Using persistent homology to quantify a diurnal cycle in hurricanes[☆]Sarah Tymochko^{a,*}, Elizabeth Munch^{a,b}, Jason Dunion^{c,d}, Kristen Corbosiero^e, Ryan Torn^e^a Dept. of Computational Mathematics, Science, and Engineering, Michigan State University, 428 South Shaw Lane, East Lansing, MI, 48824, USA^b Dept. of Mathematics, Michigan State University, 619 Red Cedar Road, East Lansing, MI 48824, USA^c Cooperative Institute for Marine and Atmospheric Studies, University of Miami, 4600 Rickenbacker Causeway, Miami, FL 33149, USA^d Hurricane Research Division, NOAA/Atlantic Oceanographic and Meteorological Laboratory, 4301 Rickenbacker Causeway, Miami, FL 33149, USA^e University at Albany - SUNY Albany, Dept. of Atmospheric and Environmental Sciences, 1400 Washington Avenue, Albany, NY 12222, USA

ARTICLE INFO

Article history:

Received 12 March 2019

Revised 10 January 2020

Accepted 21 February 2020

Available online 22 February 2020

Keywords:

Topological Data Analysis

Atmospheric science

Diurnal cycle

Image processing

ABSTRACT

The diurnal cycle of tropical cyclones (TCs) is a daily cycle in clouds that appears in satellite images and may have implications for TC structure and intensity. The diurnal pattern can be seen in infrared (IR) satellite imagery as cyclical pulses in the cloud field that propagate radially outward from the center of nearly all Atlantic-basin TCs. These diurnal pulses, a distinguishing characteristic of this diurnal cycle, begin forming in the storm's inner core near sunset each day, appearing as a region of cooling cloud-top temperatures. The area of cooling takes on a ring-like appearance as cloud-top warming occurs on its inside edge and the cooling moves away from the storm overnight, reaching several hundred kilometers from the circulation center by the following afternoon. The state-of-the-art TC diurnal cycle measurement in IR satellite imagery has a limited ability to analyze the behavior beyond qualitative observations. We present a method for quantifying the TC diurnal cycle using one-dimensional persistent homology, a tool from Topological Data Analysis, by tracking maximum persistence and quantifying the cycle using the discrete Fourier transform. Using Geostationary Operational Environmental Satellite IR imagery from Hurricanes Felix and Ivan, our method is able to detect an approximate daily cycle.

© 2020 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license.

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

1. Introduction

The field of atmospheric science has numerous observation platforms that provide high space and time resolution data, but has yet to find methods which can quantify the intuitive patterns explicitly. Meanwhile, the young field of Topological Data Analysis (TDA) encompasses methods for quantifying exactly these sorts of structural intuitions seen by atmospheric scientists. This paper merges these two fields by using persistent homology, a now well-established tool in TDA, to quantify a diurnal cycle observed in a hurricane using Geostationary Operational Environmental Satellite (GOES) infrared (IR) satellite data.

Persistent homology, and more generally TDA methods, has found significant success in rather disparate applications by finding structure in data and using this insight to answer questions from the domain of interest. For instance, Giusti et al. used the homology of random simplicial complexes to investigate the geomet-

ric organization of neurons in rat brains [1]. Persistent homology has also been used to understand periodicity in time series arising from biological [2,3] and engineering applications [4], as well as for image analysis [5–8].

This paper presents an application of the use of both time series and image analysis using TDA to study a daily cycle in hurricanes that is of interest in the field of atmospheric science. The diurnal cycle of tropical cyclones (TCs) has been described in previous studies [9–18] that provide evidence of the regularity of this cycle as well as its potential impacts. This diurnal pattern can be seen in GOES IR imagery as cyclical pulses in the cloud field that propagate radially outward from TCs at speeds of 5–10 m s⁻¹ [9–11]. These diurnal pulses, a distinguishing characteristic of the TC diurnal cycle, begin forming in the TC's core near the time of sunset each day and appear as a region of cooling cloud-top temperatures. The area of cooling then takes on a ring-like appearance as marked cloud-top warming occurs on its inside edge and it moves away from the storm overnight, reaching several hundred kilometers from the TC center by the following afternoon. Observations and numerical model simulations indicate that TC diurnal pulses propagate

[☆] Handled by Associate Editor: Ricardo Torres, PhD.

* Corresponding author:

E-mail address: tymochko@egr.msu.edu (S. Tymochko).

through a deep layer of the TC environment, suggesting that they may have implications for TC structure and intensity [9–12,14].

The current state of the art TC diurnal cycle measurement has a limited ability to analyze the behavior beyond qualitative observations. This paper presents a more advanced mathematical method for quantifying the TC diurnal cycle using tools from TDA, namely one-dimensional persistent homology to analyze the holes in a space. This research aims to detect the presence of the diurnal cycle in GOES IR satellite imagery and to track the changes through a time series.

In this paper, we present a method of automatically detecting and quantifying periodic circular structure in satellite imagery as well as the results of our method applied to GOES IR hurricane imagery for Hurricane Felix in 2007 and Hurricane Ivan in 2004. The naive combination of persistent homology with the IR imagery did not show a recurring pattern due to drastically variable values in IR brightness temperature data. Despite this, looking at the data, there is a clear circular feature that is visible in the imagery. Thus we develop a more sophisticated method using tools including the distance transform and one-dimensional persistent homology to detect the TC diurnal cycle quantitatively using maximum persistence. Using our method, we are able to detect this 24-h cycle in both hurricanes automatically, improving upon the existing qualitative methods.

