Original Article

The storm of the century! Promoting student enthusiasm for applied statistics

Lee Fawcett
School of Mathematics & Statistics, Newcastle University, Newcastle upon Tyne, UK
e-mail: lee.fawcett@ncl.ac.uk

Keith Newman
School of Mathematics & Statistics, Newcastle University, Newcastle upon Tyne, UK

Summary This article describes a hands-on activity that has been used with students aged 12–18 years to promote the study of Statistics. We believe there is evidence to suggest an increase in student enthusiasm for Statistics at school, within the Mathematics curriculum, but also within other subjects such as Geography. We also believe that the use of such activities has resulted in some students giving more serious thought to studying Statistics at University. The activity described here is supported with a web-based application to allow younger or less experienced students to engage with the material.

Keywords: Teaching statistics; extreme values; applied statistics; student motivation; *Shiny*.

BACKGROUND AND MOTIVATION

It is our experience that new undergraduate Mathematicians/Statisticians often have a rather dim view of Statistics, and it is not until they study it at University that they begin to appreciate the very practical, hands-on nature of the subject. At Newcastle University, students are offered courses in Clinical Trials, Survival Analysis, Environmental Extremes and Financial Modelling, to name but a few; they often remark that when they take such courses they finally see the relevance of Statistics and can see its place in the real world. Through the first author's role on the outreach and recruitment team at Newcastle University and associated visits to local schools, it has also become apparent that students often do not see the relevance of Statistics to other subjects, such as Geography, Biology and Psychology.

These school visits have helped to shed some light on the rather depressing scenario that such a practical subject—used in most areas of science and so having many exciting applications—is seen by many students as *dry*; shown below are comments taken directly from a questionnaire completed by 14–18 year olds on the subject of Statistics, distributed by the first author during outreach visits to local schools over the last 3 years:

"Boring boring boring. Wish this wasn't part of my Maths course at school"

"All we seem to do is flip coins and roll fair six-sided dice."

"Who cares about the chances of pulling a green sock from a drawer? Loads of rubbish examples are used and they're boring"

"Spent 3 whole classes on frequency density in histograms. That's as exciting as it gets"

"There was a question about John being late for school. How did they know [the probability of] this was 0.35?"

Some of these quotes might correspond to what Taleb (2007) refers to as the *Ludic Fallacy*, in which naive statistical assumptions underpin the modelling of complex scenarios. Students have often made comments about how unrealistic their study of Probability and Statistics is, and it is our belief that this could have a negative impact on their overall opinion of the subject and its place in the real world (and hence other subjects studied).

The classroom activities described in this article aim to dispel such concerns by bringing to life parts of the Probability and Statistics curriculum followed in schools in the UK. The activities use real-life data on annual maximum wave heights (AMWH) taken from hourly records at a location in the Gulf of Mexico, and the practical aim is very clear: to use Statistics to quantify the likelihood of extreme sea levels and hence better prepare for life-threatening flood events. The activities have been used for outreach and engagement purposes with students as young as 12 to 18 year olds considering studying Mathematics at University. *The Storm of the Century!*, as the overall activity is often advertised, has always been well-received by students, and teachers have often requested permission to use the materials in class to engage students with their Statistics curriculum. In some cases, teachers of other subjects, including Geography, have also asked to use these activities with their students.

In this paper we describe *The Storm of the Century!* activities and attempt to assess their success as tools for promoting student enthusiasm for applied Statistics. We encourage teachers to use the resources we have developed in their own classes to help engage students with the following topics: relative frequency probability, interpreting probability, basic statistical modelling and extrapolation, transposition of formulae, the Normal distribution and non-standard probability models. All materials are available to view at a dedicated webpage: www.mas.ncl.ac.uk/~nlf8/outreach.

A web-based application has also been developed to enable students and teachers to interact with the material without having to get embroiled too deeply in the mathematics. For this, we have used the *Shiny* add-on package for the (open source) R software environment for statistical computing; see Chang et al. (2015). The application itself is hosted online and can be accessed with any modern web browser. It can also be used without an internet connection, subject to the installation of R. Our supporting webpage (link above) includes full access details for the application, as well as installation details for the software should running the application locally be required. Readers are invited to take a look at this and provide any comments or feedback.

THE STORM OF THE CENTURY! ACTIVITY

In this section, we describe the five main parts of *The Storm of the Century!* activity, plus an extension to the Normal distribution for more experienced teachers and students. A two-page handout accompanies some slides for an interactive presentation that typically takes between 60 and 90 minute to complete, although this depends on the level of participation and the amount of assistance the students require. We recommend trained classroom assistants if attempting some of the more challenging parts with younger students. Parts 1–3 should be manageable with students from 12 years of age; part 4 requires some careful thought about the practical interpretation of probabilities; part 5 might require more confidence with algebra. An optional extra, probably only to be used with A level students (or equivalent), requires use of the Normal distribution (part 6). The full presentation and handout are available for readers to view on our webpage.

Part 1 - Motivation

After a five minute icebreaker, students are immediately told about the links between the statistical study of extremes and scientists who need to use such statistical methods: we talk about *hydrologists*, *seismologists* and *oceanographers*. We mention that the study of extremes, rather than averages, is a very specialized area of Statistics and one most students will not encounter until University Statistics courses. However, we explain that it is very important to these scientists, as extreme observations on variables such as rainfall, wind speeds and seismic activity (for example) are more likely to result in disasters such as floods and major earthquakes than are observations close to the average. In this activity the data we use are, by construction, extreme observations (see Table 1). Although the model we present for these data (see Part 3) originates from a rather niche area of Statistics, it is the probabilities that this model generates which provide the main focus of the activity. For older students who have studied the Normal distribution, we mention that we will be making use of the mean and standard deviation—of our extreme observations—later.

Table 1. Annual maximum wave heights (feet) taken from hourly observations at Shell Beach, Louisiana, 1955–2004

8.5	8.9	9.1	8.9	8.4	9.7	9.1
9.6	8.7	9.3	9.6	9.3	8.7	9.0
8.8	8.9	8.9	12.2	7.8	7.7	8.3
8.1	7.3	6.8	6.7	7.3	7.6	8.2
8.6	9.8	9.5	7.4	7.3	10.2	10.3
10.4	8.8	9.7	10.0	10.8	11.1	12.7
11.5	11.8	12.6	13.0	10.5	10.5	10.0
9.4						

Highlighted values are those which exceed 8.75 ft, for use in Eq. (1).

