

Predicting Halloween- and Carrington-type solar flares

Lee Fawcett, Holly Park, Thomai Tsiftsi[†], Shaun Bloomfield^{*}

[†]Universidad Nacional Autónoma de México,

^{*}Northumbria University


Structure of this talk

- Background / Motivation
 - ▶ Carrington event
 - ▶ Halloween solar storm
 - ▶ Potential consequences
- Data and typical POT analysis
- Improved analysis
 - ▶ Threshold selection
 - ▶ Dependence
 - ▶ Prior specification
 - ▶ Results / comparison

Carrington solar storm, 1859

- Powerful **geomagnetic solar storm** during solar cycle 10 (1855-1867)
- Solar **coronal mass ejection** (CME) hit Earth's magnetosphere inducing one of the largest geomagnetic storms on record
- Observed and recorded by **Richard Carrington**
- Aurorae visible in the northern hemisphere as far south as Monterrey, Cuba and southern Japan
- **Telegraph systems** failed across Europe and North America
- Peak fluxes classified as A, B, C, M, X: each letter indicates a flare an order of magnitude larger than the one preceding it
 - A1 flare has a peak X-ray flux of $1 \times 10^{-8} \text{ Wm}^{-2}$
 - X1 flare has a peak X-ray flux of $1 \times 10^{-4} \text{ Wm}^{-2}$
 - Carrington event: estimated as X45, i.e. $45 \times 10^{-4} \text{ Wm}^{-2}$

Carrington solar storm, 1859



A. STORM OF ELECTRICITY

TELEGRAPH WIRES USELESS FOR SEVERAL HOURS.

ONE OF THE MOST SEVERE DISTURBANCES FOR MANY YEARS, EXTENDING EVEN TO EUROPE—TELEPHONE WIRES ALSO OBSTRUCTED—BUSINESS DELAYED A GOOD PART OF THE DAY.

Yesterday's storm was accompanied by a more serious electrical disturbance than has been known for years. It very seriously affected the workings of the telegraph lines both on the land and in the sea, and for three hours—from 9 A. M. until noon—telegraph business east of the Mississippi and north of Washington was at a stand-still.

The Auroral Display in Boston.

Boston, Friday, Sept. 2.

There was another display of the Aurora last night, so brilliant that at about one o'clock ordinary print could be read by the light. The effect continued through this forenoon, considerably affecting the working of the telegraph lines. The auroral currents from east to west were so regular that the operators on the Eastern lines were able to hold communication and transmit messages over the line between this city and Portland, the usual batteries being discontinued from the wire. The same effects were exhibited upon the Cape Cod and other lines.

The New York Times

Published: September 3, 1859


Copyright © The New York Times



Halloween solar storms, 2003

- Series of solar flares and CMEs occurring in October/November 2003 during solar cycle 23 (1996-2008)
- Generated the largest solar flare ever recorded by the **GOES** system, estimated at X28
- **Satellite systems** and **communications** affected, numerous **spacecraft damaged**, **power outages** in Sweden and Finland
- Aurorae observed in Texas and southeast Mediterranean

Near miss: Carrington-like solar storm, 2012





What's at stake?

- Disruption to **global positioning systems**. GPS industry estimated to be a \$1 trillion business
- **Satellite communications** at high risk
- Huge disruption to the **electrical grid**: areas such as the eastern US particularly vulnerable as power infrastructure highly interconnected, large cities without power for weeks
- Key questions: When might we see the next Carrington-like solar storm? What is the **return period** of a Carrington-like event? What is the strength of solar flare we might expect to see once per century?

- Extracted from the **Space Weather Prediction Center** (SWPC) online portal
- Recorded by the **Geostationary Operational Environmental Satellite** (GOES) system, operated by the US National Oceanic and Atmospheric Administration (NOAA)
- Data represent solar flares normalised from H_{α} spectral observations and X-ray fluxes, spanning a period of 43 years: November 1975-October 2017
- Data points fairly regularly-spaced, with around four or five observations per day
- Data exhibit substantial short-term temporal dependence and cyclic variability

Data



Modelling extremes: Starting point

Let $\{X_n\}$ denote a stationary sequence of random variables with common distribution function F , and let $M_n = \max\{X_1, \dots, X_n\}$.

It is typically the case that, as $n \rightarrow \infty$,

$$\Pr(M_n \leq x) \approx F^{n\theta}(x), \quad (1)$$

where $\theta \in (0, 1)$ is known as the **extremal index**; see, for e.g., [Leadbetter & Rootzén \(1988\)](#).