2. Tropical cyclone background

Previous research has documented a clear diurnal cycle of cloudiness and rainfall in TCs: enhanced convection (i.e., thunderstorms) occurs overnight, precipitation peaks near sunrise, and upper-level cloudiness (i.e., the cirrus canopy) expands radially outward throughout the day, reaching its maximum areal coverage in the early evening hours [9–18]. To quantify the expansion and contraction of the cirrus canopy, Dunion et al. used GOES satellite IR imagery to examine the six-hour cloud-top temperature differences of major hurricanes in the Atlantic basin from 2001 to 2010 [9]. They found that an area of colder cloud tops propagated outward around 5–10 m s⁻¹ over the course of the day, with warming temperatures on its inner edge. More recently, in [11], Ditchek et al. expanded Dunion et al.'s work to include all tropical cyclones in the Atlantic basin from 1982 to 2017 and found that the diurnal pulse is nearly ubiquitous, with 88% of TC days featuring an outwardly propagating pulse.

Despite the consistent signature and documentation of this diurnal cloud signature, open questions remain as to how the diurnal cycle is linked to inner-core convective processes and whether it is a column-deep phenomenon or mainly tied to upper-level TC cloud dynamics related to incoming solar radiation [11,15,16,18]. Investigating these questions is relevant to TC forecasting as the diurnal cycle of clouds and rainfall has implications for forecasting storm structure and intensity, as evidenced by the diurnal cycle in objective measures of TC intensity and the extent of the 50-kt wind radius documented by Dunion et al. Additionally, and especially relevant to the current work, most of the papers above have identified the pulse using subjective measures of cloud-top temperature change and timing [9,11]. The current work seeks to quantify the pulse to determine its true periodicity using persistent homology, a topological tool that is particularly effective at capturing the type of patterns visible in the pulse.

3. Math background

Persistent homology is a tool from the field of TDA which measures structure in data. This data can start in many forms, including as point clouds or, as in the case of this work, as a function

on a domain. In this section, we will briefly review the necessary background to understand cubical complexes and persistent homology, and refer the interested reader to [19–22] for a more complete introduction.

3.1. Cubical complexes

In this section, we largely follow Chapter 2 of [21] with the caveat for the informed reader that because we use homology with \mathbb{Z}_2 coefficients, we can be lazy about orientations of cubes. In addition, our data consists of 2D images, so we need only define cubes up to dimension 2.

An elementary interval is a closed interval $I \subset \mathbb{R}$ of the form $[\ell, \ell + 1]$ or $[\ell]$ for $\ell \in \mathbb{Z}$, which are called nondegenerate and degenerate respectively. An elementary cube $Q \in \mathbb{R}^2$ is a product of elementary intervals $Q = I_1 \times I_2$. The dimension of Q , $\dim(Q)$, is the number of nondegenerate components of Q . Note that 0-dimensional cubes are just vertices at the points on the lattice $\mathbb{Z} \times \mathbb{Z}$ in \mathbb{R}^2 , 1-dimensional cubes are edges connecting these vertices, and 2-dimensional cubes are squares. Let \mathcal{K} denote the set of all elementary cubes in \mathbb{R}^2 and $\mathcal{K}_d \subset \mathcal{K}$ the set of d -dimensional cubes. A set $X \subset \mathbb{R}^2$ is cubical if it can be written as a finite union of elementary cubes. Then we denote the associated cubical complex as $\mathcal{K}(X) = \{Q \in \mathcal{K} \mid Q \subset X\}$, with the d -dimensional subset denoted $\mathcal{K}_d(X) = \{Q \in \mathcal{K}(X) \mid \dim(Q) = d\}$. If $Q \subseteq P$, then we say Q is a face of P , denoted $Q \leq P$. If $Q \subsetneq P$, then Q is a proper face of P , denoted $Q < P$, and is additionally a primary face of P if $\dim(Q) = \dim(P) - 1$.

A grayscale image, or more generally an $m \times n$ matrix, can be viewed as a function $M : D \rightarrow \mathbb{R}$ where $D = \{(i, j) \mid 0 \leq i < m, 0 \leq j < n\}$. We will model this as a function defined on a particularly simple cubical set $K = \mathcal{K}([0, m] \times [0, n])$. For simplicity, we denote by s_{ij} the square $[i, i + 1] \times [j, j + 1]$. So, given a matrix M , we equivalently think of this data as a function $M : K \rightarrow \mathbb{R}$ where we set $M(s_{ij})$ equal to the matrix entry M_{ij} and set $M(P) = \min_{s_{i,j} \geq P} M(s_{i,j})$ for all lower dimensional cubes P . Note that we will abuse notation and use M to denote both the original matrix and representation as a function with domain K .

3.2. Homology

Homology [20] is a standard tool in algebraic topology which provides a vector space¹ $H_k(X)$ for each dimension $k = 0, 1, 2, \dots$ for a given topological space X . The different dimensions measure different properties of the space. In particular, for this work we are interested in 1-dimensional homology; i.e. when $k = 1$. The 1-dimensional homology group measures the number of loops in the space; equivalently, we can think of this as the number of holes in the space. In particular, if we look at the black region in each of the examples in Fig. 1, the rank of the first homology for each is (1,2,1).

