Extreme Value Analysis of Tropical Cyclone Trapped-Fetch Waves

Allan W. MacAfee and Samuel W. K. Wong

National Laboratory for Marine and Coastal Meteorology, Meteorological Service of Canada, Dartmouth, Nova Scotia, Canada

ABSTRACT

Many of the extreme ocean wave events generated by tropical cyclones (TCs) can be explained by examining one component of the spectral wave field, the trapped-fetch wave (TFW). Using a Lagrangian TFW model, a parametric model representation of the local TC wind fields, and the National Hurricane Center's hurricane database archive, a dataset of TFWs was created from all TCs in the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea from 1851 to 2005. The wave height at each hourly position along a TFW trajectory was sorted into $2° \times 2°$ latitude–longitude grid squares. Five grid squares (north of Hispaniola, Gulf of Mexico, Carolina coast, south of Nova Scotia, and south of Newfoundland) were used to determine if extreme value theory could be applied to the extremes in the TFW dataset. The statistical results justify accepting that a generalized Pareto distribution (GPD) model with a threshold of 6 m could be fitted to the data: the datasets were mostly modeled adequately, and much of the output information was useful. Additional tests were performed by sorting the TFW data into the marine areas in Atlantic Canada, which are of particular interest to the Meteorological Service of Canada because of the high ocean traffic, offshore drilling activities, and commercial fishery. GPD models were fitted, and return periods and the 95% confidence intervals (CIs) for 10-, 15-, and 20-m return levels were computed. The results further justified the use of the GPD model; hence, extension to the remaining grid squares was warranted. Of the 607 grid squares successfully modeled, the percentage of grid squares with finite lower (upper) values for the 10-, 15-, and 20-m return level CIs were 100 (80), 94 (53), and 90 (16), respectively. The lower success rate of 20-m TFW CIs was expected, given the rarity of 20-m TFWs: of the 5 713 625 hourly TFW points, only 13 958, or 0.24%, were 20 m or higher. Overall, the distribution of the successfully modeled grid squares in the data domain agreed with TFW theory and TC climatology. As a direct result of this study, the summary datasets and return level plots were integrated into application software for use by risk managers. A description of the applications illustrates their use in addressing various questions on extreme TFWs.

1. Introduction

Tropical cyclones (TCs) that generate extreme ocean waves in midlatitudes (e.g., Luis, in 1995; Danielle, in 1998; and Juan, in 2003) exhibit two common features: the TC traveled in a straight line for at least 18 h and the speed exceeded 10 m s^{-1}. As explained in Bowyer and MacAfee (2005, hereinafter BM5), the continued forcing of the waves by the wind from the synchronicity of the TC's motion and the acceleration of the growing waves results in a trapped-fetch wave (TFW) situation whereby the significant wave height H_s (the average

height from trough to crest of the one-third highest waves) can become much greater than for waves generated by a stronger but slower-moving TC. While TFWs represent only one component of the spectral wave field, and optimum synchronicity is extremely sensitive to TC speed, case studies showed that observed extreme waves can be predicted by a simple Lagrangian TFW model (MacAfee and Bowyer 2005, hereinafter MB5). Such a model is used operationally at the Canadian Hurricane Centre (CHC) to predict extreme wave events with TCs affecting Atlantic Canada.

The CHC TFW model uses the significant wave method of Bretschneider and Tamaye (1976) and wind fields created by a TC parametric wind model as described in MB5. The wind model inputs are TC track position and intensity (maximum wind and central pressure). Model-specific parameters (e.g., radius of maxi-

Corresponding author address: Allan W. MacAfee, National Laboratory for Marine and Coastal Meteorology, Meteorological Service of Canada, 45 Alderney Dr., Dartmouth, NS B2Y 2N6, Canada.
E-mail: al.macafee@ec.gc.ca

FIG. 1. Marine areas in Atlantic Canada subjected to GPD analysis. The area names corresponding to each number are listed in Table 1.

TABLE 1. Atlantic Canada marine areas depicted in Fig. 1.

Identification No.	Area name	Identification No.	Area name
1	Fundy	11	East Scotian slope
2	Grand Manan	12	Fourchu
3	Lurcher	13	Banquereau
4	Browns Bank	14	Laurentian Fan
5	Georges Bank	15	Southwest coast
6	Southwestern shore	16	South coast
7	Lahave Bank	17	East coast
8	West Scotian slope	18	Northern Grand Banks
9	Eastern shore	19	Southeastern Grand Banks
10	Sable	20	Southwestern Grand Banks

mum winds) are set using statistical methods computed using the input parameters (MacAfee and Pearson 2006). Substituting position and intensity data from the National Hurricane Center's hurricane database archive (HURDAT; Jarvinen et al. 1984), the same method can be used to generate wind fields and TFWs along any historical TC. The resulting dataset, covering the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea, contains examples of extreme waves (e.g., 32-m H_s, generated by Debby in 1982). A logical question is, how often do extreme waves from TCs occur in particular areas of the Atlantic Ocean, Gulf of Mexico, or Caribbean Sea?

Modeling the occurrence of extreme events is of recognized importance because of the potentially adverse effects of the resulting circumstances (Easterling et al. 2000). A branch of statistics, namely extreme value theory, has been developed in the past century to analyze the frequency and severity of such extreme events. Early references on this subject include Fisher and Tippett (1928) and Gumbel (1958). Research in the past few decades has found more robust and sophisticated techniques for modeling extreme phenomena (e.g., Pickands 1975; Smith 1984).

The study of rare climate and weather events is an active area of research, including recent work on extreme precipitation, wind speeds, and air quality (e.g., Smith 1989; Brabson and Palutikof 2000; Beguería and Vicente-Serrano 2006; Jagger and Elsner 2006). Hindcast ocean wave datasets based on kinematically reconstructed wind fields, with particular attention to TC winds, were examined for extreme events and trends (Swail et al. 2000; Swail and Cox 2000; Wang and Swail 2002). Van Gelder et al. (2000) examined the tails of wave-height distribution statistics. Recently, return periods for high waves in the Gulf of Mexico based on

observations during Hurricane Ivan were determined by Panchang and Li (2006).

In this study, such analysis is performed on the historical synthetic TFW data generated using HURDAT information by creating two deriving datasets. First, TFW hourly data positions were sorted into 2° × 2° latitude–longitude grid squares covering the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea. The TFW hourly positions were also sorted into the marine areas in Atlantic Canada (Fig. 1 and Table 1), which are of particular interest to the Meteorological Service of Canada (MSC) because of the high ocean traffic, offshore drilling activities, and commercial fishery. MSC issues forecasts and warnings for these areas, and an assessment of the suitability of using the extreme value analysis as input during TC events was warranted. Second, for each grid square and marine area, time series spanning the June to November hurricane season were extracted.

Using these datasets a suitable statistical model was developed and tested for five separate grid squares in distinct parts of the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea with observed high-wave activity. Additional tests were performed using the marine areas. After evaluation, the model was fitted for all grid squares in the study region, so that the output inference on return periods and levels in different grid squares could be interactively compared by risk managers, using tools specifically designed for this purpose.

Creation, contents, and mining of the TFW dataset are outlined in section 2. A review of extreme value theory and selection of appropriate methods is presented in section 3. Application of these techniques to the TFW dataset appears in section 4. The results and discussion of the analysis are given in section 5. A brief description of the applications used by risk managers to

FIG. 2. Schematic diagram showing the successive steps in modeling TFWs at a track loca-
tion (circled and labeled TL): (a) construct a model grid right-of-track, oriented parallel to the
storm motion $\mathbf{V}_{st}$ (heavy arrow); (b) create a wind field using a parametric wind model; (c)
using each grid point compute a TFW, then select the dominant TFW from each row. The
dominant trajectory for the TL is denoted by a heavier arrow.

display the TFW trajectories and results of the extreme
value analysis are presented in section 6. A summary
and conclusions complete the paper in section 7.

2. TFW dataset

Several TFW datasets were used in this study, and
the construction and use of each dataset are described
in the following sections.

a. Trajectory dataset

HURDAT (information online at http://www.nhc.
noaa.gov/pastall.shtml) provided track latitude and lon-
gitude, time (UTC), maximum wind speed (kt; 1 kt =
0.5144 m s^{-1}), and minimum central pressure (mb) at
6-h intervals for all TCs in the Atlantic Ocean, Gulf of
Mexico, and Caribbean Sea from 1851 to 2005. The 6-h
data for each recorded TC was linearly interpolated to
hourly values.

As described in MB5 and illustrated in Fig. 2a, a grid
(37 × 30; 0.1° latitude resolution) aligned along the
instantaneous TC motion vector and right-of-track was
defined at each hourly TC position, as indicated by the
circle and TL label in Fig. 2a.

Surface wind observations over the ocean are limited
to buoys, ships, and, more recently, synthetic winds
constructed from satellite data (e.g., Chelton et al.
2006). Reanalysis projects (e.g., Kalnay et al. 1996)
have created surface wind fields; however, these grids
neither span the period of study nor provide the winds
at the desired resolution; hence, a parametric model
was used. MacAfee and Pearson (2006) describe sev-
eral parametric wind model formulations, various tech-

niques to set model-specific parameters, and validation against 2D wind fields (H*WIND; Powell and Houston 1996) and buoy data. Based on the validation results, the *enhanced* method and Holland model were selected to create the wind vectors illustrated in Fig. 2b.

MB5 describes the Lagrangian TFW model and shows validation of the computed wave heights against buoy data, scanning radar altimetry data, and full spectral wave model output. Operationally, the TFW model output is included in CHC forecast bulletins. Hence, we consider the model a valid approach for estimating extreme waves with TCs. Using the wind fields along the track, the dominant trajectories for each row across each hourly grid were computed as illustrated in Fig. 2c. Each grid point seeds a new set of trajectory calculations. As the wave traverses downstream grids, the wave height was computed in 1-h time steps using the local wind field, ignoring wave–wave interaction. The output data consisted of hourly values along each trajectory of the wave's latitude and longitude, H_s (m), period (s), length or distance over which growth occurred (n mi; 1 n mi = 1.852 km), the duration or elapsed time (h) during which TFW growth occurred to reach that location, and the angle (°True) of the trajectory.

The preceding grid construction and wind field and trajectory calculations were repeated for all TCs from 1851 to 2005. In total, the dataset consisted of 5 713 625 hourly wave points: 13.6% of the wave heights were 10 m or higher, 2.4% 15 m or higher, and 0.24% 20 m or higher.

Note that parametric wind models are formulated using observed data, and the wind field is considered representative of a TC with a climatological size–intensity relationship and a single eyewall: TCs with symmetric wind fields or with concentric eyewalls may not be adequately depicted (MacAfee and Pearson 2006). Different formulations may result in variations in the wind field, and TFW modeling is sensitive to these variations (BM5); however, the difference in wave heights is typically < 10% (MB5). Thus, we consider our computed TFWs to be a realistic, but synthetic, representation of the extreme waves in TCs from 1851 to 2005.

This trajectory information formed the basis of the two deriving datasets created for study and analysis: first, the sorted dataset for overall risk and data visualization; second, the time series dataset for the statistical methods in sections 3–5.

b. Sorted dataset for risk assessment

To create risk-potential datasets, the trajectories were sorted into 2° × 2° latitude–longitude grid squares, as illustrated in Fig. 3a. Using the latitude and longitude, each hourly position along a given trajectory was assigned to a grid square. The recorded TFW model outputs were then sorted into categories: H_s 1 m, period 1 s, duration 1 h, length 5 n mi, and angle 5°. If successive hourly positions on the same trajectory fell in the same grid square, the position with the highest wave growth was recorded, ignoring the earlier positions and values. In addition, the name, year, and highest H_s for each TC contributing to any given grid square were recorded on a grid square–by–grid square basis. The grid square category data (counts of H_s, period, length, duration, and angle), the contributing TC information (name, year, and grid square–specific maximum H_s), and the extracted extreme maximum wave and direction were sorted by month and recorded in a plain-text file.

Figure 3a shows the grid squares for September, color coded based on the highest TFW in the grid square. Figure 3b shows the frequency of TFW H_s for the grid square south of Newfoundland, centered at 41°N, 57°W and identified by the X in Fig. 3a. Note that the TFW tail beyond 20 m was generated by six TCs: the second TC in 1857; the fourth TC in 1870; the fifth TC in 1899; the tenth TC in 1916; the fourth TC in 1949; and Debby in 1982. This plot uses all trajectory–grid square intersections, ignoring the tendency of multiple TFWs from upstream grids for the same TC to arrive in the grid square at approximately the same time. Hence, these frequency plots acknowledge the existence of extreme waves, but inferences from this initial analysis are limited. The main value of this analysis is to identify grid squares with significant and extreme TFW activity that can be studied further using more rigorous data extraction techniques and statistical methods.

