Beach-level 24-hour forecasts of Florida red tide-induced respiratory irritation

Shane D. Ross[a,*], Jeremie Fish[b], Klaus Moeltner[c], Erik M. Bollt[b], Landon Bilyeu[d], Tracy Fanara[e,f]

[a] *Aerospace and Ocean Engineering, Virginia Tech, Blacksburg, Virginia*
[b] *Electrical and Computer Engineering and C^3S^2 the Clarkson Center for Complex Systems Science, Clarkson University, Clarkson, New York*
[c] *Agricultural and Applied Economics, Virginia Tech, Blacksburg, Virginia*
[d] *School of Plant and Environmental Sciences, Virginia Tech, Blacksburg, Virginia*
[e] *Mote Marine Laboratory and Aquarium, Sarasota, Florida*
[f] *National Oceanic and Atmospheric Administration, Washington, DC*

Abstract

An accurate forecast of the red tide respiratory irritation level would improve the lives of many people living in areas affected by algal blooms. Using a decades-long database of daily beach conditions, two conceptually different models to forecast the respiratory irritation risk level one day ahead of time are trained. One model is wind-based, using the current days' respiratory level and the predicted wind direction of the following day. The other model is a probabilistic self-exciting Hawkes process model. Both models are trained on eight beaches in Florida during 2011-2017 and applied to the large red tide bloom during 2018-2019. The Hawkes process model performs best, correctly predicting the respiratory risk level an average of 86.7% of the time across all beaches, and at one beach, it is accurate more than 91% of the time.

Keywords: Red tide, Forecast, Respiratory irritation, Modeling, Hawkes

*Email: shaneross@vt.edu

process, Aerosol, HAB, Cyanobacteria, New tools, Public health

1. Introduction

Harmful algal blooms of the toxic dinoflagellate *Karenia brevis*, referred to as "Florida red tide" (henceforth abbreviated as RT) have affected the Florida Gulf coast for centuries. There is emerging evidence that these blooms have increased in frequency, intensity, and geographic spread in recent years (e.g., [1–6]). *K. brevis* produces brevetoxin, a neurotoxin that can result in massive fish kills and mortalities to marine mammals and sea birds. Indirectly, this can lead to neurotoxic shellfish poisoning in humans from consuming contaminated shellfish [6, 7].

More directly, and of primary interest for this study, brevetoxin is released into near-shore aerosol as RT cells are lysed by wave action, or aerosolized though bubble-mediated [8] transport. If inhaled by humans, brevetoxin can produce upper and lower respiratory irritation, such as a burning sensation of eyes and nose, and a dry, choking cough. While these symptoms have been found to be relatively short-lived in healthy individuals (upon separation from the harmful aerosol), RT effects can be more severe and longer-lasting for people with chronic respiratory conditions, such as asthma [1, 4–6, 9–11].

In this study, the potential of forecasting beach-specific respiratory irritation one day ahead of time is assessed, using previous irritation reports at the same location. Two models are proposed: one based on the current respiratory irritation level and a forecast of the next day's wind direction, and the other based on RT as a self-exciting process (or Hawkes process). Both are data-driven, trained on data on from 2011-2017, and tested on data

from a severe RT bloom during 2018-2019. Both are compared with a simple persistence model, which assumes the next day's respiratory irritation level will be the same as the current day's. The Hawkes model performs the best, accurately predicting the respiratory irritation level on average about 87% of the time. Interestingly, the Hawkes process model requires only the recent history of RT respiratory irritation at a location. In particular, the method does not require water samples of *K. brevis* or even ocean or wind forecasts.

1.1. Monitoring of red tide blooms

Monitoring for *K. brevis* blooms involves the following primary components. Satellite imagery is processed to locate potential blooms. Processed images are made available to managers and state health officials in the Gulf of Mexico through the National Oceanic and Atmospheric Administration (NOAA) Harmful Algal Bloom Operational Forecast System (HAB-OFS) [12]. In Florida, where blooms tend to occur most frequently, water samples are collected weekly along the shore and from offshore transects by the Florida Fish and Wildlife Research Institute once a bloom is identified. Samples are delivered to a laboratory for cell enumeration via microscopy. Microscopic enumeration takes about one hour per sample. Typically, samples are processed within 1–2 days and can take longer for more samples from remote areas. The resulting cell counts are then used by HAB-OFS to provide broad, county-wide forecasts of brevetoxin exposure risks. The data for a particular county can be up to a week old by the time it is available to the public. In terms of forecast accuracy, Stumpf et al. [13] found that while county-wide forecasts of respiratory risk were correct 70% of the time, they were only correct 20% of the time when applied to individual beaches.

Recently, beach-level 24-hour forecasts for respiratory impacts have received the most focus from policy agencies [14]. The Gulf of Mexico Coastal Ocean Observing System (GCOOS) recently developed a beach-level risk forecast that includes more than 20 Gulf Coast beaches [15]. The forecast uses current wind forecasts as well as near real-time cell counts of *K. brevis* from water samples enabled by HABscope, a portable microscope system [16]. While the beach reporting system described below provides different data (actual beach-level respiratory impact, collected daily over several years), one can envision the potential to tie in to existing forecast frameworks, fusing the multiple data sources for more accurate forecasts or greater spatial coverage.

