Deep learning for climate and weather extremes

Jordan Richards

School of Mathematics and Maxwell Institute for Mathematical Sciences,
University of Edinburgh, UK

NORA Webinar





Climate and weather extremes

Extremes

Events with very low probability under the conditions of interest.

- **Wildfire:** how does the chance of a large burned area vary with weather and vegetation?
- **Metocean storms:** which combinations of winds and waves create dangerous conditions?
- **Alpine discharge:** how might high and low flows across catchments change together?

The relevant quantity is often a tail probability or extreme quantile.






Why extremes are difficult to learn

- **Rarity:** few observations describe the most consequential events.
- **Extrapolation:** decisions may concern events beyond the observed record.
- **Changing conditions:** tail behaviour can depend on many environmental covariates.
- **Dependence:** compound impacts involve several variables or locations.

A modelling challenge

Good predictions in the centre of a distribution do not establish accuracy in its tails.

Grey-box modelling

Extreme value theory

- Tail structure for extrapolation.
- Parameters with statistical meaning.

Deep learning

- Flexible covariate effects.
- High-dimensional environmental inputs.
- Fast!

Deep extremal regression

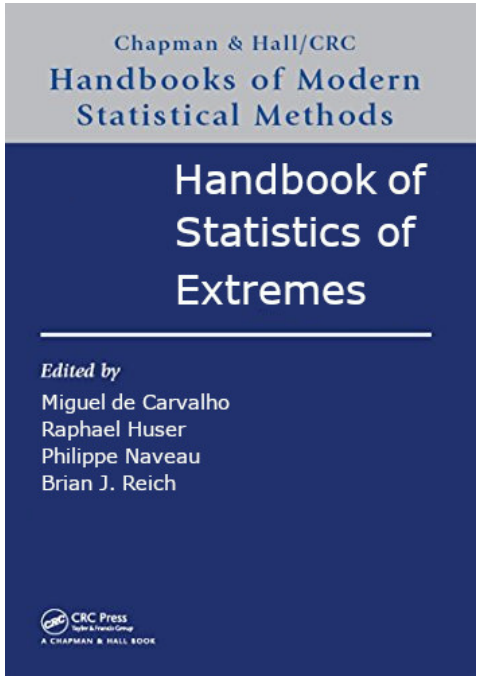
A neural network is used to approximate conditional tail functions - mappings from covariates to extremal risk measures of interest.

Interpretability can come from both the tail model and the regression structure.

Outline

1. **Conditional tails and deep extremal regression**
   Extreme quantiles, tail indices and return levels.
2. **European wildfire extremes**
   Environmental drivers and partially interpretable networks.
   Joint with Raphaël Huser, Emanuele Bevacqua, and Jakob Zscheischler
3. **UK metocean storms**
   Generative modelling with an angular–radial representation.
   Joint with Ed Mackay, Callum Murphy-Barltrop, and Phil Jonathan
4. **Alpine discharge extremes**
   Changing aggregate and concurrent risks.
   Joint with Callum Murphy-Barltrop, Ben Poschlod, Alexander Sasse, and Jakob Zscheischler

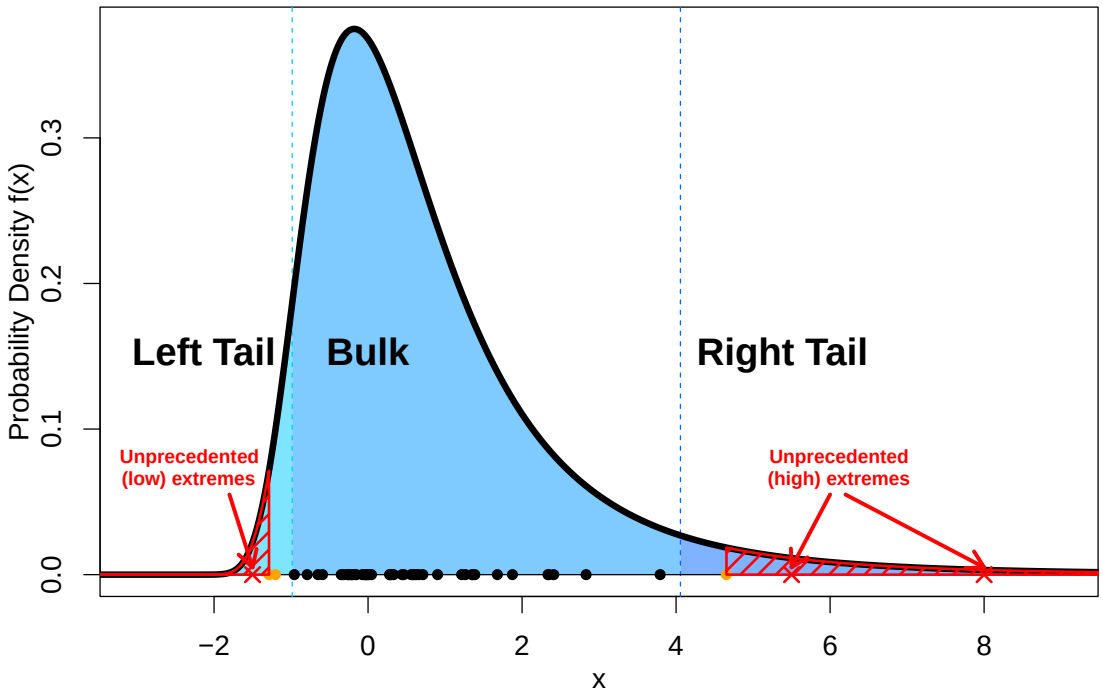
Outline



Chapter 21. Richards and Huser (2026).
Extreme Quantile Regression with Deep Learning.
Code and short courses on GitHub - `https://github.com/Jbrich95/`

Background on extreme events

- Univariate context: Observations $Y_1, \ldots, Y_n$.
  Can we estimate the probability of unprecedented extreme events of a given size (typically larger than $M_n = \max(Y_1, \ldots, Y_n)$)?



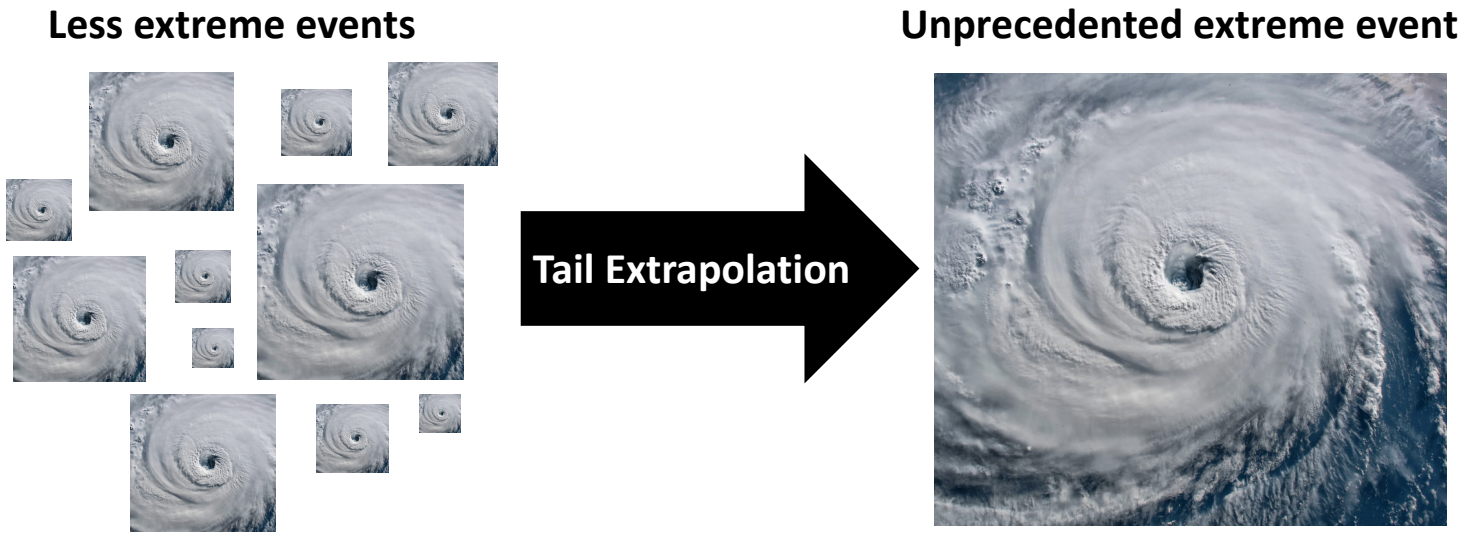
Background: Extreme-Value Theory

- Need to extrapolate into the tail, based on resilient, reliable, theoretically-justified methods

Main idea of Extreme-Value Theory (EVT)

Learn about unprecedented extreme events from less extreme events.

Requires assumptions about the regularity of the tail, or the tail decay rate



Marginal modeling of extremes — theory for peaks-over-threshold

- **Pickands–de Haan–Balkema Theorem:** high threshold excesses $Y - u \mid Y > u$ may be approximated by the generalized Pareto (GP) distribution - there exists a scaling $a(u) > 0$ such that as $u \to y_F$ (upper endpoint)

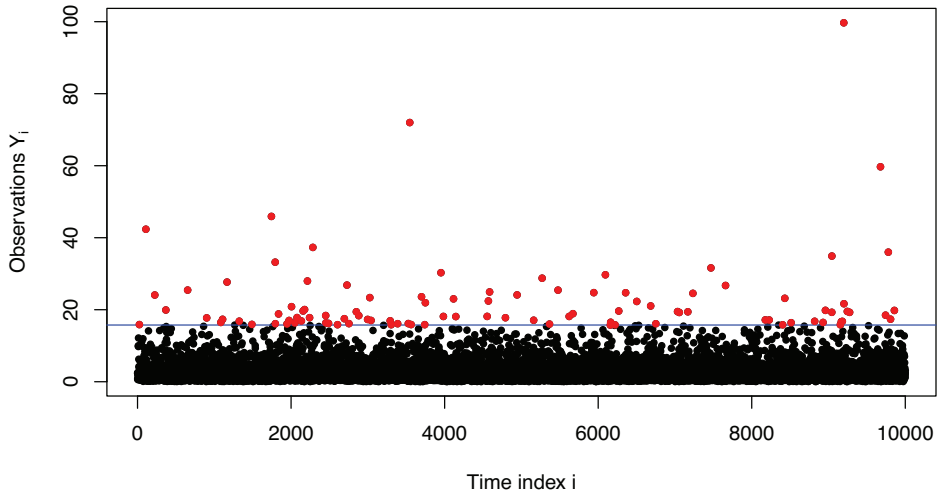
$$\Pr\left(\frac{Y-u}{a(u)} > y \mid Y > u\right) \to \begin{cases} (1+\xi y/\tau)_+^{-1/\xi}, & \xi \neq 0, \\ \exp(-y/\tau), & \xi = 0, \end{cases},$$

  where $\tau > 0$ and $\xi \in \mathbb{R}$ are a scale and shape, respectively.

- In practice, we can model excesses directly as $Y - u \mid Y > u \sim \mathrm{GP}(\tau, \xi)$.

Marginal modeling — Statistics of peaks-over-threshold

Threshold exceedance approach $Y_1, Y_2, \ldots, \overset{\text{iid}}{\sim} F$



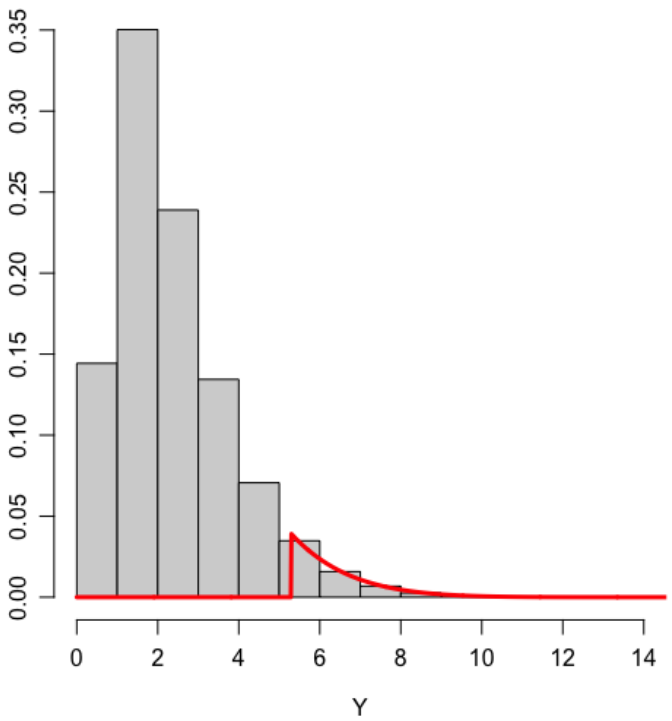
Pickands–de Haan–Balkema Theorem: *For a broad range of distributions F, we have the following large-u approximation*

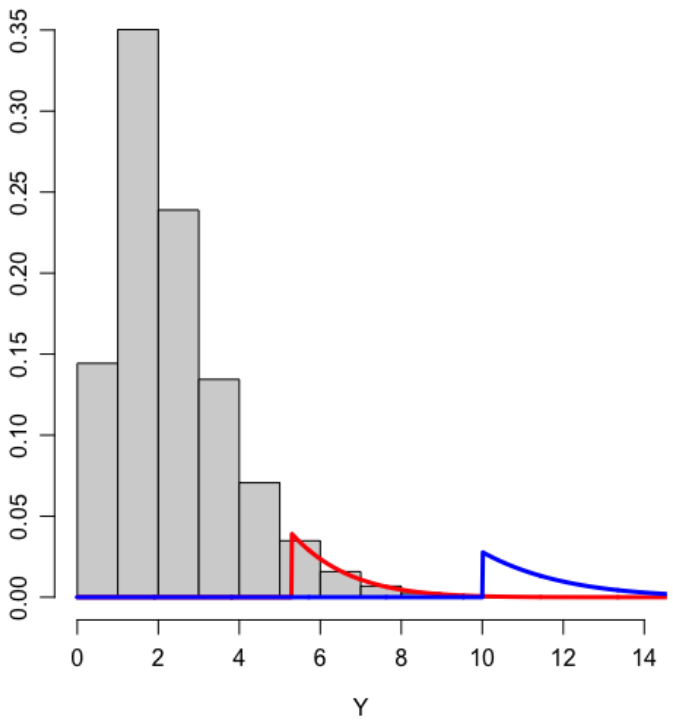
$$Y_i - u \mid Y_i > u \sim \text{GP}(\tau, \xi).$$

Extrapolation

Threshold-stability

If $Y - u \mid Y > u \sim \mathrm{GP}(\tau, \xi)$, then $Y - v \mid Y > v \sim \mathrm{GP}(\tau + \xi(u - v), \xi)$ for $v \geq u$.