Using polygonal outlines for each marine area, the sorting procedure was repeated to assess if such a dataset was useful to forecast operations. Overall, the TFW probability data for the larger marine areas were useful but lacked the detail provided by the individual grid squares within a given area.

c. Time series dataset for extreme value analysis

To create datasets suitable for the application of statistical methods, first the analysis time period had to be chosen. With the continuing reanalysis of the HURDAT database (Landsea et al. 2002; Feuer et al. 2004) to improve its precision and relevance, it seemed appropriate to make use of the entire duration of the trajectory dataset, recognizing inherent weaknesses as stated by Landsea et al. (2006): "Tropical cyclone 'best track' data sets are finalized annually by operational meteorologists, not by climate researchers, and none of
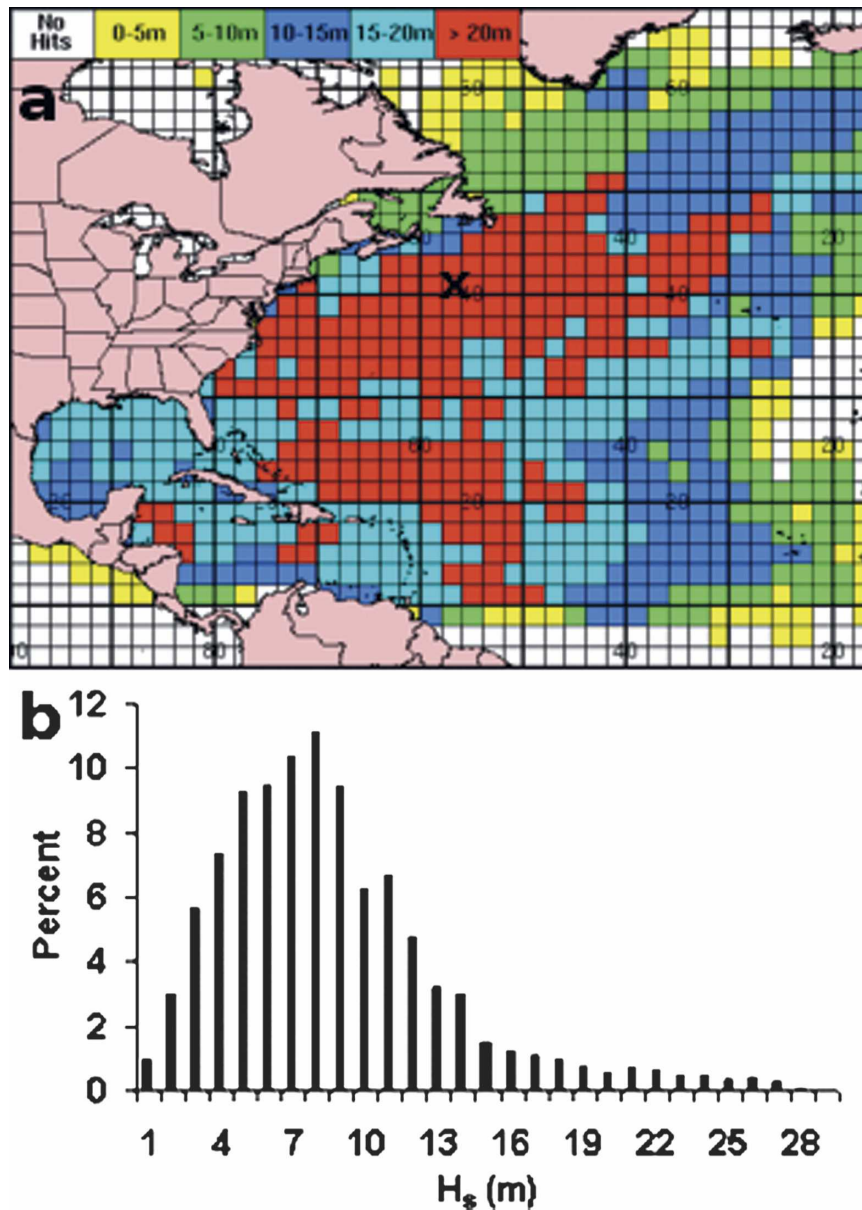
FIG. 3. (a) Sorted 2° × 2° lat–lon grid squares of TFWs for September. (b) Frequency plot of TFW heights (m) for the grid square centered at 41°N, 57°W, indicated by the X in (a).

the data sets have been quality controlled to account for changes in physical understanding, new or modified methods for analyzing intensity, and aircraft/satellite data changes." From a statistical viewpoint the longer time span provided more precise model estimates, that is, with lower variance. Landsea (1993) described the interannual variability in hurricane activity. As such, the longest data series possible is necessary to yield the most robust estimates.

Initial data mining was performed on the output trajectories to sort the wave events by time and location.

In this process, the data were reformatted into an hourly time series of H_s in each grid square. Since the majority of the time spans do not have TFW activity, zeros were used as placeholders for these hours. This gives all time series a consistent length and allows the time intervals between successive TFW events to be examined. The zeros do not add to the number of extreme waves and hence will not affect later estimates on the distribution of extreme wave heights. Further, TC activity is seasonal, with the Atlantic hurricane season running from 1 June to 30 November of each year. As

FIG. 4. Plot of hourly TFW heights for the grid square centered at 41°N, 57°W.

observations of TCs from the remainder of each year largely do not contribute to the understanding of extreme wave events (Landsea 1993, his Table 3), events outside of the hurricane season were eliminated from the time series. This resulted in 183 days (4392 hourly values) per year being included in this study.

As an example, consider the plot of wave heights over time for the grid square south of Newfoundland, (Fig. 4). Each point represents one hourly observation. In this particular raw time series, there were 1832 hourly values that exceeded a height of 5 m.

3. Extreme value theory

Extreme value theory is concerned with the distribution of rare events, rather than usual occurrences. For TFW data, the rare events are the exceptionally high waves. A method to model the frequency and severity of the waves requires some background statistical theory. The complete development is mathematically demanding, so the interested reader may refer to the referenced accounts.

Let $X_1, X_2, \ldots, X_n$ be a sequence of independent and identically distributed (iid) random variables drawn from a continuous distribution function F. The precise underlying distribution F may not be known. For example, this sequence may be a time series with n observations, where each observed value does not depend on its neighbors, with the overall distribution of the observations remaining stable over time. The observations of interest here are not the bulk of the dataset but rather the unusually large values, so let $M_n = \max(X_1, X_2, \ldots, X_n)$. This will be the maximum value observed after n periods. Then, for the distribution of M_n, for any value x,

$$
\begin{aligned}
\Pr(M_n \leq x) &= \Pr(X_1 \leq x, X_2 \leq x, \ldots, X_n \leq x) \\
&= \Pr(X_1 \leq x) \times \Pr(X_2 \leq x) \times \cdots \times \Pr(X_n \leq x) \\
&= \Pr(X_n \leq x)^n \\
&= F^n(x). \quad (1)
\end{aligned}
$$

For inference to be made on M_n, Fisher and Tippett (1928) looked for sequences of real constants $a_n > 0$, b_n such that $(M_n - b_n)/a_n$ would converge in distribution to a nontrivial distribution function H. With some rearrangement and the above, this can be written as

$$
\lim_{n \to \infty} \Pr\left(\frac{M_n - b_n}{a_n} \leq x\right) = \lim_{n \to \infty} \Pr(M_n \leq a_n x + b_n)
$$
$$
= \lim_{n \to \infty} F^n(a_n x + b_n) = H(x). \quad (2)
$$

As detailed by Gnedenko (1943), in the extremal types theorem, if H exists, H must belong to the generalized extreme value (GEV) distribution. The GEV can be parameterized using the location (μ), scale (σ), and shape (ξ) as

$$
H(x) = \exp\left\{-\left[1 + \xi \frac{(x - \mu)}{\sigma}\right]^{-1/\xi}\right\}, \quad (3)
$$

where $[1 + \xi(x - \mu)/\sigma] > 0$, and $\sigma > 0$. This one expression encompasses three families of more familiar distributions, namely the Gumbel, Frechet, and Weibull. The Gumbel case can be obtained by taking the limit $\xi \to \infty$. Gumbel (1958) showed that essentially all common continuous distributions satisfy the condition that the sequences $a_n > 0$, b_n exist. Satisfying this condition is also known as being in the domain of attraction of an extreme value distribution. For modeling purposes, maximum values may usually be acceptably fitted to the GEV, because of the convergence properties outlined above.

This classical theory on the GEV has been used widely to make inferences on data that consist of a series of maxima, such as annual maxima, on various environmental data (Hosking et al. 1985). A natural extension to this approach would involve using more data than simply the maxima. For example, with regular measurements being taken, it may be useful to incorporate a larger collection of extreme data points for modeling purposes. Pickands (1975) proposed a method for using the largest m observations from a dataset. The choice m is assumed to be much smaller than the number of observations n but typically large enough to include more data than a GEV model. Following Pickands (1975), a starting assumption was for-

mulated for the arbitrary continuous distribution function F and a real constant c that

$$\lim_{u \to \inf[x:F(x)=1]} \inf_{a>0} \sup_{x\geq0} \left| \frac{1 - F(u + x)}{1 - F(u)} \right.$$
$$\left. - \exp\left[-\int_0^{x/a} \frac{1}{(1 + ct)_+} \right] \right| = 0, \quad (4)$$

where $(1 + ct)_+ = \max(0, 1 + ct)$, and it was shown that this condition is equivalent to F being in the domain of attraction of an extreme value distribution. Note that

$$\Pr(X - u \geq x | X \geq u) = \frac{1 - F(u + x)}{1 - F(u)} \quad (5)$$

is the conditional probability that X exceeds $x + u$ given that X exceeds u. Thus, for u sufficiently large, it follows from Pickands's assumption that, for $a > 0$,

$$\frac{1 - F(u + x)}{1 - F(u)} \approx \exp\left[-\int_0^{x/a} \frac{1}{(1 + ct)_+} \right]. \quad (6)$$

Integrating the right-hand side for the three different cases for c, namely $c > 0$, $c < 0$, and $c = 0$, yields $(1 + cx/a)^{-1/c}$ for $c \neq 0$ and $\exp(-x/a)$ for $c = 0$. Letting G denote the conditional distribution of X given it exceeds a threshold u, the amount x by which an observation exceeds u is approximately distributed as

$$G(x) = 1 - (1 + cx/a)^{-1/c}, \quad c \neq 0$$
$$G(x) = 1 - \exp(-x/a), \quad c = 0. \quad (7)$$

The observations in the dataset that are greater than u are exceedances over the threshold. This distribution function is known as the generalized Pareto distribution (GPD) where c and a are the shape and scale parameters, respectively. Therefore, for u sufficiently large, the GPD is a valid approximation for threshold exceedances in data drawn from any common continuous distribution. In applications, rather than choosing m directly, m is often inferred by choosing the threshold u. The value m is then the number of observations greater than u. In the GPD, the $c = 0$ case is simply the exponential distribution, $c > 0$ yields the ordinary Pareto distribution, and $c = 1$ reduces to the uniform distribution. The parameter c controls the shape of the distribution, with the tail weight increasing with c. For $c \geq 0$, the probability density is positive for all $x > 0$. The tail is light when $c < 0$, and the density has a finite upper endpoint, being nonzero only for $0 \leq x < |a/c|$.

The GPD has been used frequently to analyze the severity of extreme events that exceed a threshold, also known as "peaks-over-threshold" (POT) modeling (e.g., Smith 1989; Beguería and Vicente-Serrano 2006).

The m threshold exceedances $X_i - u$ when $X_i > u$ are fitted to a GPD by methods such as maximum likelihood or probability-weighted moments (Hosking and Wallis 1987). The larger portion of data used to build the model, relative to fitting maxima, allows more precise parameter estimates for a time series of a given length, as illustrated by Madsen et al. (1997). This basic GPD fit assumes that the model parameters do not vary over time; that is, all the exceedances come from the same distribution.

Returning to the original sequence of random variables, let $Y_1, Y_2, \ldots, Y_m$ denote the times of m observations that exceeded the threshold, retaining the order in which they occurred. The times between successive events are $T_i = Y_{i+1} - Y_i$, for $i = 1, 2, \ldots, m - 1$. The sequence of T_i can be modeled as the event interarrival times of a Poisson process (e.g., Beguería and Vicente-Serrano 2006). The simplest case is to use a homogeneous Poisson process and to assume that the exceedance sizes and event times are completely independent. If time dependence in the arrival intensities is present, a nonhomogeneous Poisson process may be more appropriate. Given a fitted GPD model for the threshold exceedances $X_{Yi} - u$ and a fitted homogeneous Poisson process with a constant intensity λ for the T_i, quantities such as the return period may be calculated. The return period is defined as the expected time between two consecutive events of a given magnitude, or return level. Let T be the event when an extreme value exceeding u occurs. Suppose that a return level of some $x > u$ is chosen. Given an event exceeding u, it exceeds x with probability $1 - G(x - u)$. Using the independence assumptions, the interarrival times of events exceeding x thus form a thinned Poisson process with intensity $\lambda[1 - G(x - u)]$. The time between two consecutive events is then exponentially distributed, with mean $1/\lambda[1 - G(x - u)]$. Therefore, the return period t as a function of the desired return level x is

$$t(x) = \frac{1}{\lambda[1 - G(x - u)]}. \quad (8)$$

4. Application to the TFW dataset

The extreme value theory as described above requires a number of assumptions. By using one set of parameters for the GPD of wave-height extremes in each grid square, any particular extreme is assumed to be independent of the time at which it occurs as well as wave extremes that precede or follow it. If true, the frequency distribution of extreme wave heights at each particular location should remain stable over time and not be subject to significant variation. This presents

problems for climatological data such as TFWs. As described by BM5, the synchronicity of the winds and the wave growth is sustained by the TC over many hours and long distances. The resulting series of waves can thus impact a given location over a period of hours. In the present time series, this translates to large wave heights being recorded for a string of consecutive hour values. Yet, this prolonged period of impact originates from one TC, and hence nearby hour values will in fact be statistically dependent. It would be erroneous to consider these as separate independent events.

Assuming we have treated the TFW data for hourly dependence, a homogeneous Poisson process for event arrival times will still require the time between TC events that drive the TFWs to be iid. Landsea (1993) showed that some fluctuations and variability in TC occurrence and intensity can be expected, both over different months in the same year and between different years. Thus, these same fluctuations may be mirrored in the TFW model output, as it is driven by these very winds. The effect of the annual cycle should be somewhat alleviated as a result of our restriction to the Atlantic hurricane season. This iid assumption will be checked carefully with plots and diagnostics. If a constant intensity assumption fails, further study and a more complex model may be required.

The approach taken here to address the time dependence is known as *declustering*. This procedure has been applied for a variety of environmental variables because of similar persistence in natural processes (e.g., Smith 1989; Beguería and Vicente-Serrano 2006). Suppose a threshold u has been chosen. The first time an observation exceeds u is marked as the beginning of a cluster. The cluster continues until a specified number of consecutive values fall below u. All observations within that cluster will then be treated as one distinct event in analysis. Typically, the maximum value recorded within the cluster is used as the representative value. The next cluster then begins when a subsequent threshold exceedance is encountered. This clustering interval should thus be chosen large enough to span the longest TFW event but not so long that separate TC events generating TFWs become indistinguishable. This restriction was satisfied by examining the numbers of distinct clusters found over different cluster sizes.

The choice of threshold is also an important consideration and a subject of research (e.g., Lang et al. 1999; Beguería 2005). From the theory presented above, it is necessary for u to be high enough for the GPD approximation to be valid. However, should u be too large, there will be insufficient data to fit a model. One method to determine the threshold is the mean excess plot, constructed by computing the average exceedance

TABLE 2. Characteristics of the grid squares chosen for sample analysis. N_{raw} is the number of raw TFW exceeding 5 m, and H_s is the maximum TFW height (m). The letter identifiers for each grid square correspond to panels in the figures.