We show lots of motivating pictures from recent earthquake, tsunami and hurricane events; we update these pictures every year to make sure there are some that the students recognize from recent news stories. Our experience of delivering many outreach and engagement activities has told us that a non-mathematical introduction to such an activity, using topical, thought-provoking and often dramatic visual stimuli helps to engage and enthuse the audience about any forthcoming hands-on activities. Figure 1 shows some of the pictures used, including hypothetical scenarios in New York and London. The class are then given some facts and figures about Hurricane Katrina and are asked to write some of these down in the space provided on their handout (see supporting webpage), including:

- AMWH during Katrina reached 14.4 ft
- Katrina was billed as the *Storm of the Century*

The remainder of the presentation focuses on AMWH data collected at a location in the Gulf of Mexico not far from New Orleans, Louisiana (Table 1). Notice the data span the 50 years up to, and including, 2004—the year before Katrina struck. Throughout the talk students are told to imagine themselves as the mathematician/statistician working as part of a scientific team investigating the design of a new sea wall being built to protect the city of New Orleans. The main premise of the activity is to think about how we can use historical data on extremes to estimate the likelihood of future AMWH larger than those ever recorded before—notice that the largest height in Table 1 is 13 ft, 1.4 ft lower than that observed during Katrina. On the *Data Preview* page of the Shiny application there is a drop-down menu from which various built-in datasets can be selected, one of which is the AMWH data shown in Table 1. A map showing the geographical

Fig. 1. Visual stimuli used to motivate the study of environmental extremes. Top row: Hurricane Katrina; middle row: Hypothetical flooding in New York and London as a result of climate change; bottom row: the *Boxing Day Tsunami* in the Indian Ocean, 2004. The top-right photograph is the first author's own, taken during a research visit to New Orleans (2011)

location, with the raw data and various graphical/numerical summaries, is automatically displayed, along with a dataset description; see Figure 2.

Part 2 - Basic activity using relative frequencies

The first statistical activity requires students to estimate the probability that the AMWH for any randomly selected year in the future–possibly 2005–exceeds 8.75 ft, using a relative frequency approach. From the data in Table 1, we can see that 33 observations exceed 8.75 ft (highlighted) and so, expressing this as a proportion of the total number of AMWH we have, gives:

$$P(\text{AMWH} > 8.75\text{ feet}) = \frac{33}{50} = 0.66. \quad (1)$$

Students are also asked to find, in the same manner, exceedance probabilities for 11.25 and 14 ft, respectively giving:

$$P(\text{AMWH} > 11.25\text{ feet}) = \frac{6}{50} = 0.12;$$
$$P(\text{AMWH} > 14\text{ feet}) = \frac{0}{50} = 0. \quad (2)$$

After engaging with the class about the practical interpretations of these probabilities, with reference to the probability scale, the students are then asked to critique their interpretation of these probabilities—especially that given by Eq. (2). Using past data alone implies that an AMWH greater than 14 ft is impossible, although we know that wave heights *did* exceed this level during Katrina—and students are asked to recall this fact from earlier. The task of estimating such an event seems impossible, unless we extrapolate using a suitable *probability model* (see part 3 of this activity).

Although obtaining these relative frequencies is very simple we find it is an extremely accessible way to begin the activity. Almost all students can complete this task without assistance and—more importantly—can see the relevance of simple ideas of probability from their school Statistics curriculum to the real world. To finish the first task, students are then asked to complete, by hand, the graph shown in Figure 3. Here, they must plot exceedance probabilities for AMWH of 6.5, 7.0, …,12.5 ft, these probabilities being obtained using the relative frequency approach as before.

Clicking on the *Relative frequency* tab of our *Shiny* application reproduces the plot shown in Figure 3 within the application itself, as well as a table of results showing relative frequency exceedance probabilities across a range of values of the variable being studied; see the screenshot in Figure 4. For the AMWH at the location being studied here, the table produced covers the exceedance probabilities the students are asked to calculate for themselves in the handout and gives the probabilities shown in the plot in Figure 3.

Part 3 - Moving on: A probability model for extremes

We now discuss simple ideas of modelling and, given the aim of estimating probabilities of very rare events, we attempt to justify the need for a well-fitting model from which to extrapolate. We discuss that this requires a *leap of faith* in that a model which describes our observed data well can be extended beyond the reach of our data, and such uncertainty means we are more reliant than ever on the model we choose (and, of course, we do always choose a model, there is no *correct* model).

After discussing some history surrounding the development of *Extreme Value Theory*, students are introduced to *Gumbel's model for exceedance probabilities*. All notions of density and distribution functions are avoided, the *Gumbel model* being simply presented as the survival function of the Type I extreme value distribution (Coles, 2001), this function giving model-based estimates of the relative frequencies obtained empirically from the data (e.g. Eq. (1) and Eq. (2)). See the Appendix for full details of the Gumbel model (and in particular Eq. (4)).

At this point, the teacher/facilitator has two options, depending on their own level of confidence and the student audience: (i) work directly with the formula in Eq. (4) in the Appendix, demonstrating the use of Gumbel's model for the AMWH data with a scientific calculator and allowing the students to try this out for themselves; (ii) use our *Shiny* web application to automate the calculations in Gumbel's model, allowing students and facilitators alike to engage with the ideas behind the model without getting embroiled in the mathematics. In the discussion below, we focus mainly on approach (ii) as we believe many teachers/facilitators and students would be most comfortable with this; however, we also briefly discuss how we have used approach (i) with older students at recent school visits.

