted.

The tracking algorithm (Algorithm 2) is sensitive to its two input parameters $U_{\max}$ and $t_{\min}$. Choosing a value of $t_{\min}$ too low may not eliminate enough false positives, but choosing a value of $t_{\min}$ too high may eliminate weak cyclones that did exist but did not intensify enough to last long. The effects of choosing a too high or too low value for $U_{\max}$ are more subtle. A value of $U_{\max}$ that is too low can cause storm tracks to fragment into several disjoint pieces. By contrast, choosing an unrealistically high value for $U_{\max}$ could cause the tracking algorithm to merge two storms that are in reality separate entities.

Storm tracks provide a natural mechanism to count storms and to dismiss false positives. Tracks that consist of only one point, indicating a storm whose duration was only 1 time step may be dis-

4

Algorithm 2 Temporal correlation algorithm

Input: List L of potential storms from spatial search output, max storm travel speed $U_{\max}$, minimum duration
    $t_{\min}$, other storm identification criteria (temporal)

Output: List T of storm tracks

 1: set T = empty list

 2: N_T = total number of time steps in data set

 3: **for** $i = 1$ **to** N_T **do**

 4:    **for all** elements $l_i \in L$ at time step i **do**

 5:       start new track t at l_i

 6:       continue = `True`

 7:       $j = i$

 8:       **while** continue **do**

 9:          examine all $l_{j+1} \in L$ at time step $j + 1$ for possible successors to storm l_j

10:          **if** successor found **then**

11:             add l_{j+1} to track t

12:             $j = j + 1$

13:          **else**

14:             continue = `False`

15:          **end if**

16:       **end while**

17:       **if** track t meets or exceeds identification criteria **then**

18:          add t to T

19:       **end if**

20:    **end for**

21: **end for**

---

missed as noise. Tracks that persist for many time steps but do not move may possibly be regarded as topographic effects, particularly if the identification criteria use data that are sensitive to topography, such as geopotential height surfaces.

Storm tracks also provide a straightforward method of applying additional identification criteria. A study concerned with identifying regions of cyclogenesis may reject any storms that do not intensify along their track. Temporal criteria may also be used to perform more detailed classification of candidate storms. For example, a vorticity criterion or a vertical wind speed criterion may detect strong convection due to thunderstorms. To make a distinction between typical summertime afternoon thunderstorms and more persistent mesoscale convective complexes, a temporal criterion may be used to neglect storms that do not persist for longer than 12 hours.

5

2.1 Grid point search

Most storm detection studies of climate data use software that implements Algorithm 1 as a grid point
search. The Geophysical Fluid Dynamics Laboratory's TSTORMS (1997) software, for example,
has become the prevalent tool for tropical cyclone detection in high resolution climate data (Vitart
et al., 1997; Vitart and Stockdale, 2001; Knutson et al., 2008; Zhao et al., 2009; Prabhat et al., 2012;
Zarzycki and Jablonowski, 2014). Grid point searches have also been employed by other studies of
mid-latitude extratropical cyclones (Blender et al., 1997; Geng and Sugi, 2001; Wernli and Schwierz,
2006; Kleppek et al., 2008; Raible et al., 2008).

Grid point searches are designed for the common case where data are given on a uniform latitude-
longitude grid with resolution $\Delta\lambda$ (in radians) so that grid points (λ_j, θ_i) are located at

$$\lambda_j = j\Delta\lambda, \quad j = 0, \ldots, n_{lon} - 1, \tag{1a}$$

$$\theta_i = -\frac{\pi}{2} + i\Delta\lambda, \quad i = 0, \ldots, n_{lat} - 1, \tag{1b}$$

where λ is longitude, θ is latitude, $n_{lon} = 2\pi/\Delta\lambda$ is the number of longitudinal grid points, and
$n_{lat} = n_{lon}/2 + 1$ is the number of latitude grid points.

In a grid point search algorithm, each grid point (λ_j, θ_i) in the search domain is a search sector
center. Sector K_{ij} centered at grid point (λ_j, θ_i) is defined as

$$K_{ij} = \{(\lambda_{j\pm k}, \theta_{i\pm l})\} \qquad k, l = 0, \ldots, n - 1, \tag{2}$$

where n is a user-specified parameter that corresponds to the scale of a storm in latitude-longitude
space. Thus, each sector is a $(2n + 1) \times (2n + 1)$ square in grid point space.

To set up a grid point search, users define the search domain by defining a minimum and maximum
latitude, $\theta_{\min}$ and $\theta_{\max}$, and a maximum and minimum longitude $\lambda_{\min}$ and $\lambda_{\max}$. In this work we
assume $\lambda_{\min} = 0$ and $\lambda_{\max} = 2\pi$, while $\theta_{\min}$ and $\theta_{\max}$ can vary by application. Users must also
select a value for n that relates the spatial scale of the storms they wish to detect to the resolution of
the data $\Delta\lambda$. Figure 1(a) shows grid point search sectors along the equator with $n = 2$ for data with
resolution $\Delta\lambda = 10°$. The sectors are 5×5 boxes in grid point space and span approximately 5600
km $\times$ 5600 km on the Earth.

For each sector, the software collects data from the $(2n + 1) \times (2n + 1)$ points centered at (λ_j, θ_i).
The collected data are compared against the storm identification criteria. If the criteria are met or
exceeded in the sector, the algorithm checks if (λ_j, θ_i) is the location of the storm within that sector.
If so, the algorithm records the storm to its output list. If not, the algorithm cycles to the next grid
point, say $(\lambda_{j+1}, \theta_i)$, and begins again. In Figure 1(a) the blue, horizontally striped sector is cen-
tered at $(\lambda_j, \theta_i) = (150E, 0N)$ and the next two consecutive sectors are shown by the red, vertically
striped sector centered at $(\lambda_{j+1}, \theta_i) = (160E, 0N)$ and black, diagonally striped sector whose center
is $(\lambda_{j+2}, \theta_i) = (170E, 0N)$.
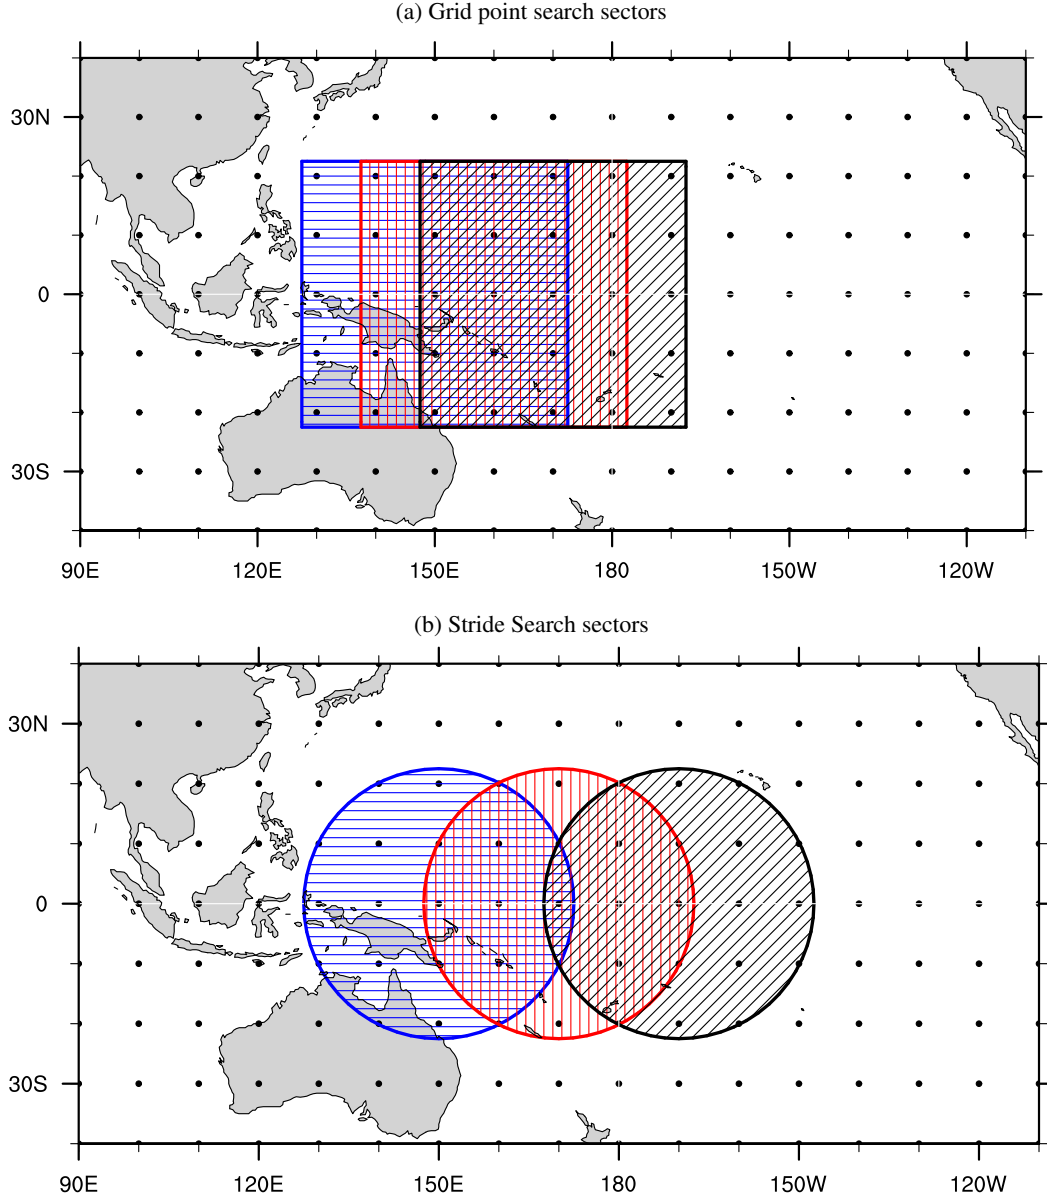
Figure 1. Adjacent search sectors along the equator. Black dots represent data points with resolution $\Delta\lambda = 10°$. Blue sector (horizontal striping) center is $\lambda = 150E$, $\theta = 0N$. Red (vertical striping) and black (diagonal striping) sectors are the next two consecutive searches; (a) grid point search, $n = 2$; (b) Stride Search, $s = 2220$ km.