Marginal modelling of extremes — GP density

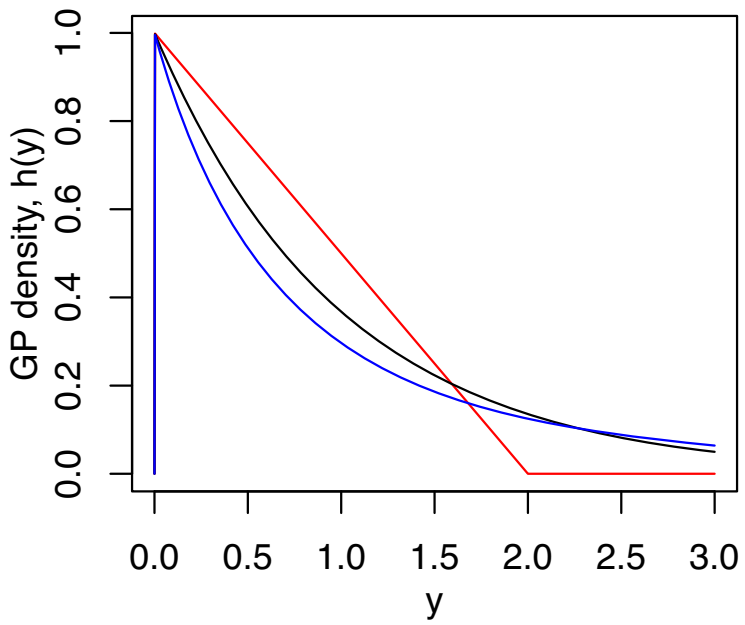


Figure: GP density; The shape parameter determines the distribution's tail decay and support.

Meaning of the shape parameter

$\xi > 0$: heavy tail. $\xi = 0$: exponential tail.
$\xi < 0$: finite upper endpoint under the fitted GP model.

Extremal regression: environmental conditions and risk

An environmental question

How does the severity of extreme wildfires vary with drought, temperature and vegetation?

Let Y denote burnt area and $\mathbf{x}$ describe the environmental conditions.

Extremal regression links these conditions to the upper tail of burnt area:

$$Y - u(\mathbf{x}) \mid \{Y > u(\mathbf{x}), \mathbf{X} = \mathbf{x}\} \sim \mathrm{GP}\{\tau(\mathbf{x}), \xi(\mathbf{x})\}.$$

- $u(\mathbf{x})$: what counts as unusually large under these conditions.
- $\tau(\mathbf{x})$: the scale of exceedances above this threshold.
- $\xi(\mathbf{x})$: how quickly the probabilities of even larger events decrease.

Goal: identify conditions associated with more severe extremes.

Regression

Standard choices for $(\tau(\mathbf{x}), \xi(\mathbf{x}))$:

- Linear: e.g., $\tau(\mathbf{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots$.
- Additive (typically utilising smoothing splines): i.e., $\tau(\mathbf{x}) = \beta_0 + \sum_{i=1}^{p} \sum_{k=1}^{K_i} \beta_{ik} \psi_{ik}(x_i)$.

Recent uptick in replacing $(\tau(\mathbf{x}), \xi(\mathbf{x}))$ with a non-parametric alternative from supervised ML, e.g., **Neural Networks**:

$$(\tau(\mathbf{x}), \xi(\mathbf{x})) = \mathbf{m}_1 \circ \mathbf{m}_2 \circ \cdots \circ \mathbf{m}_L(\mathbf{x}),$$

but this idea goes back to Carreau and Bengio (2007)!

Carreau, Julie and Bengio, Yoshua. (2007) A Hybrid Pareto Model for Conditional Density Estimation of Asymmetric Fat-Tail Data. AIaS.

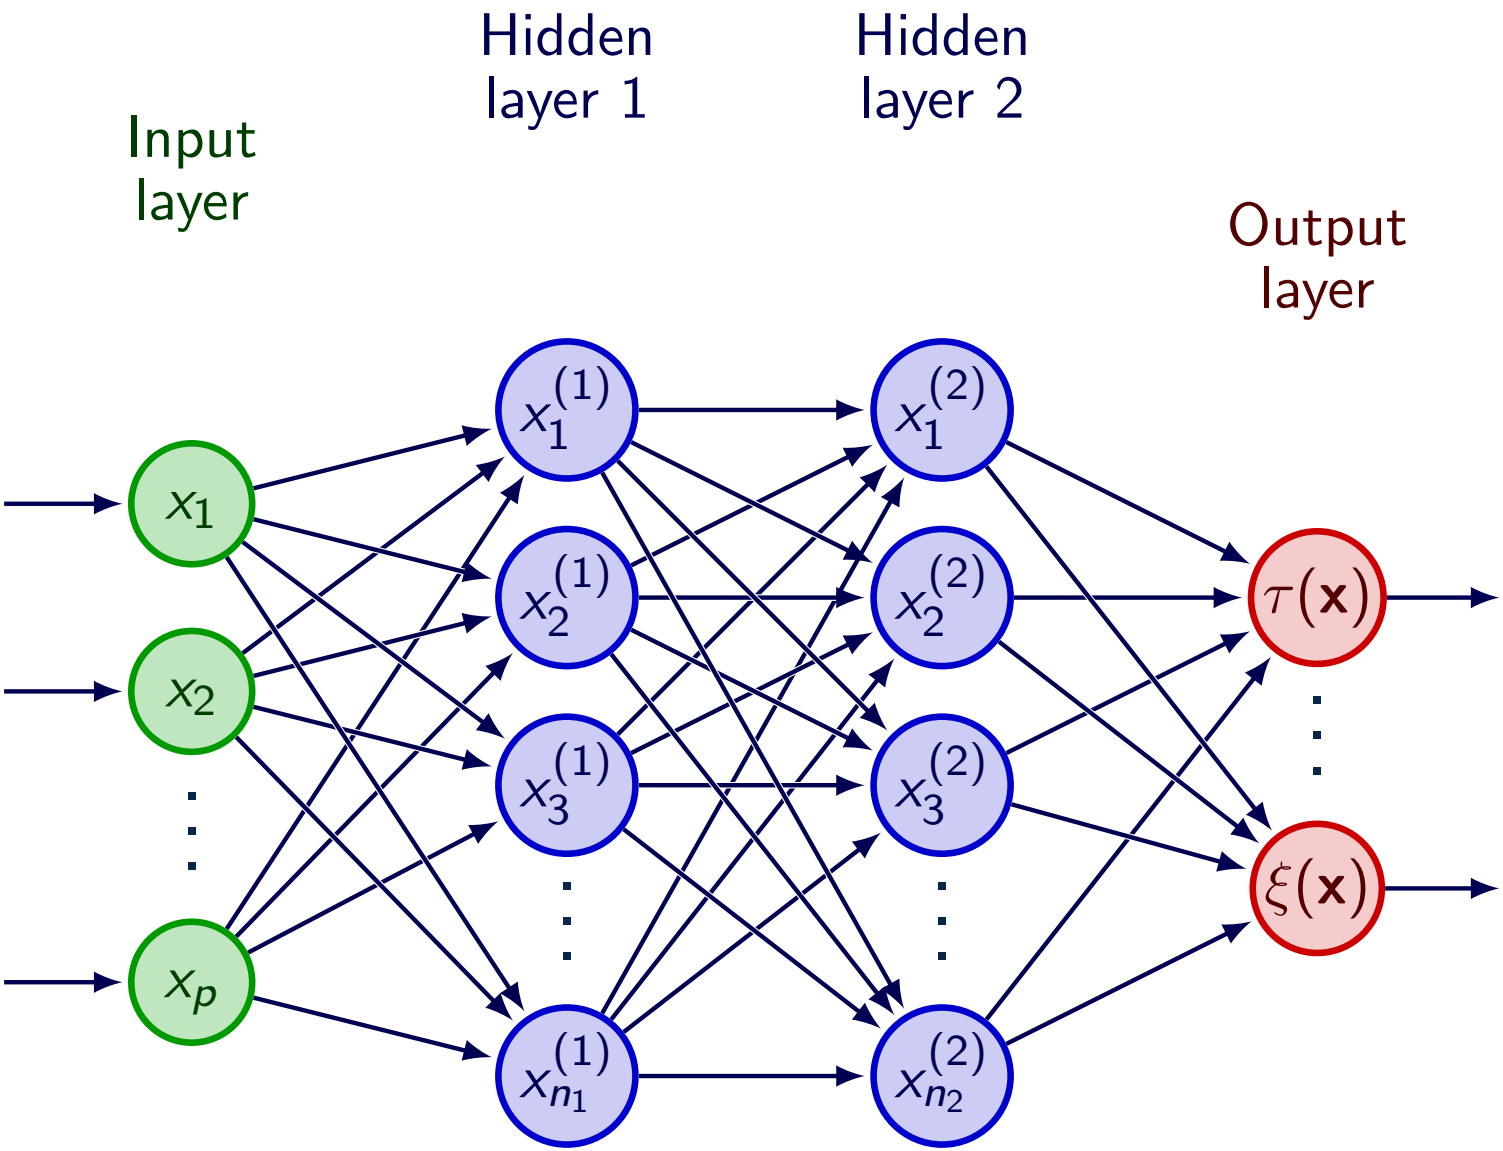
Deep regression for extremes

For covariates $\mathbf{x}$, a network outputs parameters:

$$\mathbf{h}_0 = \mathbf{x}, \qquad \mathbf{h}_\ell = a_\ell(W_\ell \mathbf{h}_{\ell-1} + \mathbf{b}_\ell), \qquad (\tau(\mathbf{x}), \xi(\mathbf{x})) = g(\mathbf{h}_L).$$

- Hidden layers learn nonlinear relationships and interactions.
- Output links enforce constraints, e.g., $\tau(\mathbf{x}) > 0$.
- The GP distribution retains control of the tail form.

Deep regression for extremes



Learning the threshold and tail

1. Estimate a conditional high quantile $u(\mathbf{x})$ using quantile loss.
2. Fit the excesses $e_i = y_i - u(\mathbf{x}_i) > 0$ using GP likelihood:

$$\mathcal{L} = \sum_{i:e_i>0} \left[\log \tau(\mathbf{x}_i) + \left(1 + \frac{1}{\xi(\mathbf{x}_i)}\right) \log\left(1 + \frac{\xi(\mathbf{x}_i)e_i}{\tau(\mathbf{x}_i)}\right)\right].$$

The support requires $1 + \xi(\mathbf{x}_i)e_i/\tau(\mathbf{x}_i) > 0$ for all $\mathbf{x}_i$.

- Regularisation and validation control network complexity.
- Threshold sensitivity and tail diagnostics assess the statistical approximation.
- Easy to couch in existing deep learning software - **just change the loss function.**

Extreme quantiles and return levels

Write $\zeta(\mathbf{x}) = \Pr\{Y > u(\mathbf{x}) \mid \mathbf{x}\}$. For $p > 1 - \zeta(\mathbf{x})$ and $\xi(\mathbf{x}) \neq 0$,

$$q_p(\mathbf{x}) = u(\mathbf{x}) + \frac{\tau(\mathbf{x})}{\xi(\mathbf{x})}\left[\left\{\frac{\zeta(\mathbf{x})}{1-p}\right\}^{\xi(\mathbf{x})} - 1\right].$$

For $\xi(\mathbf{x}) = 0$, the increment is $\tau(\mathbf{x}) \log\{\zeta(\mathbf{x})/(1-p)\}$.

Return-level interpretation

For a stationary sequence with m observation periods per year, set $1 - p = 1/(mT)$ for a T-year mean recurrence level.

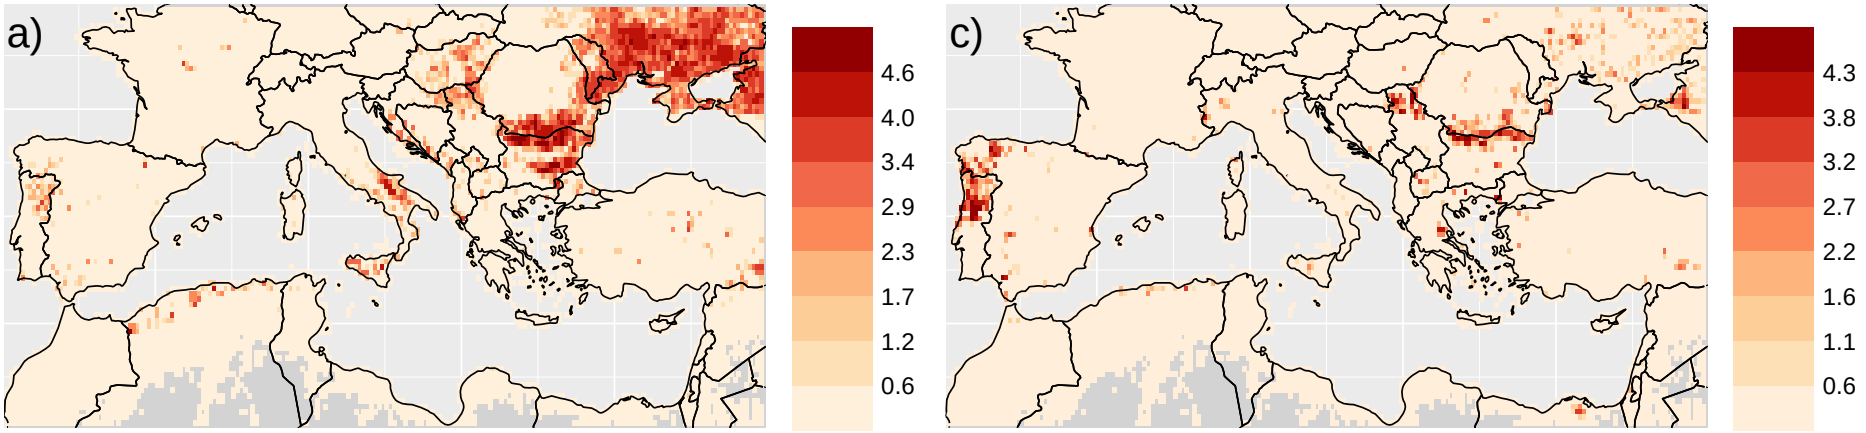
Changing covariates require a reference climate or exposure history.