1.2. Beach Conditions Reporting System

To address the need for location-specific conditions, a Beach Conditions Reporting System (BCRS) was initiated in 2006 [13, 17–19] and was redeveloped in 2015 when it began gaining public usership [20]. From 2017-2019, the site gained approximately 1.5 million users. The BCRS provided smartphones to (professional) lifeguards and park rangers with an app designed for reporting beach conditions. Twice each day (10:00 and 15:00 local time), lifeguards and park rangers report occurrence of coughing as described further below (and other conditions such as presence of dead fish). While lacking the quantification and precision of microscopy, the reports provide beachgoers with useful real-time information for adequate planning (i.e., severity of aerosolized toxins, potential risks to asthmatics, presence of dead fish, etc.). The BCRS is managed by Mote Marine Laboratory with lifeguards and park rangers as the primary reporters in several counties [20]. The BCRS data compilation is automated, with timely sharing of data with agencies, includ-

ing Florida Fish and Wildlife Conservation Commission and NOAA.

Though the BCRS provides more timely information about beach conditions than the weekly sampling described earlier, it does not provide key information needed for consistent forecasts. It provides no information on *K. brevis* cell presence, but rather an assessment of water color. Because many other factors can change water color, this information has limited value. If winds are blowing offshore in the morning, the public will not know whether the beach will be usable in the afternoon, when there may be a sea breeze onshore.

While RT outbreaks can occur throughout the entire year, aerosolized RT impacts have shown substantial variability both in a temporal and spatial sense. They can last from a few hours to multiple days or even weeks at a given site (e.g., beach), and vary in intensity across sites at a given point in time, with heavily impacted areas at times alternating with completely unaffected shoreline segments [2]. It is still somewhat unclear to what extent human activities, e.g., via nutrient flows into the ocean, contribute to the build-up and/or sustained duration of a RT bloom [1, 3, 21, 22].

While efforts are ongoing to curb RT blooms via prevention and control methods, the predominant management strategy to date has been mitigation, via early detection and avoidance of human contact [1, 6].

1.3. Respiratory irritation

Starting in August 2006, the professional lifeguard corps in Sarasota County began twice-daily reports (approx. 10:00 and 15:00 local time) of the presence of respiratory irritation at six sites, as part of the BCRS. In January 2007, two additional lifeguard sites were added in Manatee County.

Respiratory irritation is defined by the amount of coughing observed in addition to the personal conditions experienced by the lifeguard. The presence of people coughing is used as a proxy for respiratory irritation (cough, nasal congestion, throat irritation, chest tightness, wheezing, and shortness of breath). Coughing has been documented as a response to *K. brevis* aerosols in studies involving occupationally exposed workers, recreationally exposed beachgoers, and asthmatics [9, 23–25].

Lifeguards 'listen' to beachgoers for the presence and/or frequency of coughing. The symptoms observed by the lifeguards are reported at various levels of respiratory irritation as shown in Table 1. Irritation levels are given on a four-tiered scale from 'none' (no coughing noted nearby), 'slight' (a few coughs and sneezes within 30 seconds), 'moderate' (A cough/sneeze heard every 5 seconds), and 'high' (continuous coughing and sneezing in nearby surroundings). For portions of the analysis below, moderate and high classes are together as 'high' risk, as these were the level at which impacts affect the general public [13]. The 'none' and 'slight' are grouped as 'low' risk. The two-tiered (binary) classification of RT respiratory irritation risk was used in

Irritation level	**Risk level**	**In a 30 s audio sample**
None	Low	No coughing/sneezing heard in 30 s
Slight	Low	A few coughs/sneezes heard in 30 s
Moderate	High	A cough/sneeze heard every 5 s
High	High	Coughing/sneezing almost continuously

Table 1: The four-tiered red tide respiratory irritation levels reported in the Beach Conditions Reporting System, and the corresponding two-tiered risk level defined for this study.

this study for an initial analysis of forecasting methods, based on historical data.

1.4. Statistics describing the data set

Since December 2011, the BCRS has monitored RT conditions at over 40 beaches in nine counties along the Gulf coast via citizen scientists, in most cases local lifeguards. Eight beaches were the focus of this study (see Figure 1); six in Sarasota County and two in Manatee County, which had a sufficient number of reports for developing a model, as discussed below. Figure 2 shows the annual aggregate sum of RT irritation reports (level 'slight' or higher) from the 6 beaches in Sarasota County. It is clear that in 2018, the county was especially hard-hit by RT.

For the purposes of the model, only an average daily irritation level is

Figure 1: The eight beaches along the Florida Gulf Coast which are the focus of this study.

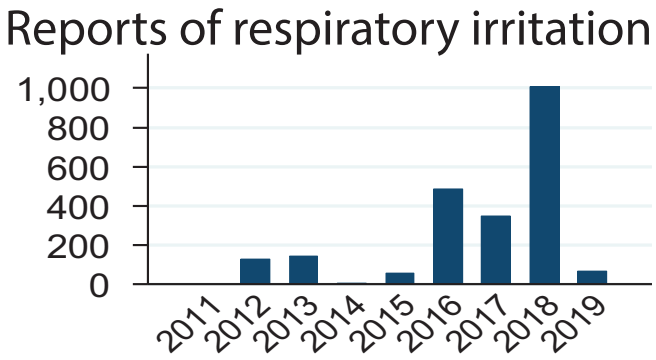

Figure 2: Reports of respiratory irritation caused by red tide for Sarasota County, from December 2011 to April 2019.

considered. Reports are given twice per day, separated by only a few hours, and are often the same. When they differ, the average of the four-tiered irritation level is used. For the two-tiered irritation risk of high and low, there was near unanimous agreement between the two daily reports.

Correlation of irritation level across distance and time. From the BCRS database, correlation of the respiratory irritation level with nearby beaches is considered. For the four-tier irritation levels described as 'none' to 'high' (see Table 1), numerical values 1 through 4 were assigned, respectively.