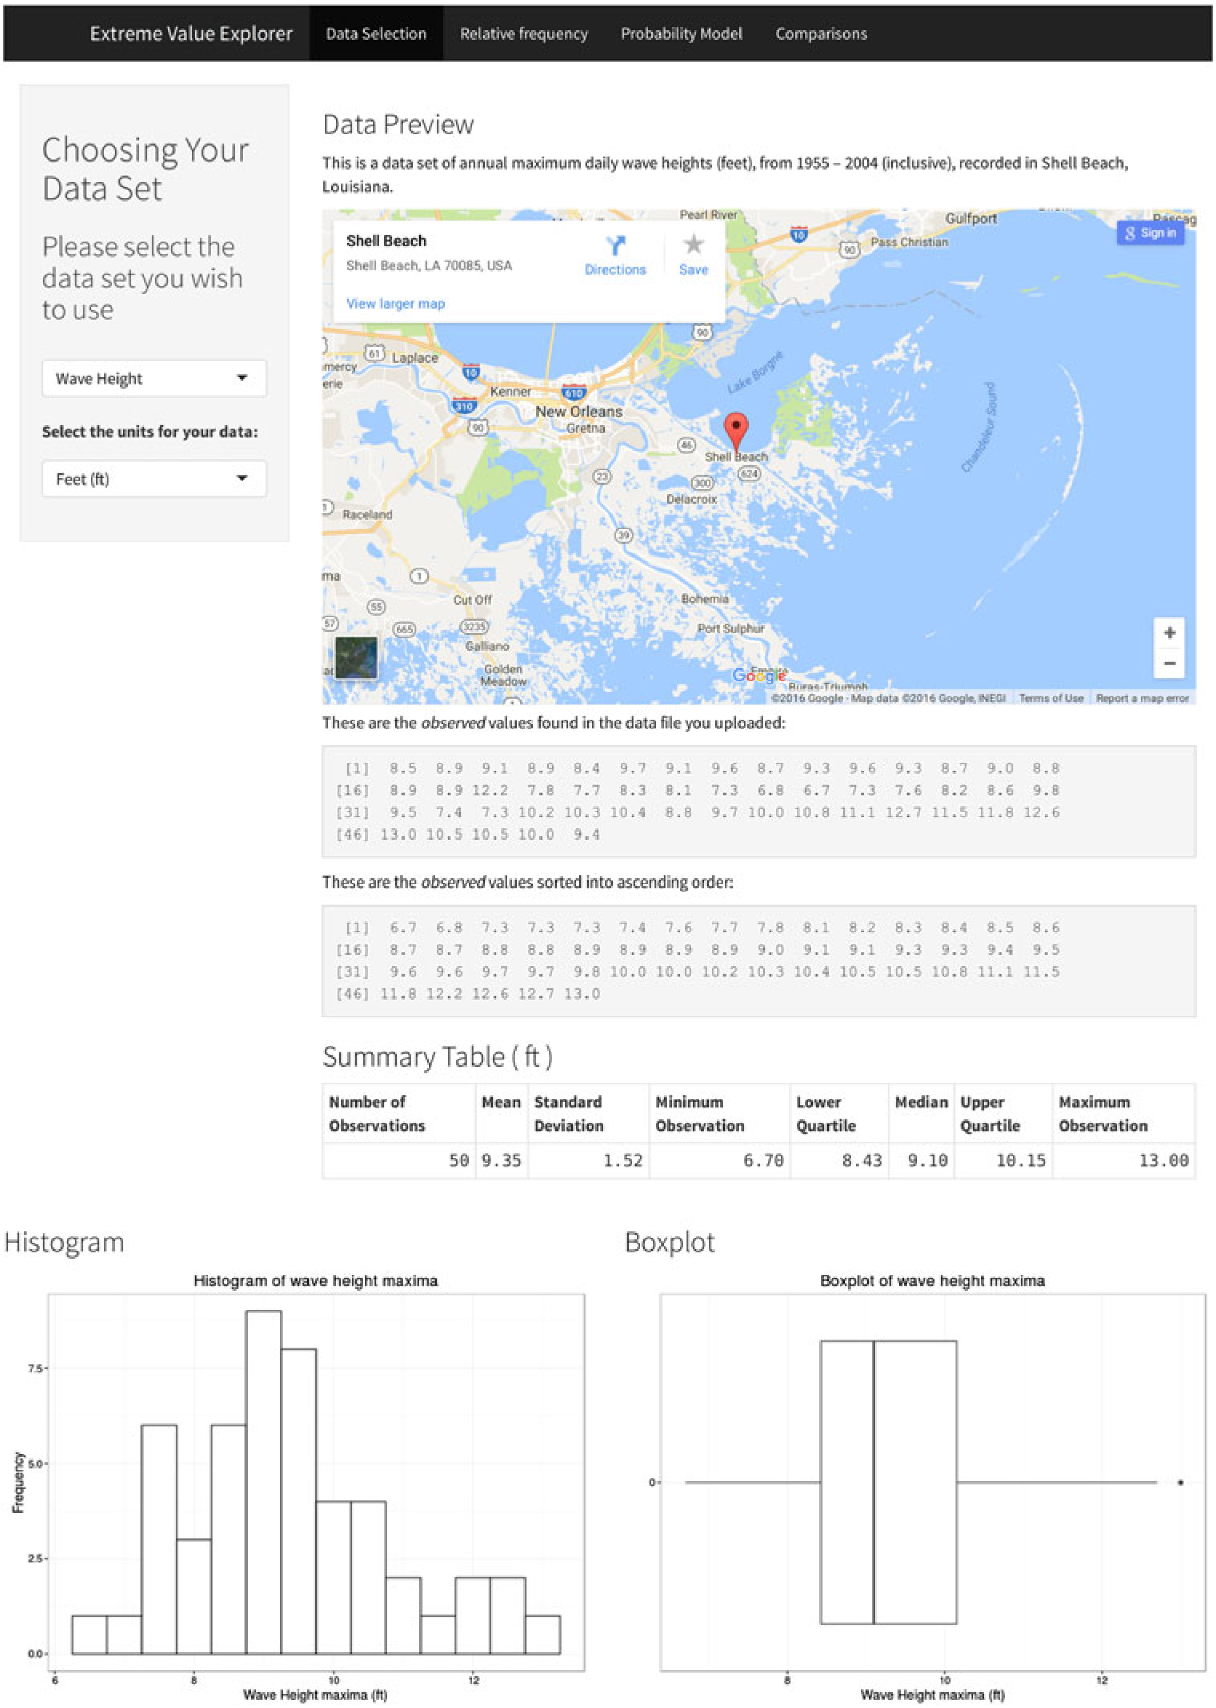

Fig. 2. Data preview page of the *Shiny* web application, showing the data selection menu, geographical location of the data collection site, the raw data and associated numerical and graphical summaries