The exact definition of homology is as follows. For any cubical set L (which for the purposes of this discussion will always be a subset of K), we have sets giving the cubes of different dimensions: $\mathcal{K}_i(L)$ for $i = 0, 1, 2$. An i -chain is a formal linear combination of i -simplices in L , $c = \sum_{Q_j \in \mathcal{K}_i(L)} a_j Q_j$, with coefficients $a_j \in \mathbb{Z}_2$. We can of course add these objects by setting $(\sum a_j Q_j) + (\sum b_j Q_j) = \sum (a_j + b_j) Q_j$ and multiply by a constant. Thus, the collection of all i -chains forms a vector space $C_i(L)$.

We define a linear transformation $\delta_i : C_i(L) \rightarrow C_{i-1}(L)$ called the boundary map, by setting $\delta_i(Q) = \sum P$ where the sum² is over the

¹ Normally a group, however, we are working with field coefficients.

² Again, notice that because we are working with \mathbb{Z}_2 coefficients, the book-keeping normally needed for orientation is unnecessary.




Fig. 1. An example matrix, M (top left) and corresponding persistence diagram (top right). Second row: The black portions are sublevel sets, M_r , where $r = 0.25, 0.35$, and 0.7 . The existence of a point far from the diagonal in the persistence diagram shows that there is a prominent circular structure; while the other points are caused by the smaller circles.

primary faces $P < Q$. The kernel of δ_1 , $\text{Ker}(\delta_1)$, (that is, the set of elements of $C_1(L)$ which map to 0) is generated by closed loops in L . The image of δ_2 , $\text{Im}(\delta_2)$, is generated by boundaries of 2-cells. Then the 1-dimensional homology group is defined to be $H_1(K) = \text{Ker}(\delta_1)/\text{Im}(\delta_2)$. An element of this group $\gamma \in H_1(L)$, represents an equivalence class of loops which can differ by collections of 2-cells.

3.3. Persistent homology

For a *static* space L , $H_1(L)$ measures information about the number of loops. Persistent homology takes as input a *changing* topological space, and summarizes the information about how the homology changes.

Let an $m \times n$ \mathbb{R} -valued matrix M be given. Fix a function value $r \in \mathbb{R}$ and let $M_r = f^{-1}(-\infty, r]$. That is, M_r is the subset of squares in K which have value at most r in the matrix, along with all edges and vertices which are faces of any included square. M_r is often called a sublevel set of M . See Fig. 1 for M_r regions corresponding to an example image.

These spaces have the property that $M_r \subseteq M_s$ for $r \leq s$, thus we can consider the sequence $M_{r_1} \subseteq M_{r_2} \subseteq \dots \subseteq M_{r_k}$ for any set of numbers $r_1 < r_2 < \dots < r_k$. This sequence of spaces is called a filtration. For each of the spaces, we can compute the homology group $H_p(M_{r_i})$. The inclusion maps give rise to linear maps $H_p(M_{r_1}) \rightarrow H_p(M_{r_2}) \rightarrow \dots \rightarrow H_p(M_{r_k})$.

It is these maps that we study to understand how the space changes. In particular, when we are focused on 1-dimensional homology ($k = 1$) as in this study, a loop is represented by an element $\gamma \in H_1(M_{r_i})$. We say that this loop is born at r_i if it is not in the image from the previous space; that is, $\gamma \notin \text{Im}(H_1(M_{i-1}) \rightarrow H_1(M_i))$. This same loop dies at r_j if it merges with this image in M_{r_j} ; that is, $\gamma \in \text{Im}(H_1(M_{i-1}) \rightarrow H_1(M_j))$ where we abuse notation by using γ to both refer to the class in $H_1(M_{r_i})$ and the image of this class under the sequence of maps in $H_1(M_{r_j})$. We refer to $r_j - r_i$ as the lifetime of the class.

A persistence diagram, as seen in the right of Fig. 1, is a collection of points where for each class which is born at r_i and dies at r_j is represented by a point at (r_i, r_j) . The intuition is that a class which has a long lifetime is far from the diagonal while a class with a short lifetime is close. In many cases, a long lifetime loop implies that there is some sort of inherent topological feature being found, and thus that this point far from the diagonal is important, while short lifetime loops are likely caused by topological noise due to sampling or other errors in the system. In the example of Fig. 1, there is a prominent off-diagonal point which

shows that the function defined by the matrix has a circular feature. Thus, a common measure for looking at the persistence diagram when investigating a single, circular structure is the maximum persistence, defined as

$$\text{MaxPers}(D) = \max_{(r_i, r_j) \in D} r_j - r_i \quad (1)$$

for a given persistence diagram, D .