As one can see in the left panel of Figure 3, the correlation of nearby beaches (within a few km) is high, but drops off to about 50% for beaches from 5 km up to approximately 25 km away, decreasing approximately monotonically with distance. One can also consider the correlation of irritation level across time for an individual beach. The average auto-correlation of a beach's respiratory irritation level versus time-lag is reported. Notice there is a moderate correlation for 1 day, but the correlation drops to about zero

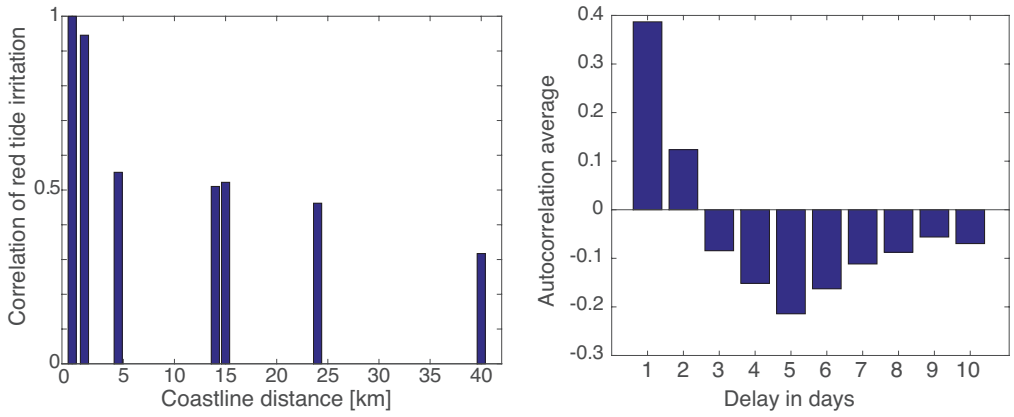

Figure 3: (left) Correlation of red tide respiratory irritation level between the different beaches in Figure 1 as a function of coastline distance. (right) The average auto-correlation of a beach's red tide respiratory irritation level with time.

after 2 days.

2. Data-driven models of respiratory irritation from red tide

2.1. Wind-based model

The wind direction reported from the BCRS at each beach was given as one of 8 octants of width 45° (N, NE, E, SE, S, SW, W, and NW). From the orientation of the coastline for the beaches considered, offshore winds (winds from the land to the ocean) are those coming from the N, NE, E, and SE and onshore winds (ocean to land) are those coming from the S, SW, W, and NW. See Figure 4.

For a wind-based model, the following statistical analysis from the BCRS database is performed for the years 2011-2017, restricted only to periods of time when a RT bloom was present. For a respiratory irritation level of r_t on day t, where $r_t \in \{\text{none, slight, moderate, high}\}$, the probability of r_{t+1} on

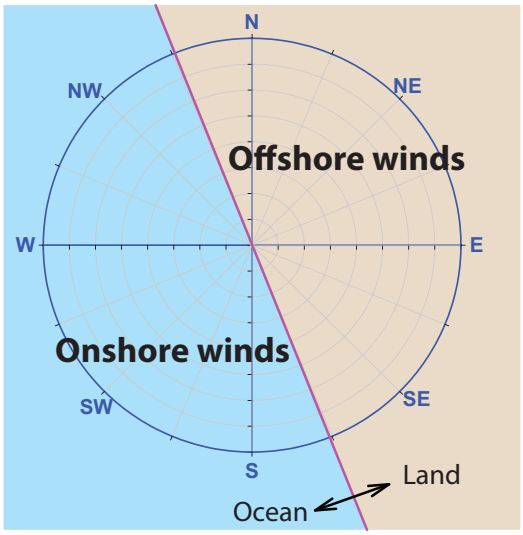

Figure 4: Definition of onshore and offshore wind directions for the beaches studied. The standard wind direction convention is used, where direction denotes where the wind is coming from.

day $t+1$ is calculated, based on the frequencies of such occurrences on the database. Two conditions are considered: onshore winds on day $t+1$ and offshore winds on day $t+1$. The results are shown in Figure 5. For instance, if today's respiratory irritation level is moderate, and there will be onshore winds tomorrow, then with over 60% probability, tomorrow's irritation level will be moderate or high (high risk). On the other hand, if tomorrow's winds are offshore, then with about 60% probability the irritation level will improve, to slight or none (low risk).

If this probability distribution is assumed to hold for future events, then one has a straightforward forecast model for the respiratory irritation level based on the weather forecast, in particular, the wind direction forecast. To get a deterministic model in place of a probabilistic model, one can assume

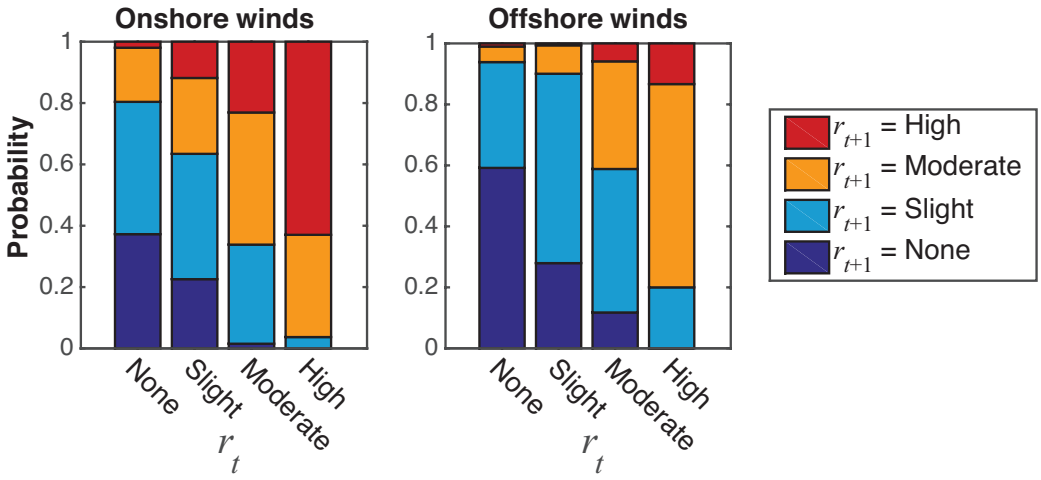

Figure 5: Respiratory irritation level on day $t+1$ as a function of the level on day t and the wind condition on day $t+1$; onshore or offshore winds.

