# Recalibration of FEH13 rainfall model for Cumbria

## Final report

Gianni Vesuviano and Elizabeth Stewart

Client Ref: ENV6000806R
Date 23/02/2021

| | |
|---|---|
| **Title** | Recalibration of FEH13 rainfall model for Cumbria |
| **Client** | Environment Agency/UKCEH |
| **Client reference** | ENV6000806R |
| **UKCEH reference** | NEC06121/NEC06813/Issue 1.2 |
| **UKCEH contact details** | Elizabeth Stewart<br>t: 01491 692229<br>e: ejs@ceh.ac.uk |
| **Author** | Vesuviano, Gianni & Stewart, Elizabeth |
| **Approved by** | James Miller |
| **Signed** | J Miller |
| **Date** | 02/03/2021 |

# Executive Summary

This report presents the results of a study to re-calibrate the UK's current rainfall depth-duration-frequency model (FEH13) over Cumbria with more recent data that were not included in calibration of the FEH13 model currently available through the FEH Web Service. These new data include the record-breaking November 2009 and December 2015 storms.

The study includes a re-mapping of the median annual rainfall for the whole UK, for different durations from 1 hour to 8 days. It then focuses on the depth-duration-frequency (DDF) relationships estimated at ten case study sites. The steps undertaken to generate the final DDF estimates from the raw annual maximum rainfall data are described in detail. As the FEH13 method was almost unchanged during this study, this report describes in detail the steps undertaken to generate the data available on the FEH Web Service.

The new model is used to estimate spatial return periods for the maximum 36-hour total of the November 2009 event, and point return periods for the November 2009 and December 2015 events for selected locations and durations. The rarest return period within the November 2009 storm's spatial field is found to decrease from almost 8000 to just over 500 years, while return periods at the periphery of the storm are practically unaffected. Point rainfall return periods for the November 2009 event are found to decrease considerably relative to FEH13 but to be similar to those estimated by the original Flood Estimation Handbook model (FEH99). However, the return period of the 24-hour record breaking rainfall at Honister Pass in December 2015 is found to be reduced approximately sevenfold compared to both the FEH13 and FEH99 models.

Finally, model outputs for the whole of Cumbria, consisting of mapped rainfall depth for specified durations and return periods, are presented and compared to estimates from the current FEH13 model, indicating that the vast majority of duration-frequency combinations result in greater modelled depths, or that specified depth-duration combinations are increasing in frequency. However, peak 1-hour rainfalls of specified return periods (here 100 and 1000 years) may be getting smaller in the north of Cumbria.

Although the new estimates for rainfall DDF relationships presented in this report are theoretically less uncertain than those from previous UK-wide rainfall models, including the FEH13, FEH99 and FSR, it is important to note that uncertainty in extreme events is still very high, and can only be lowered through the observation and incorporation of more extreme events from a longer monitoring period.

# Contents

# 1 Introduction

The FEH13 rainfall depth-duration-frequency (DDF) model is the current UK method for relating the rarity and severity of rainfalls of different durations. It supersedes all previous UK-wide DDF models for all purposes except estimation of PMP (probable maximum precipitation). The calibration data set for the FEH13 model includes events from approximately 6500 daily recording gauges and 1000 hourly recording gauges, most of which start after 1961 and none of which ends after 2006.

Since 2006, new UK records for storm depths of durations from 24 hours to 4 days have been set in Cumbria: first in November 2009 (24 hours, 3 and 4 days) and December 2015 (24 hours and 2 days). Because of their recency, they were not included in the FEH13 model calibration data set, and the current FEH13 model is unable to consider them when generating rainfall estimates. There is also some concern that the majority of calibration data in the FEH13 (as well as the previous FEH99 and FSR models) represent a climate in which winter rainfall was less intense than it is now (e.g. Watts *et al.*, 2015) and therefore that more calibration data from the recent past are required to represent the current climate. Given evidence that larger rainfall events may become more common in the future (Osborn & Maraun, 2008), it was deemed a priority to understand how the occurrence of several major recent flooding events in Cumbria may change the expected probabilities of similar events occurring in the future.