Return-level interpretation

The quantiles $q_p(\mathbf{x})$ can be interpreted as a measure of hazard for risk level p and conditions $\mathbf{x}$.

Mediterranean Wildfire Extremes

- $Y(s, t)$ is monthly aggregated burnt area [km^2] in a $0.25^\circ \times 0.25^\circ$ grid cell, from FireCCI (Lizundia-Loiola et al., 2020).
- June–Nov., 2001–2020: $\approx 10\,000$ spatial locations and 1.2 million grid-cell months, including 102 240 positive observations.
- Most observations are zero. We distinguish **whether any burning occurs** from **how much area burns, given occurrence**.

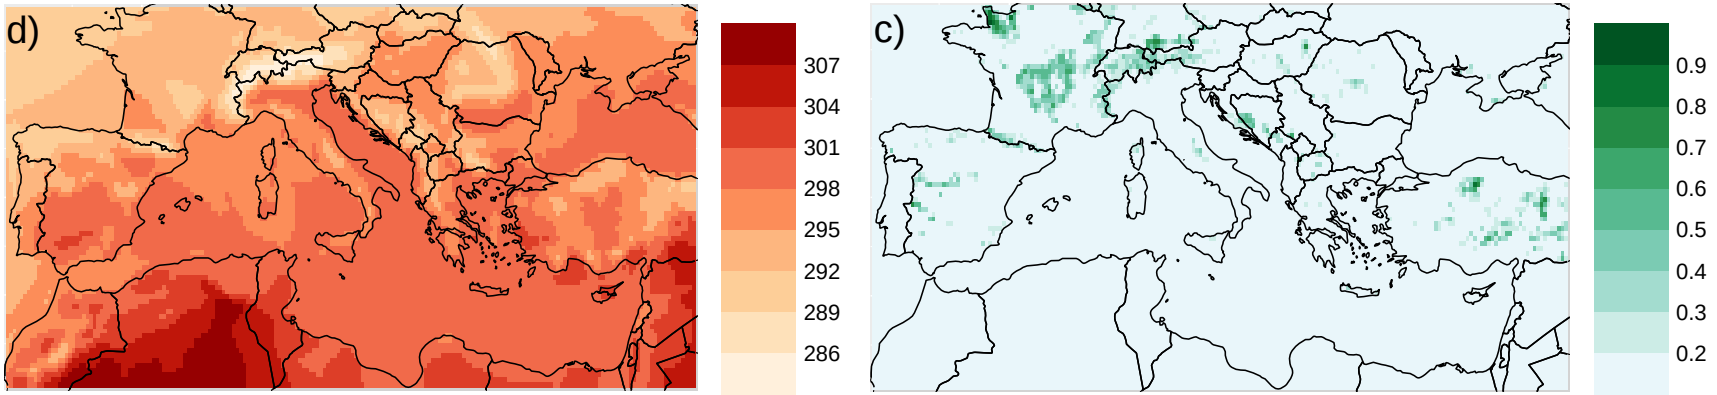


Maps of $\log(1 + \text{BA})$: August, 2001 (left) and October, 2017 (right).

Mediterranean Wildfire Extremes

- Land cover maps (**COPERNICUS**) with proportion of grid cells consisting of one of 21 types, e.g., water, urban areas, **grassland**

- 13 meteorological variables from **ERA-5 reanalysis on single levels**, e.g., **temperature**, wind-speed components, surface pressure



Temperature (left) and grassland proportion (right) for August, 2001.

Wildfire occurrence and conditional spread

For a grid-cell month, write Y for burnt area and $\mathbf{x}$ for its environmental covariates.

Occurrence

$$p_0(\mathbf{x}) = \Pr\{Y > 0 \mid \mathbf{X} = \mathbf{x}\}.$$

Does any burning occur?

Conditional spread

$$F_+(y \mid \mathbf{x}) = \Pr\{Y \leq y \mid Y > 0, \mathbf{X} = \mathbf{x}\}.$$

How much area burns, given that burning occurs?

Modelling extreme burnt area

Conditional on some burning occurring (p_0), we model the upper tail of monthly burnt area.

- **Where does the tail begin?** A neural network estimates a threshold $u(\mathbf{x})$ that varies with environmental conditions.
- **How often is that threshold exceeded?** A second network estimates the exceedance probability $p_u(\mathbf{x})$.
- **How large can exceedances become?** An extreme-value model describes their magnitude, with a scale $\tau(\mathbf{x})$ that depends on environmental conditions.

Tail behaviour

One common shape parameter across space and time:
$\widehat{\xi} = 0.32(0.28, 0.35)$.

Separating interpretable effects from flexible prediction

We model both **fire occurrence** and a **scale component of extreme burnt area** using two complementary components.

Interpretable component

Smooth curves describe associations with:

- air temperature T;
- vapour-pressure deficit VPD;
- precipitation deficit or surplus SPI.

Neural-network component

Learns complex relationships among the remaining environmental predictors.

Partially interpretable wildfire regression

Let $\mathbf{x}_{\mathcal{I}} = (T, \mathrm{VPD}, \mathrm{SPI})$ contain the predictors we wish to interpret. The remaining predictors form $\mathbf{x}_{\mathcal{N}}$ (Richards et al., 2023).

$$\mathrm{logit}\{p_0(\mathbf{x})\} = \underbrace{\sum_{j=1}^{3} f_j(x_{\mathcal{I},j})}_{\text{interpretable splines}} + \underbrace{m_0(\mathbf{x}_{\mathcal{N}})}_{\text{CNN}},$$

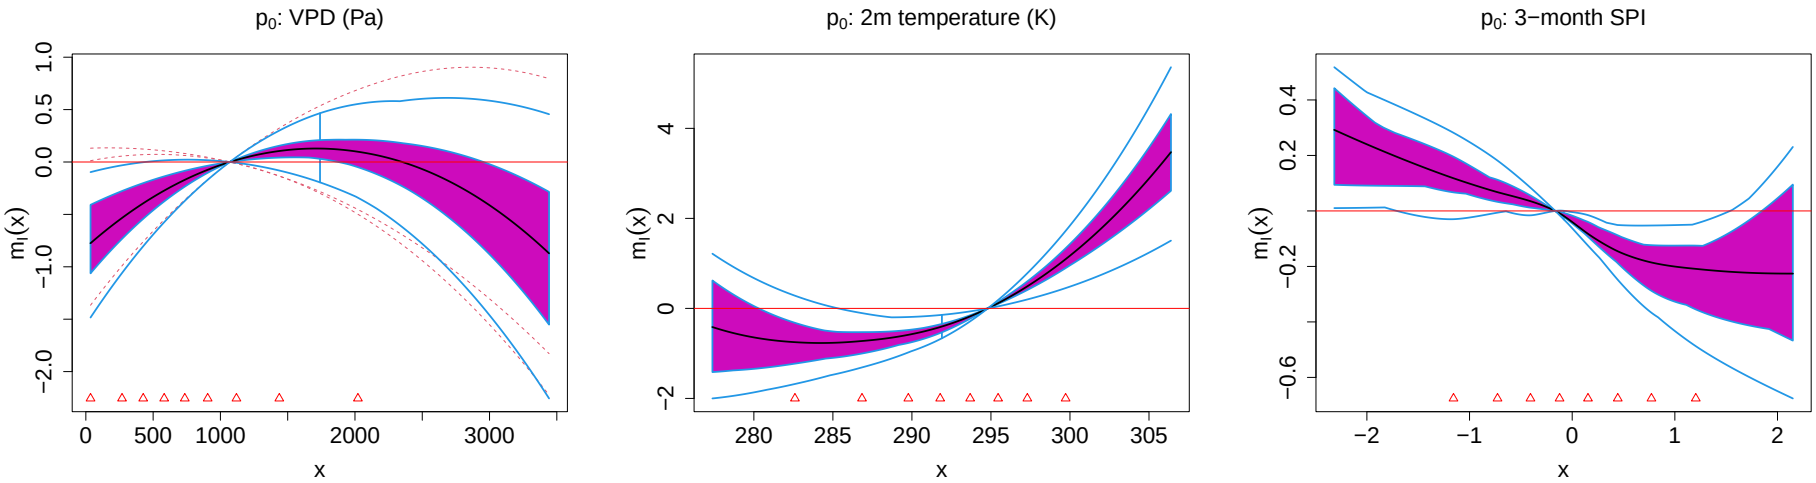
$$\log\{\tau(\mathbf{x})/\lambda(\mathbf{x})\} = \underbrace{\sum_{j=1}^{3} g_j(x_{\mathcal{I},j})}_{\text{interpretable splines}} + \underbrace{m_\tau(\mathbf{x}_{\mathcal{N}})}_{\text{dense neural network}}.$$

- $\lambda(\mathbf{x})$ is available burnable area, used as an offset in the scale model.

Richards, Huser, Bevacqua, Zscheischler (2023). Insights into the drivers and spatiotemporal trends of extreme Mediterranean wildfires with statistical deep learning. AIES.

Mediterranean Wildfire Extremes

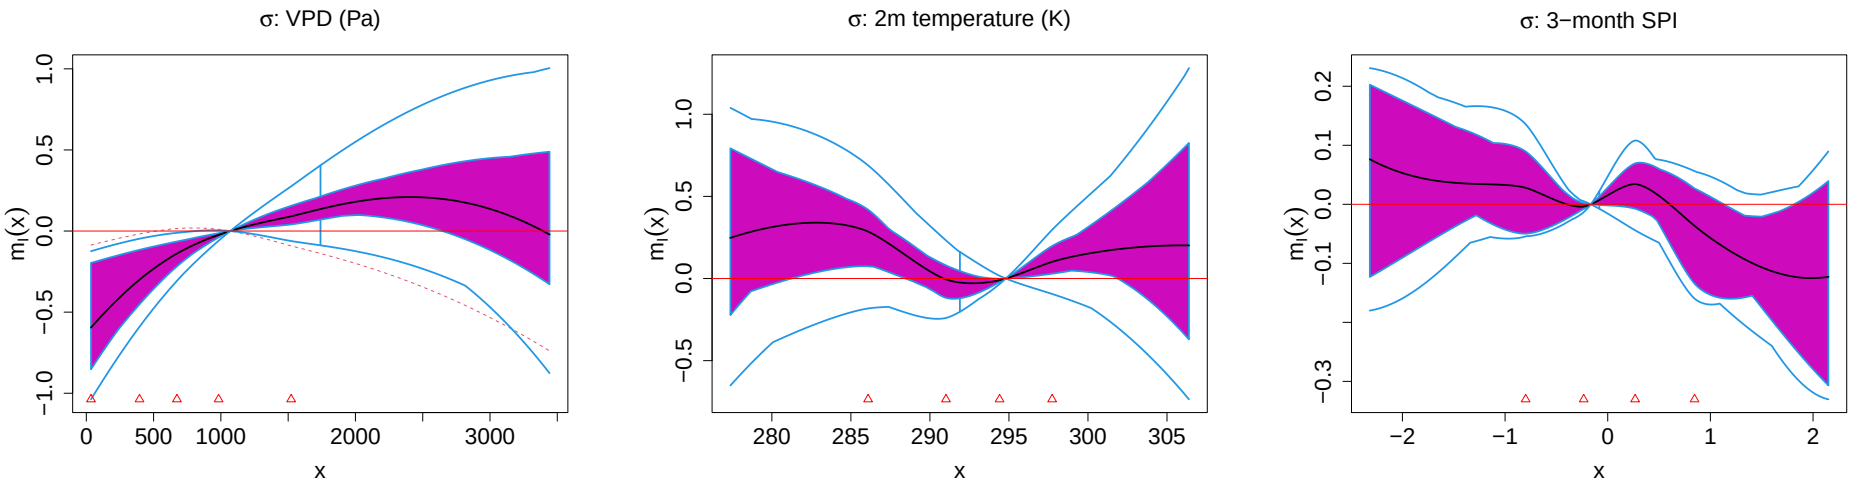
What are the non-linear drivers of wildfire occurrences?



- Higher temperature and lower SPI are associated with greater occurrence probability.
- The VPD effect rises initially, then declines above approximately 1500 Pa.

Effects are on the **log-odds** scale and centred at the predictor median. Magenta shows the central 50% of bootstrap curves.

Occurrence and severity have different drivers



Additive contributions to $\log \tau$, centred at the predictor median.

- VPD shows the clearest nonlinear relationship with conditional spread severity.
- Temperature and SPI effects on τ are less pronounced than their effects on occurrence.