As $\theta \rightarrow 0$ there is increasing dependence in the extremes of the process; for an independent process, $\theta = 1$.

Modelling extremes: Block maxima

The **Generalised Extreme Value** (GEV) distribution is the limiting model for F^n , with distribution function (d.f.)

$$\mathcal{G}(y; \mu, \varsigma, \xi) = \begin{cases} \exp \left[- (1 + \xi(y - \mu)/\varsigma)^{-1/\xi} \right], & \xi \neq 0, \\ \exp \left[-\exp \left(-(y - \mu)/\varsigma \right) \right], & \xi = 0, \end{cases}$$

defined on $\{y : 1 + \xi(y - \mu)/\varsigma > 0\}$, where $-\infty < \mu < \infty$, $\varsigma > 0$ and $-\infty < \xi < \infty$ are parameters of location, scale and shape.

- It can be shown that $F^{n\theta}$ is also GEV with d.f. $\mathcal{G}(y; \mu^*, \varsigma^*, \xi)$, provided long-range dependence is negligible¹
- In practice, the GEV is used to model maxima over some convenient calendar unit – usually years

¹Technically, this requires **Leadbetter's $D(u_n)$ condition** to hold; see **Leadbetter *et al.* (1983)**.

Modelling extremes: Threshold excesses

Pickands (1975) showed that for $\theta = 1$ and u large, $(X - u | X > u)$ follows a **Generalised Pareto** (GP) distribution with d.f.

$$\mathcal{H}(y) = \begin{cases} 1 - (1 + \xi y / \sigma)^{-1/\xi}, & \xi \neq 0, \\ 1 - \exp[-y/\sigma], & \xi = 0, \end{cases}$$

defined on $\{y : y > 0 \text{ and } (1 + \xi y / \sigma) > 0\}$, where $\sigma = \varsigma + \xi(u - \mu) > 0$ and ξ are the scale and shape.

- Graphical diagnostics available for **selection of threshold** u , but often difficult to interpret
- Extremal index θ can't be absorbed into the GP parameters; careful consideration of **extremal dependence** required

Modelling extremes: Return level estimation

The **r -year return level z_r** can then be obtained by inversion of $\mathcal{G}(z_r)$ or $\mathcal{H}^\theta(z_r)$.

For threshold excesses, on equating to $1 - r^{-1}$ this gives

$$z_r = \begin{cases} u + \sigma \xi^{-1} \left[(\lambda_u^{-1} w_r)^{-\xi} - 1 \right], & \xi \neq 0, \\ u - \sigma \log(\lambda_u^{-1} w_r), & \xi = 0, \end{cases} \quad (2)$$

where $w_r = 1 - [1 - (rn_y)^{-1}]^{1/\theta}$, λ_u is the rate of threshold excess and n_y is the (average) number of observations per year.

For seasonally varying extremes, we usually solve

$$\prod_{m=1}^M \mathcal{H}_m^{\theta_m}(z_r) = 1 - r^{-1},$$


where $\mathcal{H}_m$ is the GP distribution function in season m with parameter set $(\lambda_{u_m}, \sigma_m, \xi_m)$, and θ_m is the extremal index in season m .

Cyclic variation

- Close inspection of the time series suggests non-stationarity **inherited from the solar cycle**, affecting both the frequency and the severity of extreme flares
- After trial-and-improvement, fit sinusoidal waves of the form

$$\sum_{i=1}^3 A_i \sin \omega_i t + B_i \cos \omega_i t + C$$

to the series of moving averages, subtracting the estimated function before performing extreme value analyses



Typical peaks over threshold (POT) analysis

- Identify independent clusters of solar flare extremes over some **suitably high threshold** u
- Runs declustering, with **cluster termination interval** κ , most commonly used
- Fit the GP distribution to the set of independent cluster peak excesses
- Substitute $(\hat{\lambda}_u, \hat{\sigma}, \hat{\xi}, \theta = 1)$ into (2) to obtain $\hat{z}_r/\hat{r}$:

(u, κ)		$\hat{z}_{50}$		$\hat{r}$ years	
		Post. mean	95% CI	Halloween	Carrington
$u = 90\%$ quantile	$\kappa = 10$	X35	(X31, X50)	38	110
	$\kappa = 15$	X35	(X31.5, X52)	37.5	100
	$\kappa = 30$	X37	(X32, X65)	33	96.5
$u = 95\%$ quantile	$\kappa = 10$	X35	(X30, X46)	38	100
	$\kappa = 15$	X36	(X31, X47.5)	37	97
	$\kappa = 30$	X38.5	(X30, X72.5)	30	79

(As reported in Tsiftsi and De La Luz, 2018)

Typical peaks over threshold (POT) analysis

- Choice of threshold difficult:
 - ▶ Just choose a high quantile?
 - ▶ Graphical diagnostics – e.g. **mean residual life plot**, **Hill plots**, **parameter stability plots**...
 - ▶ Once u is chosen, it is treated as fixed
- Choice of declustering scheme difficult:
 - ▶ Runs declustering: Choice of κ ?
 - ▶ Other filtering schemes available
 - ▶ Wasteful – **large uncertainty** on return level estimates
- Prior specification – objective priors, missed opportunity?