Location	Figure panel identifier	Grid square center Lat (°N)	Lon (°W)	H_s (m)	N_{raw}
North of Hispaniola	a	21.0	69.0	20.66	2895
Gulf of Mexico	b	29.0	87.0	20.64	3199
Carolina coast	c	31.0	73.0	23.21	3989
South of Newfoundland	d	41.0	57.0	28.28	1834
South of Nova Scotia	e	41.0	63.0	27.13	1936

over the threshold over a series of threshold values. The resulting plot should appear approximately linear after an appropriate minimum threshold satisfying the above criteria is reached. Another method used to verify the selected threshold is to plot the fitted parameter values over a series of threshold values. When an appropriate threshold is reached, the parameter estimates should become relatively stable.

For estimating model parameters, the method of maximum likelihood (ML) will be used. Smith (1984) showed that when $c > -0.5$, the ML estimators for the GPD will be asymptotically efficient under certain regularity conditions. For estimating λ in the Poisson process, the ML estimate is simply the number of events divided by the time period.

The statistical calculations were performed on a desktop computer using the software package R with the evd library.

5. Analysis results and discussion

The methodology described in section 4 was first applied to a sample of representative grid squares that exhibited high TFW activity. The five grid squares selected from general areas in the data domain are listed in Table 2. The procedure used to fit GPDs to these grid squares is outlined below.

First, an appropriate clustering interval was chosen for each grid square. Figure 5 shows the number of TFW events present in the data plotted against the length of cluster intervals, at four different threshold values: 1, 3, 6, and 9 m. The plots confirm that there are many more events with a smaller threshold value. The number of separate events identified by the different cluster sizes decreases rapidly as cluster size increases; high waves close in time become grouped into single events. This is particularly true for the higher thresholds. It appears that the two grid squares from Atlantic Canada require a longer cluster interval for the number of events to stabilize. From TFW theory (BM5), this is
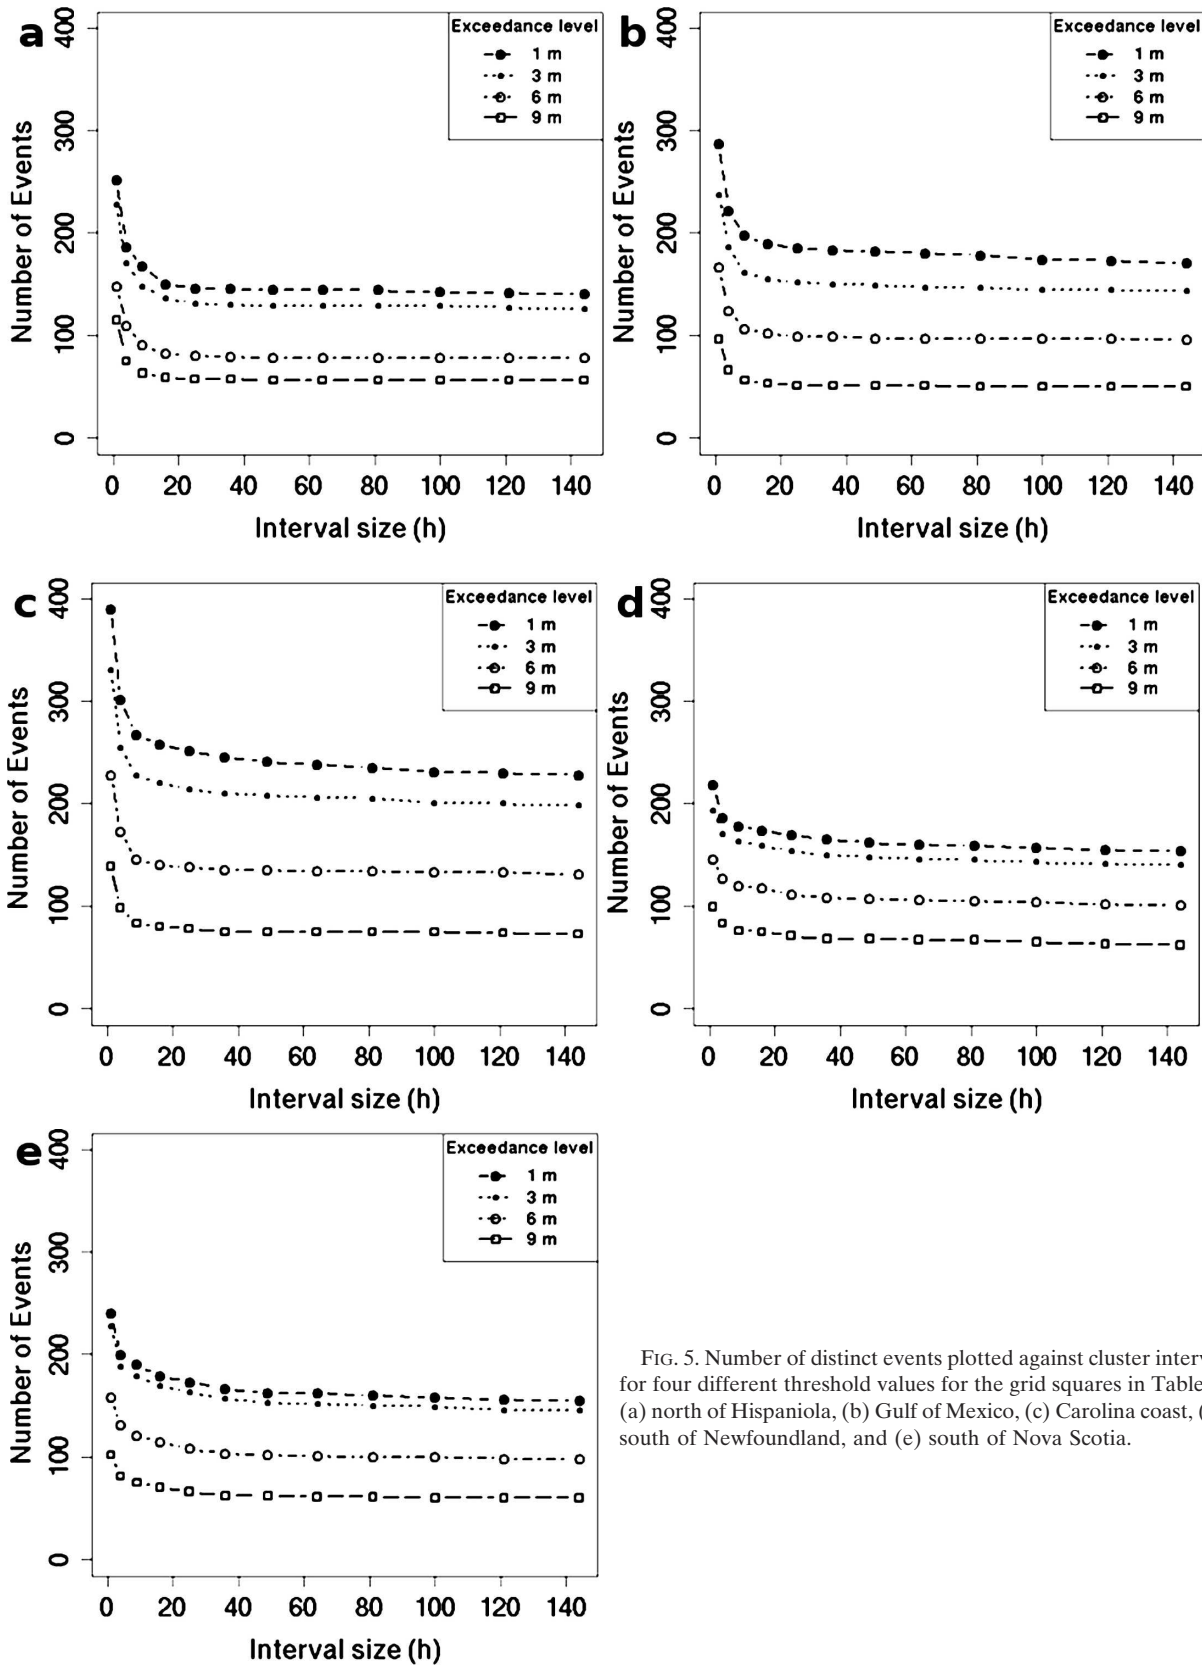
FIG. 5. Number of distinct events plotted against cluster interval for four different threshold values for the grid squares in Table 2: (a) north of Hispaniola, (b) Gulf of Mexico, (c) Carolina coast, (d) south of Newfoundland, and (e) south of Nova Scotia.
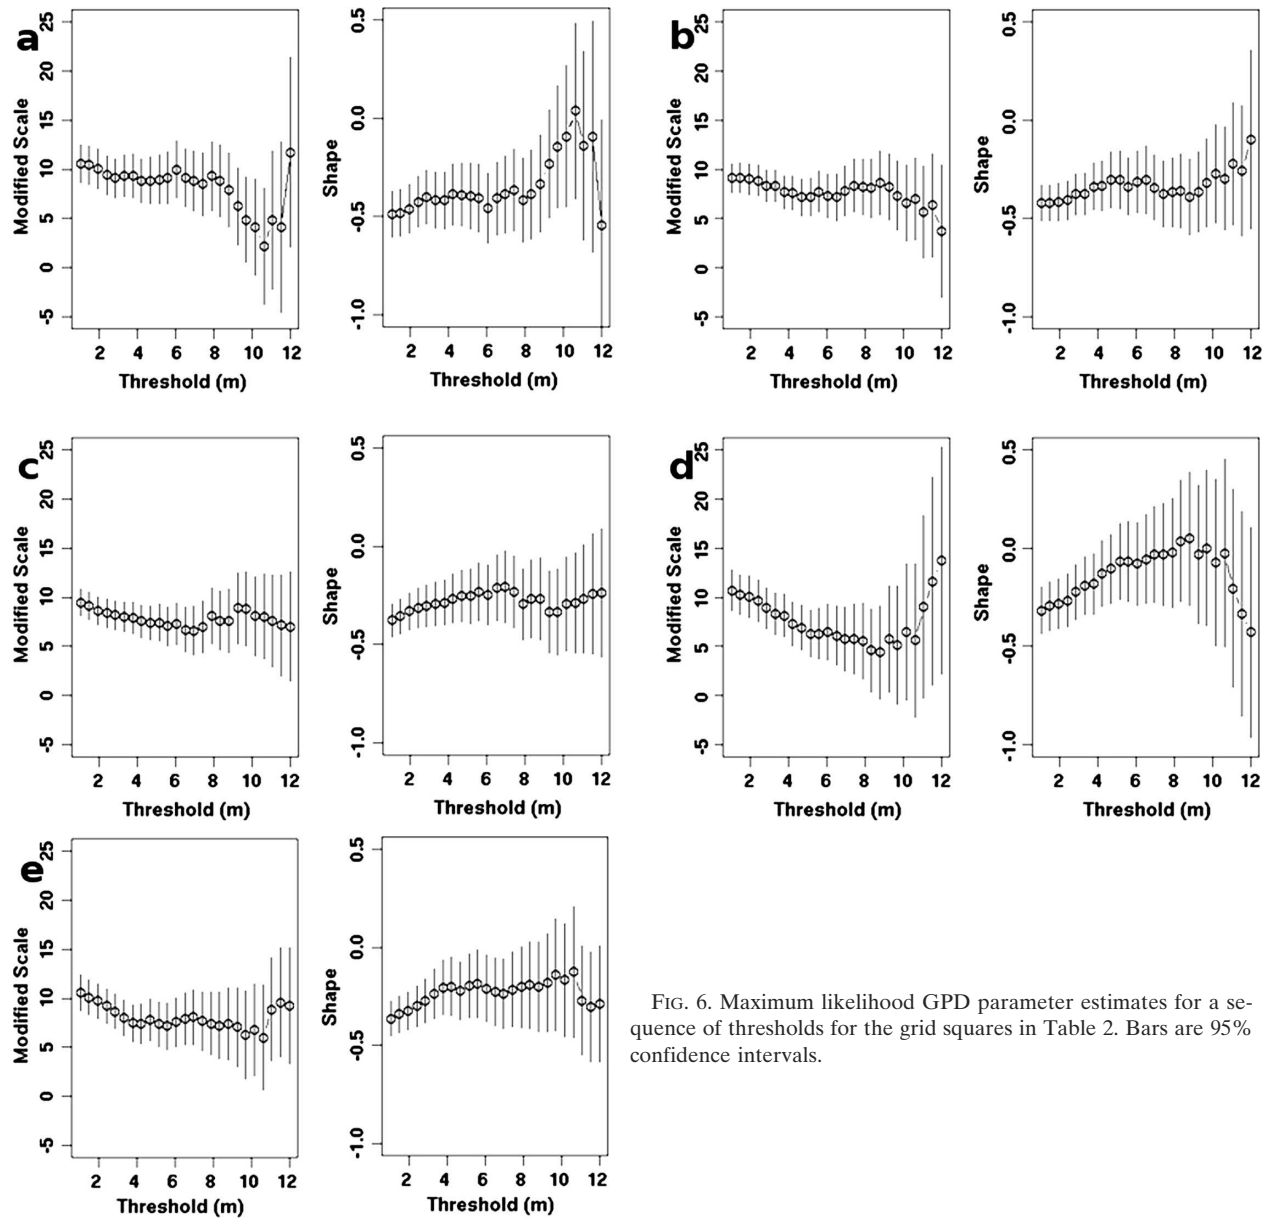
FIG. 6. Maximum likelihood GPD parameter estimates for a sequence of thresholds for the grid squares in Table 2. Bars are 95% confidence intervals.

likely due to the longer straight distance that the waves can travel during their approach to these higher latitudes. For all five grid squares, an interval of 120 h appears sufficient for declustering the TC occurrences that generated the TFWs. The number of events identified does not change when the interval is further increased to 144 h.

Second, a similar approach was used to obtain an appropriate threshold choice, given a cluster interval of 120 h. Figure 6 shows estimates of scale and shape parameters for different threshold values for the five grid squares. In the figure, scale estimates have been modified by subtracting the product of threshold and esti-

mated shape, to allow comparison across different thresholds.