In this study, we re-calibrate the FEH13 model to an increased data set of daily and hourly annual maxima, ending in 2016. The method is almost unchanged from that used to generate rainfall DDF estimates for the FEH Web Service, hence this report can also be read as a description of the model underlying the FEH13 data available on the FEH Web Service.

The FEH13 method assumes stationarity in the input data and a deterministic relationship between depth, duration and frequency at every location. This recalibration does not modify the underlying FEH13 method, so it does not allow estimation of uncertainty and cannot account for any potential effects of climate change. Additionally, as the original FEH13 method is not used for PMP estimation, this study does not attempt to supplant the current method of PMP estimation.

# 2 Background to FEH13 model

The FEH13 rainfall model is the Flood Estimation Handbook's latest UK-wide statistical model for rainfall depth-duration-frequency (DDF) estimation. The model is based on an analysis of over 170,000 station-years of data from daily rain gauges throughout the UK together with about 17,000 station-years of hourly data. It was developed to allow the estimation of rainfall depths falling over durations from 1 hour to 192 hours (8 days) for return periods ranging from 2 years to over 10,000 years. The FEH13 model is delivered via the FEH Web Service (https://fehweb.ceh.ac.uk) and has been extrapolated to provide depth estimates for durations as short as 5 minutes. It is widely used for applications such as drainage design, flood risk assessment and reservoir flood safety appraisal throughout the UK.

The key components of the FEH13 rainfall model were developed during the joint Defra/Environment Agency research project 'Reservoir Safety – Long Return Period Rainfall (FD2613)' (Stewart *et al.*, 2013). The main aim of the project was to replace the original FEH rainfall DDF model (Faulkner, 1999), now referred to as the FEH99 model. It was commissioned in response to concerns about the apparently high estimates produced by FEH99 when applied to return periods in excess of its recommended upper limit of 1,000 years. Due to the assumed straight line log-log relationship between rainfall depth and return period in the FEH99 model, it is not uncommon for the model to estimate a return period of less than 10,000 years for a PMP (probable maximum precipitation) event (Babtie Group *et al.*, 2000; MacDonald & Scott, 2001). This caused difficulties in the practical assessment of reservoir flood safety in the UK, led Defra to recommend the use of the earlier FSR rainfall model (NERC, 1975) for 10,000-year rainfall estimates – a model based largely on rain gauge records from the period 1961-1970 – pending further research (Defra, 2004). An independent assessment of FEH99 by Cox (2003) formed the basis for development of an updated FEH DDF model via the Reservoir Safety – Long Return Period Rainfall (FD2613) research project.

The FD2613 research produced a revised FEH rainfall DDF model based on rainfall durations from 1 hour to 8 days applicable to return periods from 2 to over 10,000 years. The revised model retained the basic *index-flood* approach of the FEH99 model but its main advances were:

- increased availability of rainfall maxima, particularly for sub-daily durations
- a revised standardisation that uses standard-period average annual rainfall (*SAAR*) and northing in addition to the index variable *RMED* (the median annual maximum rainfall) to remove more of the location-dependent variation in rainfall before combining maxima from networks of rain gauges
- a revised spatial dependence model
- improvements to the FEH FORGEX method of deriving growth curves
- a more flexible depth-duration-frequency (DDF) model structure

The FD2613 project provided indicative results from the new model for 71 sites in the UK but did not go as far as to generalise the model to allow its application at any point of interest. Generalisation to the whole UK was carried out during a follow-on project funded by CEH, which reviewed the spatial consistency of the model results for different durations and return periods, made minor modifications to the model, and finally generated the results that are now available as FEH13 outputs through the FEH Web Service. Rainfall estimates for sub-hourly durations of *x* minutes are still partly based on FSR procedures, as they are found by multiplying the ratio between the *x*-minute and 60-minute FSR rainfall depths by the 60-minute FEH13 rainfall depth.