that the state with the maximum likelihood is the one which occurs. This leads to the simple model given in Figure 6. For instance, if the irritation level today is at 'none', but tomorrow has onshore winds, the maximum likelihood is that it will be slight tomorrow, since the light-blue bar in the first stack of the left panel of Figure 5 is longer than all others. In place of the

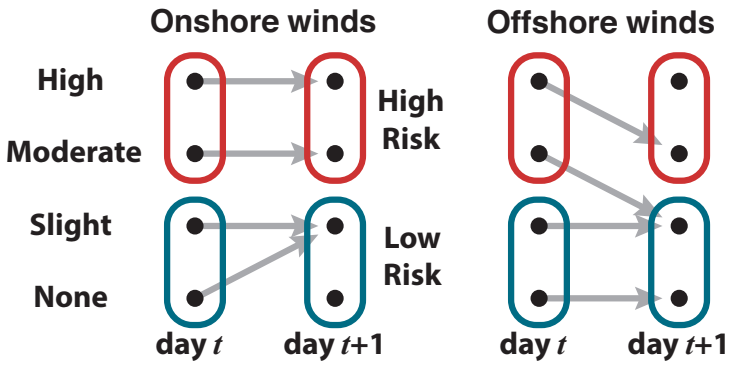

Figure 6: Wind-based respiratory irritation model.

four-tiered irritation level, one can use the simpler two-tiered irritation risk level, low and high, as given in Table 1. This coarser, binary description of the irritation level is used as it is more amenable to modeling and validation.

2.2. Hawkes process model

Hawkes processes have been used to model a remarkable range of phenomena. From earthquakes [26, 27], to gang related violence [28, 29] to econometrics [30, 31]. This type of model performs well in situations where there is evidence for a clustering of events in time.

Hawkes processes are a type of temporal point process (TPP). The classic TPP is the (homogeneous) Poisson process, an example of which is the number of asteroids striking earth above a certain size. Presumably the probability of an asteroid strike is independent of whether or not there was a previous strike. Poisson processes are thus "memoryless" and events are roughly equally spaced in time as seen in the bottom of Fig. 7. The Hawkes process by contrast is a type of non-homogeneous Poisson process, in which the rate of events are dependent upon the history of arrivals and the time which has passed between events. The probability of seeing a new event, e.g., a high risk day, increases when a previous event has occurred. This leads to the temporal clustering seen in the middle panel of Fig. 7. Below, a more detailed description of a Hawkes process is provided.

A *homogeneous* Poisson process is a stochastic point process in which events happen with constant rate λ and with probability,

$$P_{\text{Poisson}}(k|\lambda) = \frac{(\lambda t)^k}{k!} e^{-\lambda t}, \tag{1}$$

where λ has units of inverse time. A *non-homogeneous* Poisson process is

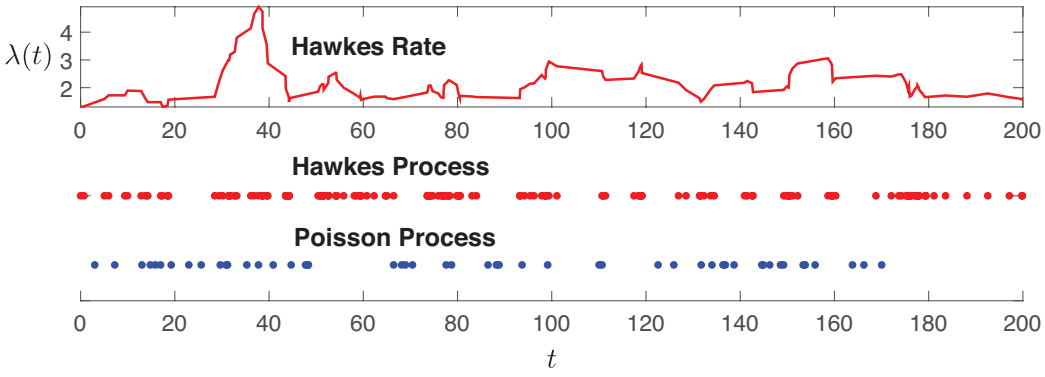

Figure 7: The varying rate of an example Hawkes process can be seen (top). A Hawkes process (middle) tends to have events which cluster in time, whereas a Poisson process with the same initial rate (bottom) produces events that are roughly equally spaced in time. The Hawkes process proves to be a good model for many natural processes, which also tend to have events which cluster in time.

one in which the rate parameter $\lambda(t)$ is a function of time. A linear Hawkes process with an exponential kernel, [32] is a non-homogeneous Poisson process where,

$$\lambda(t) = \lambda_0 + \sum_i \alpha e^{-\beta(t-\tau_i)}. \tag{2}$$

where $\lambda_0 \geq 0$ is the background intensity rate, $\alpha > 0$ is the excitation level, $\beta > 0$ is the reversion level, and $\{\tau_1, \tau_2, ...\tau_i\}$ is the observed sequence of past event times.