Centering a sector at each grid point in the search domain yields a robust algorithm. It ensures that the entire search domain will be covered and that the same storm will not be recorded twice. Even though a single storm may trigger the identification criteria in several sectors, only the sector whose center corresponds to the storm location will be recorded to output. While the robustness of the grid point search algorithm is an advantage, it comes at the cost of redundant work. The same data points are accessed multiple times because the algorithm only advances one grid point at a time and the overlap of adjacent sectors is considerable.

The data access required by a grid point search and the overlap of adjacent sectors is also illustrated by Figure 1(a). All three sectors read the data from grid points in the region $\{(\lambda, \theta) : 150E \leq \lambda \leq 170E, 20S \leq \theta \leq 20N\}$. For visual clarity we have not plotted the sectors at $(\lambda_{j-1}, \theta_i)$ or $(\lambda_j, \theta_{i\pm 1})$, which would also overlap a majority of the same grid points.

2.2 Stride Search

Instead of squares in grid point space, Stride Search sectors are circles on the surface of an Earth-sized sphere. The sectors are defined using the geodesic distance function

$$\text{dist}_G\big((\lambda_1, \theta_1), (\lambda_2, \theta_2)\big) = a \arccos\big(\cos\theta_1 \cos\theta_2 \cos(\lambda_2 - \lambda_1) + \sin\theta_1 \sin\theta_2\big), \tag{3}$$

where a is the radius of the Earth. Users select an application-dependent spatial scale s in units of distance such that a maximum of one storm can be located within any spherical circle of radius s. The search domain is divided into a collection of circles on the sphere, each with the same geodesic radius s. Stride Search sectors, illustrated in Figure 1(b) for $s = 2220$ km, are defined by the following procedure.

The number $S_{lat} = s/a$ defines the arc length corresponding to the user-specified scale s. We refer to S_{lat} as the *latitude stride* and use it to define lines of constant latitude

$$\theta_I = \theta_{min} + S_{lat}I, \quad I = 0, \ldots, N_\theta, \tag{4}$$

where $N_\theta = \lfloor (\theta_{max} - \theta_{min})/S_{lat} \rfloor + 1$. The set θ_I divides the search domain into latitudinal strips of width $\approx s$. We also define a *longitude stride* for each θ_I,

$$S_{lon}^{(I)} = \min\left(\frac{s}{a\cos\theta_I}, 2\pi\right), \quad I = 0, \ldots, N_\theta, \tag{5}$$

so that $S_{lon}^{(I)}$ are the arc lengths along each latitude circle θ_I that approximately span a geodesic distance s in the longitudinal (zonal) direction. The minimum function in (5) accounts for the case where θ_I is either pole. The longitude stride defines points λ_{IJ} along each latitude line θ_I, where $J = (j-1)S_{lon}^{(I)}$ for $j = 1, \ldots, N_\lambda^{(I)}$, creating $N_\lambda^{(I)} = \lfloor 2\pi/S_{lon}^{(I)} \rfloor$ longitude points along each θ_I.

Each point (λ_{IJ}, θ_I) defines the center of search sector K_{IJ}, where K_{IJ} is the set of all points on the sphere lying within a distance s of (λ_{IJ}, θ_I),

$$K_{IJ} = \big\{(\lambda, \theta) : \text{dist}_G\big((\lambda, \theta), (\lambda_{IJ}, \theta_I)\big) \leq s\big\}. \tag{6}$$

8

190   We note that the definition of the Stride Search sectors is determined entirely by the application-related spatial scale s and is therefore independent of resolution of the data, $\Delta\lambda$. By construction, sector K_{IJ} overlaps its neighbors $K_{I\pm1,J}$ and $K_{I,J\pm1}$ by approximately one radius s in physical space. This is a sufficient condition to ensure that the entire search domain will be covered by the circular search sectors.

195   Figure 1(b) shows three consecutive circular sectors with $s = 2220$ km. This value of s corresponds to an arc length of $\approx 20°$, and was chosen to match the scale of the sectors in Figure 1(a). The blue, horizontally striped circle is centered at $(\lambda_{IJ}, \theta_I) = $ (150E, 0N). The red, vertically striped sector is centered at $(\lambda_{I,J+1}, \theta_I) = $ (170E, 0N) and the center of the black, diagonally striped sector is located at $(\lambda_{I,J+2}, \theta_I) = $ (170W, 0N). Since each Stride Search sector is separated from its imme-

200   diate neighbors by either a latitude stride or a longitude stride, the overlap between adjacent sectors is much reduced compared to a grid point search. The Stride Search algorithm therefore covers a much larger geographic area with the same number of search sectors.

Stride Search setup is completed by defining the sectors K_{IJ} in terms the available data. Computationally, this involves creating a class and/or methods that identify and link each sector to the

205   data points enclosed by its geographic boundary. Sectors used with high resolution data sets automatically link to more grid points than the same-sized sectors used with lower resolution data. For uniform longitude-latitude grids, Equations (1) apply and the process is straightforward. For unstructured grids the mesh's connectivity information may include a node adjacency list or topological data structures such as edges and faces. Any of these may be used to determine a sector's enclosed data

210   points. In the absence of such connectivity information a *kd*-tree algorithm (e.g., Samet (2006)) may be used.

Reduced overlap between adjacent sectors leads to improved performance by decreasing the number of redundant data accesses and by reducing the total number of sectors. However, it also creates a new issue. Since each sector is searched independently, several sectors may detect and record the

215   same storm to the linked-list, as illustrated by Figure 2. Before the storms detected by Stride Search can be saved to output, duplicate entries must be removed.

Duplicates are removed by again referencing the user-specified scale s. Each pair of entries in the linked-list are compared; if a pair are separated by a distance less than s, they are considered duplicates and the less intense entry is deleted. In Figure 2, this is demonstrated with pressure data.

220   Each of the three search sectors have exceeded the storm identification criteria and independently locate their minimum pressure. The blue (left) sector finds a minimum of 988 hPa, the red (middle) finds 982 hPa, and the black (right) finds 984 hPa. These three entries are clearly separated by a distance less than s, as they are all contained within the red (middle) circle. The duplicate removal procedure will delete the blue (left) and black (right) entries because compared to the red (middle)

225   entry they have higher pressures and are less intense. Only the red 982 hPa entry will be saved to output. In general the list of detected storms at a particular time step, duplicates included, is much
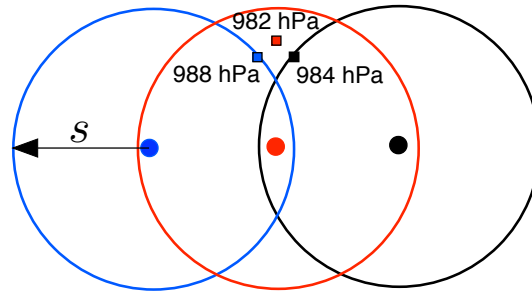
Figure 2. Duplicate detections of the same storm using the Stride Search algorithm. Each sector (blue/left, red/middle, and black/right) locates its minimum sea level pressure (corresponding squares). Duplicate entries are removed prior to the output stage; only the red 982 hPa entry will remain.

smaller than the spatial size of the data and the time required by the duplicate removal procedure is negligible.

3 Data description

230   To demonstrate and test Stride Search we use data produced by the the spectral element dynamical core of the Community Atmosphere Model, CAM-SE (Neale et al., 2012; Dennis et al., 2012). The model uses a cubed sphere grid and high resolution experiments set 240 elements per face of the cube for a total of 3,110,402 horizontal grid points. This results in a horizontal resolution of $\Delta\lambda \approx 0.125°$ (Worley et al., 2011; Dennis et al., 2012). Due to well-known issues regarding the tuning of physical

235   parameterizations within climate models, this high resolution simulation may produce high intensity storms with unrealistically high frequencies (Reed and Jablonowski, 2011; Dennis et al., 2012).

    The original goals of the high resolution experiments of Worley et al. (2011) and Dennis et al. (2012) were to demonstrate the parallel scaling of CAM-SE, to document its required run time and related statistics in various high performance computing environments, and to demonstrate the

240   model's capability to produce well-resolved features like tropical cyclones that cannot be represented well in low resolution experiments. Here, we choose this data because its high resolution ensures that small-scale storms will exist, which provides a good testing environment for storm detection algorithms. The fact that there may be an unrealistically high number of storms in the data is a benefit in this case.