4. Methods

4.1. Data preparation

The data was given in the form of storm-centered GOES IR ($10.7\text{-}\mu\text{m}$) satellite imagery. The $10.7\text{-}\mu\text{m}$ long channel detects IR energy emitted from the Earth and is not strongly affected by atmospheric water vapor. Thus, this particular channel is useful for detecting clouds at all times of the day and night and is ideal for tracing the diurnal evolution of the TC cloud fields. For our analysis, we used two different types of data sets. These two data sets have the same native spatial resolution, but differ in temporal resolution. The first type (hereafter the GOES-12 data sets), utilizes brightness temperatures derived directly from GOES-12 4-km IR satellite imagery and consists of data in hourly increments, with the exception of 0415 and 0515 UTC each day in Hurricane Felix, and 0445 and 0545 UTC each day in Hurricane Ivan (due to the GOES-12 satellite eclipse period). Imagery was remapped such that each pixel has a spatial resolution of 2 km^2 and each image covers a total area of approximately $1500\text{ km} \times 1500\text{ km}$, represented as a 752×752 matrix. This remapping was performed using the Man Computer Interactive Data Access System (McIDAS; [23]) in order to generate storm-centered satellite images that were focused on the relevant TC environment. The McIDAS 4 km to 2 km remapping procedure replicates the original 8-bit grayscale values such that none of the original pixel information is lost. The second type is the GridSat-GOES [24] data set and consists of data in 3-h increments with the exception of 0600 UTC each day. Each pixel has a resolution of approximately 8 km and each image covers a total area of approximately $2400\text{ km} \times 2400\text{ km}$, represented by a 301×301 matrix. This data is cropped to a 191×191 matrix to approximately match the area covered by the first set of data. The cropped version covers a total area of approximately $1530\text{ km} \times 1530\text{ km}$. For an example of the GOES-12 and cropped GridSat-GOES satellite imagery for Hurricane Felix, see Fig. 2.

The GridSat-GOES data set requires some additional processing. The data is stored using a different format, using short numbers rather than floats, so the following equation is applied to the GridSat-GOES brightness temperatures to do the number type conversion: $[(\text{Original} \cdot 0.01 + 200.0) - 22.858]/0.919565$. Some images in the GridSat-GOES data set also contain missing values where the brightness temperature for certain pixels was not recorded and is instead assigned a fill value. In order to prevent these values from impacting our results, we interpolate values for these pixels. For a given pixel with a missing value, we compute




Fig. 2. Original satellite imagery from Felix GOES-12 data set (left) and Felix GridSat-GOES data set (right) at approximately the same time.

the average value of a 5×5 grid centered at the pixel, not including the pixels in this range that also have missing values.

For Hurricane Felix, we studied both types of data sets to test the flexibility of the method across spatial and temporal resolution. The Felix GOES-12 data set spans 2 to 4 September 2007, while the Felix GridSat-GOES dataset spans spanning 31 August to 6 September 2007. For Hurricane Ivan, we used the GOES-12 data set, which spans 30 August to 1 September 2004.

4.2. Method of detection & quantification

We present a method of detecting and quantifying periodic circular structure representing the TC diurnal cycle in IR satellite imagery. Our method combines existing methods from the fields of image processing, topological data analysis, and signal processing.

Initially, we have a time series of IR satellite images, represented as a matrix of pixel values $S(t)$ for time t . The TC diurnal pulse is propagating outward through the day; thus, in order to see the movement and changes in the GOES satellite brightness temperature, we consider the difference in matrices six hours apart [9,11]. For all times t , given the original brightness temperature image $S(t)$, we compute the six-hour differences, $M(t) = S(t+6) - S(t)$. While circular features are visually prominent in the data, simply using persistence on the difference data did not show any relevant features. This discrepancy is due to the extreme differences in the function values between the circular sections which prevents the sublevel sets from containing the full circular structure until very late in the filtration. Thus, we are unable to detect the circular structure from $M(t)$ and we define a new function on the difference matrix using the following method. Fix a threshold μ and let $M(t)_\mu$ be the subset of $M(t)$ which has function value less than μ . This method results in a binary matrix defined entry-wise, with $M(t)_\mu[i, j] = 1$ if $M(t)[i, j] < \mu$, and $M(t)_\mu[i, j] = 0$ otherwise. We will address this choice of threshold in Section 5.2; however, we will focus on the case $\mu = 80$ degrees for most of our analysis. Note that because $M(t)$ is a difference of two images, the threshold is not isolating all pixels above a certain temperature, but rather those pixels that increase in value by at least 80 degrees over the six hours.

After thresholding, we now have binary images; however, the persistent homology is uninteresting, as there are only two possible sublevel sets. In order to create a grayscale image that maintains the visually apparent topological structure of the image, we apply the distance transform, a method from the field of image processing [25,26], which gives a new matrix $D(t)$. The distance transform is calculated as follows: given any pixel s_{ij} in an image $M(t)$ represented as a matrix of pixels, $D(t)_{i,j} = \min d(s_{i,j}, x)$ where x is a 0-valued pixel and d is any distance metric. In this application, we specifically use the L_∞ distance, also known as the chessboard distance; however, the distance transform is defined for any distance metric, the Euclidean distance transform being the most commonly used. Given two pixels, $s_{i_1, j_1}, s_{i_2, j_2}$ the L_∞ distance between them is calculated as $d(s_{i_1, j_1}, s_{i_2, j_2}) = \max\{|i_2 - i_1|, |j_2 - j_1|\}$. This defines a distance on the pixels, which are the 2-cells in the cubical complex. The distance can be extended to the lower dimensional cells in the same manner as described in Section 3.1.