A Hawkes process is known as a *self-exciting process* because when an excitation happens the rate increases before decaying to the natural unexcited rate λ_0. The rate initially increases by the amount α when an event arrives but exponentially decays with rate β towards λ_0. The Hawkes process maintains a finite rate $\lambda(t)$ so long as $\alpha < \beta$ [33]. Estimation of the parameters for the Hawkes process is done via maximum likelihood estimation (MLE); de-

tails of MLE for a Hawkes process can be found in Appendix A. It is known that for some data, the MLE method is quite sensitive to the random initial guess [34]. That was, however, not the case here. For all the beaches considered, there was a sufficient number of events that, even for different random initial conditions, the same parameter values were estimated, to within the reported significant digits in Table 2. The eight beaches on which this study focused were in fact chosen for this reason, that the MLE parameter estimation was stable, converging to the same vales independent of the initial guess.

Once the parameters have been estimated from MLE, the probability of observing no events at each time is determined by substituting $\hat{\lambda}(t)$ into eq. (1), where $\hat{\lambda}(t)$ is the rate estimated through MLE. This gives, on each day t, the probability of the next day, $\Delta t = 1$ d, being a high risk day, which is $1 - P_{\text{Poisson}}(0|\hat{\lambda}(t))$, that is,

$$P_{\text{highrisk}}(t) = 1 - e^{-\hat{\lambda}(t)\Delta t} \tag{3}$$

To calibrate the model, one trains on the same data from years 2011-2017 used to train the MLE to obtain a threshold value for when an event is likely to occur. The set of parameters are calculated independently for each beach, though the parameters are similar for all beaches, as seen in Table 2. For example, in all beaches examined the parameter $\beta \approx 0.45$ d^{-1}, which suggests a correlation time-scale of $\beta^{-1} \approx 2$ days, as also seen in the auto-correlation data (Figure 3). A longer probabilistic "memory" timescale can be estimated,

$$T_m = -\log(0.05)/\beta + 1 \quad \text{(in days)}, \tag{4}$$

Beach	λ_0	α	β
Manatee	0.0028	0.4079	0.4978
Coquina	0.0039	0.4592	0.5497
Lido	0.0057	0.2577	0.4100
Siesta	0.0063	0.4135	0.5933
Nokomis	0.0068	0.2686	0.3492
Venice North	0.0068	0.2601	0.3351
Venice	0.0056	0.2817	0.4266
Manasota	0.0051	0.3512	0.4674
Overall	**0.0054**	**0.3375**	**0.4536**
(Standard Deviation)	(0.0014)	(0.081)	(0.091)

Table 2: Estimated values of the Hawkes parameters for each of the beaches considered, listed north to south. All are in units of inverse days, d^{-1}. The bottom rows shows the average parameter value (across all beaches) along with the standard deviation.

which is about 1 week, since a high risk day which happened 7 days ago have a contribution of about 5% compared to that of high risk day which occurred 1 day ago. Also of interest is that the estimates show that $\alpha \approx 0.3$ d^{-1} for all beaches. This is roughly 10^3 times larger than the base rate, λ_0, for each beach, meaning that a single high risk day increases the probability of more high risk days by several orders of magnitude.

3. Results of 1-day forecast of red tide respiratory risk

The wind-based and Hawkes process models described above are applied to the red tide blooms during the 2018-2019 time-frame, at each of the

beaches considered (Figure 1), which was roughly from mid-2018 to late January 2019. At each beach, there are about 200 days of observations. An example of the Hawkes process is shown for a particular beach, Lido Beach, in Figure 8. The percentage of days with a correct 1-day forecast is reported (e.g., a high risk is predicted, and the actual observed risk is high). For comparison, a persistence model is included as a null hypothesis. The persistence model assumes that tomorrow will be like today. The results are reported in Table 3.

The Hawkes process model performs the best, overall and for each of the 8 beaches analyzed. The Hawkes process model correctly predicts the risk level an average of 86.7% of the time, and at one beach (Venice), it is higher

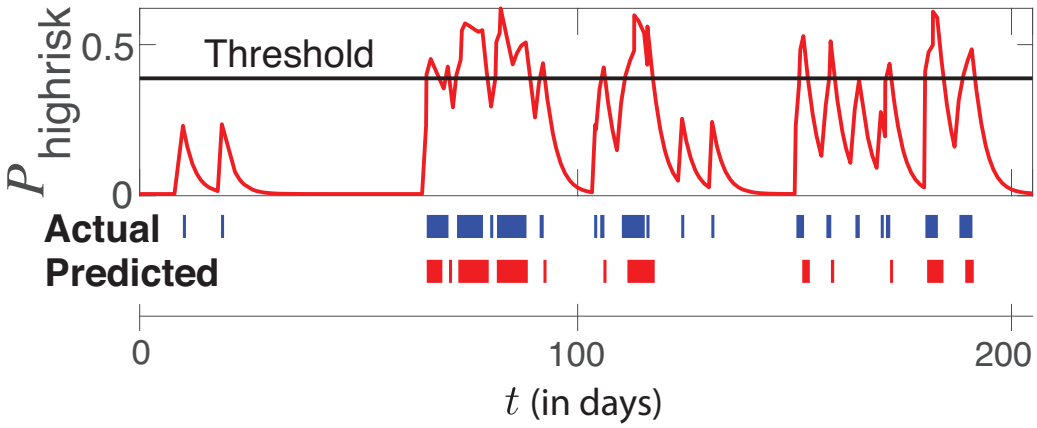

Figure 8: The red line gives the probability of a high risk day at Lido Beach, Sarasota County, Florida, using the Hawkes process model. The flat black line represents the threshold. When the probability of a high risk day exceeds the threshold, the day is forecast as a high risk day. Otherwise, it will forecast as a low risk day. The time axis is the number of days since the red tide bloom began in 2018. The actual and predicted high risk days are shown. During the clustering events the Hawkes model matches nicely.