245   The data set contains 5 years of simulated data that used CAM5 physics and pre-industrial (year 1850) initial conditions. Instead of additional model components, the CAM-SE atmospheric dynamical core is coupled to a set of land, ocean, and sea ice data that also correspond to the year 1850 to provide its boundary conditions (Dennis et al., 2012). The land, ocean, and sea ice boundary conditions are periodic, with period 1 year, and simply repeat throughout the 5-year atmospheric

250   simulation.

10

The model's native cubed sphere data were interpolated to a uniform latitude-longitude grid with $n_{lon} = 1024$ for a resolution of $\Delta\lambda = 0.35°$ using the regridding software provided by the Earth System Modeling Framework (Balaji et al., 2014). To facilitate a timing experiment (presented in the next section), we also interpolate 3 months of data to resolutions of $\Delta\lambda = 2°, 1°, 0.5°$ and $0.25°$.

4 Tropical cyclones

In this section we apply both Stride Search and TSTORMS to the problem of tropical cyclone detection. Our aim is to validate Stride Search by comparison with TSTORMS, which has a proven record. We also discuss the subtle differences between the two algorithms that lead to differences in their final results that may be of importance to climate researchers.

While there are many different combinations of variables available to define a tropical cyclone (Walsh et al., 2007; Horn et al., 2014), to provide both codes with a common set of identification criteria we choose the TSTORMS default. A tropical cyclone is identified within search sector K_{ij} if the following four criteria are met (Vitart et al., 1997):

1. There is a cyclonic vorticity maximum greater than a threshold value, τ_ζ:

$$\max_{i,j \in K_{ij}} \left[\mathrm{sgn}(\theta_i) \cdot \zeta_{850}(\lambda_j, \theta_i) \right] > \tau_\zeta, \tag{7}$$

where ζ_{850} is the relative vorticity at the 850 hPa level.

2. The distance between the cyclonic vorticity maximum and the sector's sea level pressure minimum is less than a threshold value τ_{D_1} :

$$\mathrm{dist}\left((\lambda_\zeta, \theta_\zeta), (\lambda_P, \theta_P) \right) < \tau_{D_1}, \tag{8}$$

where $(\lambda_\zeta, \theta_\zeta)$ and (λ_P, θ_P) are the locations of sector K_{ij}'s vorticity maximum and sea level pressure minimum, respectively.

3. The difference between the vertically averaged temperature's maximum value and its sector average exceeds a threshold τ_T:

$$\max_{i,j \in K_{ij}} \overline{T}(\lambda_j, \theta_i) - \mathrm{Avg}_{K_{ij}}(\overline{T}) > \tau_T, \tag{9}$$

where $\overline{T}$ is defined as

$$\overline{T}(\lambda, \theta) = \frac{1}{2}\left(T_{500}(\lambda, \theta) + T_{200}(\lambda, \theta) \right), \tag{10}$$

and T_{500} and T_{200} are the temperatures the 500 hPa and 200 hPa pressure levels, respectively. To maintain consistency between both codes, the sector average $\mathrm{Avg}_{K_{ij}}(\overline{T})$ is approximated as a simple arithmetic average,

$$\mathrm{Avg}_{K_{ij}}(\overline{T}) = \frac{1}{N_{K_{ij}}} \sum \{\overline{T}(\lambda_j, \theta_i) : (\lambda_j, \theta_i) \in K_{ij}\}, \tag{11}$$

11

	τ_ζ	τ_{D_1}	τ_T	τ_{D_2}
TSTORMS	8.5×10^{-4} s^{-1}	4°	2 K	2°
Stride Search	8.5×10^{-4} s^{-1}	450 km	2 K	225 km

Table 1. Threshold values used for tropical cyclone detection.

where $N_{K_{ij}}$ is the number of data points in sector K_{ij}.

4. The distance between the maximum vertically averaged temperature and the sea level pressure minimum is less than a threshold value τ_{D_2}:

$$\mathrm{dist}\left((\lambda_T, \theta_T), (\lambda_P, \theta_P)\right) < \tau_{D_2}, \tag{12}$$

where (λ_T, θ_T) is the location of the sector maximum of $\overline{T}$.

Differences between the two algorithms' detections arise due to the differences in the algorithms themselves. For computing the collocation criteria, Equations (8) and (12), TSTORMS uses the angular distance function

$$\mathrm{dist}_A\left((\lambda_1, \theta_1), (\lambda_2, \theta_2)\right) = \sqrt{(\lambda_2 - \lambda_1)^2 + (\theta_2 - \theta_1)^2}. \tag{13}$$

For TSTORMS, whose intended application is in tropical regions, this is a simple and effective strategy because angular distance is a reasonable proxy for geodesic distance near the equator. Stride Search uses the geodesic distance function (3). Users of TSTORMS must specify τ_{D_1} and τ_{D_2} in angular units, while users of Stride Search must use units of length. The arithmetic averages of the vertically averaged temperature (11) will be different from one algorithm to the other, because their sectors will contain a different number of data points. As a result criteria 2, 3, and 4 may behave differently for each algorithm.

4.1 Spatial search results

We apply both algorithms to the data described in Section 3. We set TSTORMS $n = 12$, Stride Search $s = 450$ km, and use the threshold values shown in Table 1 for Equations (7), (8), (9), and (12).

Results from an arbitrarily chosen 3 months of data, July 18 to October 18 of simulation year 4, are plotted in Figure 3. Each dot represents a storm detected at one time step. All 6-hourly time steps over the entire three months are shown, colored by the windspeed-dependent hurricane categories defined by the Saffir-Simpson intensity scale.

Both algorithms produce qualitatively similar results. Visually they appear to agree nearly perfectly on the identifiable storm tracks and intensities. They both have false positives over land and in the Southern Ocean. Stride Search produces more false positives, particularly in the Southern Ocean, than TSTORMS. This is due to the fact that Stride Search sectors – especially at higher latitudes –

12
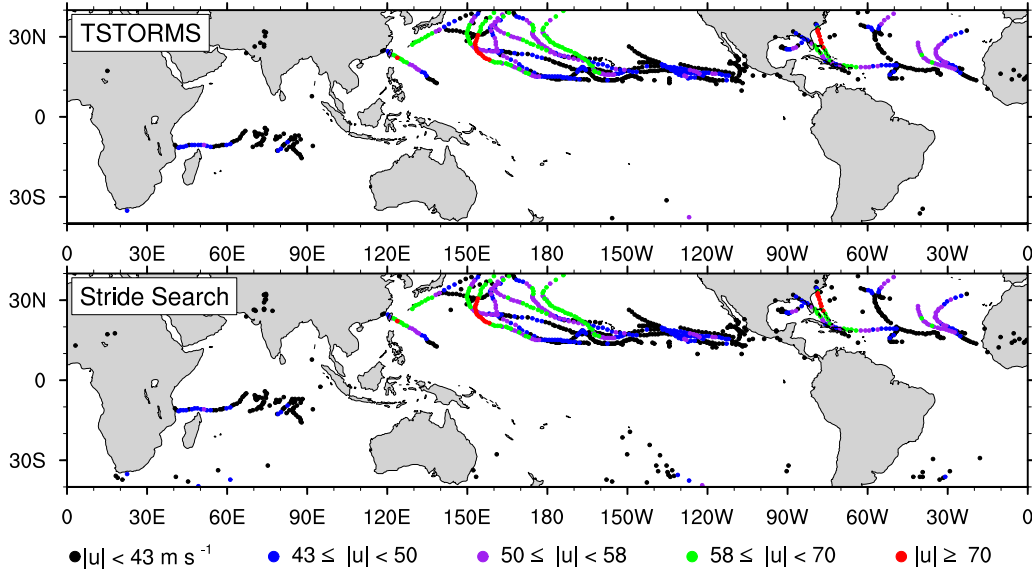
Figure 3. Spatial search results, northern hemisphere summer simulation year 4. Each dot is a storm detected at one time step; colors correspond to the categories of the Saffir-Simpson hurricane scale.

typically contain more data points than TSTORMS sectors. The larger number of points per sec-
tor reduces the sector average of the vertically averaged temperature, Equation (11), compared to a
TSTORMS sector at the same location. Thus, the warm core temperature excess criterion (9) is more
easily achieved using Stride Search than TSTORMS for the same value of τ_T. For both codes, false
positives are eliminated by the temporal correlation algorithm discussed in the next section. How-
ever, the consequences of this different behavior of criterion (9) between the codes will propagate
into the storm tracks and final output.

4.2 Temporal correlation results

In this section we apply the temporal correlation algorithm, Algorithm 2, to the spatial search results.
Since tropical cyclones are inherently maritime events (Cotton and Anthes, 1989), at this stage we
also apply a land mask to remove any tracks whose origins are not over water.

In Figure 4 we show the storm tracks that correspond to the spatial search output of Figure 3 with
parameters $U_{\max} = 15 \text{ m} \cdot \text{s}^{-1}$ and $t_{\min} = 2$ days. These results show that the temporal correlation
algorithm succeeds in eliminating false positives and gives a better representation of the storms
within the data set than the raw output from the spatial search algorithm. Table 2 presents the final
storm counts for each algorithm for the three month data set separated by hurricane category. Again,
both algorithms produce nearly identical results which validates the present work.