There are six stages to the FEH13 procedure:

- Abstraction of annual maxima (AMAX) of different durations from continuous hourly and daily rain gauge data;
- Estimation of UK-wide *RMED* (equivalent to the 2-year rainfall) for the same durations, using the median values of the abstracted AMAX series through a combined approach of modelling and error-correction;
- Standardisation of AMAX for each site and duration to a reference distribution, using estimated *RMED*, *SAAR* and northing;
- Pooling of standardised AMAX records, based on concentric circles, and fitting segmented lines to the maximum values of progressively larger pools of network maxima (the FOcused Rainfall Growth EXtension or FORGEX methodology);
- Fitting a consistent DDF model to the FORGEX lines, to ensure a monotonic relationship between rainfall depth, duration and rarity.
- Spatial smoothing of modelled rainfalls, to avoid sudden changes from point-to-point.

All six of these stages were implemented in the current research project, with only minor alterations to re-parameterise the equations for standardisation of rainfall maxima and correct a very minor error in the growth curve derivation routine. The FORGEX methodology and DDF model were applied at each point on a regular 1-km grid for durations based on rainfall accumulations over 1, 2, 4, 6, 12, 18 and 24 hours, and 1, 2, 4 and 8 rainfall days, where a rainfall day starts at 0900 UTC.

## 2.1 Uncertainty and non-stationarity

Estimates of rainfall depth or return period obtained from the original FEH13 model, this recalibration, as well as the earlier FEH99 and FSR models, are given as single values without error bounds or indications of uncertainty. It is important to know how uncertain an estimate is in order to answer questions such as “how likely is it that the estimate given is 10% or 20% too low?”. However, while considering and estimating uncertainty becomes more important as event rarity increases, defining uncertainty bounds and confidence intervals becomes increasingly difficult as there are fewer observations on which to base these. Methods such as bootstrapping (Efron, 1979) can be used to estimate uncertainty in a sample of data. However, modifying the several stages of the FEH13 method to allow this or other uncertainty estimation methods was outside the scope of this project. Estimation of UK-wide *RMED* considers uncertainty to some extent, as the gridded estimate of the true *RMED* value is allowed to deviate further from gauged estimates at rain gauges with shorter records, and less at rain gauges with longer records. Gauges located within 300 metres of each other have their records concatenated prior to extraction of AMAX, unless there is a valid reason not to do so in any individual case, helping to reduce uncertainty.

The original FEH13 model, as well as the FEH99 and FSR models, cannot account for potential non-stationarity in the rainfall data, and modifying the several stages of the FEH13 model to allow it was outside the scope of this project. However, an analysis of flood peaks in the USA by Luke *et al.* (2017) finds that “updated stationary" models (i.e. regularly recalibrated stationary models) can outperform non-stationary models, even when trends are observed to persist. Serinaldi *et al.* (2018) conclude that a clear, physical, deterministic cause for non-stationarity must be specified before a non-stationary model can be justified or configured. From a practical perspective, non-stationary models always have more parameters than equivalent stationary models, as each "stationary" parameter is expressed as a constant and at least one time-varying component. Non-stationary models therefore increase the scope for inaccurate parameterization (Faulkner *et al.*, 2020), particularly increasing the risk of generating models that overfit to their calibration data, then behave inappropriately when extrapolated outside of the range of their calibration data or validation space. Although it is a stationary model, the performance of the FEH99 model at return periods longer than 1000 years clearly demonstrates the potential dangers of extrapolation.

# 3 Study area

Cumbria is an upland county in north-west England. With approximately 500,000 residents in its 6767 km$^2$ area, it is one of England's most sparsely populated places (Cumbria County Council, 2017). Carlisle, in the north of the county, is Cumbria's only city.

Much of the land in Cumbria is rural and mountainous: The Lake District National Park lies entirely inside Cumbria, as do significant portions of the Yorkshire Dales National Park. The Lake District comprises 35% of the total area of Cumbria, or 2362 km$^2$, and contains England's highest mountain, Scafell Pike (978 metres), England's largest lake, Windermere, and England's deepest lake, Wast Water (Lake District National Park Authority, 2018).