To calculate the distance transform, we use the python submodule `scipy.ndimage`, specifically the function `distance_transform_cdt` with the chessboard metric. Therefore, each entry in $D(t)$ corresponds to the minimal distance to an entry where $M(t)[i, j] \geq \mu$. The distance transform $D(t)$ is then scaled by the resolution for each data set in order to convert the distance units to kilometers instead of pixels. For the GridSat-GOES data set, we scale by a factor of 8 km/pixel while for the GOES-12 data sets, we scale by a factor of 2 km/pixel. We then compute

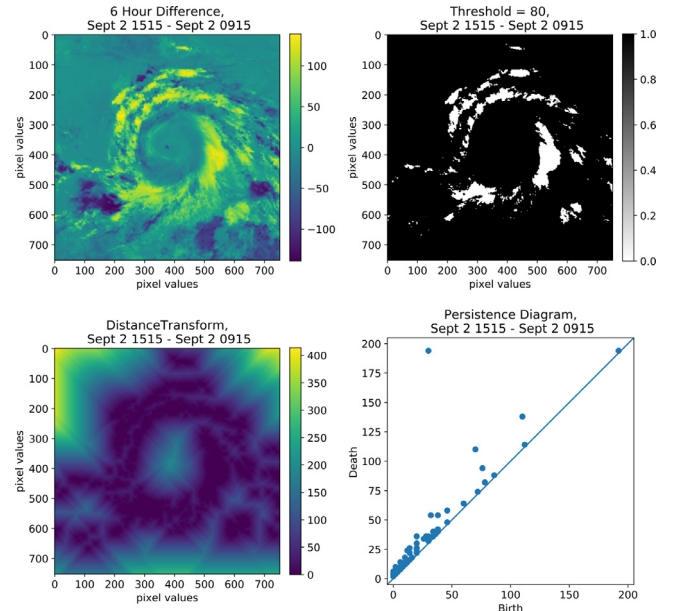


Fig. 3. Top left: Example of 6 h difference, $M(t)$, from the Felix GOES-12 data set; Top right: thresholded subset, $M(t)_\mu$ where $\mu = 80$; Bottom left: distance transform function; Bottom right: corresponding persistence diagram.

sublevel set persistence on the function $D(t)$ using the cubtop method in Perseus [27,28], which calculates persistent homology for cubical complexes using concepts from discrete Morse theory. Note that Perseus requires an integer value filtration function on the cubical complex; thus, we chose to use the chessboard metric for the distance transform. Fig. 3 shows an example of each step described so far.

For each six-hour difference in each data set, we apply the steps described above, then calculate maximum persistence as defined in Eq. 1. By plotting maximum persistence over time, we can see how the most prominent circular feature changes through the progression of the day and life of the TC. This plot should show an oscillatory pattern, detecting the change in the diurnal cycle throughout the day. In order to quantify this oscillatory pattern, we use the Fourier transform. In general, the Fourier transform is a commonly used method for investigating periodicity of time series [29] by decomposing a wave into a sum of sinusoids with different frequencies. Since we are working with discrete data, we will work with the discrete Fourier transform (DFT). Let T be the time between discrete samples, then let $t_k = kT$ where $k = 1, \dots, N-1$. Then, the discrete Fourier transform is $F_n = \sum_{k=0}^{N-1} f(t_k) e^{-2\pi i n k / N}$. This converts a function from the time domain to the frequency domain. The power spectrum of F_n can be estimated by calculating the square of the absolute value of the discrete Fourier transform, $|F_n|^2$.

Using the DFT, we calculate the most prominent frequency in the data in order to determine how often the cyclic behavior repeats. Note, to use the discrete Fourier transform the time steps must be equal; however, because of the missing times in our data, this is not the case. Therefore, we approximate the maximum persistence at these values by adding a point along the line between the times immediately before and after the missing time. Additionally, we must truncate the maximum persistence to only include the days where we have the data for the entire day. This means truncating the Felix GridSat-GOES data set to include only 1-4 September 2007, the Felix GOES-12 data set to include only 1-2 September 2007, and the Ivan GOES-12 data set to include only 30-31 August 2004. The discrete Fourier transform was calculated using the python submodule `numpy.fft`.

We first calculate the Fourier transform using the function `fft`, then calculate the frequency bins using `fftfreq`. Using this information, we plot the approximate power spectrum for each data set. Note, if working with real data as we are in this application, the power spectrum will be symmetric for positive and negative frequencies; therefore, we only need to look at the positive frequencies. Picking the frequencies corresponding to the highest peaks in each power spectrum gives us the frequency of the most prominent periodic signal in the data. Using this frequency, we can calculate the period of the oscillatory pattern and quantify the signal we are detecting. To verify the detected signal matches the visual oscillatory pattern we see, we can reconstruct the sinusoid corresponding to the most prominent periodic signal using the inverse discrete Fourier transform, using `ifft`. Then plotting this reconstructed sinusoid over the original data, we can see how closely this signal matches the patterns in the data.

5. Results

5.1. Experimental results

After the data is prepared, we apply the steps described in Section 4.2 to each data set. For the two Hurricane Felix data sets, as they are from the same hurricane, we would expect the results to be similar despite the temporal and spatial resolution differences. Plotting the calculated maximum persistence over time, we get the time series plotted as solid lines in Fig. 4. The plots show an oscillatory pattern for all three data sets which appears to repeat approximately daily.