Beach	**Persistence**	**Wind**	**Hawkes**	**Red Tide Bloom Dates**
Manatee	81.9%	83.2%	85.0%	2018-Aug-03 to 2019-Jan-30
Coquina	77.4%	79.2%	82.1%	2018-Aug-03 to 2019-Jan-30
Lido	83.5%	85.4%	89.6%	2018-Jun-06 to 2019-Jan-30
Siesta	83.0%	82.1%	88.2%	2018-Jun-08 to 2019-Jan-30
Nokomis	79.2%	81.0%	83.7%	2018-Jun-06 to 2019-Jan-30
Venice North	78.5%	80.7%	84.8%	2018-Jun-06 to 2019-Jan-30
Venice	76.2%	82.7%	91.1%	2018-Jun-05 to 2019-Jan-22
Manasota	79.3%	85.4%	87.8%	2018-Jun-07 to 2019-Jan-30
Overall	**79.8**%±2.6%	**82.6**%±2.2%	**86.7**%±3.1%	Across all beaches

Table 3: For each beach, the percent of days for which the respiratory irritation risk was correctly forecast by each model during the dates shown is reported. Bottom row, the average correct forecast percentage across all eight beaches, along with ± the standard deviation.

than 90% of the time. The wind-based model is a marginal improvement over the persistence model overall, but at one beach (Siesta), it performed worse.

The self-excitation behavior can be seen in the example of Lido Beach, Figure 8. The probability of the following day being high risk is a function of all actual high risk days in the recent past. When the probability exceeds the threshold, the model predicts the following day will be high risk day. A visualization of the comparison between the three forecast models and the actual risk for Venice Beach is shown in Figure 9.

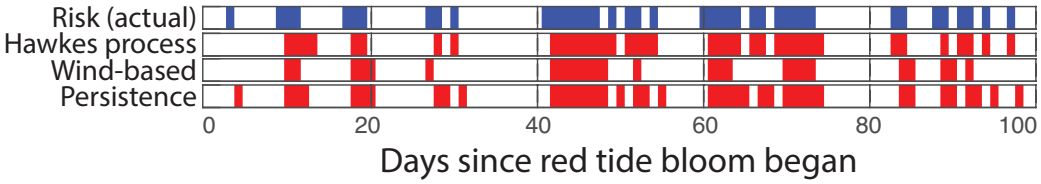

Figure 9: For Venice Beach, Sarasota County, Florida, the days of a high respiratory risk level is shown, and several 1-day forecast models, during the red tide bloom that began June 2018. The time axis shows days since the bloom began on 5-June-2018. For illustration purposes, only the first 100 days are shown.

4. Discussion

The BCRS has accumulated a significant amount of data since its launch in 2006. In the absence of a more sophisticated first-principles-based forecasting system, the historical data could be utilized to provide a near-term beach-level forecast at the locations considered. For forecasting the beach-specific respiratory irritation risk level one day ahead of time, the simple persistence model (tomorrow will be like today) does rather well, near 80% accuracy.

Interestingly, the probabilistic self-exciting Hawkes process model (about 87% accuracy), which does not contain wind as an input, outperforms the more intuitive wind-based model (about 83% accuracy). This may be an artifact of the relatively small number of events (high-risk days) in the data, even over a $\sim$10 year period. Fortuitously, there were enough events to get numerical estimates for the parameters of a univariate Hawkes process, but not enough for a multivariate analysis (i.e., including wind data). With additional data containing high-risk events, adding wind data to the Hawkes process model may be possible. That is, a hybrid of the two models consid-

ered here could potentially improve accuracy, but the amount of historical data available to train the model would need to increase significantly.

Note that the given models do not predict when a red tide bloom will occur. However, when one does occur, the models provide a method for forecasting the respiratory irritation risk for the next day, based on the recent risk history. The accuracy of the Hawkes process model suggests it may be fruitful to consider self-excitation-based approaches in larger-scale models of harmful algal blooms.

5. Conclusions

The Beach Condition Reporting System which has been operational for over decade, has accumulated a wealth of data. In particular, red tide-induced respiratory irritation levels at individual beaches have been reported daily over this time period. The analysis performed here considered eight beaches, providing one of the first reports of the statistics of this data set. Beach-level next-day forecasts of the respiratory irritation risk are also developed. Training only on data from red tide blooms during the time period 2011-2017, both models provided a forecast of the respiratory level day-by-day during the extensive red time bloom of 2018-2019. One model was wind-based, using the current days' respiratory level and the predicted wind direction of the following day. The other model was a probabilistic self-exciting Hawkes process model, which used as input the record of the recent risk history. The Hawkes process model performed the best, correctly predicting the respiratory risk level an average of 86.7% of the time across all beaches, and at one beach, was accurate more than 91% of the time. Interestingly,

the Hawkes process model does not require water samples of *K. brevis*, nor ocean or wind forecasts.

These results suggest that beach-level on-site reports of respiratory irritation are a valuable data source, providing an excellent means to forecast the following day's beach-specific respiratory irritation risk at the same location. Moreover, the efficacy of the BCRS suggests that timely and regular reports of red tide-induced respiratory irritation level should continue to be supported and should be incorporated in operational forecasts used by resource managers and the public.

Acknowledgements

We thank the Mote Marine Laboratory & Aquarium, the Florida Fish and Wildlife Conservation Commission, the Sarasota County Beach Patrol, the Manatee County Department of Public Safety, the Marine Rescue Division, and all of the beach lifeguards, park rangers, and citizen scientists responsible for providing data into the Beach Conditions Reporting System. This project was partially supported by the Global Change Center, the Fralin Life Sciences Institute, and the Institute for Society, Culture, and Environment at Virginia Tech. Moeltner also acknowledges partial support by the USDA/NIFA Multi-State project #VA-136344. Ross acknowledges partial support by the National Science Foundation (NSF) under grant number 1922516.