Improved analysis: Aims

Here, we have three main aims:

- 1 Estimate the **extremal index** for use in return level estimation via a Markov chain model for threshold excesses
 - ▶ Avoids declustering altogether, and the problems it brings
 - ▶ Presses into use information on *all* extremes
- 2 Capture uncertainty in the **choice of threshold**, and enable a more 'automatic' selection of u
- 3 Get more out of working within the Bayesian framework
 - ▶ Where possible, formulate **informative priors**
 - ▶ Make use of the **posterior predictive distribution**

Improved analysis: Mixture model

Suppose our data comprise a sequence of n identically distributed observations $\mathbf{x} = \{x_i : i = 1, \dots, n\}$. We assume the following mixture model:

$$f(x) = (1 - \lambda_u) \cdot f_1(x) + \lambda_u \cdot f_2(x),$$

where

$$f_1(x) = \frac{h(x|\phi, \mathbf{x})}{H(u|\phi, \mathbf{x})} I_{(-\infty, u]}(x) \quad \text{and}$$

$$f_2(x) = g(x|u, \sigma_u, \xi) I_{(u, \infty)}(x).$$

- Observations sub-threshold follow a non-parametric density h
- Excesses above the threshold follow a GP distribution with parameter set (u, σ_u, ξ) and extremal index θ
- Possibility of a discontinuity at u

Improved analysis: Bulk distribution

For $f_1(x)$ we use a kernel estimator of the form

$$\hat{f}(x) = n^{-1} \sum_{i=1}^n K_{\phi}(x - x_i), \quad (3)$$

where $K_{\phi}(y) = \phi^{-1}K(y/\phi)$ is chosen to be a unimodal probability density function that is symmetric about zero.

- Kernel function used in (3) not critical; we use a zero mean Gaussian kernel
- GP distribution is used for the upper tail – extrapolation into the upper tails is of no concern
- **Kernel bandwidth likelihood** $L_{\kappa}(\phi, u|\mathbf{x})$, for $A = \{i : x_i \leq u\}$:

$$\left\{ \frac{(1 - \lambda_u)}{n^{-1} \sum_{k=1}^n \Phi\left(\frac{u - x_i}{\phi}\right)} \right\}^{|A|} \prod_A \frac{1}{n-1} \sum_{j=1; j \neq i}^n K_{\phi}(x_i - x_j)$$

Improved analysis: Threshold excess distribution

We assume a first-order Markov structure, giving the likelihood:

$$\prod_{i=1}^{n-1} f(x_i, x_{i+1}; \psi) \bigg/ \prod_{i=2}^{n-1} f(x_i; \psi). \quad (4)$$

- Transformation from GP to unit Fréchet margins gives a range of bivariate models to use for the numerator in (4)
- Most commonly-used is the **logistic family** with $\alpha \in [0, 1]$:

$$G(x_i, x_{i+1}) = \exp \left\{ - \left(x_i^{-1/\alpha} + x_{i+1}^{-1/\alpha} \right)^\alpha \right\}. \quad (5)$$

- **Extreme value Markov chain likelihood** $L_{\text{EV}}(u, \sigma_u, \xi, \alpha | \mathbf{x})$:
 - ▶ Differentiation of (5), with careful censoring, gives pairwise contributions to the numerator in (4)
 - ▶ Univariate contributions to the denominator are given through the GP distribution

Aside 1: Estimating the extremal index?

$$z_r = \begin{cases} u + \sigma \xi^{-1} \left[(\lambda_u^{-1} w_r)^{-\xi} - 1 \right], & \xi \neq 0, \\ u - \sigma \log (\lambda_u^{-1} w_r), & \xi = 0, \end{cases}$$

where $w_r = 1 - [1 - (rn_y)^{-1}]^{1/\theta}$, λ_u is the rate of threshold excess and n_y is the (average) number of observations per year.

Aside 1: Estimating the extremal index?

Define (arbitrarily) x_n such that $F^n(x_n) = 1/2$ in (1).