The corresponding mean excess plots appear in Fig. 7. Based on Figs. 6 and 7, the model estimates appear stable in the range of 6–10 m for all five grid squares. Beyond 10 m, fluctuations are evident as the sparse available data yields large variances in the estimates, as evident by the widening 95% confidence intervals in Fig. 6. The generally downward-sloping plots in Fig. 7 indicate a relatively light tail in the data, confirmed by the negative shape parameter estimates from Fig. 6. With the possible exception of the Hispaniola grid square (Fig. 7a), the mean excesses are approximately
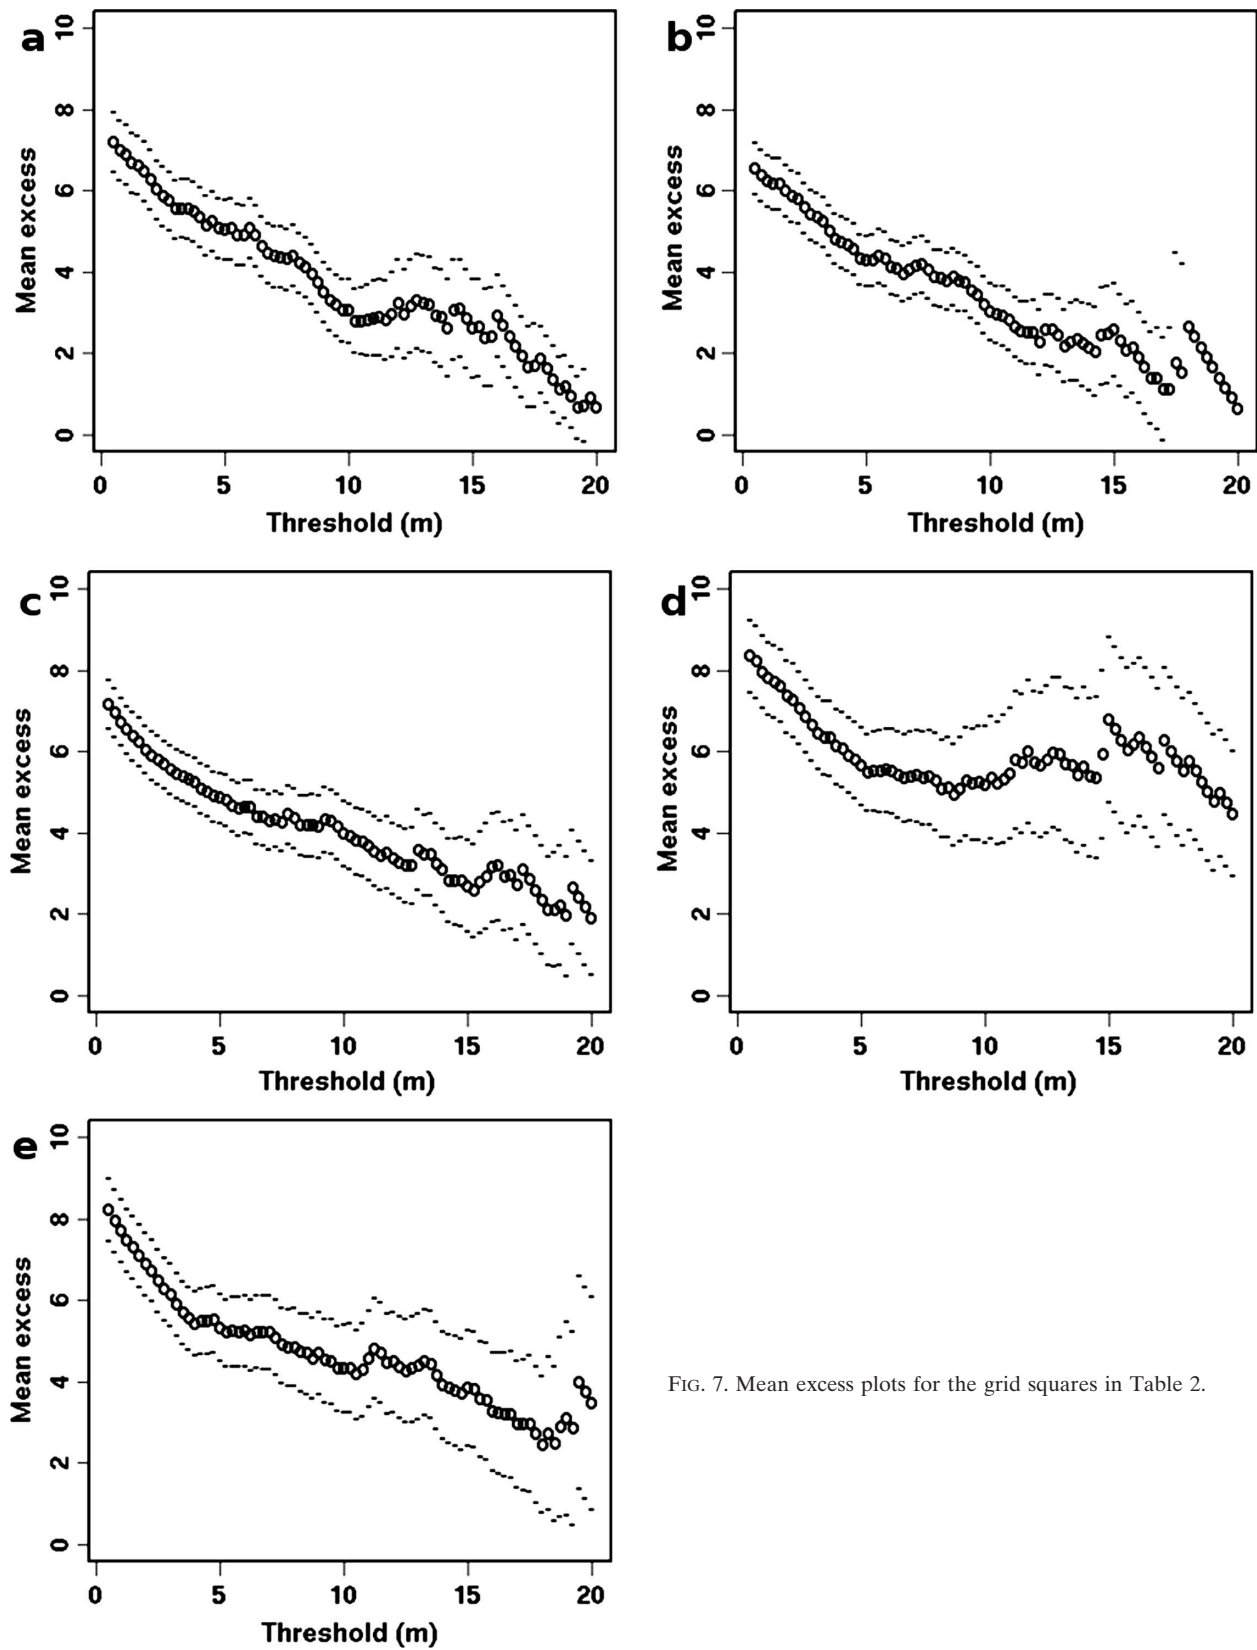
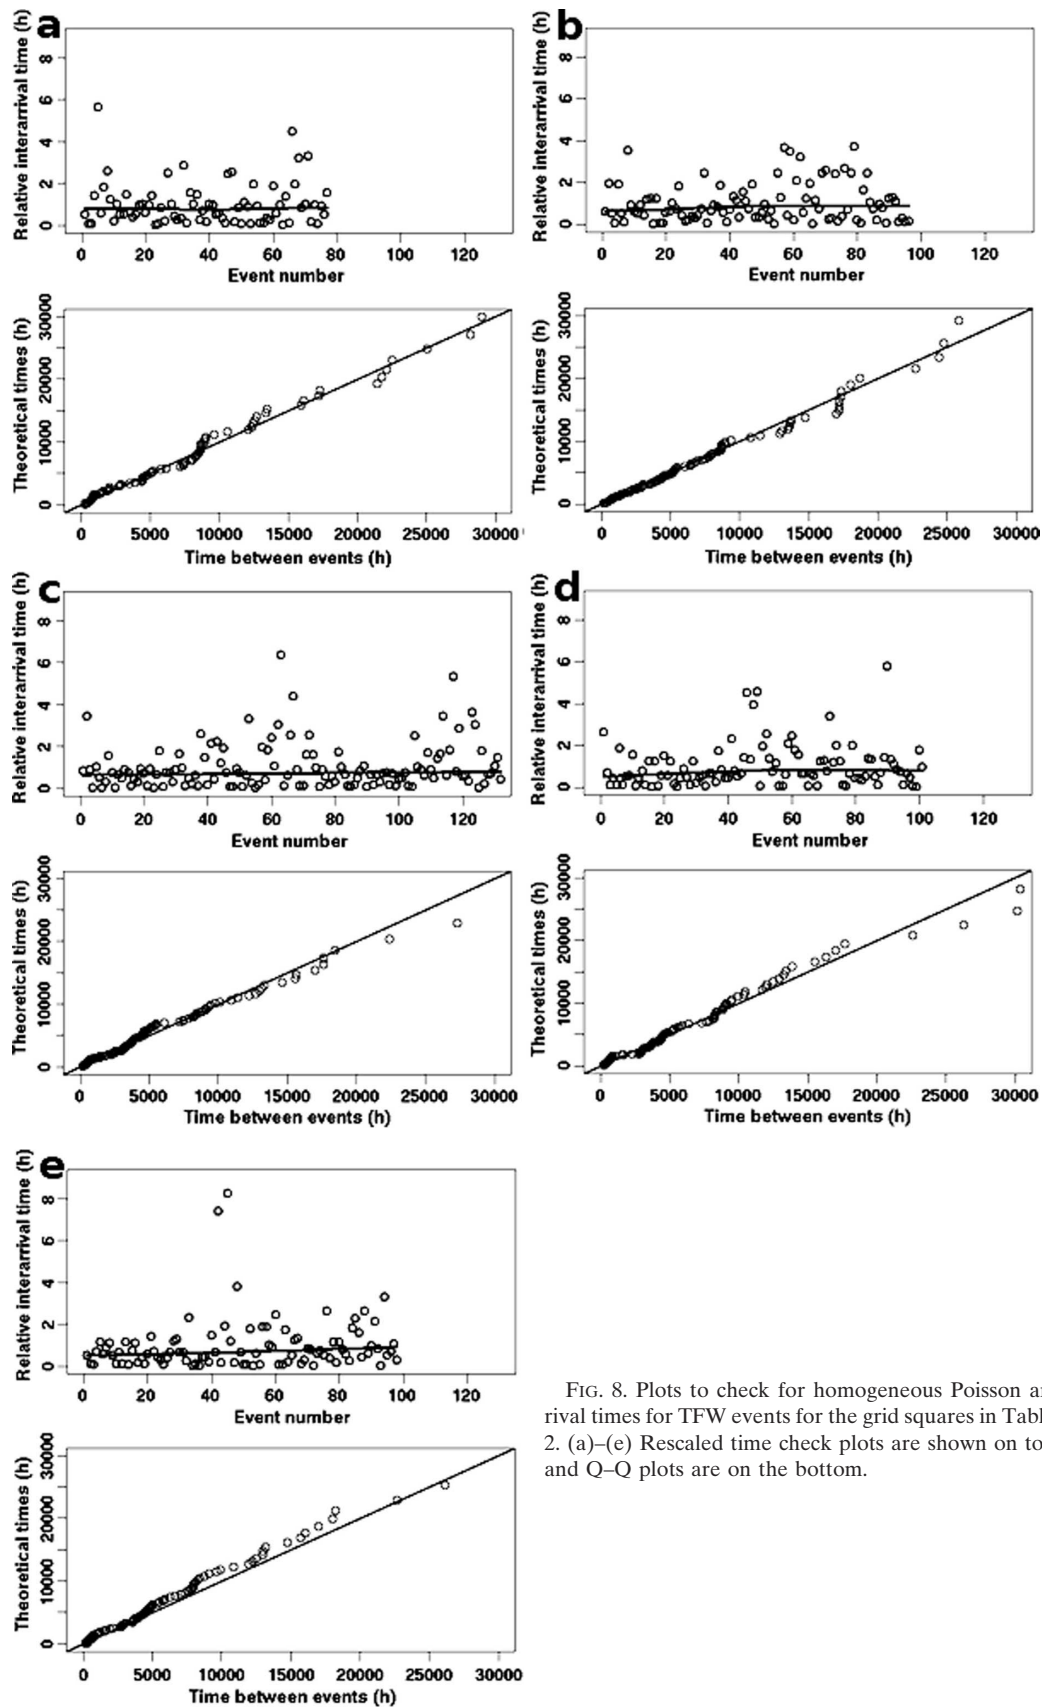FIG. 7. Mean excess plots for the grid squares in Table 2.

FIG. 8. Plots to check for homogeneous Poisson arrival times for TFW events for the grid squares in Table 2. (a)–(e) Rescaled time check plots are shown on top and Q–Q plots are on the bottom.

FIG. 9. Same as Fig. 4 except after setting the threshold ($u = 6$) and declustering.

linear beyond a threshold of 5 m, suggesting that a GPD fit should be plausible with a threshold of 5–6 m. Hence, the remainder of the fitting procedure will use 6 m as the threshold.