Table 2 shows that Stride Search detects two fewer category 1 storms and one additional category
2 storm than TSTORMS. Looking for differences between the panels of Figure 4, we see that the

13
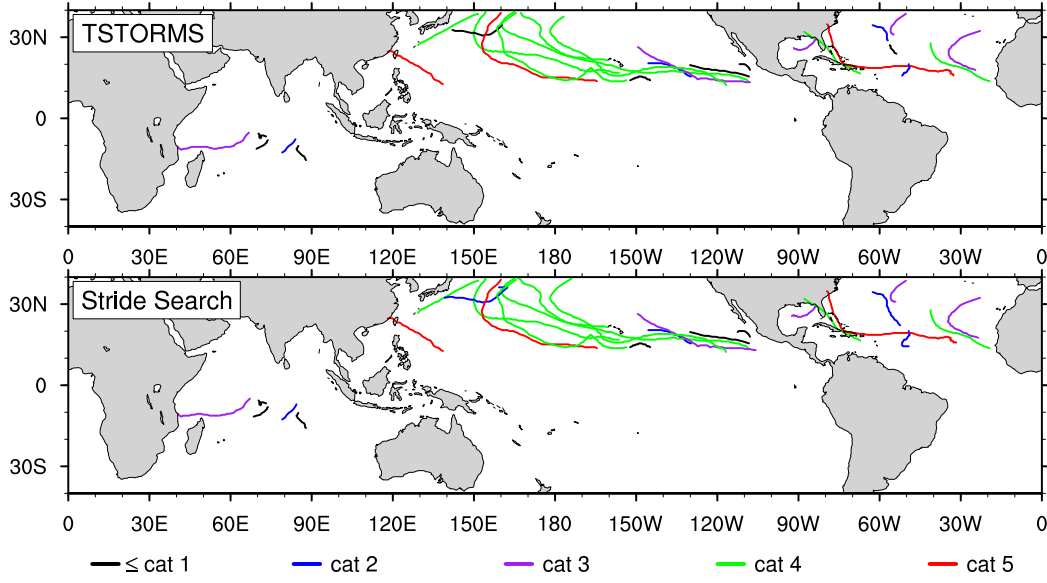
Figure 4. Storm tracks, northern hemisphere summer, simulation year 4. Each track is colored by the hurricane category corresponding to the maximum wind speed achieved during its lifetime.

	cat 0	cat 1	cat 2	cat 3	cat 4	cat 5	total
max. wind (m·s^{-1})	$\|u\| < 33$	$33 \leq \|u\| < 43$	$43 \leq \|u\| < 50$	$50 \leq \|u\| < 58$	$58 \leq \|u\| < 70$	$\|u\| \geq 70$	
TSTORMS	1	7	4	5	8	3	28
Stride Search	1	5	5	5	8	3	27

Table 2. Storm count by hurricane category for northern hemisphere summer simulation year 4; this table corresponds to the storm tracks shown in Figure 4.

category 1 storms correspond to a storm in the western Pacific off the east coast of Japan near (150E, 050W) and a storm in the north central Atlantic near (050W, 20N).

Comparing the panels, we see that Stride Search classified the western Pacific storm as category 2, which accounts for two of the three differences between rows in Table 2. Viewing the data, we note that the Stride Search track for this particular storm is 4 timesteps longer than the corresponding TSTORMS track. During its final data points in the Stride Search output its temperature excess was decreasing and very close to the detection threshold. Since the TSTORMS sector average of $\overline{T}$ is higher than the Stride Search sector average of $\overline{T}$, the storm did not pass criterion 3 (Equation 9) at the end of its life cycle in TSTORMS. We also find that this storm only achieved a category 2 wind speed at the very end of its life cycle, after the point in time where TSTORMS had stopped detecting it, which explains why Stride Search counts the storm as a category 2 and TSTORMS does not.

A similar explanation holds for the Atlantic storm. We see that TSTORMS counts the same storm twice, once as a category 1 and once as a category 2, while Stride Search shows only one longer

14
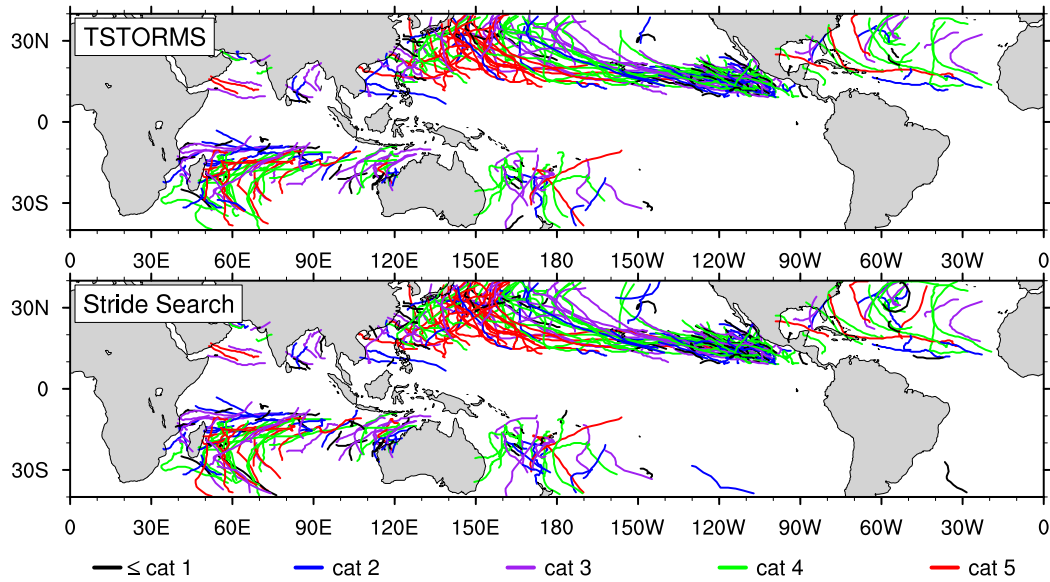
Figure 5. Storm tracks. The ultimate output of each algorithm for the entire 5-year data set. Coloring as in Figure 4.

category 2 track. This is again due to the difference in the computation of the temperature excess between the two codes. This particular storm weakened in its early days to the point where its temperature excess was not sufficient for TSTORMS to detect it before finally intensifying into a category 2 storm. This creates a hole in the TSTORMS track that does not show up in the Stride

345 Search results because the temperature excess criterion is not as strict in Stride Search as it is in TSTORMS.

We point out that this is not a weakness of the TSTORMS software – our choice of $\tau_T = 2\text{K}$ was somewhat aribtrary and choosing a lower threshold value τ_T for TSTORMS would remedy this problem for this particular cyclone. Rather, we stress that these differences arise simply due to

350 the differences in the definition of both algorithm's search sectors. To investigate these differences further we tested Stride Search using a midpoint rule quadrature approximation of the sector average of the average vertical temperature, Equation (11), and found similar results. We therefore chose to use the arithmetic average to keep the tropical cyclone identification criteria the same between the two codes.

355 Unfortunately, differences in specific storm tracks between each algorithm become more difficult to sort by cause as the size of the data set grows. Figure 5 shows the storm tracks produced by each algorithm for the entire data set, using the same identification criteria and threshold values as our previous discussion. Since the temperature excess criterion is more easily achieved by Stride Search, we would expect Stride Search to identify more storms than TSTORMS, particularly in the lower-

360 intensity storm categories. We might also expect TSTORMS to count too many storms because, for

max. wind (m·s^{-1})	cat 0 $\|u\| < 33$	cat 1 $33 \le \|u\| < 43$	cat 2 $43 \le \|u\| < 50$	cat 3 $50 \le \|u\| < 58$	cat 4 $58 \le \|u\| < 70$	cat 5 $\|u\| \ge 70$	total
TSTORMS	4	59	78	87	87	41	356
Stride Search	7	75	82	86	83	43	376

Table 3. Storm counts by hurricane category for the entire 5-year data set; this table corresponds to the storm tracks shown in Figure 5.

these threshold values, some storms may be split into multiple tracks. Both of these predictions may be born out in the data, which is tabulated in Table 3. Stride Search does indeed detect more category 0, 1, and 2 storms. TSTORMS also finds a higher number of high intensity (category ≥ 3) storms than Stride Search, possibly due to track splitting. Without investigating differences in individual tracks, which is impractical for large data sets, one can only state that since the identification criteria used by both codes were identical, the different results can only be due to differences at the level of the algorithms' design.

4.3 Performance and timing

As discussed previously, climate data sets are large and expected to increase in size as climate models run at high resolutions with $\Delta\lambda < 0.5°$. Storm detection algorithms must be both accurate and efficient. In this section we document the dependence of each search algorithm's run time on the resolution of its input data set.