Appendix A. Hawkes Maximum Likelihood Estimation

The parameters α, β, λ_0 may be estimated via maximum likelihood estimation (MLE). The log-likelihood function of the Hawkes process is given by [35]:

$$\mathcal{L}(\tau_1, ..., \tau_n | \alpha, \beta, \lambda_0) = -\lambda_0 T + \sum_{i=1}^{n} \frac{\alpha}{\beta}[e^{-\beta(T-\tau_i)} - 1] + \sum_{i=1}^{n} \log(\lambda_0 + \alpha A(i)), \tag{A.1}$$

where T is the total time which has been recorded,

$$A(i) = \sum_{\tau_i > \tau_j} e^{-\beta(\tau_i - \tau_j)} \ (\forall i \geq 2) \tag{A.2}$$

and τ_n is the time of the last recorded *event*. One may find the partial derivatives of the log likelihood function given in eq. (A.1) in [35]. However, estimation of the maximum (by setting the partial derivatives equal to zero and solving) is challenging. So as an alternative, standard numerical techniques for nonlinear optimization can be used, in this case the Nelder-Mead direct search technique [36] to estimate the parameters.

References

[1] F. Alcock, An assessment of Florida red tide: Causes, consequences and management strategies, 2007. Technical report #1190. Mote Marine Laboratory.

[2] K. Nierenberg, A. Reich, R. Currier, B. Kirkpatrick, L. C. Backer, R. Stumpf, L. Fleming, G. Kirkpatrick, Beaches and HABs: successful expansion of the Florida red tide reporting system for protection of public health through community education and outreach, Florida Journal of Environmental Health (2009) 18–24.

[3] K. Nierenberg, M. M. Byrne, L. E. Fleming, W. Stephan, A. Reich, L. C. Backer, E. Tanga, D. R. Dalpra, B. Kirkpatrick, Florida red tide perception: Residents versus tourists, Harmful Algae 9 (2010a) 600–606.

[4] K. Nierenberg, K. Kirner, P. Hoagland, S. Ullmann, W. G. LeBlanc, G. Kirkpatrick, L. E. Fleming, B. Kirkpatrick, Changes in work habits of lifeguards in relation to Florida red tide, Harmful Algae 9 (2010b) 419–425.

[5] L. E. Fleming, B. Kirkpatrick, L. C. Backer, C. J. Walsh, K. Nierenberg, J. Clark, A. Reich, J. Hollenbeck, J. Benson, Y. S. Cheng, et al., Review of Florida red tide and human health effects, Harmful Algae 10 (2011) 224–233.

[6] A. Corcoran, M. Dornback, B. Kirkpatrick, A. Jochens, A primer on Gulf of Mexico harmful algal blooms, 2013. Technical Report. Gulf of Mexico Alliance & Gulf of Mexico Coastal Ocean Observing System.

[7] B. Kirkpatrick, L. E. Fleming, L. C. Backer, J. A. Bean, R. Tamer, G. Kirkpatrick, T. Kane, A. Wanner, D. Dalpra, A. Reich, et al., Environmental exposures to Florida red tides: effects on emergency room respiratory diagnoses admissions, Harmful Algae 5 (2006) 526–533.

[8] R. B. Pietsch, H. Grothe, R. Hanlon, C. W. Powers, S. Jung, S. D. Ross, D. G. Schmale III, Wind-driven spume droplet production and the transport of *pseudomonas syringae* from aquatic environments, PeerJ 6 (2018) e5663.

[9] L. C. Backer, L. E. Fleming, A. Rowan, Y.-S. Cheng, J. Benson, R. H. Pierce, J. Zaias, J. Bean, G. D. Bossart, D. Johnson, R. Quimbo, D. G. Baden, Recreational exposure to aerosolized brevetoxins during Florida red tide events, Harmful Algae 2 (2003) 19–28.

[10] B. Kirkpatrick, J. Bean, L. E. Fleming, L. C. Backer, R. Akers, A. Wanner, D. Dalpra, K. Nierenberg, A. Reich, D. Baden, Aerosolized red tide toxins (brevetoxins) and asthma: a 10 day follow up after 1 hour acute beach exposure, in: Proceedings of the 12th International Conference on Harmful Algae, 2009, pp. 297–299.

[11] B. Kirkpatrick, L. E. Fleming, J. A. Bean, K. Nierenberg, L. C. Backer, Y. S. Cheng, R. Pierce, A. Reich, J. Naar, A. Wanner, W. Abraham, Y. Zhou, J. Hollenbeck, D. Baden, Aerosolized red tide toxins (brevetoxins) and asthma: continued health effects after 1 h beach exposure, Harmful Algae 10 (2011) 138–143.

[12] National Oceanic and Atmospheric Administration, Harmful

Algal Bloom Operational Forecast System, 2021. Web site: https://tidesandcurrents.noaa.gov/hab/gomx.html/; last accessed May 19, 2021.

[13] R. P. Stumpf, M. C. Tomlinson, J. A. Calkins, B. Kirkpatrick, K. Fisher, K. Nierenberg, R. Currier, T. T. Wynne, Skill assessment for an operational algal bloom forecast system, Journal of Marine Systems 76 (2009) 151–161.

[14] K. Moeltner, T. Fanara, H. Foroutan, R. Hanlon, V. Lovko, S. Ross, D. S. III, Harmful algal blooms and toxic air: The economic value of improved forecasts, Marine Resource Economics under revision (2021).