Third, the assumption that the event arrival time is a homogeneous Poisson process was verified. Let λ' be the ML estimate for the intensity of the process, which is simply the number of event clusters divided by the length of the time period. From probability theory, the sequence of interarrival time gaps in a homogeneous Poisson process is iid exponential with rate $1/\lambda'$. Using T_i as defined in section 3, the rescaled quantities T_i/λ' will be iid unit-exponential if the assumption holds. The rescaled values plotted in the top panel of each of Figs. 8a–e were used to check for trend over time. Since the abscissa of these plots is the event indices in time order, if there is a trend present, the distribution of the scatter points may shift from left to right across the plot. A smooth fit curve was overlaid to detect such shifts. Apart from a few visible outliers, in general, the slope of each fitted curve is very flat: the maximum percentage difference from the mean of the curve to each point of the curve does not exceed 12%, except for the Nova Scotia grid square (Fig. 8e), which exhibits a maximum of 31% difference. Apart from the Nova Scotia grid square, this suggests that the time gaps between successive TFW events remain relatively stable throughout the 155-yr period. Adding to the suspicion of the Nova Scotia grid square (Fig. 8e), it also has two extreme outliers. These indicate two very long periods in which there was no TFW activity. Further, the points in this grid square seem clustered more tightly in the lower part of the distribution near the beginning of the 155-yr period. These anomalies indicate potential model difficulties for this grid square. The quantile–quantile (Q–Q) plots of the interarrival times against the modeled exponential distribution with rate $1/\lambda'$ are shown in the

bottom panels of each of Figs. 8a–e. There is some degree of departure from the line of slope 1, most notably in the upper region of the time intervals, where the actual times between events are mostly longer than expected. Any significance of this will be left to future study. In particular, for the Nova Scotia grid square, the two outliers (two rightmost points) are very influential, pulling the Q–Q line away from the central portion of the data. The bulk of the data from the five grid squares fit adequately, so the homogeneous Poisson assumption overall seems tenable and should not need to be seriously questioned.

As an illustration of the dataset after the above handling, consider the plot, shown in Fig. 9, of wave heights over time for the Newfoundland grid square. Each circle represents one distinct event. There are 1615 hourly values exceeding 6 m, which are grouped into 102 distinct clusters.

Fourth, a GPD model was fitted by ML to the grid squares. An immediate check on the plausibility of the GPD fit to the clustered data as a whole was done using test statistics. The null hypothesis (H_0: the data came from a GPD) was tested using two standard test statistics. Specific application of the selected tests to the GPD is described in Choulakian and Stephens (2001). The Cramér–von Mises (CM) statistic is defined by

$$\sum_{i=1}^{n} [z_i - (2i-1)/(2n)]^2 + 1/(12n), \qquad (9)$$

and the Anderson–Darling (AD) statistic by

$$-n - (1/n)\sum_{i=1}^{n} (2i-1)[\ln z_i + \ln(1 - z_{n+1-i})], \qquad (10)$$

where z_i is the ith-order statistic of the dataset, transformed to a uniform (0, 1) distribution. Choulakian and Stephens (2001) include a table of critical values of these tests for a GPD. When both parameters must be estimated, the critical values are dependent on the shape parameter. For example, at $c = 0$, the acceptable range of the test statistics at the 95% level are CM $<$ 0.153 and AD $<$ 0.974, while at $c = -0.4$, the limits are CM $<$ 0.201 and AD $<$ 1.221. The fit summary for the five grid squares is shown in Table 3. The standard errors for the estimated parameters in Table 3 were computed using the numerical approximate ML information matrix, the variance of the score vector. The shape parameter c appears to be significant and negative for all grid squares except Newfoundland. This indicates a generally light tail with a finite probability density endpoint $|a/c|$. The AD and CM statistics indicate no problems for all except the Hispaniola grid square, which is slightly over the acceptable range at the 95% level. This indicates that the GPD model with

TABLE 3. GPD fit summary for the selected analysis grid squares of Table 2. The columns are threshold (u), cluster interval (r), number of exceedances in the raw dataset (N_{raw}), number of clusters (n), and estimates of scale (a) and shape (c) parameters (standard error in brackets); AD and CM refer to the test statistics described in section 5.

Location	u (m)	r (h)	N_{raw}	n	a (m)	c	AD	CM
North of Hispaniola	6	120	2490	78	7.32 (0.99)	−0.46 (0.09)	1.43	0.25
Gulf of Mexico	6	120	2526	97	5.40 (0.68)	−0.31 (0.08)	0.45	0.08
Carolina coast	6	120	3228	133	5.79 (0.66)	−0.25 (0.08)	0.31	0.04
South of Newfoundland	6	120	1615	102	5.95 (0.87)	−0.08 (0.11)	0.36	0.05
South of Nova Scotia	6	120	1679	99	6.36 (0.83)	−0.22 (0.09)	0.21	0.03

the present parameters may be a poor fit for the Hispaniola time series.

Fifth, similar to plots checking for the Poisson process, the iid distribution of wave heights over time must be examined. Following the development by Davison (1984), an empirical Laplace transform can be applied to the excess amounts over the threshold, $X_{Yi} - u$. If these values are in fact iid from the GPD with the estimated parameters, the transformation

$$z_i = \frac{1}{c} \ln\left[1 + \frac{c(X_{Yi} - u)}{a} \right] \qquad (11)$$

should yield iid unit-exponential random variables representing their relative size. The top panels of each of Figs. 10a–e show the resulting relative size values z_i plotted against the time indices, which can be interpreted in the same way as Fig. 8. These show that the slopes of the smooth curves are very flat, except for a slight downward trend in the Nova Scotia grid square (Fig. 10e). Any significance or causes of this apparent aberration will not be investigated here but left to future work. Q–Q plots of the wave-height quantiles against the unit-exponential distribution appear in the bottom panels of each of Figs. 10a–e. The Q–Q plots also show that the data as a whole fit well to the GPD model. Overall, this series of plots provides support for iid in the GPD fit of wave heights over time.

Last, return period and return level calculations were determined for the fitted models using Eq. (8) and standard bootstrap techniques. TFWs were extracted from the set with repetition to obtain new datasets. For each grid square, 1000 bootstrap samples were taken with the same size as the clustered dataset, and each of the new samples were used to generate a corresponding return level. Figure 11 shows plots of expected return periods against wave heights with goodness-of-fit points: for a good fit the points should lie close to the curve. Insets in each panel of Fig. 11 are the 95% confidence interval (CI) for 10-, 15-, and 20-m return levels, which are listed in Table 4. The 95% CI shown is the empirical 0.025 and 0.975 quantiles. The CI for 10-m events is very narrow, indicating good certainty for the

prediction. As 15- and 20-m events are rare, the resulting CIs are much wider. In grid squares where these extreme wave heights almost never occur, the resulting CI may not be useful. The computed upper endpoints for such squares are extremely large and exceed the software limit; these are shown as infinity in the results (Fig. 11 and Table 4). The fitted model provides little information and high uncertainty in these cases, as there is no practical return period estimate possible when none or very few such events occurred in the past. The Carolina coast grid square (Fig. 11c) exhibited the most frequent 10-m wave occurrence, but 20-m waves become very sparse in that grid square. The grid squares with the most frequent 20-m waves are the two selected in Atlantic Canada (Figs. 11d,e), especially south of Newfoundland (Fig. 11d). These results are not unexpected, given known TC trajectory data (Landsea 1993), TFW theory (BM5), and TFW case studies (MB5). Easterly winds in the Tropics push TCs westward south of an area of high pressure (i.e., the Bermuda high). On the west side of the high pressure, near the coast of North America, the circulation around the high directs TCs slowly northwest, where they come under the influence of increasingly strong westerly winds. These westerly winds ultimately result in TCs accelerating in a northeast direction. The Carolina coast has a high frequency of passing TCs, but, because of the prevalence for recurvature as they approach North America, conditions favorable for TFW growth are limited (BM5). Similarly, near Hispaniola (Fig. 11a) and in the Gulf of Mexico (Fig. 11b) restricted open-water distance and limits on TC speed and wave growth synchronicity, due to the slower and often steady TC motion, explain the infrequent extreme heights attributed to the TFW growth mechanism. However, in Atlantic Canada, approaching accelerating TCs track northeast for long distances, providing adequate time and suitable conditions for extreme TFW formation. This is especially true for the grid square south of Newfoundland, as waves that entered this grid square have typically been generated from TCs traveling northeast for an extended period after recurvature (e.g., Charlie,
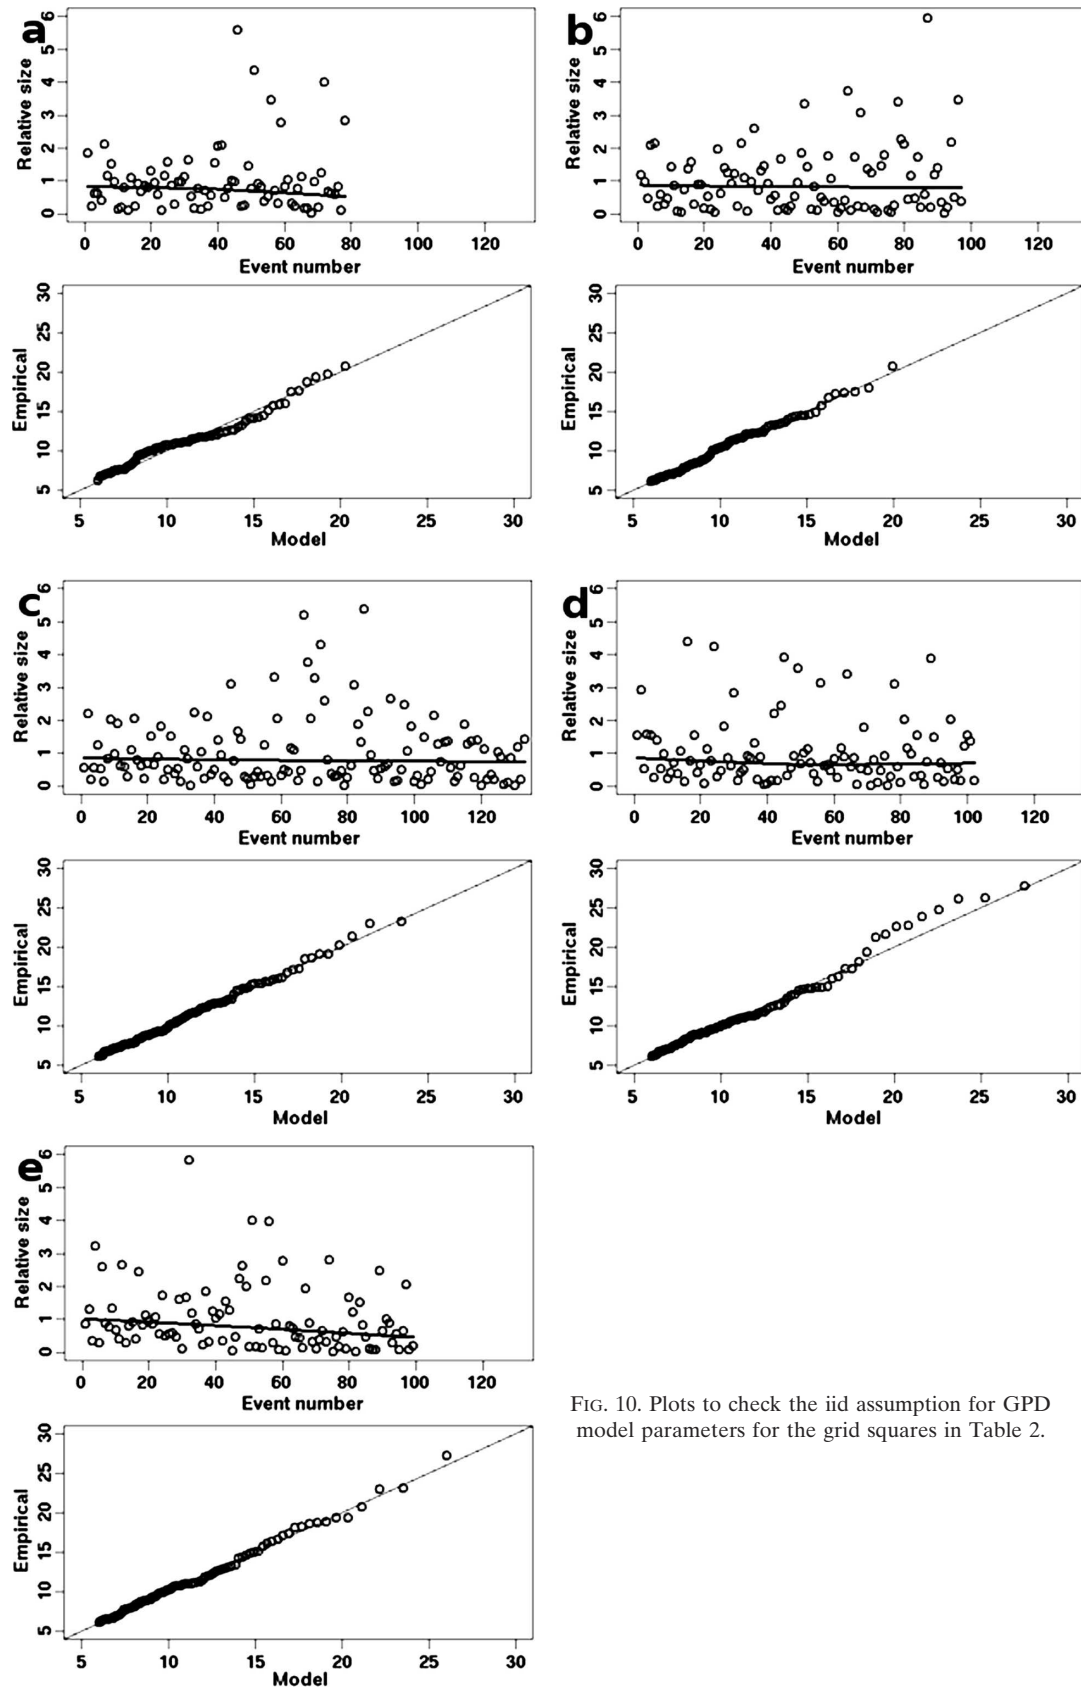
FIG. 10. Plots to check the iid assumption for GPD model parameters for the grid squares in Table 2.
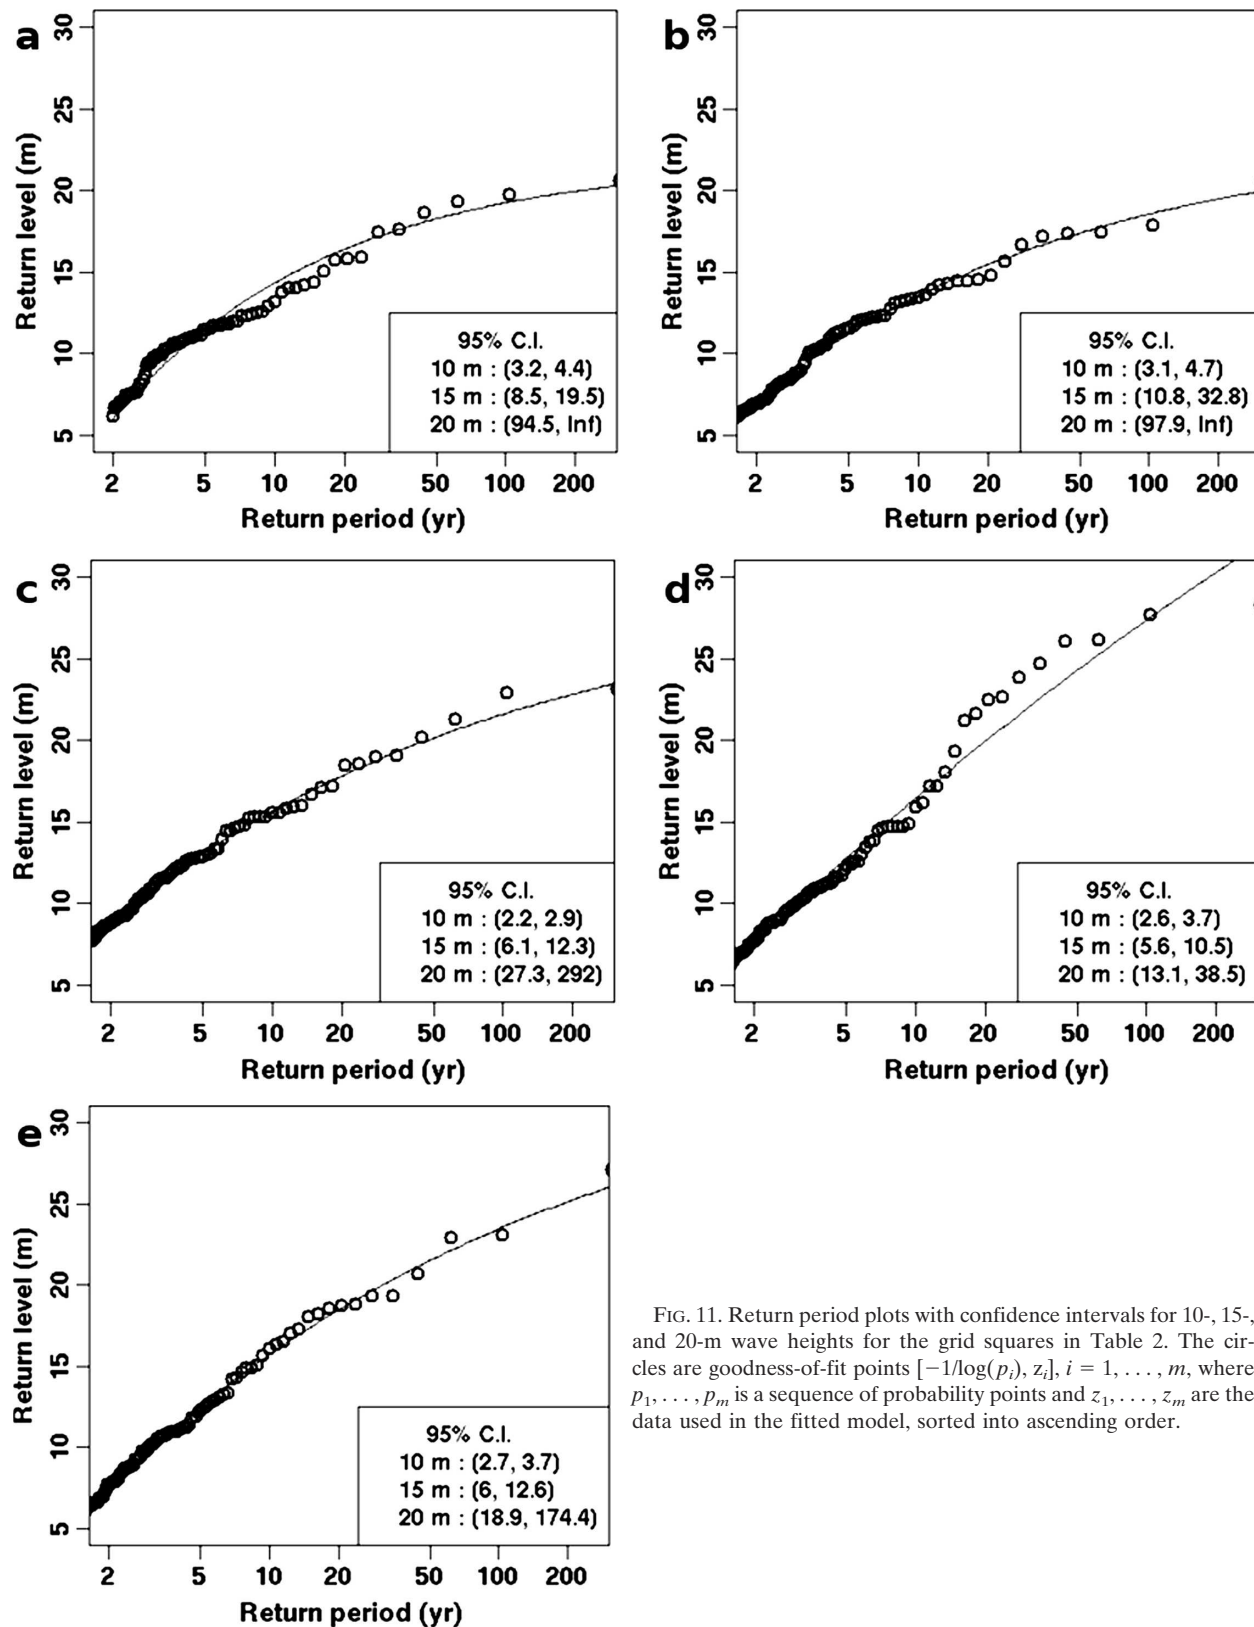
FIG. 11. Return period plots with confidence intervals for 10-, 15-, and 20-m wave heights for the grid squares in Table 2. The circles are goodness-of-fit points $[-1/\log(p_i), z_i]$, $i = 1, \ldots, m$, where $p_1, \ldots, p_m$ is a sequence of probability points and $z_1, \ldots, z_m$ are the data used in the fitted model, sorted into ascending order.