Differences in the structure of the two algorithms result in notable differences in the number of search sectors and the number of times each grid point is accessed in memory per time step. In the top row of Table 4 we present the total number of search sectors required by each algorithm to search the tropical domain, $\theta_{\min} = 40S$, $\theta_{\max} = 40N$ for each data resolution. For TSTORMS this is equivalent to the number of grid points in the domain, hence the number of search sectors increases by a factor of four as the data resolution is halved. The number of Stride Search sectors remains constant across all data resolutions. For data on a uniform latitude-longitude grid, the number of points in a Stride Search sector grows as a function of latitude. The maximum points per sector listed for Stride Search are an upper bound that depends on the search domain, specifically $\theta_{\min}$ and $\theta_{\max}$ and the spatial scale s. For TSTORMS, the max points per sector are a function of the user-specified parameter n and are equal to $(2n+1) \times (2n+1)$.

In the last row of the table the total number of per time step data accesses required by each algorithm are given; these numbers are the product of the number of sectors and the number of points per sector. This number provides an indication of the cost of each algorithm. The smaller number of sectors and reduced overlap between sectors in the Stride Search algorithm result in many fewer data accesses (by orders of magnitude) than the grid point search used by TSTORMS.

16

	TSTORMS				Stride Search			
$\Delta\lambda$	$2°$	$1°$	$0.5°$	$0.25°$	$2°$	$1°$	$0.5°$	$0.25°$
number of sectors	7.29E3	2.92E4	1.17E5	4.66E5	1616	1616	1616	1616
max points per sector	25	49	169	625	49	143	437	1575
max data accesses	1.82E5	1.43E6	1.97E7	2.92E8	7.91E4	2.31E5	7.06E5	2.55E6

Table 4. Numbers of sectors, maximum data points per sector, and total number of data accesses per time step to search tropical domain $\theta \in [40S, 40N]$ vs. data resolution. TSTORMS used $n = \{2, 3, 6, 12\}$ for the corresponding $\Delta\lambda$, and Stride Search used $s = 450$ km.

Both TSTORMS and the current implementation of Stride Search are written in Fortran and run in serial on a single thread. As a timing experiment, we run both codes on three months of data that were interpolated to resolutions of $\Delta\lambda = 2°, 1°, 0.5°, 0.25°$, as described in Section 3. Each data set contains 360 time steps evenly spread across three NetCDF files. We choose $s = 450$ km for Stride Search and set the TSTORMS parameter $n = 2, 3, 6, 12$ for the corresponding $\Delta\lambda$. The total wall clock time required for each algorithm to search each data set, including I/O, is recorded and the average time per time step is computed by dividing the total by 360. We repeat each experiment three times for each resolution and average the results to produce Figure 6.

Figure 6(a) shows the average wall clock time required by each algorithm to search one time step of data as a function of the data resolution $\Delta\lambda$. The plot shows results from experiments run on a standard desktop workstation using GNU's Fortran compiler. The same tests were run on one node of Sandia's Red Sky High Performance Computing cluster using the Intel Fortran compiler and produced similar results. Both algorithms appear to scale at the expected rate of $O(\Delta\lambda^{-2})$ as $\Delta\lambda \to 0$, but for each resolution Stride Search is faster. The speed advantage of Stride Search over TSTORMS improves as resolution increases. Figure 6(b) shows the speedup due to Stride Search, defined as the ratio of the TSTORMS average wall clock time per data time step to the Stride Search average wall clock time per data time step. Stride Search is approximately 15% faster for the $\Delta\lambda = 2°$ data and this ratio increases as $\Delta\lambda \to 0$. Stride Search is 2.5 times faster for the $\Delta\lambda = 0.25°$ data.

5 Polar search

A key motivation for developing Stride Search was to provide a detection algorithm capable of searching all latitudes, including polar regions. The Arctic and Antarctic climates become increasingly frequent subjects of study due to recent significant changes in these environments (Stocker et al., 2013), and a detection algorithm capable of searching data near the poles is necessary. Grid point searches have been used at midlatitudes up to $\approx 60N$ and 60S (König et al., 1993; Raible and Blender, 2004), but users must exercise care when choosing the sector size parameter n at high lati-

17
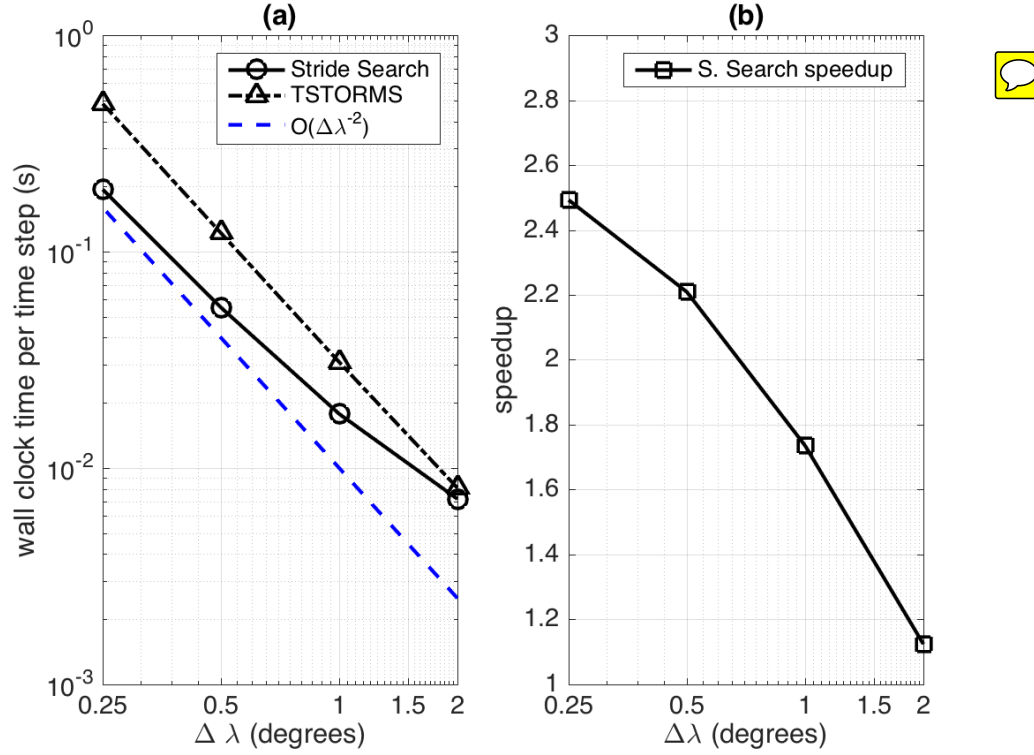
Figure 6. Timing results. (a) Average wall clock time required to search one timestep of data. (b) Speedup due to Stride Search.

tudes. A grid point search near the poles may have to use sectors whose physical size are no longer representative of the physical features of the storms it is meant to locate.

Figure 7(a) shows three consecutive grid point search sectors along $\theta = 60N$. As in Figure 1, the grid of data points has resolution $\Delta\lambda = 10°$ and we have used the same $n = 2$ to set up 5×5 grid point search sectors. The blue (horizontal stripes) sector is centered at $(\lambda_j, \theta_i) = (150E, 60N)$ and the red (vertical stripes) and black (diagonal stripes) sectors are at $(\lambda_{j+1}, \theta_i) = (160E, 60N)$ and $(\lambda_{j+2}, \theta_i) = (170E, 60N)$, respectively. Each grid point search sector spans a distance of 5 grid points in latitude, or approximately 5600 km south to north.

The square grid point search sectors may appear correct in the left plot of Figure 7(a), a Mercator projection, but the problem with them is clear in the polar stereographic projection to the right. The southern boundary of each sector lies along the 40N latitude circle, where 5 grid points in longitude span 4000 km east to west. But the northern boundary of each sector, along 80N, spans only about 1000 km east to west. Users can choose higher values of n at higher latitudes to ensure the sectors will have sufficient zonal extent to capture the desired feature, but the square sectors will still have different spatial scales in the longitude direction compared to the latitude direction. As we have