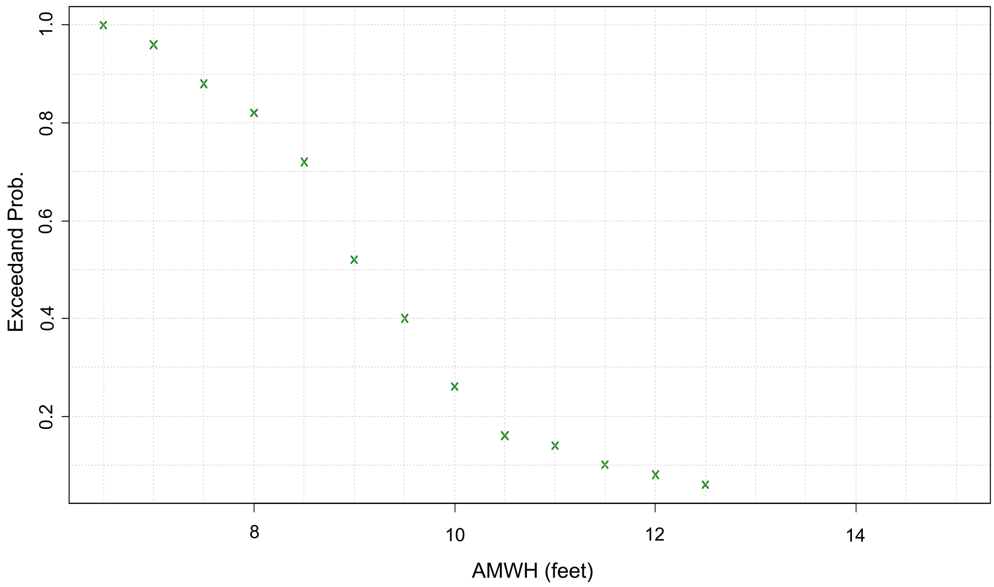

Fig. 3. Plot of relative frequencies completed by students, with annual maximum wave heights on the *x*-axis and the associated relative frequency exceedance probabilities on the *y*-axis

The *Probability Model* page of our *Shiny* application allows the user to obtain model-based estimates of relative frequency exceedance probabilities from the Gumbel model by selecting the option *Two-parameter Gumbel Model*. A box appears, giving the functional form of the model and a description of the parameters in the model with their (maximum likelihood) estimates. However, this feature can be ignored, if desired, and attention immediately given to the *Table of probabilities* and *Plot of probabilities* underneath; see the screenshot in Figure 5. The model-based estimates shown here can be compared directly to those estimates obtained from the data on the *Relative Frequency* page of the application (Figure 4) or, analogously, to those obtained by hand by the students (e.g. Figure 3). Notice from Figure 5 that the user has the option to switch between a variety of commonly-used probability models; the *Normal Model*, in particular, is considered in the extension activity (see part 6 below). Underneath the table and plot another feature of the application is a slider which can be adjusted to allow model-based exceedance probabilities for *any* value of interest to be returned.

If the teacher/facilitator chooses to work with Gumbel's formula directly with the students, using a scientific calculator instead of our *Shiny* application, the estimates of the model parameters (as provided by the application) should be given. We recommend the use of trained classroom assistants to help the students perform the calculations on a scientific calculator. With a more experienced audience discussion of the exponential function can be made here. We often allow students to work in pairs at this point, and although they find the calculator work challenging they are often extremely satisfied when they realize they can do it!

Discussion surrounding best-fitting models is usually made, supported by the *Comparisons* page of the *Shiny* application; see the screenshot in Figure 6. Students are reminded of the importance of a good-fitting model as a basis for extrapolation. An informal assessment of the Gumbel model, relative to other probability models, can be made by visually comparing the curves shown in Figure 6 to the relative frequencies from the data. More formally, the *goodness-of-fit* table in the application gives the sum of the squared vertical distances between the model-based estimates and the relative frequencies in the plot—the smaller this value, the better.

Students are asked to return model-based estimates of the three exceedance probabilities considered in the first part of the activity and complete Table 2. For the Gumbel model, they are advised to take a reading from the curve shown in the plot in Figures 5 or 6, or to use the slider in Figure 5 to obtain a more accurate estimate. We explain the relevance of the estimates based on the *Normal Model* in the extension activity (part 6).

Part 4 - Practical interpretation

Students are asked to think about why, according to our analysis and results in Table 2, Katrina might be considered the *Storm of the Century*. The ensuing discussion is often very interesting. Some students make the connection between the

Extreme Value Explorer | Data Selection | Relative frequency | Probability Model | Comparisons

Relative Frequency

How probability is calculated using relative frequency:

There are $n = 50$ observations in the dataset you provided. Therefore the probability of exceeding x is

$$\Pr(X > x) = \frac{\text{Number of observations exceeding } x}{\text{Total number of observations } (n = 50)}.$$

Table of probabilities

x	Number of observations exceeding x	Probability of exceeding x
5.50	50	1.00
6.00	50	1.00
6.50	50	1.00
7.00	48	0.96
7.50	44	0.88
8.00	41	0.82
8.50	36	0.72
9.00	26	0.52
9.50	20	0.40
10.00	13	0.26
10.50	8	0.16
11.00	7	0.14
11.50	5	0.10
12.00	4	0.08
12.50	3	0.06
13.00	0	0.00
13.50	0	0.00
14.00	0	0.00
14.50	0	0.00