To verify the periodicity of the oscillatory pattern, we apply the discrete Fourier transform and calculate the power spectrum for each data set. Each power spectrum is shown in Fig. 5. Picking the frequencies corresponding to the highest peaks in each power spectrum gives us the frequency corresponding to the most prominent sinusoidal signal in the data. The maximum peaks in Fig. 5 give a frequency of 0.976 cycles per day for the Felix GridSat-GOES data set, 0.979 cycles per day for the Felix GOES-12 data set,




Fig. 4. Maximum persistence plotted over time for all data sets using threshold $\mu = 80$ in addition to the reconstructed versions, created using inverse Fourier transform. Gray vertical lines separate days according to UTC.




Fig. 5. Power spectrum for each data set. The highest peak on the Felix GridSat-GOES power spectrum occurs at a frequency of approximately 0.976 cycles/day, the peak on the Felix GOES-12 power spectrum occurs at a frequency of approximately 0.979 cycles/day, and the peak on the the Ivan GOES-12 power spectrum occurs at a frequency of 1.0 cycles/day.

and 1.0 cycles per day for the Ivan GOES-12 data set. We use this frequency, f , to calculate the periodicity of the cycle by calculating using $1/f$, giving the period of the sinusoid in days per cycle, then multiply by 24 to rescale to hours per cycle. Doing so gives the result that the cycle is repeating every 24.6 h for the Felix GridSat-GOES data set, every 24.5 h for the Felix GOES-12 data set, and 24.0 h for the Ivan GOES-12 data set.

Using the most prominent frequency for each data set, we calculate the inverse Fourier transform, and plot these reconstructed sinusoids over the original data. These sinusoids, plotted as the lighter dashed lines in Fig. 4, closely resemble the patterns exhibited by the original maximum persistence versus time plots; therefore, these approximately 24 h patterns visible in the plots are also detected mathematically, which verifies the claim that our method is detecting a daily cycle in each data set. Additionally, since for both Hurricane Felix data sets, the plots of maximum persistence against time seem to match and both have similar detected periodicity from the DFT, our method seems robust to the temporal and spatial resolution differences in these two data sets.

5.2. Choice of threshold

The method described involves a choice of threshold, so we used a variety of thresholds, $\mu \in \{25, 30, \dots, 100\}$, to test the sensitivity of our method to the parameter choice. For both data sets of Hurricane Felix, our method is very robust to the choice of threshold. In Fig. 6, the top row are plots that represent maximum persistence versus time for the Hurricane Felix data sets using a variety of thresholds. There is a clear periodic pattern for both data sets across all the thresholds shown. In fact, for all thresholds tested $\mu \in \{35, 40, \dots, 90\}$, the period is consistent at 24.6 h for the Felix GridSat-GOES data set and 24.5 h for the Felix GOES-12 data set. For $\mu < 35$ and $\mu > 90$ the Fourier transform is unable to pick up the daily pattern in the Felix GridSat-GOES data set.

For Hurricane Ivan, the plot is shown on the bottom row of Fig. 6 for thresholds $\mu \in \{80, 85, \dots, 100\}$. For all of the threshold value shown, our method consistently detects a 24.0 h period. This is a smaller range of threshold values than those that detect a daily cycle in Hurricane Felix, but for thresholds $\mu \in \{80, 85, 90\}$, our method detects a daily cycle in all three data sets. Thus, the method may require some parameter tuning, but our analysis of these three data sets gives a range of values to start with when testing new data sets.

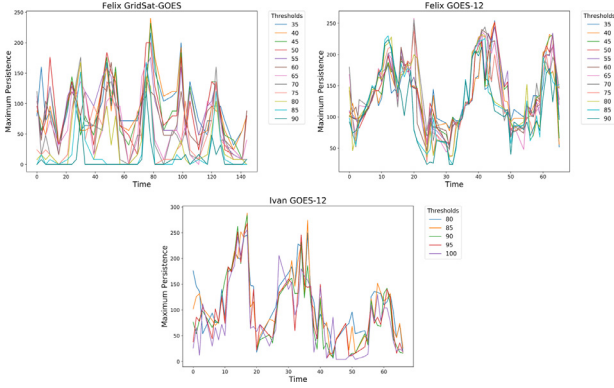


Fig. 6. Maximum persistence vs time plot for Hurricane Felix (top row) and Hurricane Ivan (bottom row). Hurricane Felix results are shown for all thresholds $\mu \in \{35, 40, \dots, 90\}$ while Hurricane Ivan results are shown for $\mu \in \{80, 85, \dots, 100\}$.

5.3. Removal of noise

While the above method detects a daily cycle, there are some instances where the six-hour differencing introduces noise because of varying behavior in the center of the hurricane. The left image of in Fig. 7 shows an example of how this noise can appear in the distance transform for the Felix GOES-12 data set. A small area of pixels above the threshold cause the distance transform to fill in the center of the circular region, thus potentially changing the value of maximum persistence. Therefore, before applying the distance transform, we use a method from mathematical morphology [30] called opening to de-noise the image and see how this impacts the detected periodicity.