Combining occurrence and severity into risk

For $y > u(\mathbf{x})$,

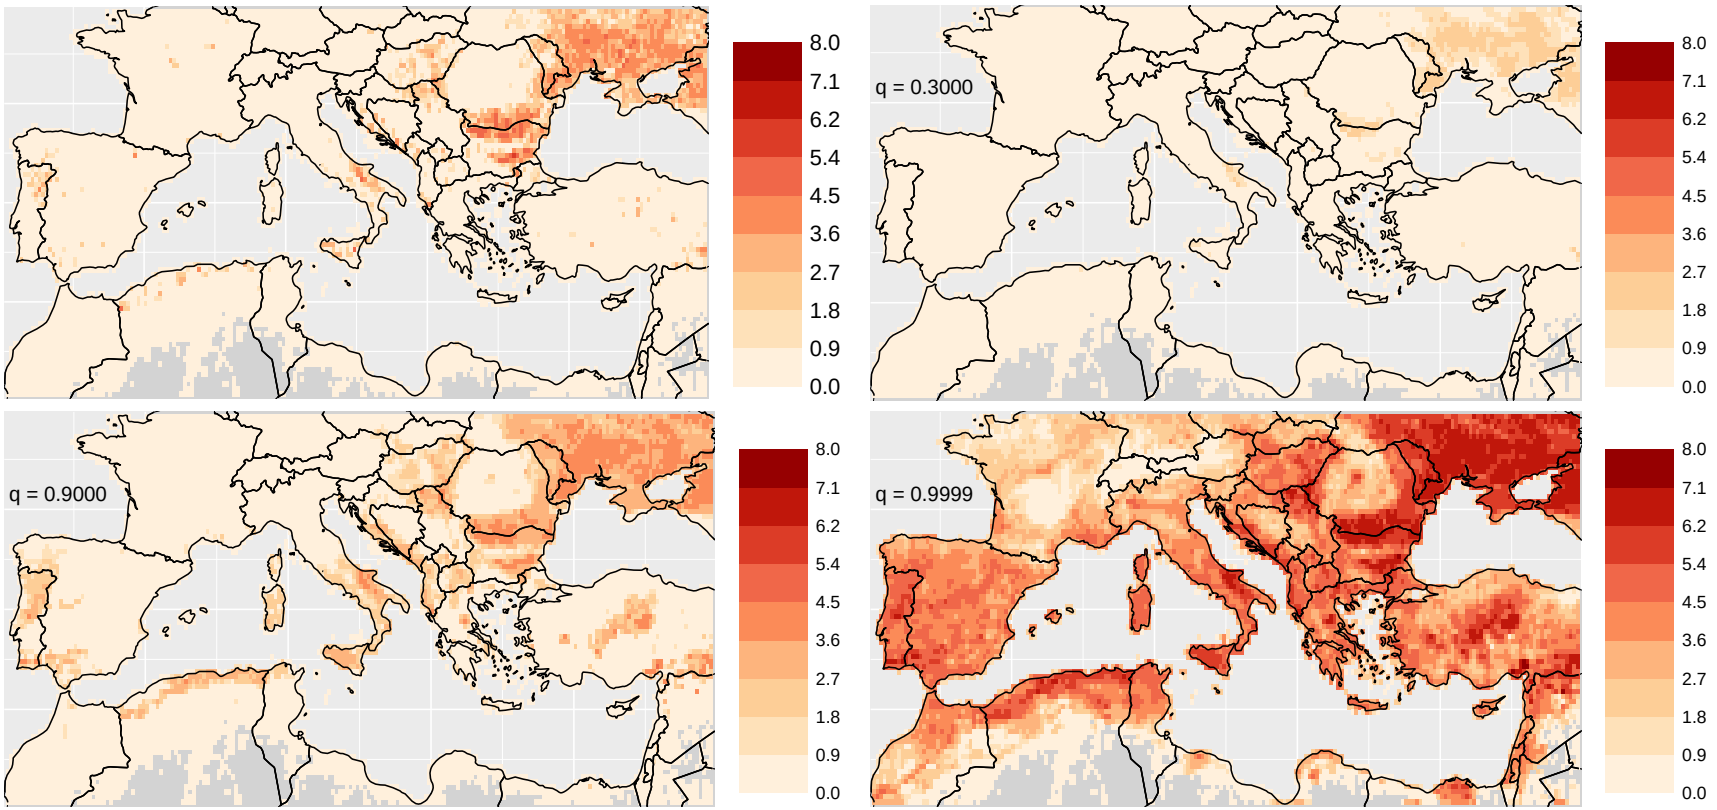
$$\Pr(Y > y \mid \mathbf{x}) \approx \underbrace{p_0(\mathbf{x})}_{\text{occurrence}} \underbrace{p_u(\mathbf{x})}_{\text{threshold exceedance}} \underbrace{\left[1 + \xi \frac{y - u(\mathbf{x})}{\tau(\mathbf{x})}\right]^{-1/\xi}}_{\text{excess magnitude}}.$$

- A large positive-burnt-area quantile describes severity **conditional on occurrence**.
- A quantile of the full distribution of Y combines occurrence probability and conditional severity.
- Locations with potentially severe fires need not have the highest overall burnt-area quantiles.

An example

If $p_0(\mathbf{x}) \leq 0.05$, the full-distribution 95% quantile is zero, even if positive burnt areas can be large.

Mediterranean Wildfire Extremes



Observed $\log(1 + \text{BA})$ (top-left) and extreme q-quantiles for August 2001.

A temperature-trend perturbation experiment

Question

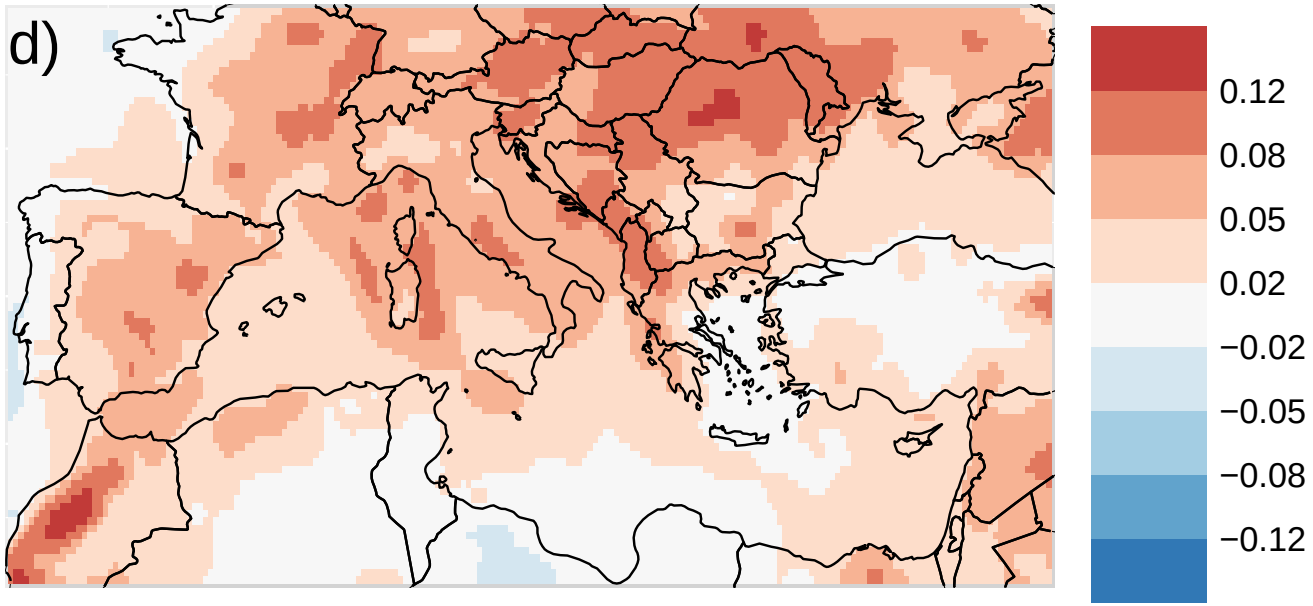
How would estimated wildfire risk in August 2001 change if temperature experienced its fitted 2001–2020 trend?

1. Estimate a local linear trend in August temperature using ERA5 data from 2001–2020.
2. Add the fitted 2001–2020 temperature change to the observed August 2001 temperature field.
3. Re-evaluate the fitted wildfire model, retaining the other covariates and the fitted p_u.

Compare changes in

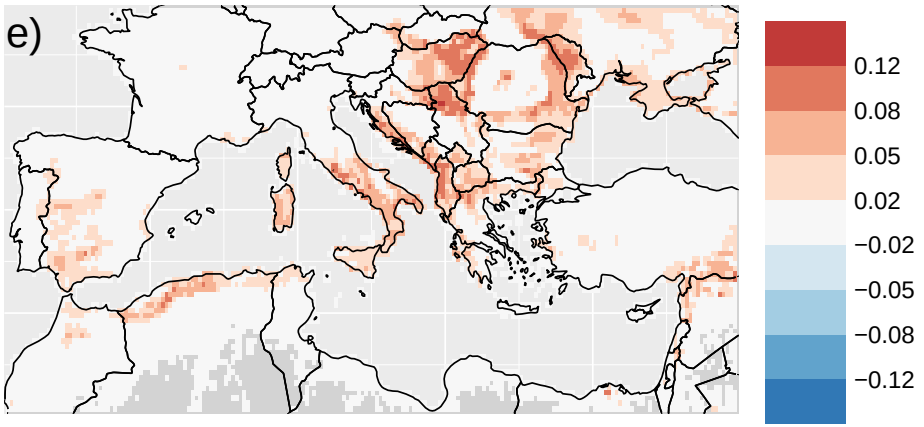
$$p_0(\mathbf{x}) \qquad \text{and} \qquad q^{+}_{0.9}(\mathbf{x}) = Q_{0.9}\{Y \mid Y > 0, \mathbf{x}\}.$$

Mediterranean Wildfire Extremes

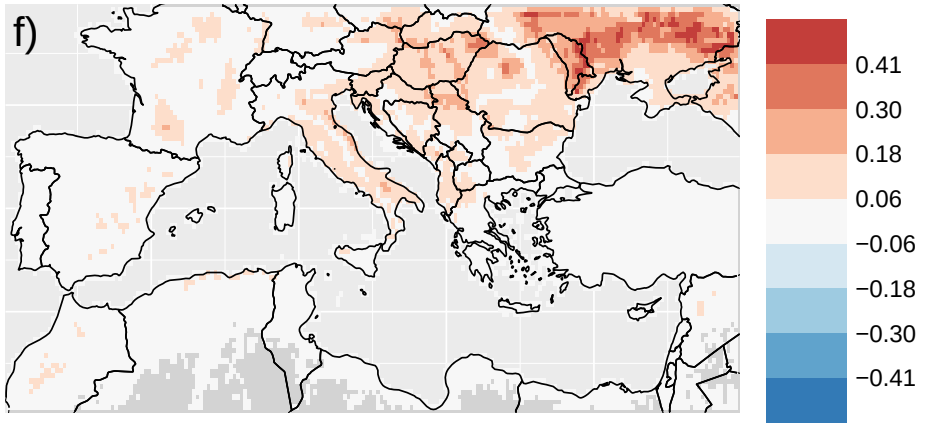


Estimated trends in August 2m air temperature (K).

Temperature trends affect occurrence and severity differently



Change in p_0



Change in $\log\{1 + q^{+}_{0.9}\}$

- Median relative site-wise change in occurrence probability: $+17.1\%$.
- Median relative site-wise change in conditional 90% burnt-area quantile: $+1.6\%$.

Maps show bootstrap-median changes under the temperature perturbation. Effects vary spatially.

Extremal regression for multivariate extremes

Wildfires: the tail depends on environmental conditions

For burnt area Y and environmental covariates $\mathbf{x}$,

$$Y - u(\mathbf{x}) \mid \{Y > u(\mathbf{x}), \mathbf{X} = \mathbf{x}\} \sim \mathrm{GP}\{\tau(\mathbf{x}), \xi(\mathbf{x})\}.$$

Neural networks describe how tail behaviour varies with the environment.

Metocean storms: the tail depends on the storm configuration

For a vector $\mathbf{Z} \in \mathbb{R}^d$ of wind and wave variables, write

$$R = \|\mathbf{Z}\|, \qquad \mathbf{W} = \mathbf{Z}/R.$$

R measures overall magnitude; $\mathbf{W}$ describes the direction.

$$R - u(\mathbf{w}) \mid \{R > u(\mathbf{w}), \mathbf{W} = \mathbf{w}\} \sim \mathrm{GP}\{\tau(\mathbf{w}), \xi(\mathbf{w})\}.$$

The same extremal regression framework, now with angle as the covariate.

Metocean multivariate extremes

- 31-year hindcast of hourly values for a site in Celtic sea ($n \approx 270000$).
- $d = 5$, x- and y-components of wind speed (U) and wave height (H), and (log)-wave period (T_m).