Plot of probabilities

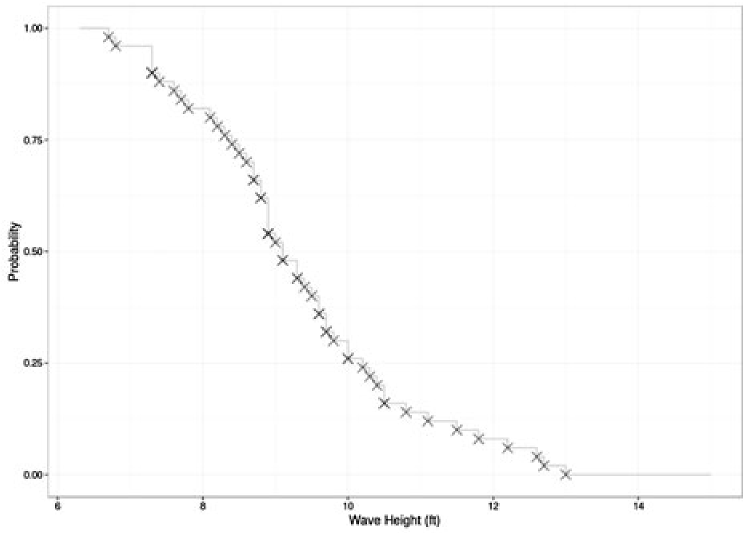

Calculate a probability from this model

Choose an extreme wave height amount to find the probability of exceeding it:

We have seen 0 instances when the wave height has exceeded 14.4 ft in our 50 observations. Therefore, the probability of observing a wave height greater than $x = 14.4$ ft every is given by

$$\Pr(X > 14.4) = \frac{0}{50} = 0 \text{ (to 4 significant figures)}.$$

How extreme could it get?

Choose how rare the event should be:

How extreme would we expect the wave height to be once every 100 years?

A once in a 100 year event corresponds to an exceedance probability $p = 0.01$ (to 4 significant figures).

The required height x of the wall can be calculated as,

$$x = 13 \text{ ft (to 2 decimal places)}.$$

Fig. 4. Plot of relative frequency exceedance probabilities in the *Shiny* web application, along with a slider bar, which can be used to obtain these probabilities. A slider bar is also included, which uses the inverse function to obtain empirical quantiles given a particular exceedance probability