Then, using (1) $[\Pr(M_n \leq x) \approx F^{n\theta}(x)]$, we can define

$$\theta_n = -\frac{\log \Pr(\max\{X_1, \dots, X_n\} \leq x_n)}{\log 2}, \quad (6)$$


and so $\theta_n \rightarrow \theta$ as $n \rightarrow \infty$.

We can use (6) to investigate the relationship between θ and α in the logistic model:


- Simulate N first-order Markov chains, each of length n , with logistic dependence α
- The probability in the numerator of (6) is estimated as the proportion of simulated chains whose maximum does not exceed x_n
- We find values for θ_n for $\alpha = \{0.05, 0.10, \dots, 0.95\}$, with $N = n = 10^6$

Aside 1: Estimating the extremal index?


Logistic



Negative logistic



Bilogistic




Giving, for example,


$$\theta \approx 0.013 - 0.092\alpha + 1.833\alpha^2 - 0.756\alpha^3$$

Aside 1: Estimating the extremal index?


$AR(1)$ with $p = 0.9$



$AR(1)$ with $p = 0.95$



$AR(1)$ with $p = 0.99$



Improved analysis: Likelihood

Combining the Kernel density model for sub-threshold observations with the GP extreme value Markov chain for threshold excesses gives the **extreme value mixture model** likelihood:

$$L(\boldsymbol{\psi}|\boldsymbol{x}) = L_K(\phi, u|\boldsymbol{x}) \cdot L_{\text{EV}}(u, \sigma_u, \xi, \alpha|\boldsymbol{x}),$$

where $\boldsymbol{\psi} = (\phi, u, \sigma_u, \xi, \alpha)$.

Improved analysis: Prior specification

The joint prior for ψ is expressed as

$$\pi(\phi, u, \sigma_u, \xi, \alpha) = \pi(\phi) \cdot \pi(u) \cdot \pi(\sigma_u, \xi) \cdot \pi(\alpha)$$

1. For the **kernel density bandwidth parameter** ϕ we follow **Brewer (2000)** by specifying the prior for the precision $1/\phi^2$ as inverse gamma:

$$\pi(\phi|d_1, d_2) = \frac{1}{d_2^{d_1} \Gamma(d_1)} \left(\frac{1}{\phi^2} \right)^{d_1-1} \exp \left(-\frac{1}{\phi^2 d_2} \right),$$

with d_1 and d_2 chosen to give a largely non-informative prior specification.

2. In the absence of any expert prior information regarding the **logistic dependence parameter** α , we assume $\alpha \sim \mathcal{U}(0, 1)$.

3. Following Behrens et al. (2004) we assume a truncated Normal prior for the **threshold** u (truncated below e_1):

$$\pi(u|\mu_u, v_u^2) = \frac{1}{\sqrt{2\pi v_u^2}} \frac{\exp\{-0.5[(u - \mu_u)/v_u]^2\}}{\Phi[-(e_1 - \mu_u)/v_u]},$$

where μ_u is set at some high data percentile; v_u^2 is set sufficiently large to represent a diffuse prior.

Improved analysis: Prior specification

4. For the **GP parameters** (σ_u, ξ) we attempt to formulate informative priors.

- ▶ We work with a re-parameterised GP distribution:
 $\eta = \log(\sigma_u - \xi u)$
- ▶ Following **Coles & Tawn (2005)** we assume

$$\begin{aligned}\tilde{z}_1 &= z_{11} - e_1 \sim \text{Ga}(g_1, h_1) \\ \tilde{z}_2 &= z_{100} - z_{11} \sim \text{Ga}(g_2, h_2),\end{aligned}$$

where (g, h) reflect the opinions of an expert²

- ▶ Then:

$$\pi(z_{11}, z_{100}) \propto z_{11}^{g_1-1} e^{-h_1 z_{11}} (z_{100} - z_{11})^{g_2-1} e^{-h_2(z_{100}-z_{11})}$$

- ▶ Multiplying by the Jacobian of the transformation
 $(z_{11}, z_{100}) \longrightarrow (\eta, \xi)$ gives a prior in terms of (η, ξ)

²Shaun Bloomfield, Department of Mathematics, Physics and Electrical Engineering, Northumbria University