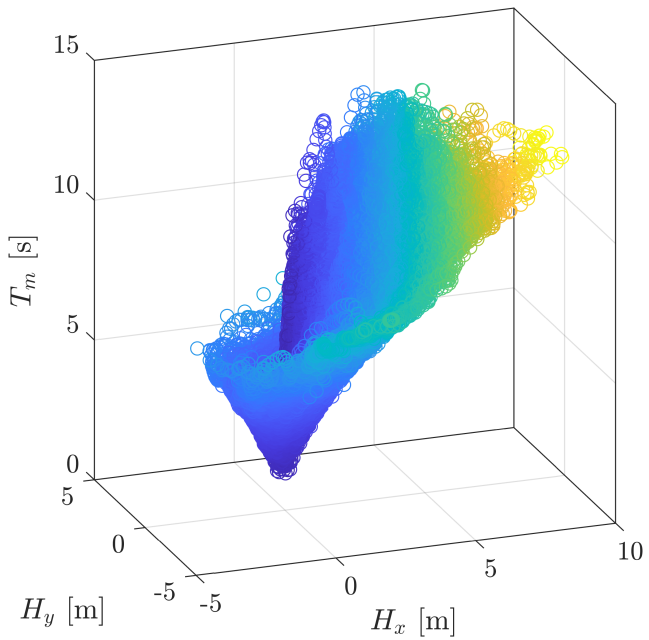
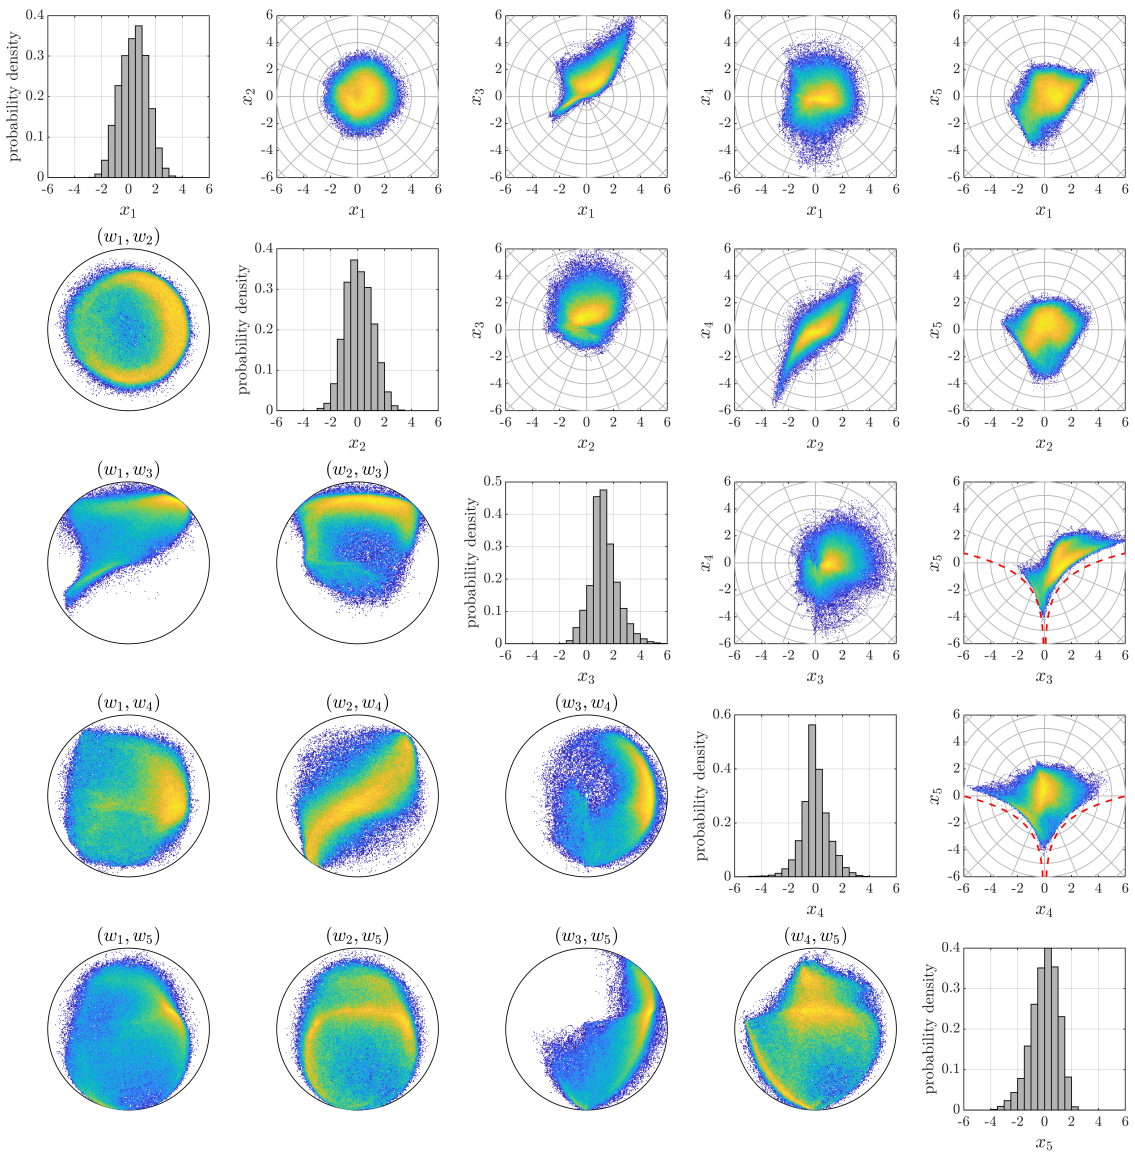


Figure: x-y components of H_s and T_m. Colours indicate the value of $H_s = \sqrt{H_x^2 + H_y^2}$ (low, high), which relates to wave steepness $s = 2\pi H_s/(gT_m^2)$.

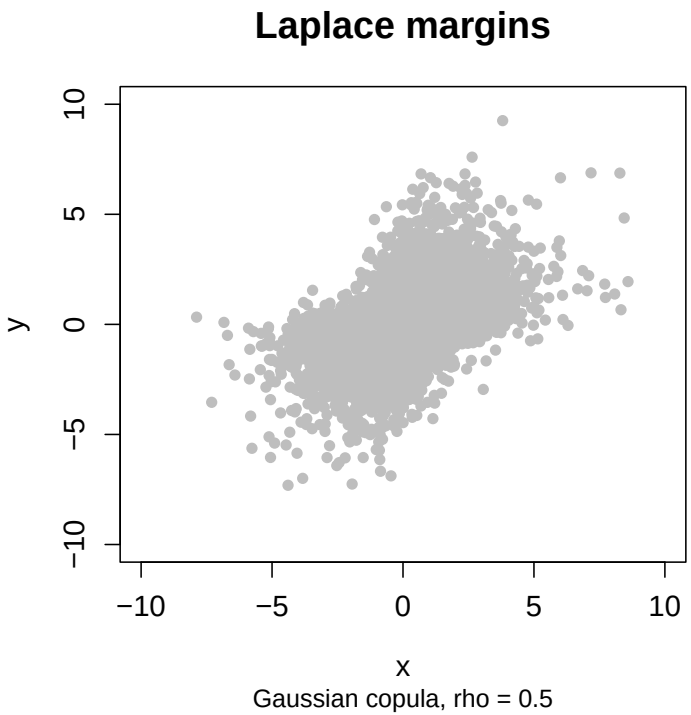
Metocean multivariate extremes

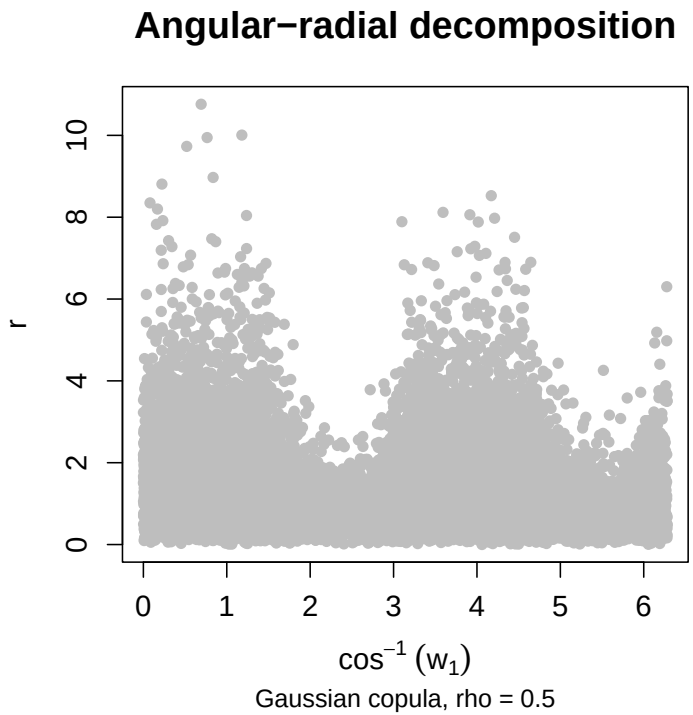


Angular-radial representation

Let $\mathbf{Z} \in \mathbb{R}^d$ and

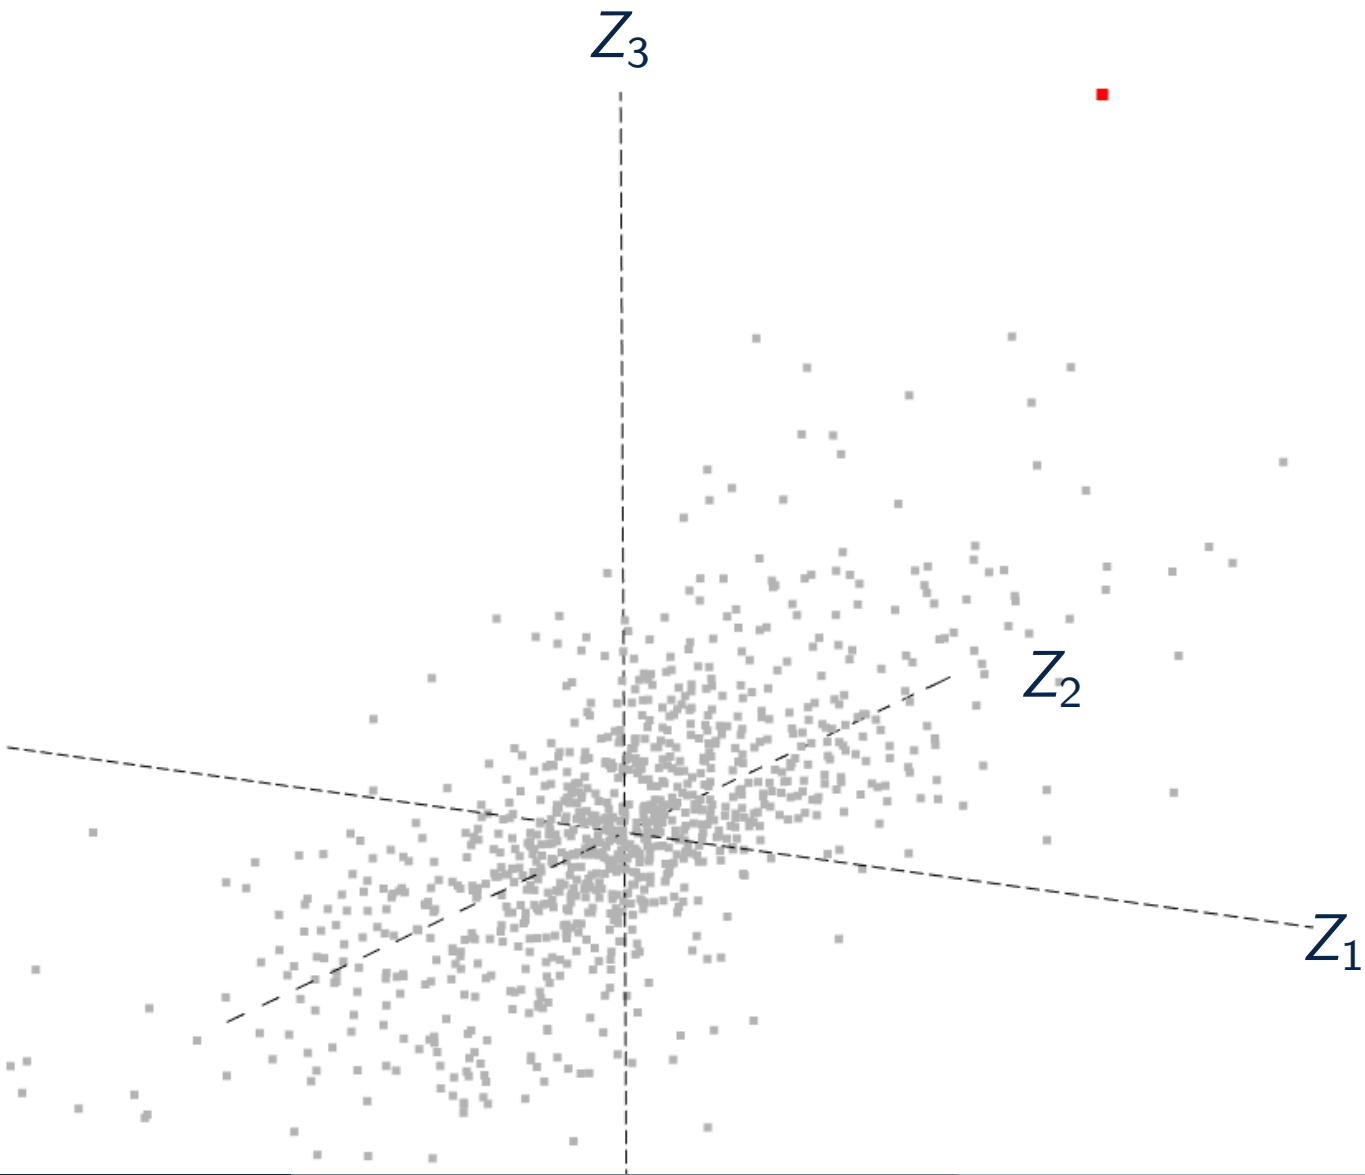
$$R := \|\mathbf{Z}\| > 0, \qquad \mathbf{W} := \frac{\mathbf{Z}}{\|\mathbf{Z}\|} \in \mathcal{S}^{d-1}.$$