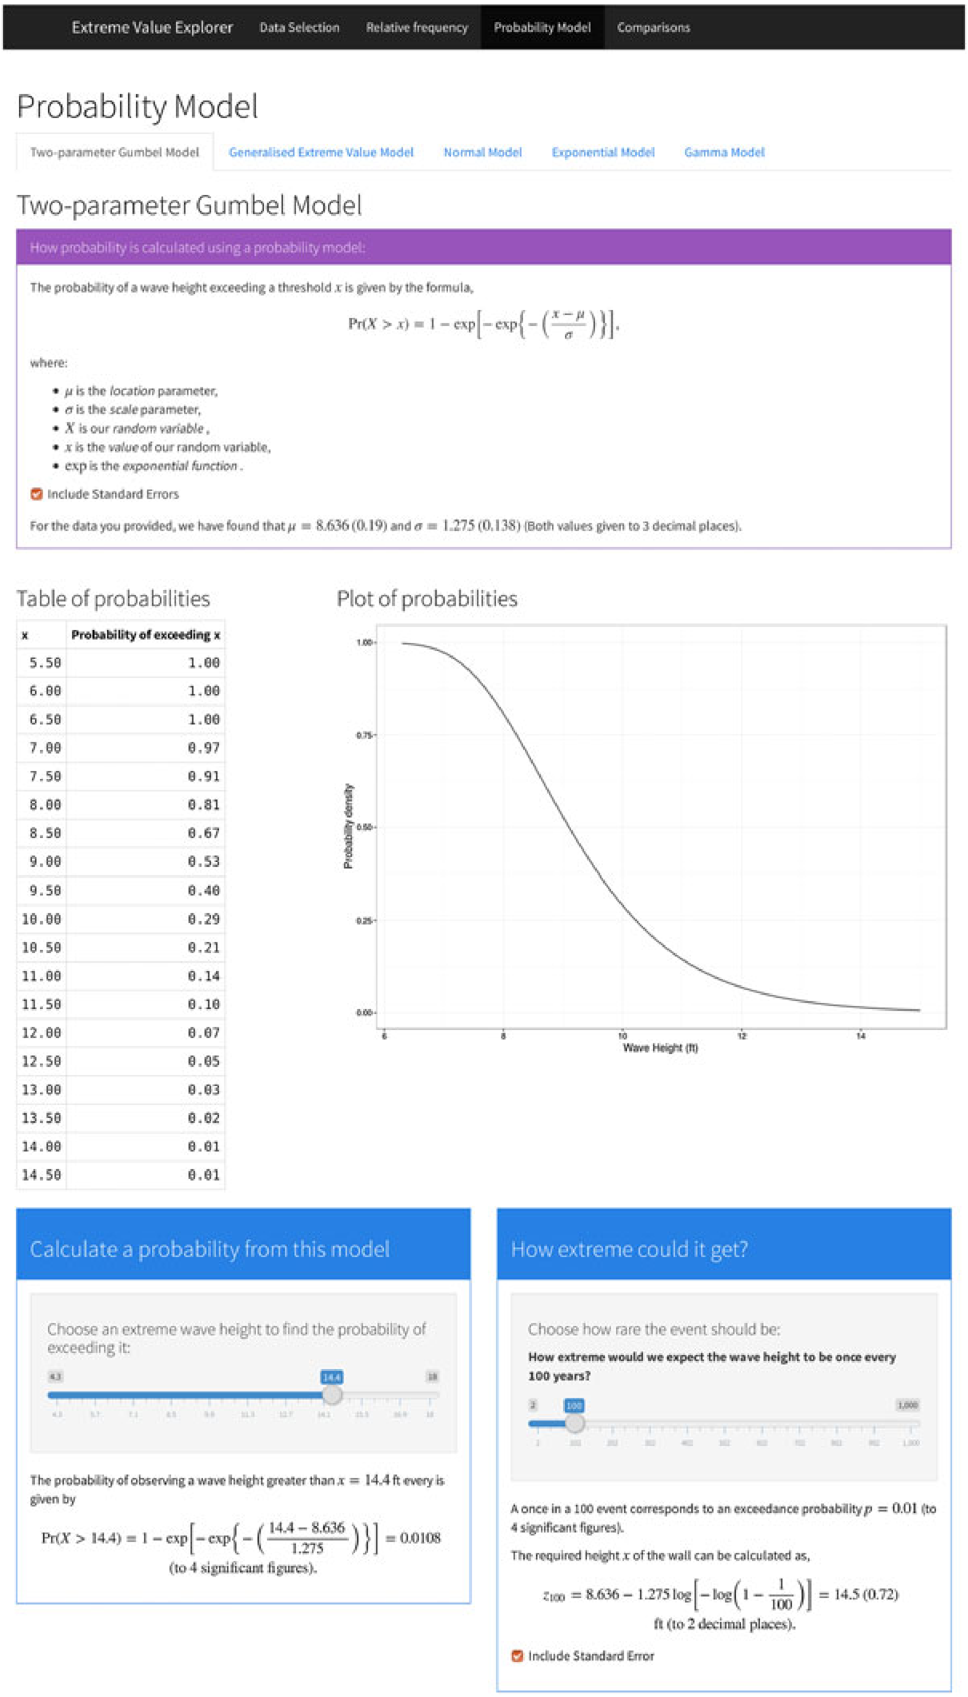

Fig. 5. Exceedance probabilities based on the fitted Gumbel model as shown in the *Shiny* web application. The slider bars allow model-based exceedance probabilities for any chosen value, as well as quantiles obtained on inversion of the fitted Gumbel model