Opening is the combination of two tools from mathematical morphology, erosion and dilation. Both involve a kernel moving through a binary image. In erosion, a pixel in the original image will remain a 1 only if all pixels under the kernel are 1's, otherwise it becomes a 0. Dilation is the opposite of erosion, A kernel moves through the binary image and a pixel is assigned a 1 if at least one pixel under the kernel is a 1, otherwise it is assigned a 0. Opening is erosion followed by dilation, which will remove noise and rebuild the area around the boundary.

We apply opening to the binary thresholded image using a 8×8 pixel kernel for the GOES-12 data sets and a 2×2 pixel kernel for the GridSat-GOES data set to remove noise such as these center pixels. Note, the difference in size of the kernel is due to the differences in spatial resolution between the two data sets. We use the python module cv2 for these computations. Opening is specifically implemented using the function `cv2.morphologyEx` using `cv2.MORPH_OPEN` as the second input. The right image in Fig. 7 show the result when opening is used on the thresholded




Fig. 7. A example comparison of distance transform images with and without opening for the Felix GOES-12 data set. At left, the noise in the center of the hurricane causes the distance transform to fill in. At right, performing opening gets rid of the small noisy point, and the distance transform does not get filled in.




Fig. 8. Maximum persistence plotted over time for all data sets using threshold $\mu = 80$ in addition to the versions using opening to remove noise. Gray vertical lines separate days according to UTC.

matrix and then the distance transform is applied. Since the distance transform is no longer filled in, the opening process has removed the noisy pixels causing the issue.

Using this extra step in the method, we recalculate maximum persistence for all times and compute the estimated period of the new maximum persistence values using Fourier transforms.

Fig. 8 shows maximum persistence plotted versus time for using our original method, and the method including the additional opening step. While the new maximum persistence values vary a little from the originals, the general oscillatory behavior seems similar. For both the Felix and Ivan GOES-12 data sets, the Fourier transform still detects a 24.5 and 24.0 h cycle respectively. Thus the presence of noise in these data sets is not impacting the results. However, for the Felix GridSat-GOES data set, the Fourier transform now detects a 15.375 h cycle, likely due to the difference in spatial resolution. The GOES-12 data has higher spatial resolution, so applying opening to remove noise does not impact the global circular structure. The GridSat-GOES data has lower spatial resolution, and is therefore more sensitive to noise in the image. Thus, our method is more reliable when applied to higher spatial resolution data, and should be used with caution on lower quality data.

6. Discussion

This paper presents a novel method for detecting and analyzing the diurnal cycle of tropical cyclones using methods from TDA. Current state of the art TC diurnal cycle measurement in the satellite imagery is mostly qualitative; our method provides a mathematically advanced method for automatic detection and measurement. While our method involves a choice of a parameter for the threshold, we present evidence that a range of threshold values yield the same results.

Here, we show that using two sets of GOES satellite data for Hurricane Felix and one set of data for Hurricane Ivan, our method is able to detect almost identical patterns across all three data sets. Our method performs more consistently and robustly on higher spatial resolution data sets, represented by the GOES-12 data for

both Hurricanes Felix and Ivan. While the method does detect a daily cycle in the lower resolution Felix GridSat-GOES data, when the images are blurred to remove noise, the cycle is no longer detected by the Fourier transform.

This is a novel application of methods from TDA and image processing to the TC diurnal cycle. We believe this method could be used to study additional atmospheric phenomena exhibiting circular structure. A future direction of this project is to apply this analysis to more TCs, other satellite channels and other atmospheric data to further test our method.

Declaration of Competing Interest

None.

Acknowledgments

The authors would like to thank William Dong, who was instrumental in creating code for the first iteration of this project. The authors also thank the two anonymous reviewers whose suggestions improved the quality of the paper. ST and EM were supported in part by NSF grants [DMS-1800446](#) and [CMMI-1800466](#); EM was also supported in part by NSF [CCF-1907591](#). JD was supported in part by the Office of Naval Research Grant N000141410132 and by NOAA Grant [NA14OAR4830172](#) from the Unmanned Aircraft Systems (UAS) Program. KC was supported by NSF Grant [AGS-1636799](#).