Major rivers and catchments in Cumbria include the Eden, passing through Carlisle with a catchment area of 2300 km$^2$, the Derwent, passing through Cockermouth and Workington with a catchment area of 700 km$^2$, and about half of the Lune, passing through Lancaster in Lancashire with a total catchment area of 1000 km$^2$.

Cumbria, with all of the places and features mentioned in this report, is mapped in Figure 1.

## 3.1 Extreme rainfalls in Cumbria

Cumbria has been the site of several major rainfall and flooding events in the 21st century, in January 2005, November 2009, summer and autumn 2012, and December 2015. The floods in January 2005 severely affected Carlisle, causing water levels in the town to rise at least one metre higher than at any time since 1771 (Environment Agency, 2006). The extreme rainfall of November 2009 set new 24-hour, 3-day and 4-day depth records for the UK: 316.4 mm, 456.4 mm and 495.0 mm respectively, at Seathwaite Farm. The 24-hour rainfall total surpassed the previous UK record, set in July 1955 at Martinstown, Dorset, by 37 mm. Although record-breaking at several longer durations, from 24 hours to 4 days, the rainfall rate over the main body of the storm was consistent, with no significant peaks (Met Office, 2012), and so the event was unremarkable at shorter timescales. Sediment analysis from Bassenthwaite Lake, in the Derwent catchment, shows that the peak flow into the lake during this event, as well as the frequency of large floods since 1990, has had no precedent at any time since at least 1460 (Chiverrell *et al.*, 2019).

A series of floods punctuated the period from June 2012 to November 2012, before the UK's 24-hour rainfall record was broken again in December 2015, when 341.4 mm of rain was recorded at Honister Pass, less than 2 km from Seathwaite Farm. Nearby, 405 mm of rain was recorded at Thirlmere Reservoir in 38 hours, setting a new UK-wide 2-day rainfall record.

Extreme orographic enhancement to rainfall results from the mountainous topography of Cumbria. Each of the January 2005, November 2009 and December 2015 rainfalls were driven by warm, moist, westerly or south-westerly airstreams associated with deep Atlantic low pressure systems being forced onto the high ground of Cumbria (Stewart et al., 2012; Met Office, 2019). Generally wet conditions prevail throughout winter in Cumbria, exacerbating winter floods (Blöschl *et al.*, 2019).