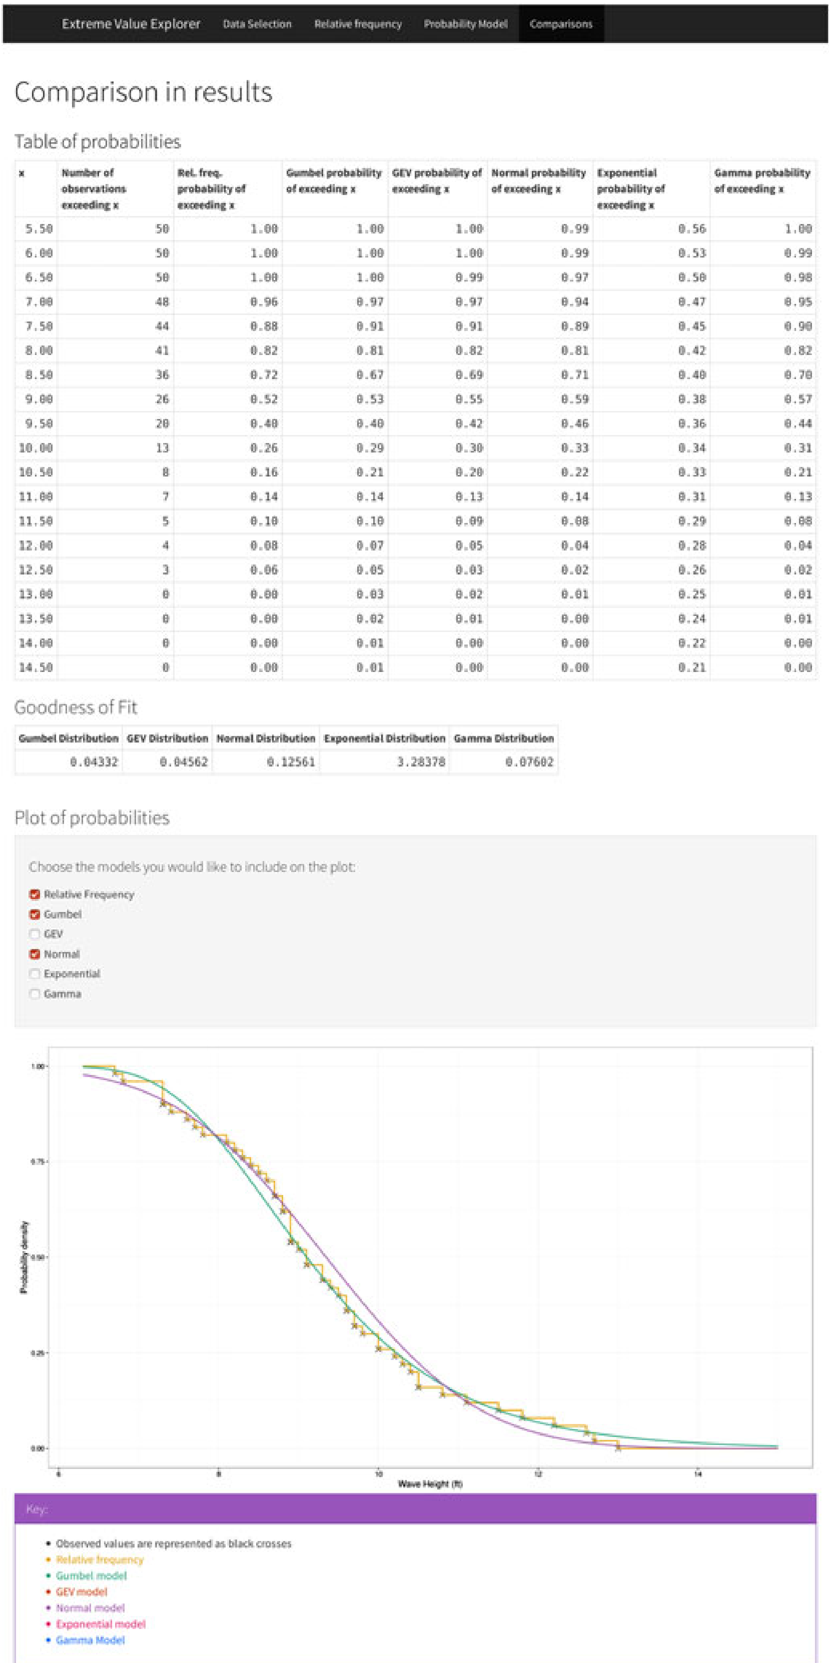

x	Number of observations exceeding x	Rel. freq. probability of exceeding x	Gumbel probability of exceeding x	GEV probability of exceeding x	Normal probability of exceeding x	Exponential probability of exceeding x	Gamma probability of exceeding x
5.50	50	1.00	1.00	1.00	0.99	0.56	1.00
6.00	50	1.00	1.00	1.00	0.99	0.53	0.99
6.50	50	1.00	1.00	0.99	0.97	0.50	0.98
7.00	48	0.96	0.97	0.97	0.94	0.47	0.95
7.50	44	0.88	0.91	0.91	0.89	0.45	0.90
8.00	41	0.82	0.81	0.82	0.81	0.42	0.82
8.50	36	0.72	0.67	0.69	0.71	0.40	0.70
9.00	26	0.52	0.53	0.55	0.59	0.38	0.57
9.50	20	0.40	0.40	0.42	0.46	0.36	0.44
10.00	13	0.26	0.29	0.30	0.33	0.34	0.31
10.50	8	0.16	0.21	0.20	0.22	0.33	0.21
11.00	7	0.14	0.14	0.13	0.14	0.31	0.13
11.50	5	0.10	0.10	0.09	0.08	0.29	0.08
12.00	4	0.08	0.07	0.05	0.04	0.28	0.04
12.50	3	0.06	0.05	0.03	0.02	0.26	0.02
13.00	0	0.00	0.03	0.02	0.01	0.25	0.01
13.50	0	0.00	0.02	0.01	0.00	0.24	0.01
14.00	0	0.00	0.01	0.00	0.00	0.22	0.00
14.50	0	0.00	0.01	0.00	0.00	0.21	0.00

Gumbel Distribution	GEV Distribution	Normal Distribution	Exponential Distribution	Gamma Distribution
0.04332	0.04562	0.12561	3.28378	0.07602

Fig. 6. The *comparisons* page of the *Shiny* web application, in this case comparing relative frequency exceedance probabilities to those obtained from the Gumbel and Normal models. The *goodness-of-fit* figures show the squared vertical discrepancies between the model-based and empirical exceedance probabilities for each model

Table 2. Model-based estimates of some exceedance probabilities, with associated empirical estimates

	Exceeds		
Probabilities	8.75 ft	11.25 ft	14 ft
Relative frequency	0.66	0.12	0
Gumbel model	0.575	0.1	0.01
Normal model	0.653	0.105	0.001

probability 0.01 and *once in a hundred years* straight away, whilst others do not make such a direct link but correctly remark that a probability of 0.01 "*... means this sort of event is extremely unlikely"*. Once explained, most students are often extremely satisfied with the title of the activity, and some are even excited by the fact that they can see the place of Statistics in the formation of such headlines *and* that they have managed to do the calculations themselves! At this point, we discuss the process of extrapolation that we referred to earlier, and the reliance—more than ever—on a well-fitting model for past observations on extremes.

Part 5 - Structural design: transposition of formulae

With older students, we now usually discuss a potential application of the fitted Gumbel model: its use as a tool for assisting the design of a new sea wall. We discuss the trade-offs between safety and cost—the higher the sea wall, the greater the level of safety afforded to a town or city, but also the greater the construction costs incurred. The following question is posed:

> *How tall should the sea wall be to protect against the AMWH we might expect to see, on average, once every 500 years?*

Students are now left to ponder how they might be able to use the fitted Gumbel model to help answer this question. Classroom assistants often wander round the room, asking students if they would know where to start with this. Usually, very few do. However, the following prompt is usually enough for some students to begin to tackle the problem:

$$P(\text{AMWH} > x) = \frac{1}{500}. \tag{3}$$

Replacing the left-hand-side of Eq. (3) with the fitted Gumbel model (see Eq. (4) in the Appendix) and then solving for x is usually manageable for students who are both algebraically confident *and* familiar with exponentials/natural logarithms. For those who are not—especially younger students—we refer to the box at the bottom-right of the *Probability Model* page in the *Shiny* application (Figure 5). Here, the calculations are performed automatically depending on the exceedance probability determined by the value selected on the slider (500 in this example). Thus, the height of the sea-wall offering protection against the AMWH we might expect to see once (on average) every 500 years is 16.56 ft (to 2 d.p.). This is known as the *500 year return level estimate* (the screenshot in Figure 5 shows the estimated 100 year return level).

Part 6 - Extension task: Comparison to the Normal distribution

The extension task here might prove useful for older students who have some experience of working with the Normal distribution. Interested teachers should read on; otherwise, the activities can end with the tasks in parts 4 or 5.

Students who are taking/have taken Statistics at a more advanced level might already be familiar with some basic models for probability. To complete this activity and so students can contextualize this work with their own study of probability models, we compare estimates of exceedance probabilities and quantiles, such as those given by the Gumbel model in Table 2, with those from a model they are more familiar with: the Normal $N(\mu, \sigma^2)$ distribution. This requires estimation of the mean μ and variance σ^2 from the data in Table 1—an exercise in summary statistics in its own right—but also the use of tables of cumulative probabilities from the standard normal distribution and quantiles from this distribution. Doing so gives the exceedance probabilities shown in the bottom row of Table 2, and an estimate of the 500-year return level of 13.71 ft, considerably smaller than that suggested by the Gumbel model. The exceedance probability associated with 14 ft is also 10 times smaller than that suggested by the Gumbel model. Of course, probabilities and quantiles from the Normal distribution can be obtained automatically from the *Probability Model* page of the Shiny application without the need to perform calculations by hand using statistical tables.

Students are then asked:

- What might be the consequences of using the Normal distribution instead of the Gumbel model?
- Which model would you trust?

Of course, in a practical setting, using the Normal model relative to the Gumbel model results in an under-estimate of quantities such as the 500-year return level, possibly leading to substantial under-protection of a town or city if such a model were used to inform the design of a sea wall. Simple graphs of the data, such as histograms or boxlots, reveal the unsuitability of the Normal distribution for the AMWH data, with some positive skew; see Figure 2. More detailed discussions can lead to the consideration of standard errors for estimates of return levels, and perhaps confidence intervals, although the computation of such is often beyond the ability and experience of most students in the age range for which this activity is intended. Within the *Shiny* application, a check box can be selected if estimated standard errors are to be displayed.