References

- [1] C. Giusti, E. Pastalkova, C. Curto, V. Itskov, Clique topology reveals intrinsic geometric structure in neural correlations, *Proc. Natl. Acad. Sci.* (2015), doi:[10.1073/pnas.1506407112](#).
- [2] A. Deckard, R.C. Anafi, J.B. Hogenesch, S.B. Haase, J. Harer, Design and analysis of large-scale biological rhythm studies: a comparison of algorithms for detecting periodic signals in biological data, *Bioinformatics* 29 (24) (2013) 3174–3180, doi:[10.1093/bioinformatics/btt541](#).
- [3] J.A. Perea, J. Harer, Sliding windows and persistence: An application of topological methods to signal analysis, *Found. Comput. Math.* (2015) 1–40, doi:[10.1007/s10208-014-9206-z](#).
- [4] F.A. Khasawneh, E. Munch, Chatter detection in turning using persistent homology, *Mech. Syst. Signal Proc.* 70–71 (2016) 527–541, doi:[10.1016/j.ymssp.2015.09.046](#).
- [5] G. Carlsson, T. Ishkhanov, V. de Silva, A. Zomorodian, On the local behavior of spaces of natural images, *Int. J. Comput. Vis.* 76 (2008) 1–12, doi:[10.1007/s11263-007-0056-x](#).
- [6] V. Robins, P.J. Wood, A.P. Sheppard, Theory and algorithms for constructing discrete morse complexes from grayscale digital images, *IEEE Trans. Pattern Anal. Mach. Intell.* 33 (8) (2011) 1646–1658, doi:[10.1109/TPAMI.2011.95](#).
- [7] V. Robins, M. Saadatfar, O. Delgado-Friedrichs, A.P. Sheppard, Percolating length scales from topological persistence analysis of micro-CT images of porous materials, *Water Resour. Res.* 52 (1) (2016) 315–329, doi:[10.1002/2015wr017937](#).
- [8] M. Saadatfar, H. Takeuchi, V. Robins, N. Francois, Y. Hiraoka, Pore configuration landscape of granular crystallization, *Nat. Commun.* 8 (15082) (2017), doi:[10.1038/ncomms15082](#).
- [9] J.P. Dunion, C.D. Thorncroft, C.S. Velden, The tropical cyclone diurnal cycle of mature hurricanes, *Mon. Weather Rev.* 142 (10) (2014) 3900–3919.
- [10] J.P. Dunion, C.D. Thorncroft, D.S. Nolan, Tropical cyclone diurnal cycle signals in a hurricane nature run, *Mon. Weather Rev.* 147 (1) (2019) 363–388.
- [11] S.D. Ditchek, K.L. Corbosiero, R.G. Fovell, J. Molinari, Electrically active tropical cyclone diurnal pulses in the atlantic basin, *Mon. Weather Rev.* 147 (10) (2019) 3595–3607, doi:[10.1175/MWR-D-19-0129.1](#).
- [12] S.D. Ditchek, J. Molinari, K.L. Corbosiero, R.G. Fovell, An objective climatology of tropical cyclone diurnal pulses in the atlantic basin, *Mon. Weather Rev.* 147 (2) (2019) 591–605, doi:[10.1175/MWR-D-18-0368.1](#).
- [13] J.P. Kossin, Daily hurricane variability inferred from goes infrared imagery, *Mon. Weather Rev.* 130 (9) (2002) 2260–2270.
- [14] K.D. Leppert, D.J. Cecil, Tropical cyclone diurnal cycle as observed by trmm, *Mon. Weather Rev.* 144 (8) (2016) 2793–2808.
- [15] E.L. Navarro, G.J. Hakim, Idealized numerical modeling of the diurnal cycle of tropical cyclones, *J. Atmos. Sci.* 73 (10) (2016) 4189–4201.
- [16] M.E. O'Neill, D. Perez-Betancourt, A.A. Wing, Accessible environments for diurnal-period waves in simulated tropical cyclones, *J. Atmos. Sci.* 74 (8) (2017) 2489–2502.
- [17] J. Steranka, E.B. Rodgers, R.C. Gentry, The diurnal variation of atlantic ocean tropical cyclone aoud distributioninferred from geostationary satellite infrared measurements, *Mon. Weather Rev.* 112 (11) (1984) 2338–2344.
- [18] X. Tang, F. Zhang, Impacts of the diurnal radiation cycle on the formation, intensity, and structure of hurricane edouard (2014), *J. Atmos. Sci.* 73 (7) (2016) 2871–2892.
- [19] H. Edelsbrunner, J. Harer, *Computational Topology: an Introduction*, American Mathematical Society, 2010.
- [20] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [21] T. Kaczynski, K. Mischaikow, M. Mrozek, *Computational Homology*, 157, Springer Science & Business Media, 2006.
- [22] E. Munch, A user's guide to topological data analysis, *J. Learn. Anal.* 4 (2) (2017), doi:[10.18608/jla.2017.42.6](#).
- [23] M.A. Lazzara, J.M. Benson, R.J. Fox, D.J. Laitsch, J.P. Rueden, D.A. Santek, D.M. Wade, T.M. Whittaker, J. Young, The man computer interactive data access system: 25 years of interactive processing, *Bull. Ame. Meteorolog. Soc.* 80 (2) (1999) 271–284.
- [24] K.R. Knapp, S.L. Wilkins, Gridded satellite (GridSat) GOES and CONUS data, *Earth Syst. Sci. Data* 10 (3) (2018) 1417–1425, doi:[10.5194/essd-10-1417-2018](#).
- [25] A. Rosenfeld, J.L. Pfaltz, Sequential operations in digital picture processing, *J. ACM* 13 (4) (1966) 471–494, doi:[10.1145/321356.321357](#).
- [26] A. Rosenfeld, J. Pfaltz, Distance functions on digital pictures, *Pattern Recognit.* 1 (1) (1968) 33–61, doi:[10.1016/0031-3203\(68\)90013-7](#).
- [27] K. Mischaikow, V. Nanda, Morse theory for filtrations and efficient computation of persistent homology, *Discrete Comput. Geometry* 50 (2) (2013) 330–353.
- [28] V. Nanda, Perseus, 2015, <http://www.sas.upenn.edu/~vnanda/perseus/>.
- [29] E.O. Brigham, *The Fast Fourier Transform and Its Applications*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.
- [30] J. Serra, *Image Analysis and Mathematical Morphology*, Image Analysis and Mathematical Morphology, Academic Press, 1984.