TABLE 4. The 95% bootstrap CI for return periods of specific TFW heights for the selected analysis grid squares of Table 2.

Location	10-m CI (yr, yr)	15-m CI (yr, yr)	20-m CI (yr, yr)
North of Hispaniola	3.2, 4.4	8.5, 19.5	94.5, ∞
Gulf of Mexico	3.1, 4.7	10.8, 32.8	97.9, ∞
Carolina coast	2.2, 2.9	6.1, 12.3	27.3, 292.0
South of Newfoundland	2.6, 3.7	5.6, 10.5	13.1, 38.5
South of Nova Scotia	2.7, 3.7	6.0, 12.6	18.9, 174.4

in 1952; Debby, in 1982; Luis, in 1995; Danielle, in 1998; and Gabrielle, in 2001).

Having completed a detailed analysis of these five grid squares and finding useful results, the same GPD model was independently fitted to the marine areas in Atlantic Canada (Fig. 1 and Table 1) using the preceding steps, sorting the hourly TFW data into polygonal representations of each area. The larger marine areas span many grid squares; however, the number of events in a particular area may not be appreciably different given the method of selecting the highest TFW value from a trajectory traversing a given area. The results are summarized in Table 5. The low test statistics provide some evidence that the GPD model fits adequately to these datasets. In some areas, such as Fundy, the small number of events in the raw data and after declustering gives large variances in the parameter estimates, and the validity of the model to these areas requires further study. Nonetheless, these results provide a very complete summary of how TFWs affect (or do

TABLE 5. Same as Table 3 but for the Atlantic Canada marine areas depicted in Fig. 1 and listed in Table 1.

Area	u (m)	r (h)	N_{raw}	n	a (m)	c	AD	CM
1	6	120	22	7	0.73 (0.48)	0.35 (0.56)	0.14	0.02
2	6	120	75	14	1.94 (0.75)	−0.12 (0.28)	0.46	0.07
3	6	120	258	27	3.52 (0.85)	−0.21 (0.15)	0.27	0.04
4	6	120	524	36	5.19 (1.12)	−0.32 (0.15)	0.85	0.14
5	6	120	1258	81	5.58 (0.83)	−0.21 (0.10)	0.45	0.06
6	6	120	215	20	4.51 (1.41)	−0.02 (0.22)	0.76	0.14
7	6	120	524	45	5.49 (1.20)	−0.12 (0.16)	0.27	0.04
8	6	120	1525	88	6.02 (0.90)	−0.21 (0.11)	0.28	0.04
9	6	120	306	36	2.69 (0.74)	0.06 (0.22)	0.21	0.03
10	6	120	575	51	4.07 (0.91)	−0.02 (0.17)	0.26	0.03
11	6	120	1630	97	5.97 (0.85)	−0.24 (0.10)	0.25	0.04
12	6	120	172	23	2.73 (0.86)	0.06 (0.24)	0.20	0.03
13	6	120	718	62	3.19 (0.63)	−0.01 (0.15)	0.26	0.03
14	6	120	1393	88	4.89 (0.74)	−0.02 (0.11)	0.25	0.03
15	6	120	268	28	3.42 (1.00)	−0.01 (0.22)	0.32	0.06
16	6	120	163	21	1.55 (0.50)	0.19 (0.24)	0.28	0.04
17	6	120	211	23	2.63 (0.89)	0.30 (0.27)	0.34	0.05
18	6	120	441	42	2.78 (0.75)	0.40 (0.23)	0.49	0.07
19	6	120	940	72	3.90 (0.71)	0.15 (0.14)	0.36	0.06
20	6	120	1227	81	4.27 (0.76)	0.14 (0.14)	0.23	0.03

TABLE 6. Same as Table 4 but for the Atlantic Canada marine areas depicted in Fig. 1 and listed in Table 1.

Area	10-m CI (yr, yr)	15-m CI (yr, yr)	20-m CI (yr, yr)
1	214.1, ∞	530.2, ∞	735.1, ∞
2	31.9, ∞	780.7, ∞	4873.3, ∞
3	14.0, 38.6	82.9, ∞	724.8, ∞
4	7.6, 14.7	26.4, 380.7	325.4, ∞
5	3.4, 5.1	9.1, 24.5	36.1, 1598.2
6	6.2, 13.5	25.9, 157.3	85.3, ∞
7	5.6, 10.6	13.1, 43.7	37.9, 1163.0
8	7.6, 14.7	26.4, 380.7	325.4, ∞
9	10.8, 32.8	46.1, 1419.6	133.1, ∞
10	5.9, 12.3	17.4, 61.9	52.8, 2072.0
11	2.8, 4.0	7.1, 15.7	28.3, 198.0
12	15.6, 136.6	55.9, ∞	160.0, ∞
13	6.2, 13.5	25.9, 157.3	85.3, ∞
14	3.3, 5.1	8.3, 20.5	20.7, 175.8
15	11.0, 33.9	36.2, ∞	130.0, ∞
16	28.7, ∞	101.2, ∞	196.5, ∞
17	12.1, 50.8	34.6, 965.7	66.0, ∞
18	8.0, 17.4	18.4, 76.8	31.8, 308.7
19	4.5, 7.2	10.3, 28.8	21.3, 119.4
20	3.8, 6.0	8.4, 20.1	17.4, 69.2

not affect) parts of Atlantic Canada. Table 6 shows a return level summary. From these models, the areas with the most TFW activity are the Grand Banks (areas 18–20 in Fig. 1) and Scotian slope (areas 5, 8, 11, and 14 in Fig. 1), consistent with the right-of-track TFW buildup described by MB5.

The GPD model was then applied in the same way to all of the remaining grid squares in the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea. Plots of return period and the CI at the 95% level for 10-, 15-, and 20-m return levels were generated for each grid square, where sufficient data existed. As depicted in Fig. 12a, 607 grid squares were successfully modeled. The CI for the 10-m return level had a finite upper endpoint for 80% of the 607 grid squares (Fig. 12b). The CI for the 15-m return level had a finite lower endpoint for 94% of the 607 grid squares, but only 53% had a finite upper endpoint (Fig. 12c). The CI for the 20-m return level had a finite lower endpoint for 90% of the 607 grid squares, but only 16% had a finite upper endpoint (Fig. 12d). The region of useful model results is consistent with TFW theory and TC climatology.