18

(a) Grid point search sectors

(b) Stride Search sectors
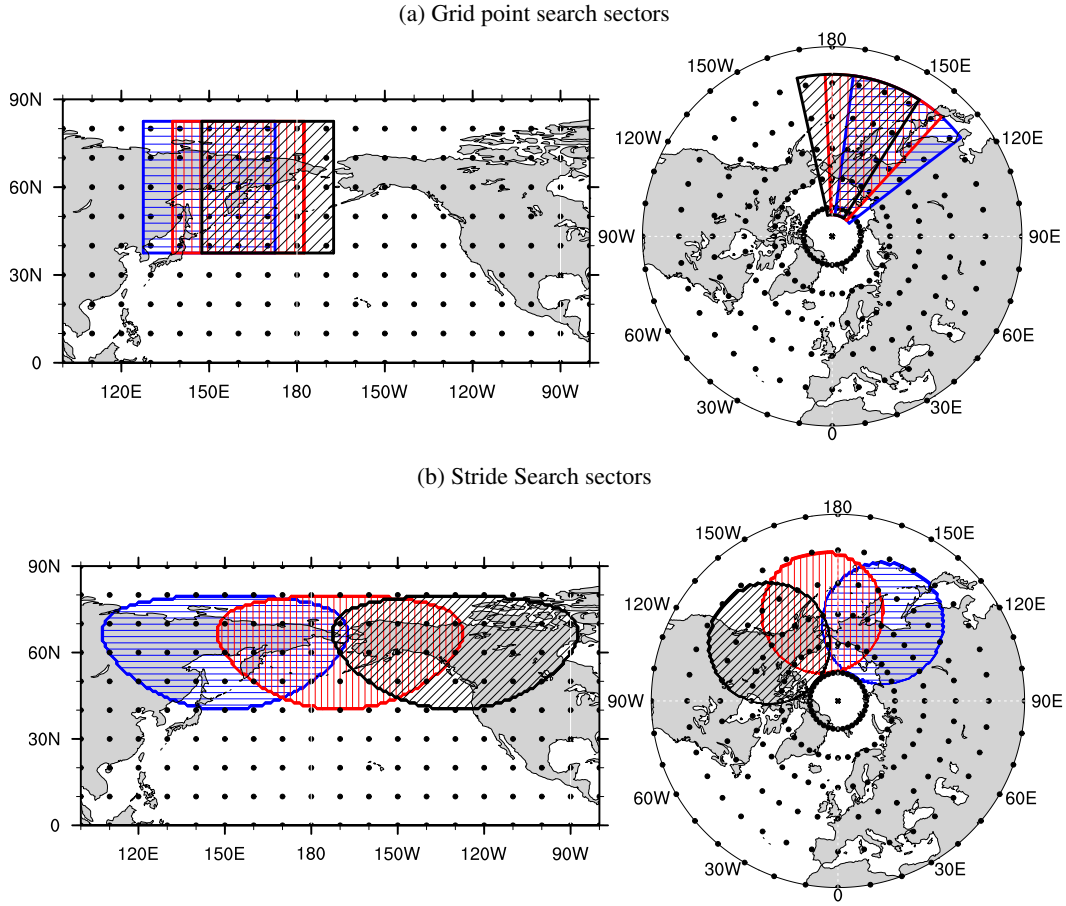
Figure 7. Search sectors along 60N. Black dots represent data points with resolution $\Delta\lambda = 10°$. Blue sector center is $\lambda = 150E$, $\theta = 60N$. Red and black sectors are the next two consecutive searches; (a) TSTORMS, $n = 2$; (b) Stride Search, $s = 2500$ km.

430   seen already in Section 4, the definition of search sector size can impact the final output of a search

algorithm, independently of the identification criteria.

     The constant geodesic radius of Stride Search sectors removes this dependence on the data. Figure 7(b) shows three Stride Search sectors along $\theta = 60N$, with $s = 2500$ km. The blue (horizontally striped) sector is centered at $(\lambda_J, \theta_I) = (150E, 60N)$. The red (vertically striped) and black (diago-

435   nally striped) sectors are at $(\lambda_{j+1}, \theta_i) = (170W, 60N)$ and $(\lambda_{j+2}, \theta_i) = (130W, 60N)$, respectively. The longitude stride along $\theta = 60N$ is twice as large as the longitude stride along the equator, so the three consecutive sectors in Figure 7(b) cover twice as many longitude lines than the three sectors in Figure 1(b). The shapes of each sector in the left plot are due to the effects of the Mercator projection. All sectors are still circles on the sphere, as shown in the polar stereographic projection (right).

440   Stride Search sectors – even one centered at the pole – have the same geographic size regardless

19

of latitude, and are therefore capable of searching polar regions as effectively as midlatitude and tropical regions.

As an example application, we consider polar lows. Polar lows are distinct from mid-latitude low pressure systems due to their different developmental forcing and a typical lack of associated fronts (Montgomery and Farrell, 1992; Rasmussen and Turner, 2003). They contribute to the break-up of sea ice which has significant implications for the polar climate.

Individual polar lows may develop a barotropic structure more similar to tropical cyclones than to baroclinic midlatitude storms (Nordeng and Rasmussen, 1992; Føre et al., 2012). These particularly interesting polar lows are typically smaller in diameter and duration than tropical cyclones, and therefore require a high resolution model such as the one described in Section 3 to resolve. A companion study will investigate the Arctic climatology of variable resolution climate models and their ability to simulate "hurricane-like" polar lows (Roesler et al., In prep.). Here, our goal is to test Stride Search's ability to locate such a small-scale high latitude storm.

A set of objective identification criteria for polar lows are given by Bracegirdle and Gray (2008), which we adapt to the Stride Search algorithm. Storm intensity is measured with vorticity and pressure and, as was the case with the tropical cyclone identification criteria, a collocation requirement is applied. New to this application is a criterion that identifies regions of increased instability associated with cold air outbreaks over relatively warm ocean water.

Stride Search records a polar low in sector K_{ij} if the following criteria are met:

1. A sea level pressure minimum of sufficient intensity exists,

$$\min_{i,j \in K_{ij}} P_{sl}(\lambda_j, \theta_i) < \tau_P, \tag{14}$$

where P_{sl} is sea level pressure and τ_P is the pressure threshold.

2. A cold air outbreak exists,

$$\min_{i,j \in K_{ij}} [\theta_{700}(\lambda_j, \theta_i) - SST(\lambda_j, \theta_i)] \leq \tau_T, \tag{15}$$

where θ_{700} is the potential temperature at the 700 hPa level, SST is the sea surface temperature, and τ_T is the cold air outbreak threshold.

3. A cyclonic vorticity maximum of sufficient strength exists,

$$\max_{i,j \in K_{ij}} [\text{sgn}(\theta_i)\zeta(\lambda_j, \theta_i)] > \tau_\zeta. \tag{16}$$

4. The vorticity maximum must be collocated with the pressure minimum,

$$\text{dist}\left((\lambda_P, \theta_P), (\lambda_\zeta, \theta_\zeta)\right) < \tau_D. \tag{17}$$

Stride Search set up uses a sector radius $s = 500$ km, and search region boundaries $\theta_{\min} = 45$N, and $\theta_{\max} = 90$N. Threshold values are set at $\tau_P = 980$ hPa, $\tau_T = 7$ K, $\tau_\zeta = 2.0 \times 10^{-4}$ s^{-1} and $\tau_D =$

200 km. To the temporal correlation algorithm we add a minimum duration $t_{\mathrm{min}} = 12$ hours and set $U_{\mathrm{max}} = 20 \text{ m} \cdot \text{s}^{-1}$.

The data include only temperature (not potential temperature) and do not include the 700 hPa pressure level. The required θ_{700} data are approximated as

$$\theta_{700} = \frac{1}{2} \left(\theta_{850} + \theta_{500} \right), \tag{18}$$

where $\theta_{850} = T_{850} \left(\frac{1000}{850} \right)^{0.286}$ and $\theta_{500} = T_{500} \left(\frac{1000}{500} \right)^{0.286}$.

Results from the entire 5-year data set are presented in Figure 8, separated by season. Storm tracks are colored by their maximum strength on the U.S. National Weather Service's maritime warning scale (Bowditch, 2002); gale force storms (black) have maximum wind speeds $17.5 \leq u_{\mathrm{max}} < 24.5$ $\text{m} \cdot \text{s}^{-1}$, storm force (blue) has $24.5 \leq u_{max} < 33 \text{ m} \cdot \text{s}^{-1}$, and hurricane force storms (red) have maximum wind speeds greater than $33 \text{ m} \cdot \text{s}^{-1}$. The results show the expected seasonal variation of storm frequencies, with the maximum number of storms and the maximum intensity of storms occurring in the winter (DJF) months. Spring (MAM) months show more activity over the pole than the fall (SON) months, and there are few storms in the summer (JJA).

Once storm tracks are built users may investigate individual storms more easily. To fulfill our goal of finding a "hurricane-like" polar low near the pole we search the storm tracks for storms that get within 5 degrees of the pole, then plot the vorticity associated with each storm. Since the size of the storm track list is much smaller than the size of the dataset, we quickly find the polar low shown in Figure 9, from 12:00 UTC December 25, simulation year 3. At the plotted time step the storm is located at (017.0E, 86.5N). Spatial scale is illustrated by the 1000 km line segment. The plot shows the 850 hPa relative vorticity of the storm and we note how similar its structure is to the typical tropical cyclone. The diameter of the storm's core is approximately 200 km and the diameter of the whole storm, including its vorticity bands, is approximately 500 km. Due to its small spatial scale this storm would not be resolved in low resolution data. Additionally, it is located very close to the pole. This storm therefore demonstrates the capability of Stride Search to find specific features in high resolution data, even in polar regions.