FURTHER DISCUSSION

Although we provide a practical motivation for the study of extremes (rather than averages, for example) and explain that filtering out a set of annual maxima might be a good way of classifying observations as *extreme*, we also explain that this approach is wasteful of data; we might have daily, or even hourly records (as is the case with the AMWH used in our activities here), and we discard all but the largest value in each year. With older students, we explain that the procedures we use assume that our observations are *independent and identically distributed*. In the case of our AMWH data, the largest hourly observations each year usually occur at some point during the hurricane season—often August or September—and so successive values in the series of annual maxima are usually far enough apart to be deemed independent. Using daily, weekly or monthly maxima *would* give us more data to work with, but in doing so, we are likely to encounter issues of dependence between consecutive maxima and other issues associated with *non-stationarity*, including seasonal variability. Studies have shown that violations of the assumption of *independent and identically distributed* observations can lead to biased estimates of return levels (e.g. Fawcett and Walshaw, 2016).

With older students, we occasionally discuss climate change. In part 3, we discuss the importance of a well-fitting model for historical observations as a basis for making predictions of future levels of AMWH; of course, any knowledge about how our variable is changing through time, perhaps as a result of climate change, should be utilized to provide more realistic estimates of return levels. Occasionally, and where appropriate, we explain how models like the Gumbel model can be adapted to account for changes in the underlying level of the extremes of our variable. For example, a simple way to account for trend might be to allow the location parameter in the Gumbel model to depend linearly on time. Recent studies examining AMWH at locations in the Gulf of Mexico extol the merits of such an approach, and there is evidence to suggest an increasing trend in the location parameter of the Gumbel model for the AMWH studied here.

EVALUATION

We believe that the activities discussed in this paper have had a positive impact on students' enthusiasm for Statistics. It is evident when we run the activities that students are generally engaged with the topic and many seem to genuinely enjoy taking part in the work. School teachers have often been even more enthusiastic, asking our permission to use the materials in class with other students and asking if any follow-up material exists. A key to the success of these activities, we have been told, is not just their demonstration of very practical applications of Statistics, but the fact that the material is directly related to our own personal research; as such, we are always extremely enthusiastic about the material. The merits of research-informed teaching and learning are discussed in, for example, Griffiths (2004) and Healey (2005), although as far as we are aware there is little in the way of evaluating the success of such methods in engagement and outreach activities.

Although rather anecdotal, teachers have told us that their students have become much more receptive to using Statistics in subjects other than Mathematics at school. After taking part in the activities discussed in this paper, some have also shown an increased enthusiasm for studying Science, Technology, Engineering and Mathematics subjects after their school study. Other evidence of the success of our activities comes from student evaluation questionnaires given out at the end of our sessions. For example, at a recent student conference held at our University, at which *The Storm of the Century!* activities were used, 65% of respondents (aged 16–17 years) said they would be more likely to study Mathematics/Statistics at University after having taken part in the sessions; 75% said they felt more enthusiastic about their school study of the subject. Other open-ended comments from recent school visits include:

"Wow! Loved the Storm of the Century. Didn't know the stuff we learn at school could be so interesting"

"… ticked the boxes for me, I like to see how this stuff can actually be used in the real world"

"Storm of Century is great as I love Maths and Geography and didn't know the two could be linked"

"Started off really easy but then we were doing high level stuff in no time, I never thought I'd be able to do that stuff."

CONCLUSIONS

We have outlined some hands-on classroom activities centred around the analysis of annual maximum wave height data to enthuse students about real-world applications of Statistics. These activities have been used with students as young as 12 years old, although extra layers of complexity can be added on to include material relevant, and challenging, for older students. The activities have always been popular, and there is some evidence to suggest they have been successful in promoting the study of Statistics and its use in other school subjects (such as Geography). Interested readers are invited to take a closer look at the materials used for the activities discussed, available to download from our webpage, and use them where they deem appropriate. Other activities are also available from this webpage, including *Speed Cameras Save Lives?*, *The Pepsi Challenge*, *The Lie Detector Test* and *The Game Show Problem (Revisited)*. The Shiny application can be used remotely by following the link from our webpage.

Appendix

In this paper, and in the *Storm of the Century!* activity, we focus primarily on modelling extremes using the Gumbel distribution. In fact, this is just one of three *extreme value distributions*, which can be shown, due to the *Extremal Types Theorem*, to be the limiting distributions for re-scaled maxima $(M_n - b_n)/a_n$, where

$$M_n = \max(X_1, X_2, \ldots, X_n)$$

and X_i, $i = 1, 2, \ldots, n$, are independent and identically distributed random variables. The other two extreme value distributions are often referred to as the Fréchet and Weibull distributions. The *generalized extreme value (GEV) distribution* unifies the three extreme value distributions, with the value of the shape parameter in this distribution controlling the tail heaviness and reducing the GEV to the Gumbel/Fréchet/Weibull models when it is zero/positive/negative. The Gumbel model, as used in this activity, has survival function:

$$P(X > x) = 1 - \exp\left\{-\exp\left[-\left(\frac{x-\mu}{\sigma}\right)\right]\right\}, \qquad (4)$$

where X is the random variable, x represents a specific value of this random variable, and μ and σ are parameters of location and scale, respectively. It is common practice to estimate these parameters via maximum likelihood; see, for example, Coles (2001, Ch. 2). In this activity, no description of maximum likelihood is given, and the estimates are simply provided by the *Shiny* application (which are, incidentally, μ=8.636 and σ=1.275).

References

Chang, W., Cheng, J., Allaire, J.J., Xie, Y. and McPherson, J. (2015). Shiny. R Package Version 0.12.1. Available at: https://cran.r-project.org/web/packages/shiny/shiny.pdf

Coles, S.G. (2001). *An Introduction to Statistical Modeling of Extreme Values*, London: Springer-Verlag London.

Fawcett, L. and Walshaw, D. (2016). Sea-surge and wind speed extremes: optimal estimation strategies for planners and engineers. *Stochastic Environmental Research and Risk Assessment*, **30**, 463–480.

Griffiths, R. (2004). Knowledge production and the research-teaching nexus: the case of the built environment disciplines. *Studies in Higher Education*, **29**(6), 709–726.

Healey, M. (2005). Linking research and teaching: exploring disciplinary spaces and the role of inquiry-based learning. In: Barnett, R. (ed.) *Reshaping the university: new relationships between research scholarship and teaching*, pp. 67–78. Maidenhead: McGraw-Hill/Open University Press.

Taleb, N. (2007). *The Black Swan*, pp. 309. New York: Random House.