6. Application

The TFW trajectories, grid square and marine area sorted data, and the results of the extreme value analysis are accessible using LINUX-based graphical user interfaces. These applications are currently used by CHC operational forecasters and MSC marine specialists to address the following questions on extreme TFWs: 1) how high are the highest TFWs likely at a
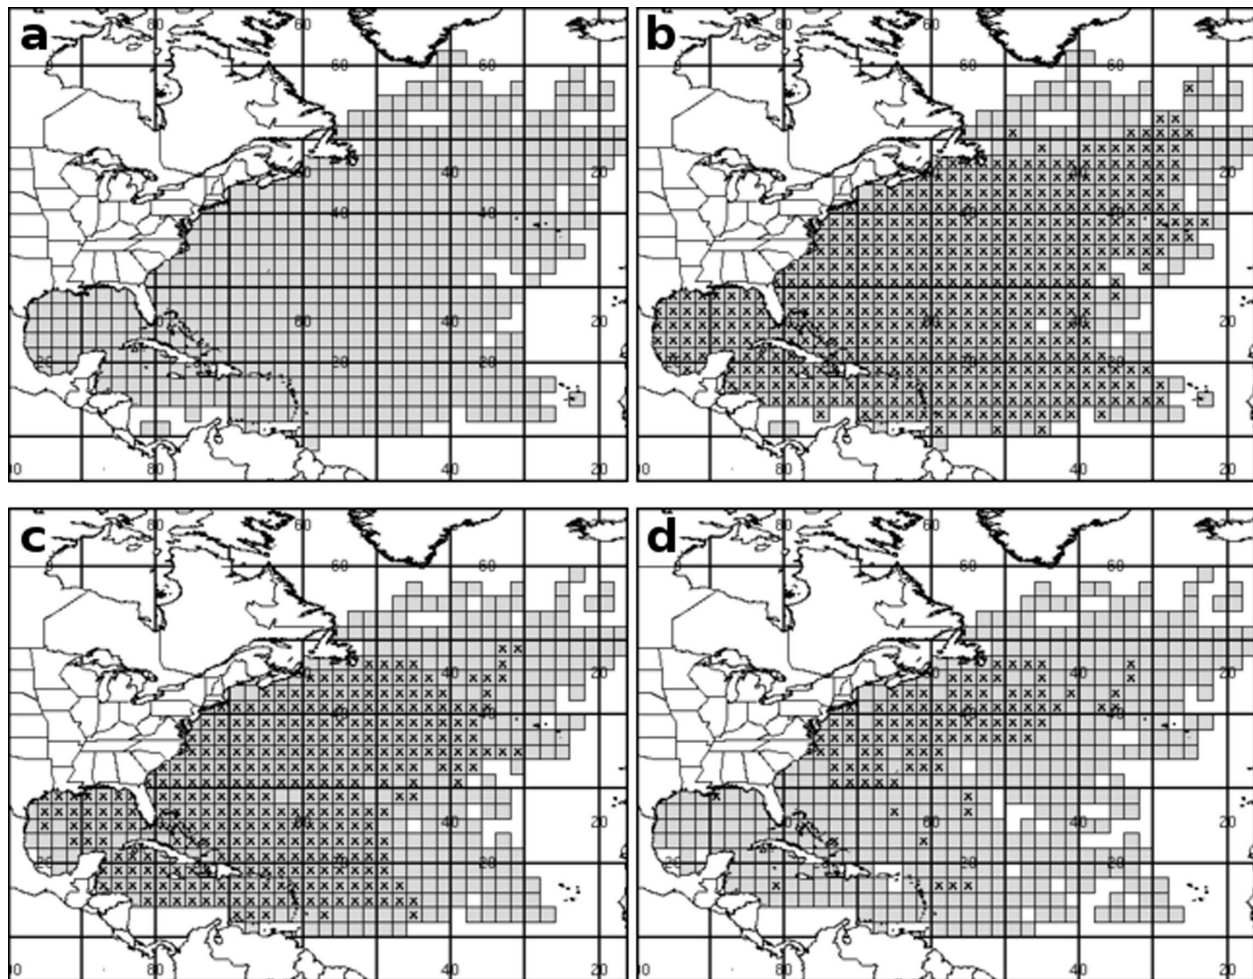
FIG. 12. (a) The $2° \times 2°$ lat–lon grid squares with valid return period plots. (b)–(d) Grid squares with finite CI endpoints for 10-, 15-, 20-m return levels, respectively. The x denotes grid squares with a finite upper CI endpoint.

particular location, 2) how realistic are the values, 3) from what direction are the TFWs most likely to approach the location, 4) how much lead time is there for preparation before the TFWs arrive, 5) what past TCs had an impact on the location, and 6) how do past TCs compare with a real-time prediction? As the datasets are exploited, the application tools will be expanded. In this context, Web-based versions are under consideration.

Figure 13 illustrates an application (CLIMVIEWER) designed specifically to display TC tracks and TFW trajectories. The full track and dominant trajectory from each hourly grid for Floyd in 1999 are shown in Fig. 13a. In the zoomed view shown in Fig. 13b, the dominant TFW from each row of the right-of-track modeling grid at the circled track point (1600 UTC 13 September 1999) is depicted. As illustrated in Fig. 13b, numeric values of H_s, period, number of hours of TFW growth, and TC HURDAT data (e.g., maximum wind) can be added to the display.

Figure 14 shows an application (CHART) used to access the grid square and marine area analyses. The upper left-hand panel shows either the grid squares color-coded based on the highest hourly wave height in a specific month or valid lower and upper CI values (e.g., Fig. 12). Clicking on a grid square (e.g., the southwestern tip of Florida) refreshes the remaining panels: the upper right-hand panel displays plots of probability and cumulative frequency for H_s, period, hours of wave growth, distance covered during wave growth, TFW direction frequency, and return periods (e.g., Fig. 11); the lower right-hand panel lists TCs contributing TFWs to the grid square, with the TC contributing the highest H_s automatically highlighted, and its HURDAT data listed in a table; and the lower left-hand panel displays the tracks of the listed TCs, with the TC contributing the highest TFW highlighted in a thicker line.

Figure 15 shows contours of the expected maximum TFWs in 5-, 10-, 15-, and 20-yr periods, generated
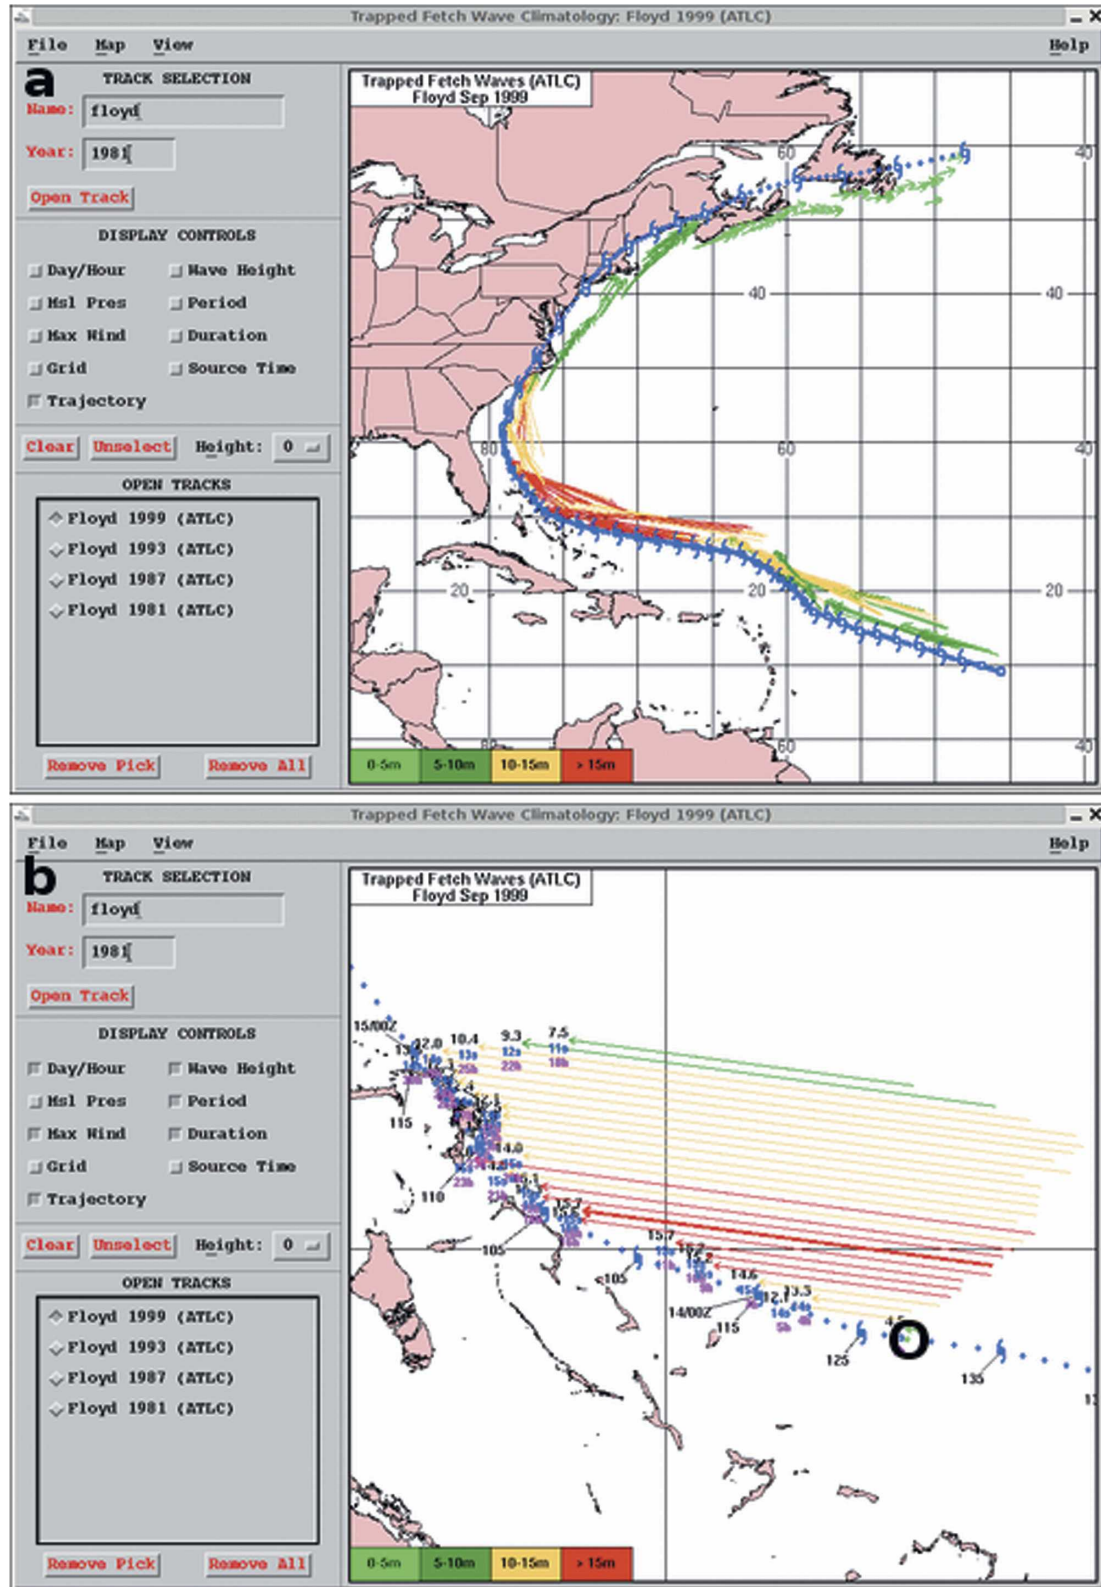
FIG. 13. Example of an application (CLIMVIEWER) used to view the TFW trajectories: (a) dominant TFWs along the track of Floyd in 1999; (b) zoomed view of TFWs computed for each row of the grid at one track point (circled) with H_s, period, and duration values appended at the end of each trajectory. The day, hour, and maximum wind are shown at HURDAT points.
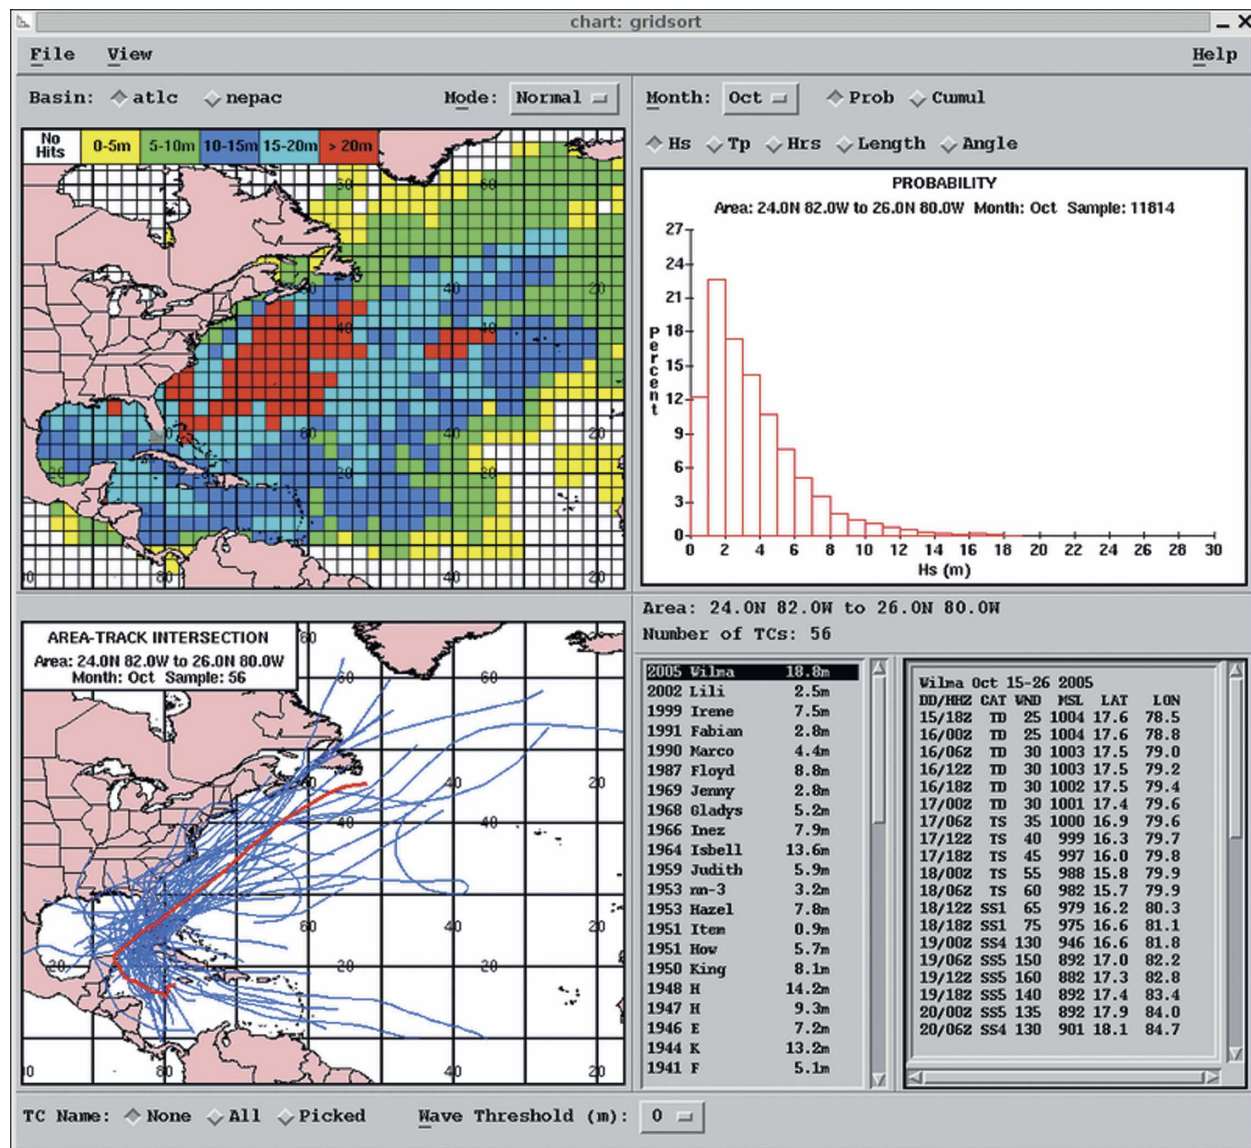
FIG. 14. Example of an application (CHART) for viewing the grid square and marine area analyses. (top left) Grid squares color-coded by maximum H_s (m) for the month of October. (top right) H_s probability plot for the user-selected grid square centered at 25°N, 81°W. (bottom right) List of TCs contributing TFWs to the selected grid square and track data for a particular TC (i.e., Wilma in 2005). (bottom left) Tracks of contributing TCs with Wilma highlighted.

from an application (RPPLOT) using the fitted GPD models.