6 Conclusions

We have introduced the Stride Search algorithm for detection of extreme events within climate data sets. The algorithm is defined independently of a data set's layout and resolution, and depends only on the spatial scale associated with a user's intended application. Stride Search was designed to be a flexible algorithm, capable of searching data sets for a variety of extreme events while treating physical space the same for different data sets. Extreme events must be described by a set of quantifiable identification criteria and a representative spatial scale. As examples we have shown detections of two different varieties of cyclonic storms: tropical cyclones and polar lows. The capability to search
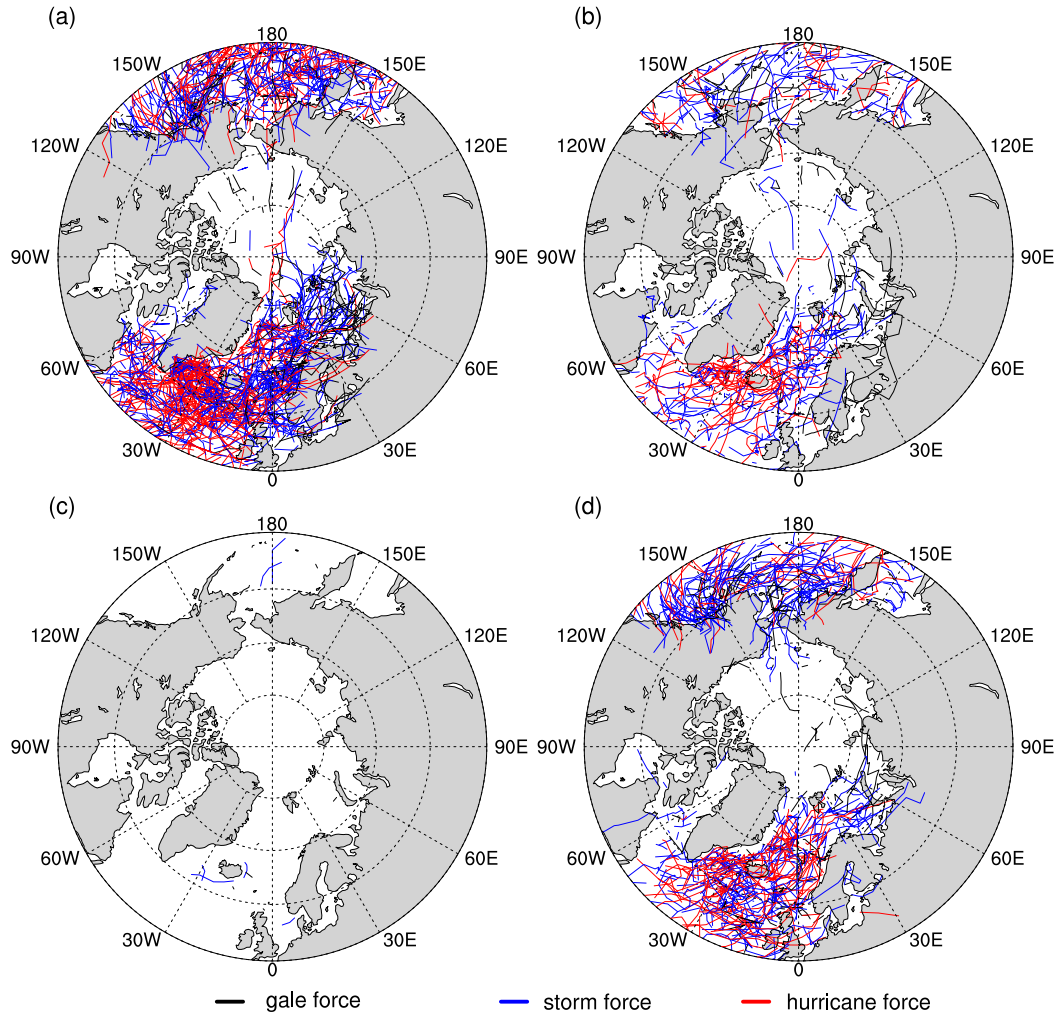
Figure 8. Northern hemisphere polar low storm tracks by season; (a) DJF, (b) MAM, (c) JJA, (d) SON. Tracks are colored by the NOAA warning category corresponding to their lifetime maximum wind speed.

polar latitudes in exactly the same manner as tropical- and mid-latitudes is a new feature introduced by Stride Search, and was the primary motivation for its development.

To validate Stride Search we compare its output to the output of TSTORMS, the current standard tool for tropical cyclone detection. We show that Stride Search performs faster than TSTORMS and its relative speed up increases as the data resolution increases. The final output of both algorithms, tropical cyclone tracks, generally agree. However, due to their different definitions of spatial search sectors, results between the two codes can differ in some cases even when both use the same identification criteria and threshold values.

Our results show that the storm track statistics associated with a particular climate data set can depend not only on the storm identification criteria, as is widely reported in the literature (e.g.,
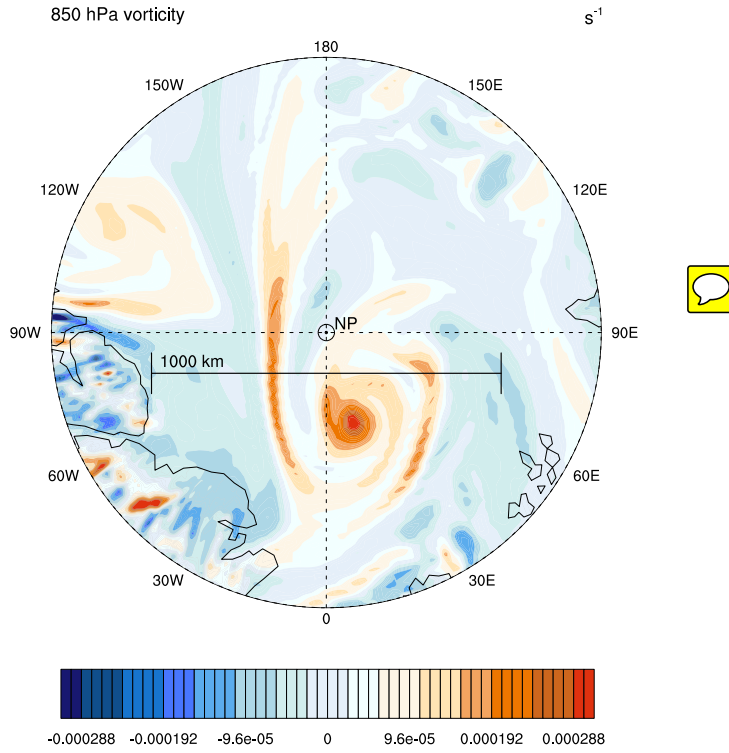
Figure 9. An example polar low on Dec. 25, simulation year 3, located at (017.0E, 86.5N) with structural similarities to a tropical cyclone. The center of the plot is the north pole and the perimiter is the 80N latitude circle.

Bracegirdle and Gray (2008); Raible et al. (2008); Horn et al. (2014)), but also on the spatial search algorithm used to produce the storm tracks. Since the Stride Search algorithm is defined independently of data layout and resolution, we posit that it may provide a more objective analysis tool and be less sensitive to differences in spatial discretizations between data sets. Further experiments are necessary to investigate this claim; they should include variable resolution data, data defined on different types of spherical meshes, and sensitivity analyses covering a range of identification criteria and threshold values.

We anticipate that extending Stride Search to other, more specialized applications such as locating multi-centered cyclones (Hanley and Caballero, 2012) and atmospheric rivers (Ralph et al., 2004; Prabhat et al., 2012) will be straightforward. The software uses the object-oriented design capabilities provided by modern Fortran (Adams et al., 2009) and is intended to allow users to extend its data types to new applications.

Finally, we note that the performance of both Stride Search and TSTORMS software may be improved via parallelization. It is already common to take advantage of temporal parallelism by applying the spatial search algorithm to multiple time steps and multiple files concurrently using several compute nodes. This may be implemented with customized run scripts or dedicated software

23

such as NASA's Portable Distributed Script (PoDS) software (Kouatchou and Oloso, 2014), GNU Parallel (Tange, 2011), and the Toolkit for Extreme Climate Analysis (Prabhat et al., 2012). However, there also remains a significant amount of unexploited parallelism in the storm detection problem, as individual search sectors at the same time step may be distributed across intranode threads. We mark the parallel development of the Stride Search software as an additional item for future work.

7 Code availability

A basic implementation of Stride Search written in Fortran for data on uniform latitude-longitude grids is available at https://github.com/pbosler/StrideSearch. The code is available open-source and distributed under the GPL-2.0 license. Development of a C++ implementation and support for unstructured grids are ongoing projects.

24

References

Adams, J. C., Brainerd, W. S., Hendrickson, R. A., Maine, R. E., Martin, J. T., and Smith, B. T.: The Fortran 2003 Handbook, Springer, 2009.

Balaji, V., Boville, B., Cheung, S., Clune, T., Collins, N., Craig, T., Cruz, C., da Silva, A., DeLuca, C., de Fainchtein, R., Eaton, B., Hallberg, B., Henderson, T., Hill, C., Iredell, M., Jacob, R., Jones, P., Kluzek, E., Kauffman, B., Larson, J., Li, P., Liu, F., Michalakes, J., Murphy, S., Neckels, D., O.Kuinghttons, R., Oehmke, B., Panaccione, C., Rosinski, J., Sawyer, W., Schwab, E., Smithline, S., Spector, W., Stark, D., Suarez, M., Swift, S., Theurich, G., Trayanov, A., Vasquez, S., Wolfe, J., Yang, W., Young, M., and Zaslavsky, L.: Earth System Modeling Framework (ESMF) version 6.3, www.earthsystemmodeling.org, 2014.

Blender, R., Fraedrich, K., and Lunkeit, F.: Identification of cyclone-track regimes in the North Atlantic, Q. J. Roy. Meteor. Soc., 123, 727–741, 1997.

Bowditch, N.: The American Practical Navigator, Pub. No. 9, National Imagery and Mapping Agency, bicentennial edn., http://www.nws.noaa.gov/om/marine/home.htm, 2002.

Bracegirdle, T. J. and Gray, S. L.: An objective climatology of the dynamical forcing of polar lows in the Nordic seas, Int. J. Climatology, 28, 1903–1919, 2008.

Cotton, W. R. and Anthes, R. A.: Storm and Cloud Dynamics, Academic Press, 1989.