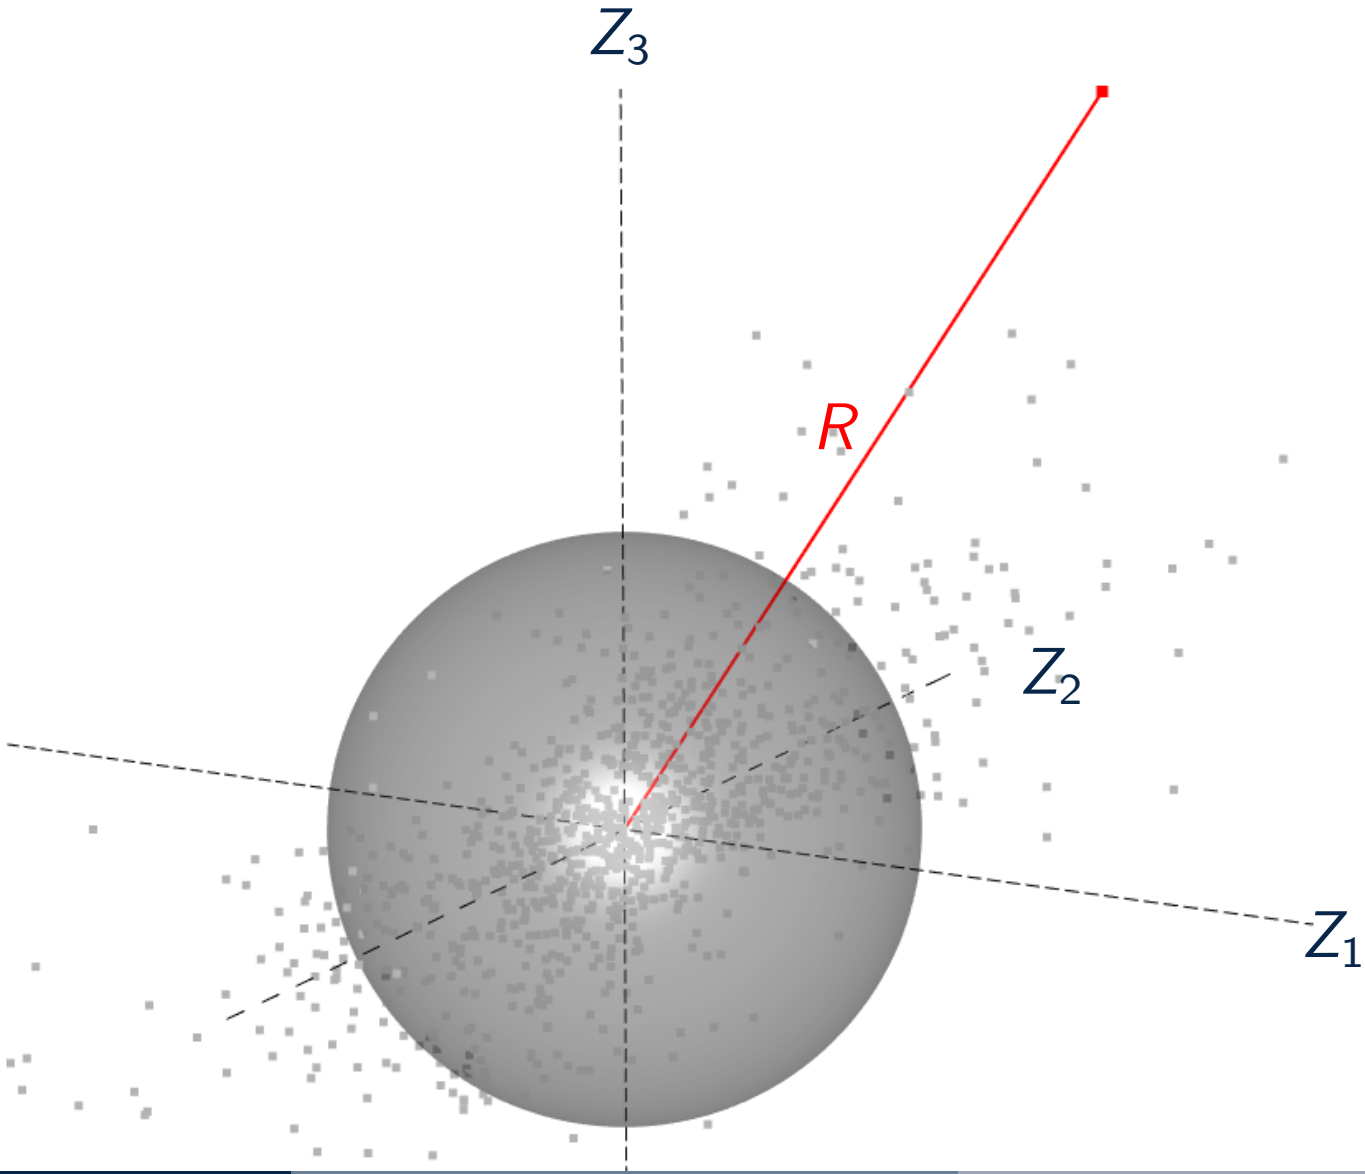


Murphy-Barltrop, C. J., Majumder, R., & Richards, J. (2024). Deep learning of multivariate extremes via a geometric representation. arXiv:2406.19936.

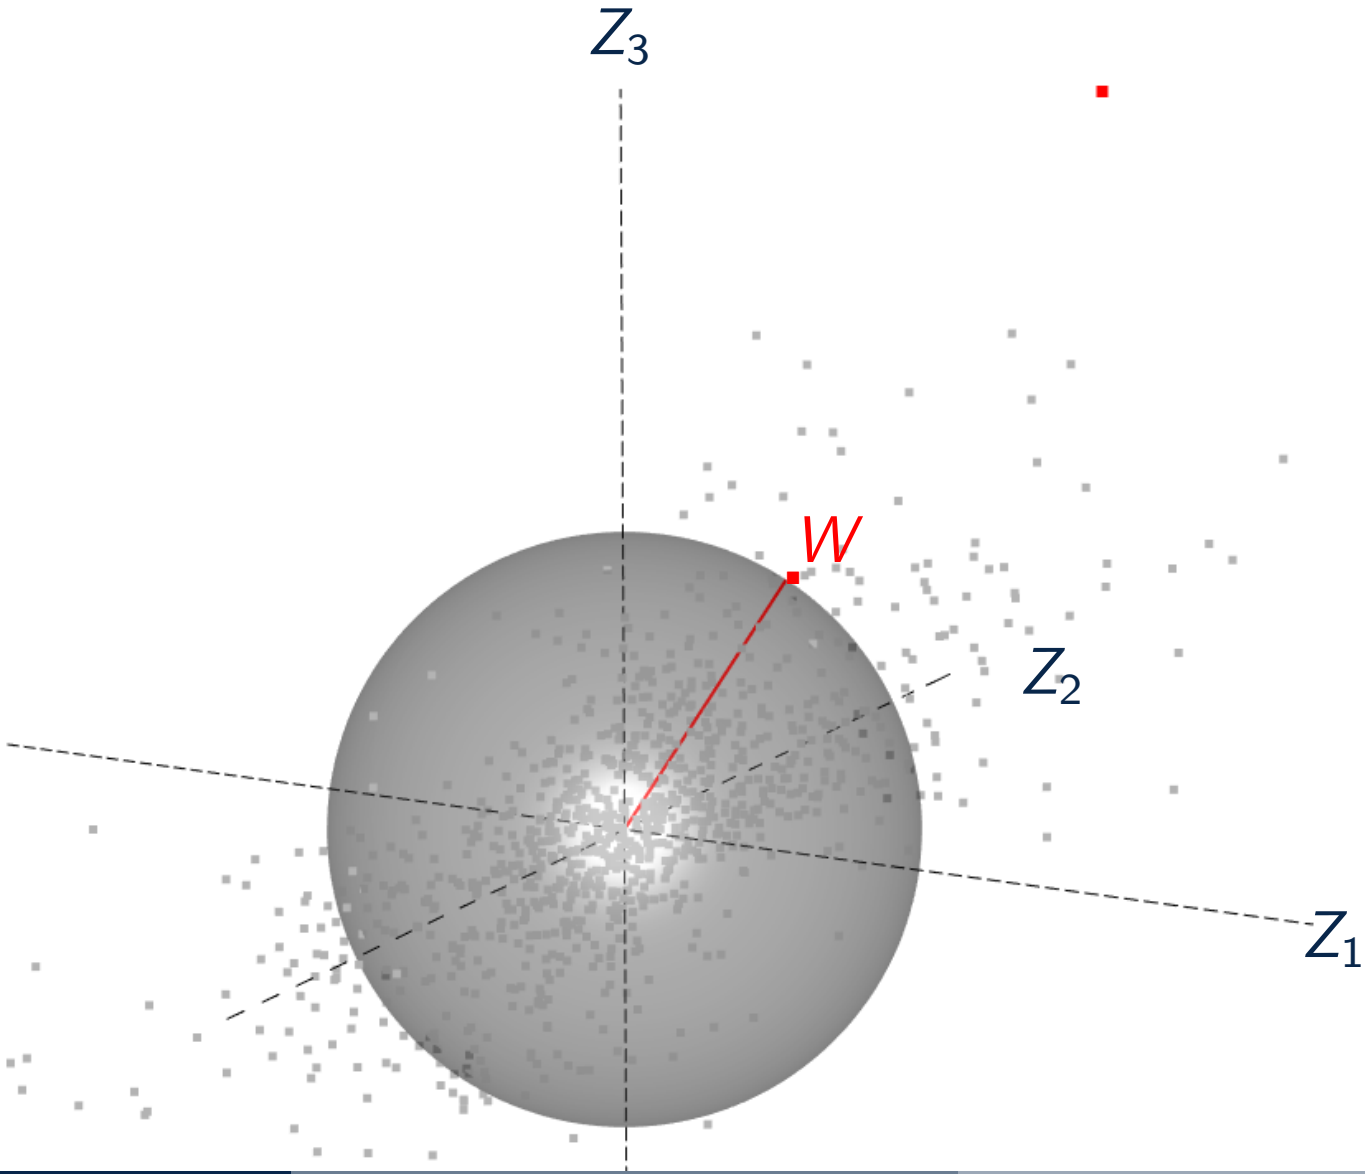
Angular-radial representations



Angular-radial representations



Angular-radial representations



Angular-radial representations

The joint density of $(R, \mathbf{W})$ can be written as

$$f_{R,\mathbf{W}}(r, \mathbf{w}) = f_{\mathbf{W}}(\mathbf{w}) f_{R|\mathbf{W}}(r|\mathbf{w}). \tag{1}$$

The Semi-Parametric Angular-Radial model (SPAR; Mackay and Jonathan, 2023) replaces $f_{R|\mathbf{W}}(r|\mathbf{w})$ with a conditional GP model **when r is big**, i.e.,

$$f_{R|\mathbf{W}}(r|\mathbf{w}) \approx \zeta f_{\mathrm{GP}}(r - u(\mathbf{w}); \tau(\mathbf{w}), \xi(\mathbf{w})), \tag{2}$$

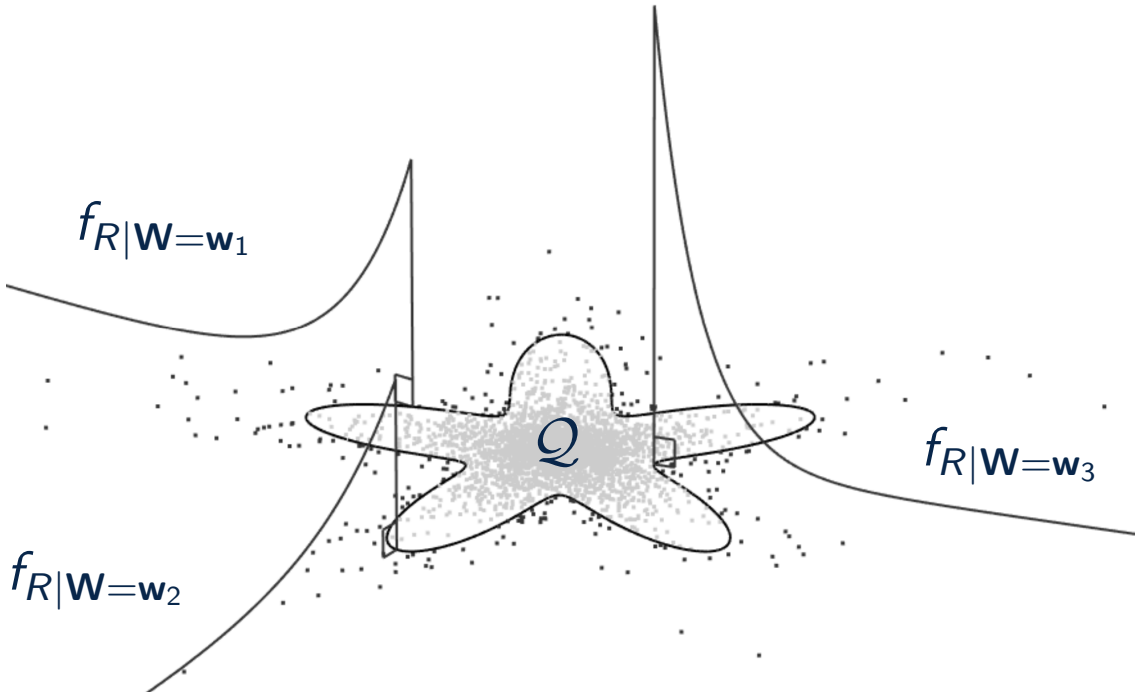
for large $r > u(\mathbf{w})$ and where $\zeta := \Pr(R > u(\mathbf{w})|\mathbf{W} = \mathbf{w})$ is fixed.

- Mackay et al. (2025) model $(u(\mathbf{w}), \tau(\mathbf{w}), \xi(\mathbf{w}))$ via **deep extremal regression**.

Mackay, Jonathan (2023). Modelling multivariate extremes through angular-radial decomposition of the density function. arXiv.

Mackay, Murphy-Barltrop, Richards, Jonathan (2025). Deep Learning Joint Extremes of Metocean Variables Using the SPAR Model. JOMAE.

Angular-radial representations



De Monte, L., Huser, R., Papastathopoulos, I., & Richards, J. (2025). Generative modelling of multivariate geometric extremes using normalising flows. arXiv:2505.02957.