As an example, consider using CHART and CLIMVIEWER to locate and evaluate historical tracks matching a real-time TC. During the real-time prediction of Wilma in 2005, the operational TFW model indicated 16-m waves near the southwest coast of Florida. Using CHART, Isbell in 1964 was identified as having a track similar to Wilma's, but with 12-m TFWs. CLIMVIEWER was used to compare the details of the two TCs: the maximum predicted TFWs for Wilma were

prior to crossing Florida, while Isbell's were over the North Atlantic after crossing Florida. The variation in location was explained by the rapid acceleration of Wilma after crossing Florida versus a more gradual acceleration for Isbell. Further examination of the frequency data and return period plots was used to qualify if the forecast heights were abnormally high and suggest different track and intensity scenarios.

Using raw TFWs from 1905 to 2004, a preliminary trend analysis was presented in MacAfee (2006); however, the results were inconclusive. The analysis of
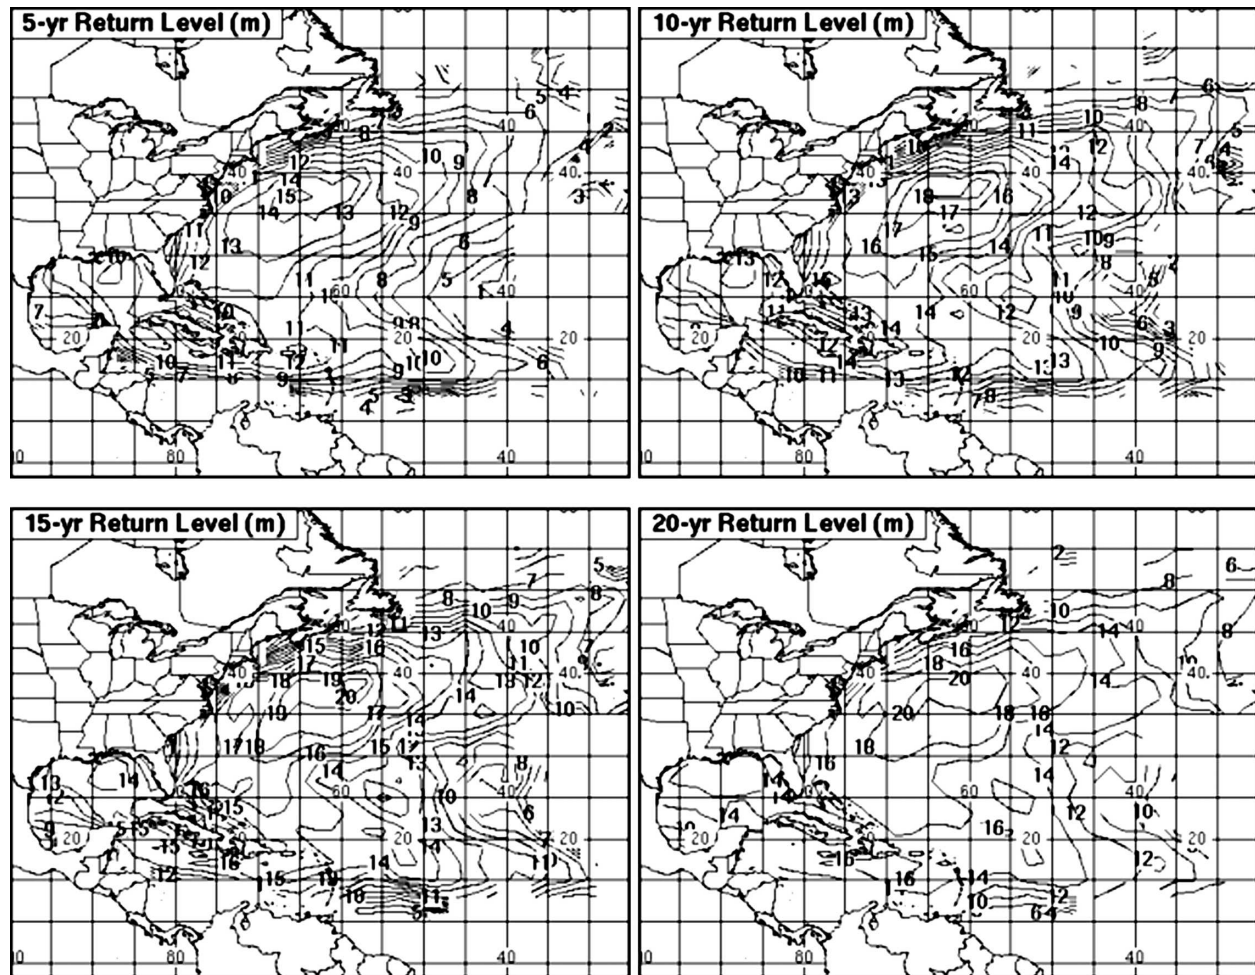
FIG. 15. Contours of the highest H_s (m) expected from TC TFWs in 5-, 10-, 15-, and 20-yr periods.

monthly trends by amalgamating the entire basin and fitting overall models to each month from June to November (hurricane season) may yield useful results; however, this paper focused on justifying that an extreme value analysis of TFWs was statistically sound. Investigation of TFW trends in the extreme value analysis and their relationship (if any) to TC trends in intensity and monthly and annual frequency is a subject of future work.

7. Summary and conclusions

A TFW dataset for all TCs in the Atlantic Ocean, Gulf of Mexico, and Caribbean Sea from 1851 to 2005 was created using a Lagrangian TFW model, a parametric model depiction of the local TC wind fields, and HURDAT track data. The hourly points along each trajectory were sorted into $2° \times 2°$ latitude–longitude grid squares and the marine areas of Atlantic Canada. The data samples in five representative grid squares and marine areas were declustered to create iid samples. These samples were subjected to statistical analysis to determine if extreme value theory could be used to model the extreme wave heights in the dataset, in particular the peaks-over-threshold approach and generalized Pareto distributions. The datasets were mostly modeled adequately by the generalized Pareto model, and much of the resulting inference was useful. Further investigation may be appropriate for grid squares in which potential anomalies were noted. Nonetheless, the results warranted extension of the generalized Pareto model to the remainder of the grid squares.

As a result of this study, return period plots for the marine areas and grid squares covering most of the climatological TC domain were successfully created and integrated into desktop applications for use by operational Canadian Hurricane Centre forecasters and marine specialists. Using illustrative examples, these applications were described and some suggestions for future development outlined.

REFERENCES

Beguería, S., 2005: Uncertainties in partial duration series modeling of extremes related to the choice of the threshold value. *J. Hydrol.,* **303,** 215–230.

——, and S. M. Vicente-Serrano, 2006: Mapping the hazard of extreme rainfall by peaks over threshold extreme value analysis and spatial regression techniques. *J. Appl. Meteor. Climatol.,* **45,** 108–124.

Bowyer, P. J., and A. W. MacAfee, 2005: The theory of trapped-fetch waves with tropical cyclones—An operational perspective. *Wea. Forecasting,* **20,** 229–244.

Brabson, B. B., and J. P. Palutikof, 2000: Tests of the Generalized Pareto Distribution for predicting extreme wind speeds. *J. Appl. Meteor.,* **39,** 1627–1640.

Bretschneider, C. L., and E. E. Tamaye, 1976: Hurricane wind and wave forecasting techniques. Preprints, *15th Conf. on Coastal Engineering,* Vol. 1, Honolulu, HI, American Society of Civil Engineers, 202–237.

Chelton, D. B., M. H. Freilich, J. M. Sienkiewicz, and J. M. Von Ahn, 2006: On the use of QuikSCAT scatterometer measurements of surface winds for marine weather prediction. *Mon. Wea. Rev.,* **134,** 2055–2071.

Choulakian, V., and M. A. Stephens, 2001: Goodness-of-fit tests for the generalized Pareto distribution. *Technometrics,* **43,** 478–484.

Davison, A. C., 1984: Modelling excesses over high thresholds, with an application. *Statistical Extremes and Applications,* J. Tiago De Oliveira, Ed., Springer-Verlag, 461–482.

Easterling, D. R., G. A. Meehl, C. Parmesan, S. A. Changnon, T. R. Karl, and L. O. Mearns, 2000: Climate extremes: Observations, modeling, and impacts. *Science,* **289,** 2068–2074.

Feuer, S. E., C. W. Landsea, L. Woolcock, and J. Berkeley, 2004: The reanalysis of Atlantic basin tropical cyclones from the 1920's: A reexamination of three catastrophic hurricanes that impacted Florida. Preprints, *26th Conf. on Hurricanes and Tropical Meteorology,* Miami, FL, Amer. Meteor. Soc., 669–670.

Fisher, R. A., and L. H. C. Tippett, 1928: Limiting forms of the frequency distributions of the largest or smallest member of a sample. *Proc. Cambridge Philos. Soc.,* **24,** 180–190.

Gnedenko, B. V., 1943: Sur la distribution limite du terme maximum d'une serie aleatoire. *Ann. Math.,* **44,** 423–453.

Gumbel, E. J., 1958: *Statistics of Extremes.* Columbia University Press, 375 pp.

Hosking, J. R. M., and J. R. Wallis, 1987: Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics,* **29,** 339–349.

——, ——, and E. F. Wood, 1985: Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics,* **27,** 251–261.

Jagger, T. H., and J. B. Elsner, 2006: Climatology models for extreme hurricane winds near the United States. *J. Climate,* **19,** 3220–3235.

Jarvinen, B. R., C. J. Neumann, and M. A. S. Davis, 1984: A tropical cyclone data tape for the north Atlantic Basin, 1886–1983: Contents, limitations, and uses. NOAA Tech. Memo. NWS NHC 22, Miami, FL, 21 pp.

Kalnay, E., and Coauthors, 1996: The NCEP/NCAR 40-Year Reanalysis Project. *Bull. Amer. Meteor. Soc.,* **77,** 437–471.

Landsea, C. W., 1993: A climatology of intense (or major) Atlantic hurricanes. *Mon. Wea. Rev.,* **121,** 1703–1713.

——, and Coauthors, 2002: The Atlantic hurricane database reanalysis project—Documentation for the 1851–1920 alterations and additions to the HURDAT database. Preprints, *25th Conf. on Hurricanes and Tropical Meteorology,* San Diego, CA, Amer. Meteor. Soc., 460–461.

——, B. A. Harper, K. Hoarau, and J. A. Knaff, 2006: Can we detect trends in extreme tropical cyclones? *Science,* **313,** 452–454.

Lang, M., T. B. M. J. Ouarda, and B. Bobée, 1999: Towards operational guidelines for over-threshold modeling. *J. Hydrol.,* **225,** 103–117.

MacAfee, A. W., 2006: A synthetic trapped-fetch wave climatology for the North Atlantic and Eastern Pacific. Preprints, *27th Conf. on Hurricanes and Tropical Meteorology,* Monterey, CA, Amer. Meteor. Soc., CD-ROM, 16C.8.

——, and P. J. Bowyer, 2005: The modeling of trapped-fetch waves with tropical cyclones—A desktop operational model. *Wea. Forecasting,* **20,** 245–263.

——, and G. M. Pearson, 2006: Development and testing of tropical cyclone parametric wind models tailored for midlatitude application—Preliminary results. *J. Appl. Meteor. Climatol.,* **45,** 1244–1260.

Madsen, H., P. F. Rasmussen, and D. Rosbjerg, 1997: Comparison of annual maximum series and partial duration series methods for modeling extreme hydrologic events. *Water Resour. Res.,* **33,** 747–758.

Panchang, V. G., and D. Li, 2006: Large waves in the Gulf of Mexico caused by Hurricane Ivan. *Bull. Amer. Meteor. Soc.,* **87,** 481–489.

Pickands, J., 1975: Statistical inference using extreme order statistics. *Ann. Stat.,* **3,** 119–131.

Powell, M. D., and S. H. Houston, 1996: Hurricane Andrew's landfall in south Florida. Part II: Surface wind fields and potential real-time applications. *Wea. Forecasting,* **11,** 329–349.

Smith, R. L., 1984: Threshold methods for sample extremes. *Statistical Extremes and Applications,* J. Tiago De Oliveira, Ed., Springer-Verlag, 621–638.

——, 1989: Extreme value analysis of environmental time series: An application to trend detection in ground-level ozone. *Stat. Sci.,* **4,** 367–393.

Swail, V. R., and A. T. Cox, 2000: On the use of NCEP–NCAR reanalysis surface marine wind fields for a long-term North Atlantic wave hindcast. *J. Atmos. Oceanic Technol.,* **17,** 532–545.

——, E. A. Ceccacci, and A. T. Cox, 2000: The AES40 North Atlantic wave reanalysis: Validation and climate assessment. Preprints, *Sixth Int. Workshop on Wave Hindcasting and Forecasting,* Monterey, CA, Environment Canada, 1–15.

van Gelder, P., J. De Ronde, N. M. Neykov, and P. Neytchev, 2000: Regional frequency analysis of extreme wave heights: Trading space for time. *Proc. 27th ICCE,* Sydney, Australia, Coastal Engineering 2000, 1099–1112.

Wang, X. L., and V. R. Swail, 2002: Trends of Atlantic wave extremes as simulated in a 40-yr wave hindcast using kinematically reanalyzed wind fields. *J. Climate,* **15,** 1020–1035.