Dennis, J. M., Edwards, J., Evans, K. J., Guba, O., Lauritzen, P. H., Mirin, A. A., St-Cyr, A., Taylor, M. A., and Worley, P. H.: CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model, Int. J. High Perform. C., 26, 74–89, 2012.

Føre, I., Kristjánsson, J. E., Kolstad, E. W., Bracegirdle, T. J., Saetra, Ø., and Røsting, B.: A 'hurricane-like' polar low fuelled by sensible heat flux: High-resolution numerical simulations, Q. J. Roy. Meteor. Soc., 138, 1308–1324, 2012.

Geng, Q. and Sugi, M.: Variability of the North Atlantic cyclone activity in winter analyzed from NCEP-NCAR reanalysis data, J. Climate, 14, 3863–3873, 2001.

Hanley, J. and Caballero, R.: Objective identification and tracking of multicentre cyclones in the ERA-Interim reanalysis dataset, Q. J. Roy. Meteor. Soc., 138, 612–625, 2012.

Hodges, K. I.: A general method for tracking analysis and its application to meteorological data, Mon. Weather Rev., 122, 2573–2586, 1994.

Horn, M., Walsh, K., Zhao, M., Camargo, S. J., Scoccimarro, E., Murakami, H., Wang, H., Ballinger, A., Kumar, A., Shaevitz, D. A., Jonas, J. A., and Oouchi, K.: Tracking scheme dependence of simulated tropical cyclone response to idealized climate simulations, J. Climate, 27, 2014.

Kleppek, S., Muccione, V., Raible, C. C., Bresch, D. N., Koellner-Heck, P., and Stocker, T. F.: Tropical cyclones in ERA-40: A detection and tracking method, Geophys. Res. Lett., 35, 5, 2008.

Knutson, T. R., Sirutis, J. J., Garner, S. T., Vecchi, G. A., and Held, I. M.: Simulated reduction in Atlantic hurricane frequency under twenty-first-century warming conditions, Nat. Geosci., 1, 359–364, 2008.

König, W., Sausen, R., and Sielmann, F.: Objective identification of cyclones in GCM simulations, J. Climate, 6, 2217–2231, 1993.

Kouatchou, J. and Oloso, A.: Portable Distributed Script (PoDS), Tech. Rep. GSC-16531-1, NASA, 2014.

Montgomery, M. T. and Farrell, B. F.: Polar low dynamics, J. Atmos. Sci., 49, 2484–2505, 1992.

25

Neale, R., Chen, C., Gettlelman, A., Lauritzen, P. H., Park, S., Williamson, D. L., Conley, A. J., Garcia, R., Kinnison, D., Lamarque, J., Marsh, D., Mills, M., Smith, A. K., Tilmes, S., Vitt, F., Morrison, H., Cameron-Smith, P., Collins, W. D., Iacono, M. J., Easter, R. C., Ghan, S. J., Liu, X., Rasch, P. J., and Taylor, M. A.:

595   Description of the NCAR Community Atmosphere Model (CAM 5.0), Tech. Rep. NCAR/TN-486+STR, NCAR, 2012.

Neu, U., Akperov, M. G., Bellenbaum, N., Benestad, R., Blender, R., Caballero, R., Cocozza, A., Dacre, H. F., Feng, Y., Fraedrich, K., Grieger, J., Gulev, S., Hanley, J., Hewson, T., Inatsu, M., Keay, K., Kew, S. F., Kindem, I., Leckebusch, G. C., Liberato, M. L. R., Lionello, P., Mokhov, I. I., Pinto, J. G., Raible, C. C., Reale,

600   M., Rudeva, I., Schuster, M., Simmonds, I., Sinclair, M., Sprenger, M., Tilinina, N. D., Trigo, I. F., Ulbrich, S., Ulbrich, U., Wang, X. L., and Wernli, H.: IMILAST: A community effort to intercompare extratropical cyclone detection and tracking algorithms, B. Am. Meterol. Soc., 94, 529–547, 2013.

Nordeng, T. E. and Rasmussen, E. A.: A most beautiful polar low. A case study of a polar low development in the Bear Island region, Tellus, 44A, 81–99, 1992.

605   Prabhat, Ubel, O., Byna, S., Wu, K., Li, F., Wehner, M., and Bethel, E. W.: TECA: A Parallel Toolkit for Extreme Climate Analysis, in: Third Workshop on Data Mining in Earth System Science (DMESS 2012), LBNL-5352E, International Conference on Computational Science (ICCS 2012), http://vis.lbl.gov/Vignettes/2012-climatePatterns-vignette/, Omaha NE, 2012.

Putnam, W. M. and Lin, S.-J.: Finite-volume transport on various cubed-sphere grids, J. Comput. Phys., 227,

610   55–78, 2007.

Raible, C. C. and Blender, R.: Northern hemisphere midlatitude cyclone variability in GCM simulations with different ocean representations, Clim. Dynam., 22, 239–248, 2004.

Raible, C. C., Della-Marta, P. M., Schwierz, C., Wernli, H., and Blender, R.: Northern hemisphere extratropical cyclones: A comparison of detection and tracking methods and different reanalyses, Mon. Weather Rev.,

615   136, 880–897, 2008.

Ralph, F. M., Nieman, P. J., and Wick, G. A.: Satellite and aircraft observations of atmospheric rivers over the eastern North Pacific ocean during the winter of 1997/98, Mon. Weather Rev., 132, 1721–1745, 2004.

Rasmussen, E. A. and Turner, J., eds.: Polar Lows: Mesoscale Weather Systems in the Polar Regions, Cambridge University Press, 2003.

620   Reed, K. A. and Jablonowski, C.: Assessing the uncertainty in tropical cyclone simulations in NCAR's Community Atmosphere Model, J. Adv. Mod. Earth Sys., 3, 16, 2011.

Roesler, E. L., Bosler, P. A., and Taylor, M. A.: Arctic storms in a regionally refined atmospheric general circulation model, Geosci. Model Devel., In prep.

Samet, H.: Foundations of Multidimensional and Metric Data Structures, Morgan Kaufman, 2006.

625   Sinclair, M. R.: An objective cyclone climatology for the Southern Hemisphere, Mon. Weather Rev., 122, 2239–2256, 1994.

Sinclair, M. R.: Objective identification of cyclones and their circulation intensity, and climatology, Weather Forecast., 12, 595–612, 1997.

Skamarock, W. C., Klemp, J. B., Duda, M. G., Fowler, L. D., and Park, S.-H.: A multiscale nonhydrostatic

630   atmospheric model using centroidal Voronoi tesselations and C-grid staggering, Mon. Weather Rev., 140, 2012.

Stocker, T., Qin, D., Plattner, G.-K., Alexander, L., Allen, S., Bindoff, N., Bréon, F.-M., Church, J., Cubasch, U., Emori, S., Forster, P., Friedlingstein, P., Gillett, N., Gregory, J., Hartmann, D., Jansen, E., Kirtman, B., Knutti, R., Krishna Kumar, K., Lemke, P., Marotzke, J., Masson-Delmotte, V., Meehl, G., Mokhov, I., Piao, S., Ramaswamy, V., Randall, D., Rhein, M., Rojas, M., Sabine, C., Shindell, D., Talley, L., Vaughan, D., and Xie, S.-P.: Technical Summary. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Tech. rep., IPCC, 2013.

Tange, O.: GNU Parallel - The command-line power tool, The USENIX Magazine, pp. 42–47, 2011.

TSTORMS: Diagnosis and Detection of Tropical Storms in High-Resolution Atmospheric Models, http://www.gfdl.noaa.gov/tstorms, accessed January 27, 2015, 1997.

Vitart, F. and Stockdale, T. N.: Seasonal forecasting of tropical storms using coupled GCM integrations, Mon. Weather Rev., 129, 2521–2537, 2001.

Vitart, F., Anderson, J. L., and Stern, W. F.: Simulation of interannual variability of tropical storm frequency in an ensemble of GCM integrations, J. Climate, 10, 745–760, 1997.

Walsh, K. J. E., Fiorino, M., Landsea, C. W., and McInnes, K. L.: Objectively determined resolution-dependent threshold criteria for the detection of tropical cyclones and reanalyses, J. Climate, pp. 2307–2314, 2007.

Wernli, H. and Schwierz, C.: Surface cyclones in the ERA-40 dataset (1958-2001). Part I: Novel identification method and global climatology, J. Atmos. Sci., 63, 2486–2507, 2006.

Williamson, D. L.: Storm track representation and verification, Tellus, 33, 513–530, 1981.

Worley, P. H., Craig, A. P., Dennis, J. M., Mirin, A. A., Taylor, M. A., and Vertenstein, M.: Performance of the Community Earth System Model, in: International Conference for High Performance Computing, Networking, Storage and Analysis (SC11), 2011.

Zarzycki, C. M. and Jablonowski, C. J.: A multidecadal simulation of Atlantic tropical cyclones using a variable resolution global atmosphere general circulation model, J. Adv. Mod. Earth Sys., 6, 805–828, 2014.

Zhao, M., Held, I. M., Lin, S. J., and Vecchi, G. A.: Simulations of global hurricane climatology, interannual variability, and response to global warming using a 50-km resolution GCM, J. Climate, 22, 6653–6678, 2009.