[15] Gulf of Mexico Coastal Ocean Observing System, Experimental Red Tide Respiratory Forecast, 2021. Web site: https://habforecast.gcoos.org/; last accessed May 19, 2021.

[16] D. R. Hardison, W. C. Holland, R. D. Currier, B. Kirkpatrick, R. Stumpf, T. Fanara, D. Burris, A. Reich, G. J. Kirkpatrick, R. W. Litaker, HABscope: A tool for use by citizen scientists to facilitate early warning of respiratory irritation caused by toxic blooms of *karenia brevis*, PLOS ONE 14 (2019) 1–17.

[17] B. Kirkpatrick, R. Currier, K. Nierenberg, A. Reich, L. C. Backer, R. Stumpf, L. Fleming, G. Kirkpatrick, Florida red tide and human health: a pilot beach conditions reporting system to minimize human exposure, Science of the Total Environment 402 (2008) 1–8.

[18] R. D. Currier, C. Boyes, A. Hails, K. Nierenberg, B. Kirkpatrick, G. Kirkpatrick, An ocean observing system for harmful algal bloom detection and tracking, in: OCEANS 2009, IEEE, 2009, pp. 1–4.

[19] K. Nierenberg, J. Hollenbeck, L. E. Fleming, W. Stephan, A. Reich, L. C. Backer, R. Currier, B. Kirkpatrick, Frontiers in outreach and education: the Florida red tide experience, Harmful Algae 10 (2011) 374–380.

[20] Mote Marine Laboratory & Aquarium, Beach Conditions Reporting System, 2021. Web site: `https://visitbeaches.org/`; last accessed Jan. 14, 2021.

[21] G. A. Vargo, C. A. Heil, K. A. Fanning, L. K. Dixon, M. B. Neely, K. Lester, D. Ault, S. Murasko, J. Havens, J. Walsh, S. Bell, Nutrient availability in support of *karenia brevis* blooms on the central West Florida shelf: what keeps Karenia blooming?, Continental Shelf Research 28 (2008) 73–98.

[22] B. Kirkpatrick, K. Kohler, M. Byrne, L. E. Fleming, K. Scheller, A. Reich, G. Hitchcock, G. Kirkpatrick, S. Ullmann, P. Hoagland, Human responses to Florida red tides: Policy awareness and adherence to local fertilizer ordinances, Science of the Total Environment 493 (2014) 898–909.

[23] L. C. Backer, B. Kirkpatrick, L. E. Fleming, Y. S. Cheng, R. Pierce, J. A. Bean, R. Clark, D. Johnson, A. Wanner, R. Tamer, Y. Zhou, D. G. Baden, Occupational exposure to aerosolized brevetoxins during Florida

red tide events: effects on a healthy worker population, Environmental Health Perspectives 113 (2005) 644–649.

[24] L. E. Fleming, B. Kirkpatrick, L. C. Backer, J. A. Bean, A. Wanner, D. Dalpra, R. Tamer, J. Zaias, Y. S. Cheng, R. Pierce, J. Naar, W. Abraham, R. Clark, Y. Zhou, M. S. Henry, D. Johnson, G. V. D. Bogart, G. D. Bossart, M. Harrington, D. G. Baden, Initial evaluation of the effects of aerosolized Florida red tide toxins (brevetoxins) in persons with asthma, Environmental Health Perspectives 113 (2005) 650–657.

[25] L. E. Fleming, B. Kirkpatrick, L. C. Backer, J. A. Bean, A. Wanner, A. Reich, J. Zaias, Y. S. Cheng, R. Pierce, J. Naar, W. M. Abraham, D. G. Baden, Aerosolized red-tide toxins (brevetoxins) and asthma, Chest 131 (2007) 187 – 194.

[26] Y. Ogata, Statistical models for earthquake occurrences and residual analysis for point processes, Journal of the American Statistical association 83 (1988) 9–27.

[27] K. Türkyilmaz, M. N. M. van Lieshout, A. Stein, Comparing the Hawkes and trigger process models for aftershock sequences following the 2005 Kashmir earthquake, Mathematical Geosciences 45 (2013) 149–164.

[28] G. O. Mohler, M. B. Short, P. J. Brantingham, F. P. Schoenberg, G. E. Tita, Self-exciting point process modeling of crime, Journal of the American Statistical Association 106 (2011) 100–108.

[29] J. Park, F. P. Schoenberg, A. L. Bertozzi, P. J. Brantingham, Investigating clustering and violence interruption in gang-related violent crime

data using spatial–temporal point processes with covariates, Journal of the American Statistical Association 0 (2021) 1–14.

[30] E. Errais, K. Giesecke, L. R. Goldberg, Affine point processes and portfolio credit risk, SIAM Journal on Financial Mathematics 1 (2010) 642–665.

[31] E. Bacry, S. Delattre, M. Hoffmann, J.-F. Muzy, Modelling microstructure noise with mutually exciting point processes, Quantitative Finance 13 (2013) 65–77.

[32] A. G. Hawkes, Spectra of some self-exciting and mutually exciting point processes, Biometrika 58 (1971) 83–90.

[33] P. J. Laub, T. Taimre, P. K. Pollett, Hawkes processes, arXiv preprint arXiv:1507.02822 (2015).

[34] A. Veen, F. P. Schoenberg, Estimation of space-time branching process models in seismology using an EM-type algorithm, Journal of the American Statistical Association 103 (2008) 614–624.

[35] T. Ozaki, Maximum likelihood estimation of Hawkes' self-exciting point processes, Annals of the Institute of Statistical Mathematics 31 (1979) 145–155.

[36] J. A. Nelder, R. Mead, A simplex method for function minimization, The Computer Journal 7 (1965) 308–313.