Improved analysis: Prior specification

Holly	Shaun
What classification of flare would you expect to occur only once per solar cycle (or once every 11 or so years)?	X20
What range of classifications would you expect the vast majority (say 95%) of extreme solar flares occurring only once in one solar cycle to lie in? (i.e in a given solar cycle, about 95% of the once-in-a-cycle flares will lie between X... class flare and an X... class flare)	Confident to give a lower limit of $\sim X5$ but an upper threshold an would be unsubstantiated
What classification of flare would you expect to occur only once every hundred years?	X40
What range of classifications would you expect the vast majority (say 95%) of extreme solar flares occurring only once in a hundred years to lie in? (i.e in a one hundred year period, about 95% of the once-in-a-hundred-year flares will lie between an X... class flare and an X... class flare)	Again, simply a lower bound would be most appropriate here, and this can be given as $\sim X30$
Could you also answer these questions for the 30-year, 150-year and 300-year scenarios	For 30 years, $\sim X30$, for 150 years, $\sim X50$ and for 300 year, $\sim X60$
How often would you expect to see an event of the same calibre as the Halloween event of 2003 (an X35 event)?	Every ~ 35 -50 years
How often would you expect to see an event of the same calibre as the Carrington event of 1859 (an X45 event)?	Every ~ 100 -150 years

Improved analysis: Prior specification

Constructing gamma priors:

i	r_i	Expected	95% Lower Bound	g_i	h_i
1	11	20	5	2.40	1012.80
2	100	40	30	36.33	9001.01

Checking coherence:

Return Period (years)	Elicited Returns Expected	Derived Returns		
		Estimate	Predictive	95% CI
30	30	25.74	25.78	(20.81, 31.31)
150	50	54.60	55.14	(41.21, 73.52)
300	60	74.90	76.30	(53.84, 105.24)

Improved analysis: Posterior sampling scheme

We use a simple Metropolis-Hastings scheme with random walk updates, tuning the innovation variances to optimise efficiency, giving successive draws

$$\boldsymbol{\psi}^{[j]} = (\phi^{[j]}, u^{[j]}, \eta^{[j]}, \xi^{[j]}, \alpha^{[j]})^T, \quad j = 1, \dots, 10^6.$$

- Chains thinned by a factor of 50 to remove substantial autocorrelation
- Convergence assessed using some standard diagnostics as discussed in [Gelman and Rubin \(1992\)](#)
- For threshold excesses, posterior draws for $(u, \sigma_u = f(u, \eta, \xi), \xi, \theta = f(\alpha))^{[j]}$ used to obtain posterior draws $z_r^{[j]}$ via (2)
- This leads to standard reporting of, e.g., **posterior mean** plus **posterior st. dev.** / **credible interval**

Improved analysis: Predictive inference

- Bayesian framework offers natural extension to prediction through the **posterior predictive distribution**
- Let Y denote a future extreme of our solar flares process. Then we can write


$$\Pr(Y \leq y|x) = \int_{\Psi} \Pr(Y \leq y|\psi) \pi(\psi|x) d\psi \quad (7)$$

- Equating to $1 - r^{-1}$ and solving for $y = z_{r,\text{pred}}$ gives a return level estimate which incorporates uncertainty due to parameter estimation
- (7) analytically intractable; approximate using

$$\frac{1}{S} \sum_{j=1}^S \Pr \left\{ Y \leq z_{r,\text{pred}} | \psi^{[j]} \right\} = \frac{1}{S} \sum_{j=1}^S \mathcal{H}^{\theta^{[j]}}(z_{r,\text{pred}} | u^{[j]} \sigma_u^{[j]}, \xi^{[j]})$$

where S is the number of MCMC iterations after burn-in

Aside 2: Summarising our return level inference?



Improved analysis: Results

	z_{50}	z_{1000}
Post. mean	X35	X57
95% CI	<i>X35</i> (X32, X41) <i>(X31, X50)</i>	<i>X54</i> (X49, X66) <i>(X40, X88)</i>
Predictive	X39	X79

	Halloween	Carrington
Post. mean	37 years	93 years
95% CI	<i>38 years</i> (31, 45) years <i>(22, 67) years</i>	<i>110 years</i> (80, 114) years <i>(83, 239) years</i>
Predictive	43 years	112 years

Conclusions

- Fitted a model which **doesn't require a choice of threshold / declustering scheme**
 - ▶ Graphical diagnostics for threshold choice are often difficult to interpret / contradictory
 - ▶ Uncertainty in threshold choice is accounted for
 - ▶ Return level estimates can be sensitive to the choice of declustering scheme
 - ▶ Using declustered extremes can result in low precision on return level/period estimates
- Is the model for dependence the 'best' ?
 - ▶ We used a model which implies asymptotic dependence
 - ▶ Even if this a valid assumption, there are several choices of bivariate extreme value model
 - ▶ Model order?
- Unsure about the effect of the bulk distribution
- **Directional modelling?**