Generating compound extremes

For a fixed conditional exceedance probability ζ:

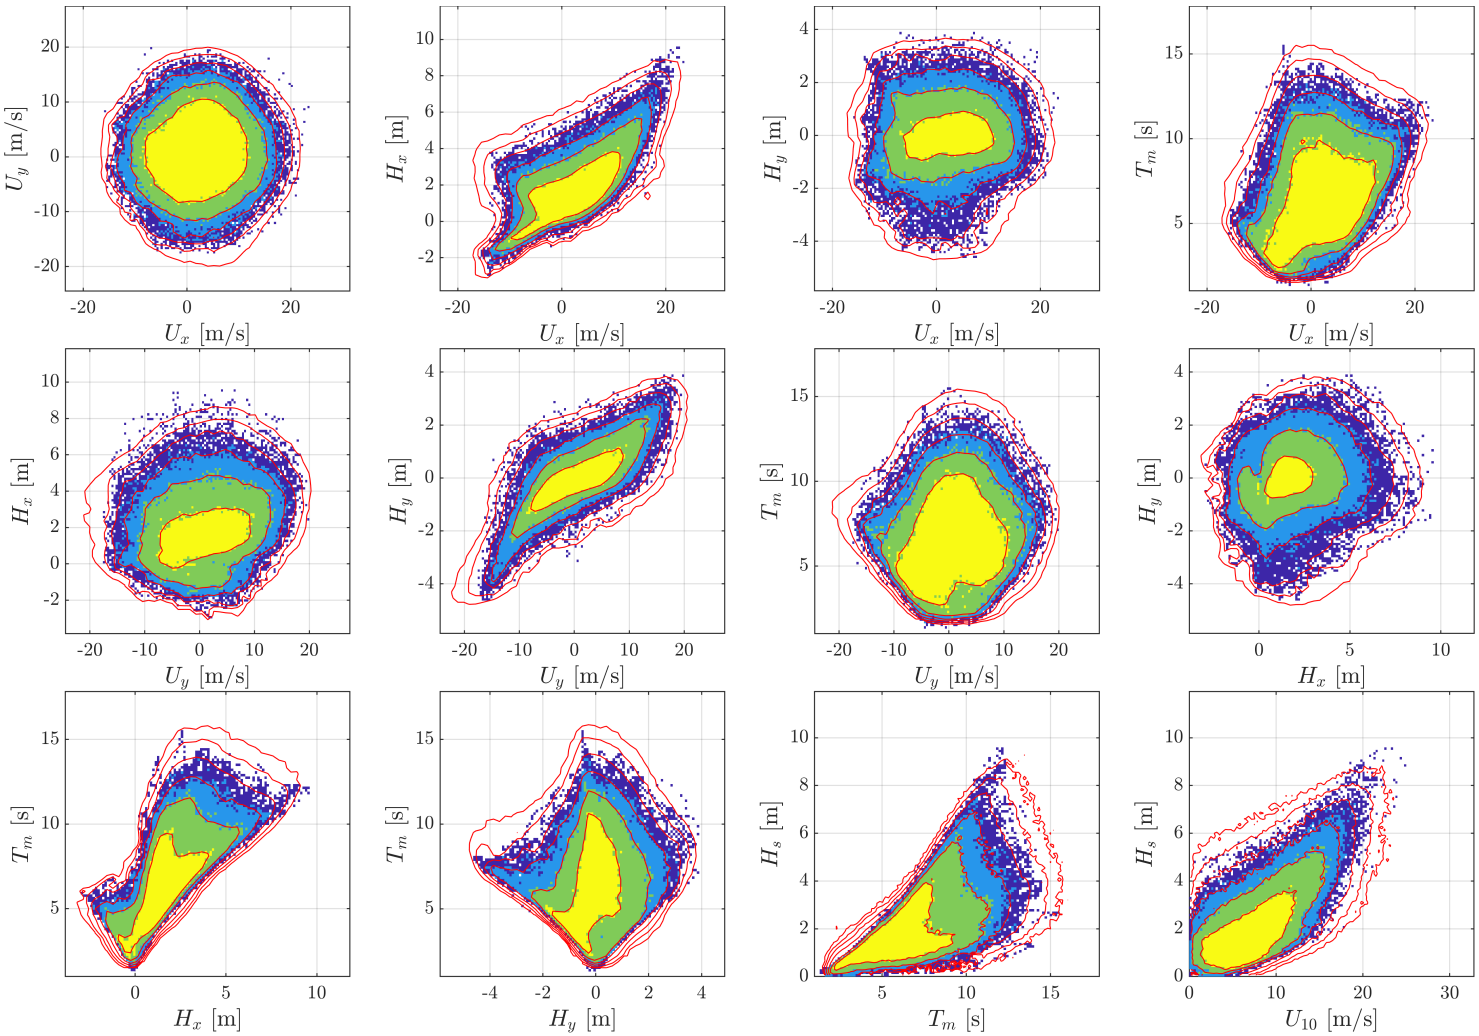
1. Draw an angle $\mathbf{W}$ from its fitted distribution.
2. Draw $E \sim \mathrm{GP}\{\tau(\mathbf{W}), \xi(\mathbf{W})\}$.
3. Set $R = u(\mathbf{W}) + E$, then $\mathbf{Z} = R\mathbf{W}$.
4. Undo preprocessing to recover physical units.

These draws represent the **joint tail**, with probability mass ζ.

Probabilities of an event A in the joint tail

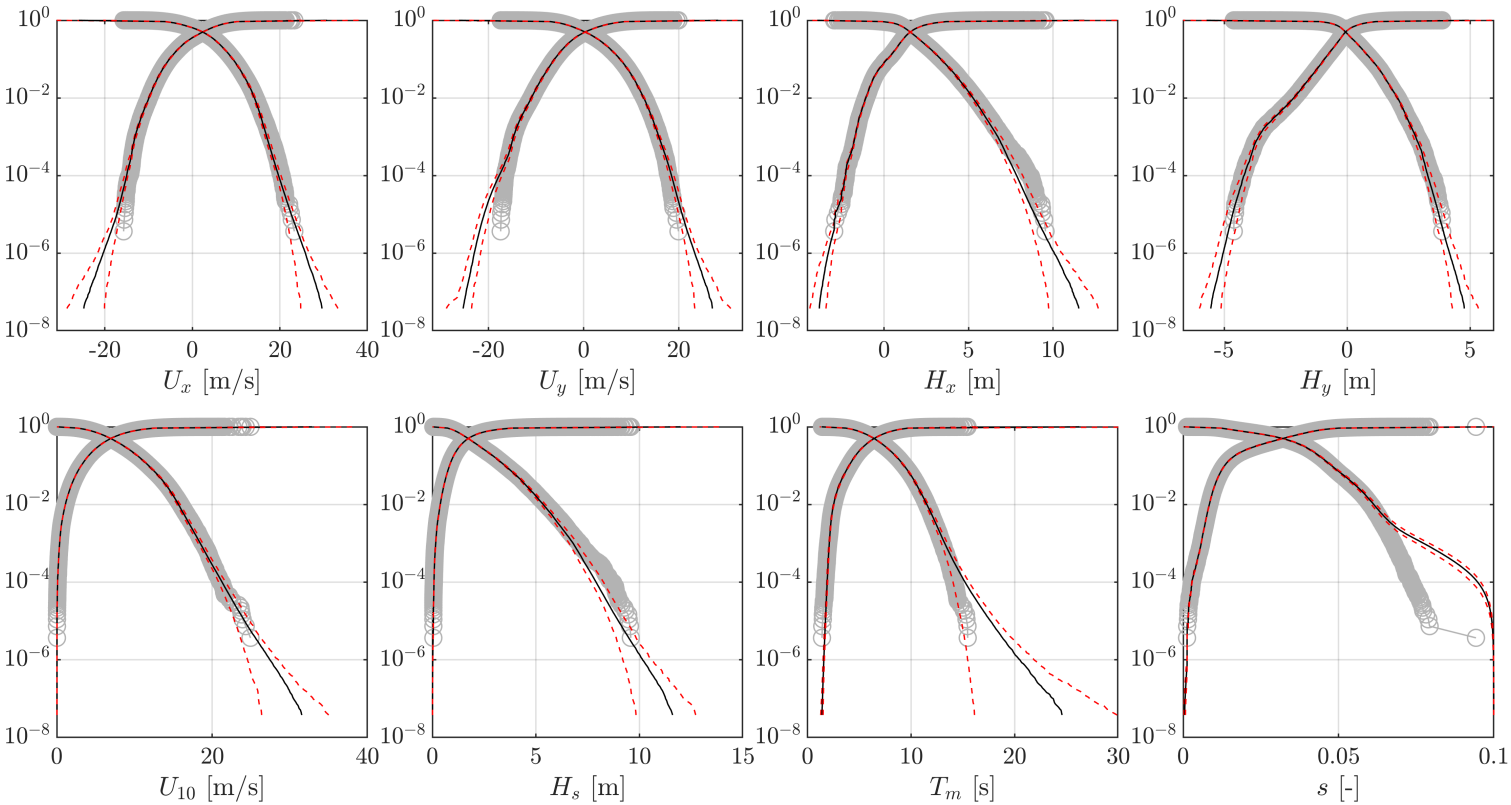
$$\Pr(\mathbf{Z} \in A) \approx \zeta N^{-1} \sum_{j=1}^{N} \mathbb{1}\{\mathbf{Z}^{(j)} \in A\}.$$

Estimates



Contours produced from deepSPAR model using 3000 years of reps.

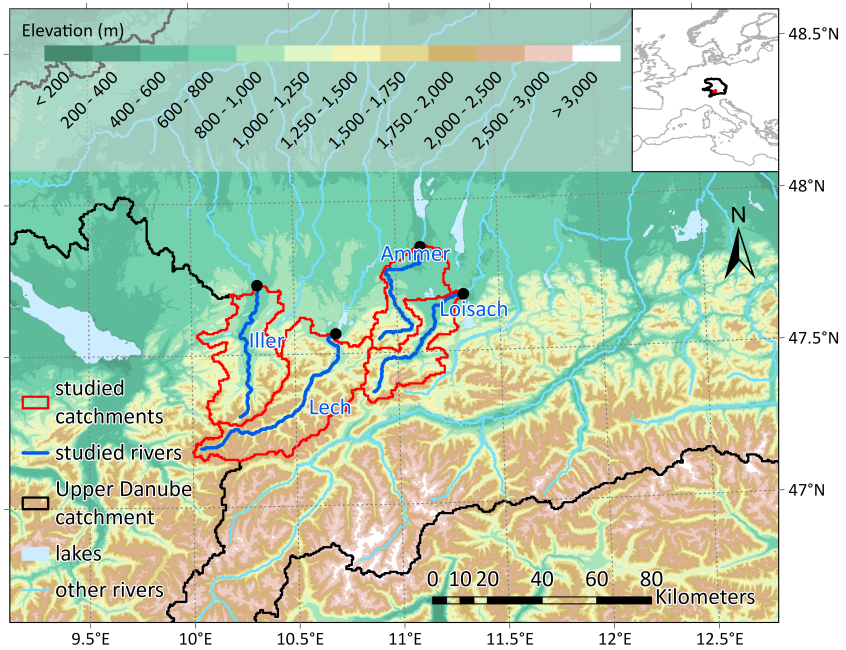
Marginal diagnostics



Wind speed: $U_{10} = \sqrt{U_x^2 + U_y^2}$. Steepness: $s = 2\pi H_s/(gT_m^2)$.

How does concurrent discharge risk change?

- Four neighbouring head catchments: Ammer, Iller, Lech and Loisach.
- Shared weather can produce extremes across several catchments, despite no flow routing between the gauged basins.
- 50 climate–hydrology ensemble members under RCP8.5, using CanESM2 / CRCM5 / WaSiM.



Scientific question

How do changes in individual catchment extremes translate into changes in aggregate and concurrent extremes?

A joint tail model for each climate window

1980–2009 2010–2039
2040–2069 2070–2099

- Fit separate four-dimensional models to **weekly maxima** and **monthly minima** in each window: eight SPAR fits.
- After centring and transforming discharge, use angle as the regression covariate:

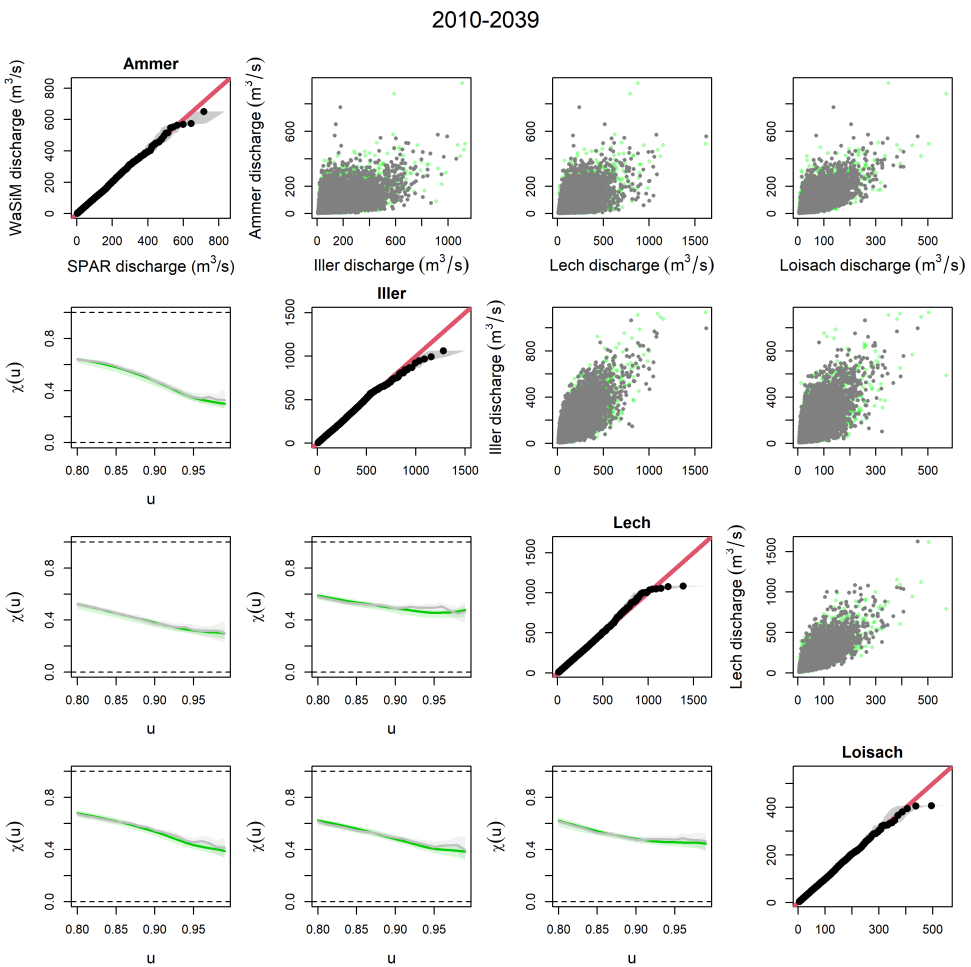
$$R - u(\mathbf{w}) \mid \{R > u(\mathbf{w}), \mathbf{W} = \mathbf{w}\} \sim \mathrm{GP}\{\tau(\mathbf{w}), \xi(\mathbf{w})\}.$$

- Neural networks estimate u, τ and ξ. Resampled angles and simulated radii generate joint extremes in physical discharge units.

Comparing climates

Estimate the same types of event probability under each fitted joint distribution.

Does SPAR reproduce the joint extremes?



Weekly maxima, 2010–2039.

Diagonal
Marginal QQ plots: extreme magnitudes.

Above diagonal
Pairwise scatter: joint configurations.

Below diagonal
$\chi(v)$: co-exceedance at high marginal quantiles.

Generally good agreement with the hydrological ensemble, with some discrepancies at the most extreme marginal quantiles.

Aggregate magnitude and concurrent rarity

Let Z_i be discharge at catchment i, summarised within a week or month, and let $S = \sum_{i=1}^{4} Z_i$.

Aggregate events: fixed physical thresholds

$$\Pr(S > 1500\,\mathrm{m}^3/\mathrm{s}) \quad \text{or} \quad \Pr(S < 30\,\mathrm{m}^3/\mathrm{s}).$$

Changes reflect both marginal distributions and their dependence.

Concurrent events: matched marginal rarity

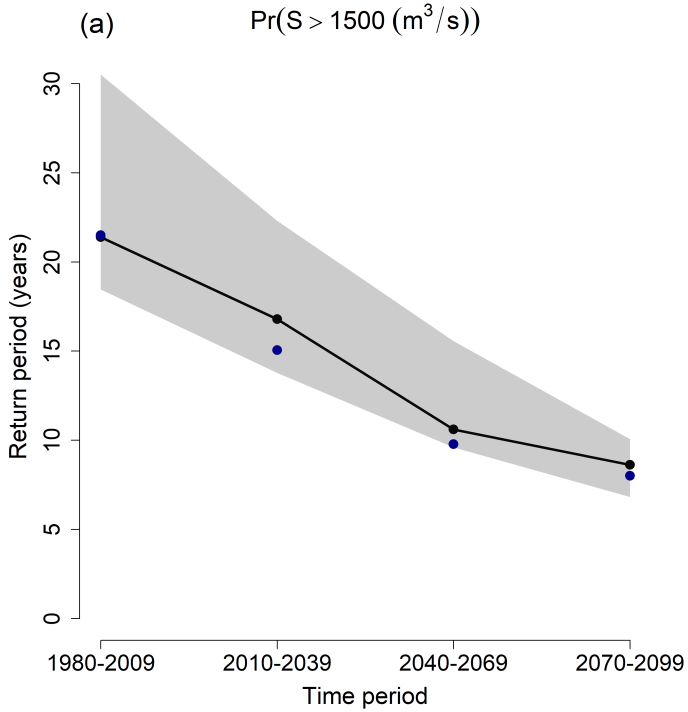
$$\Pr\{Z_i > q^H_{i,k}\ \forall i\} \quad \text{or} \quad \Pr\{Z_i < q^L_{i,k}\ \forall i\}.$$

$q^H_{i,k}$ and $q^L_{i,k}$ are marginal 10-year return levels in climate window k.

Concurrent-event comparisons therefore focus on changes in **extremal dependence**.

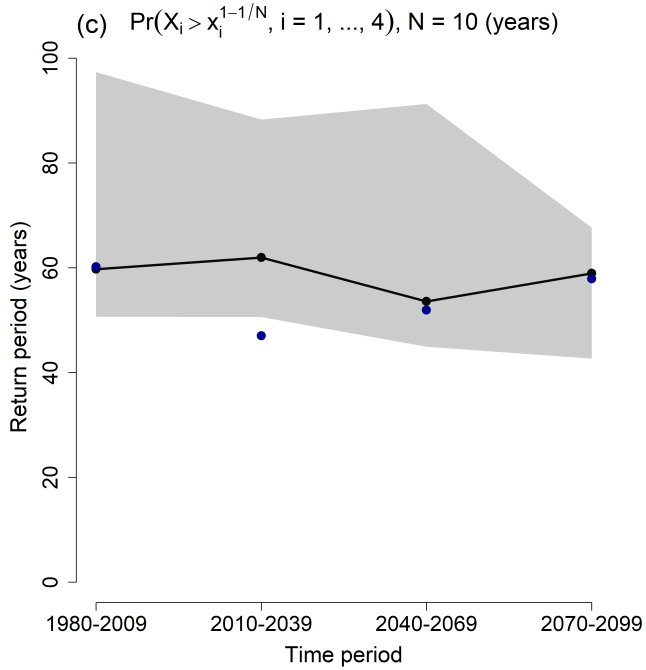
Aggregate high flows increase; concurrence changes little

Fixed aggregate threshold



$S > 1500\,\mathrm{m}^3/\mathrm{s}$: shorter return periods.

Matched marginal rarity



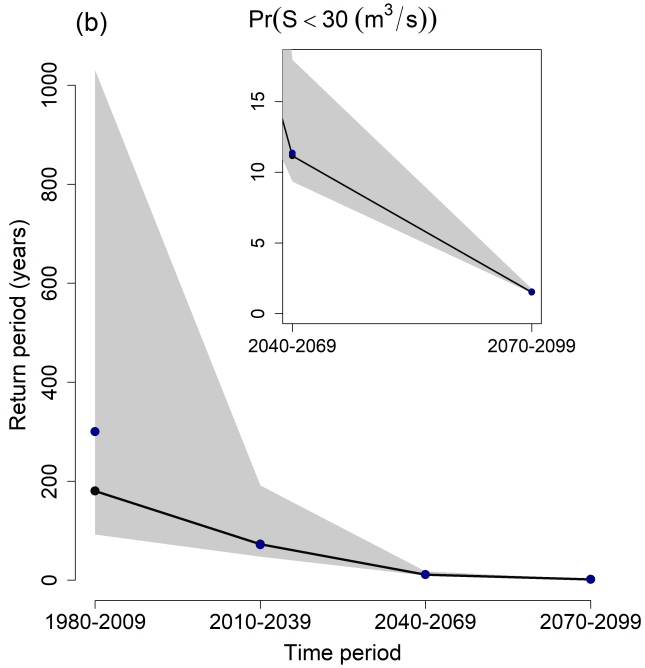
All sites above their window-specific 10-year levels: little systematic change.

Interpretation: greater aggregate high-flow risk can occur without a comparable increase in all-site upper-tail dependence.

Black: SPAR. Blue: empirical hydrological ensemble. Shorter return periods mean more frequent events.

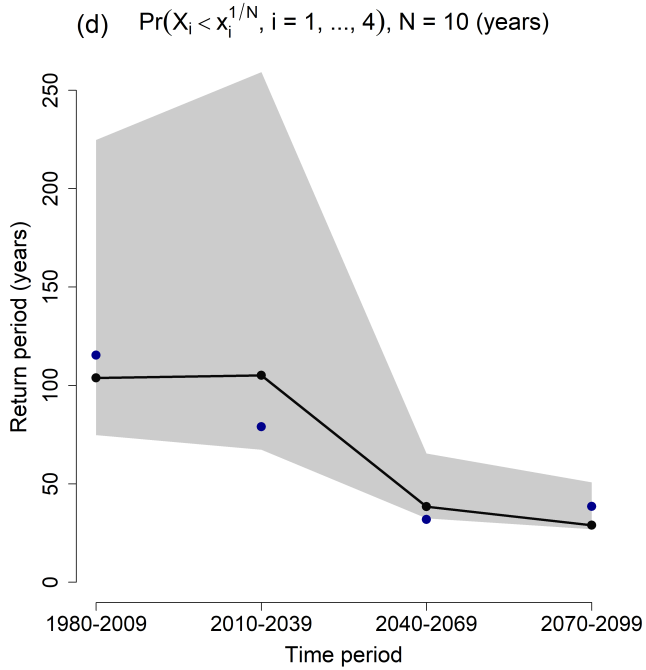
Low flows become more severe and more concurrent

Fixed aggregate threshold



$S < 30\,\mathrm{m}^3/\mathrm{s}$.

Matched marginal rarity



All sites below their window-specific 10-year low-flow levels.

- Aggregate-event return period: approximately **180 years to 1.5 years** from the first to the last window.
- Concurrent low-flow events also become more likely after accounting for changing marginal rarity.

What deep learning contributes to the risk assessment

- **Flexible joint tails:** networks learn how the radial threshold, scale and shape vary across discharge configurations.
- **Interpretable outputs:** generated events give probabilities for aggregate stress and simultaneous catchment extremes.
- **Different climate responses:** high-flow and low-flow concurrence need not evolve in the same way.

Scope of the conclusions

These are projections conditional on one model chain and RCP8.5. Bootstrap intervals quantify sampling and fitting variability, not the full uncertainty across climate models or scenarios.

How computationally demanding are these models?

Application	Available observations	Network size
Wildfire	1.21 million grid-cell months; 102 240 positive values	Four networks: 2–5 layers, with 4–16 filters/nodes per layer, plus interpretable splines.
Metocean	271 704 hourly vectors; $d = 5$	Two MLPs: 3 hidden layers of 16 neurons; approximately 650 parameters each.
Alpine discharge	Per climate window: 78 250 weekly-maxima or 18 000 monthly-minima vectors; $d = 4$	Two MLPs per fit: 3 hidden layers of 32 neurons.

Take-home messages

- Extreme value theory gives deep learning an explicit model for tail extrapolation.
- Deep extremal regression connects environmental covariates to extreme quantiles and tail parameters.
- Interpretable components help describe drivers of fitted risk.
- Angular–radial models extend the approach to generative compound extremes.

Tail diagnostics and uncertainty remain central to scientific interpretation. Mackay et al. (2026)!

Mackay, Richards, Jonathan (2026+). Diagnostic Tools for Extreme Value Regression Models. arXiv.

Discussion and resources

Jordan Richards
`https://jbrich95.github.io/`
jordan.richards@ed.ac.uk

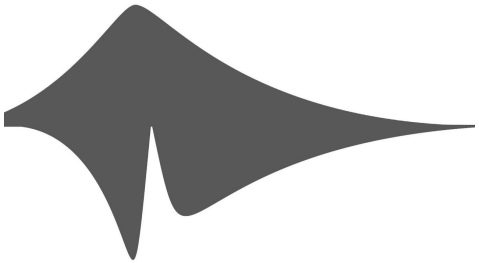
- Extreme quantile regression: Jbrich95/extQuantRegressDL
- Partially interpretable networks: Jbrich95/pinnEV
- Deep SPAR: callumbarltrop/DeepSPAR

Other uses? Neural inference.

Which extreme-event probability matters for your application?

Shameless advertisements







Bridging Heavy Tails and Artificial Intelligence

Participating journal: **Extremes**

The increasing frequency and severity of extreme events—such as catastrophic flooding, record-breaking temperatures, and unprecedented heatwaves—has highlighted the urgent need for innovative approaches in risk assessment and modeling. Modern advancements in data-collection techniques have provided increasingly large and complex datasets, which can only be processed used fast and scalable algorithms and computational software. This special issue aims to bridge the gap between Artificial Intelligence (AI) and Extreme Value Theory (EVT) to harness the strengths of both fields and address the growing challenges posed by these extreme events.



Journal

Extremes

Extremes is a dedicated platform for the publication of original research in statistical extreme value theory and its applications across various fields.

Publishing model	**Hybrid**
Journal Impact Factor	**1.1 (2023)**
Downloads	**58k (2024)**
Submission to first decision (median)	**7 days**

GAMEX

Generative AI Modeling for Extreme Events: Statistical Foundations, or simply **GAMEX**, is a **Una Europa seed-funded initiative (2026)** that brings together researchers in statistics, machine learning, artificial intelligence and data science to build an international research network at the interface of **Generative Artificial Intelligence (AI)** and **Extreme Value Theory (EVT)**, with the aim of developing reliable, interpretable, and scalable methods for modeling rare, high-impact events.

Key activities in 2026

- **Scientific Workshop**
  University of Bologna, 11–12 June 2026
- **Edinburgh Summer School on Generative AI for Extremes**
  University of Edinburgh, 8–11 September 2026
- **Online Seminar Series**
  Regular talks by GAMEX partners and invited external speakers
- **Public Lectures and Visits**
  Philippe Naveau public lecture and visit at the University of Edinburgh
- **Short-term Research Visits**
  Mobility opportunities for PhD students and early-career researchers (ECRs)

https://hecsciencesdecision.github.io/eva2027/datachallenge.html

References I

Carreau, J. and Bengio, Y. (2007). A hybrid Pareto model for conditional density estimation of asymmetric fat-tail data. In *Artificial Intelligence and Statistics*, pages 51–58. PMLR.

Lizundia-Loiola, J., Otón, G., Ramo, R., and Chuvieco, E. (2020). A spatio-temporal active-fire clustering approach for global burned area mapping at 250 m from modis data. *Remote Sensing of Environment*, 236:111493.

Mackay, E. and Jonathan, P. (2023). Modelling multivariate extremes through angular-radial decomposition of the density function. *arXiv preprint arXiv:2310.12711*.

Mackay, E., Murphy-Barltrop, C. J. R., Richards, J., and Jonathan, P. (2025). Deep Learning Joint Extremes of Metocean Variables Using the SPAR Model. *Journal of Offshore Mechanics and Arctic Engineering*, 148(021201).

Mackay, E., Richards, J., and Jonathan, P. (2026). Diagnostic tools for extreme value regression models. *arXiv preprint arXiv:2606.02676*.

Murphy-Barltrop, C., Richards, J., Poschlod, B., Sasse, A., and Zscheischler, J. (2026). Exploring climate change effects on concurrent floods and concurrent droughts via statistical deep learning. *Water Resources Research, to appear*.

Richards, J. and Huser, R. (2026). Extreme Quantile Regression with Deep Learning. In de Carvalho, M., Huser, R., Naveau, P., and Reich, B. J., editors, *Handbook on Statistics of Extremes*. Chapman & Hall/CRC, Boca Raton, FL.

Richards, J., Huser, R., Bevacqua, E., and Zscheischler, J. (2023). Insights into the drivers and spatiotemporal trends of extreme mediterranean wildfires with statistical deep learning. *Artificial Intelligence for the Earth Systems*, 2(4):e220095.

Fin.



Scan for webpage.