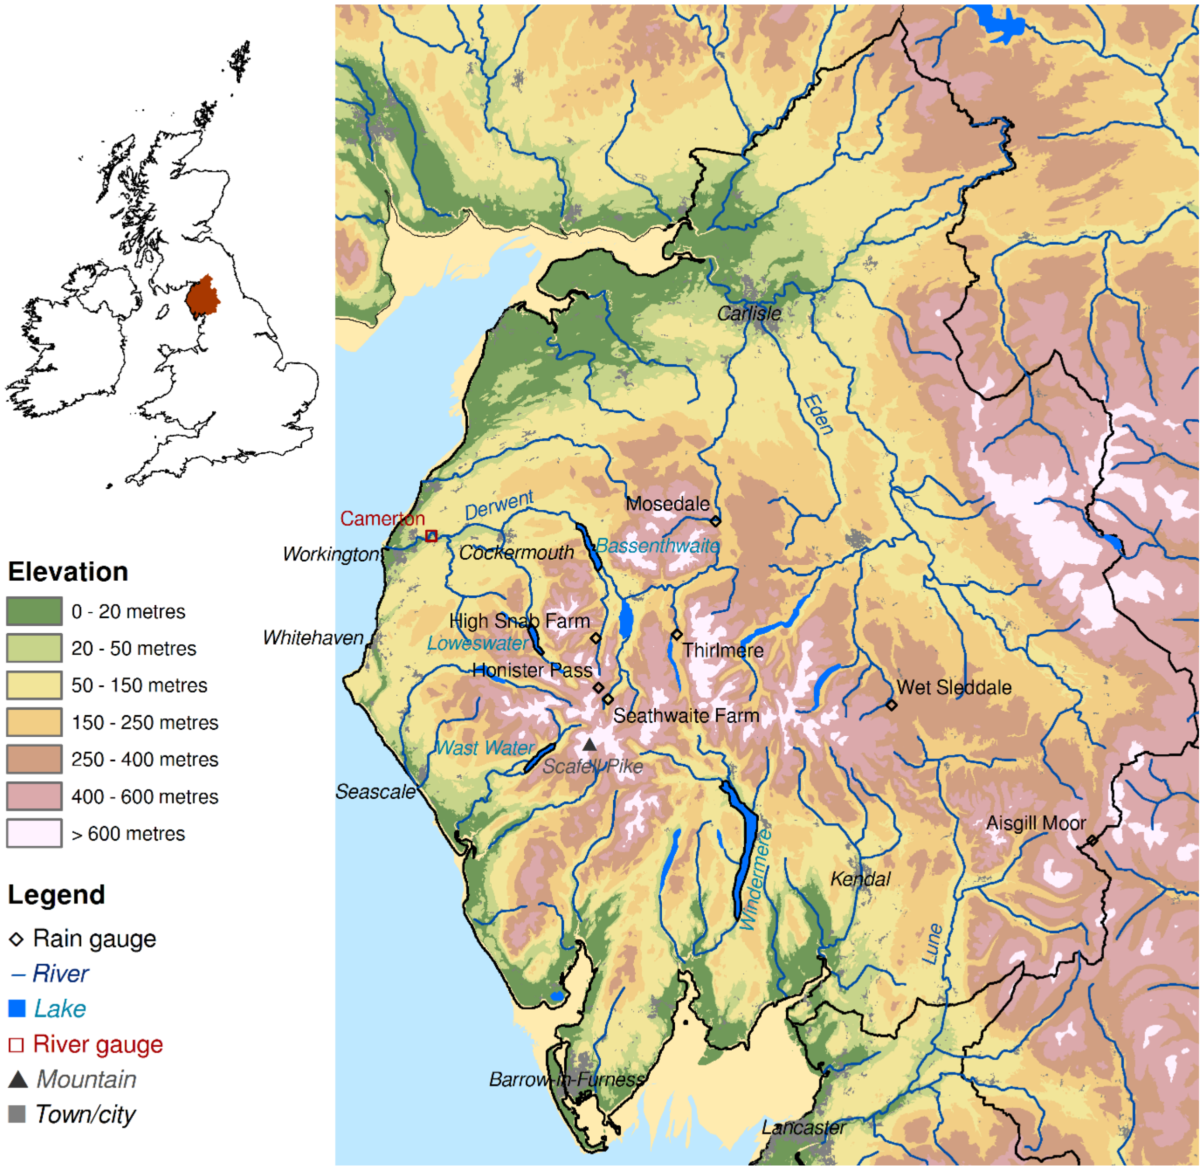


Figure 1 – Map of the British Isles with Cumbria in orange (top left). Map of Cumbria, outlined in black, with all named places and features mentioned in this report (right). Contains OS data © Crown copyright and database right 2018.

# 4 Data & quality control

This project, as with previous UK rainfall models (NERC, 1975; Faulkner, 1999; Stewart *et al.*, 2013), is based on rainfall data recorded at a large number of sites across the UK. As before, the majority of these sites use storage gauges, read daily at 0900 UTC, to produce a record of 1-day rainfall accumulations with fixed start and end times. A smaller group of sites (including some of those with storage gauges) use recording gauges, which record the time-of-occurrence of fixed-depth small rainfall quantities. Interpolation between these times is used to derive a record of 1-hour rainfall accumulations, starting and ending on-the-clock-hour.

## 4.1 Daily rainfall data

All daily rainfall data for this project were supplied by the Met Office as a single data set covering the period 1853-2016, including data already available to the FEH13 study and previous studies.

Longer rainfall records benefit from reduced uncertainty when they are used for depth-duration-frequency (DDF) estimates. Hence, daily rainfall records belonging to gauges located within 300 metres of each other were concatenated in order to produce the longest and most complete series possible at each site. The Met Office gives all of its gauges a site ID, which corresponds to a location rather than a piece of equipment. Different pieces of equipment at the same location have the same site ID, allowing Met Office gauge records to be concatenated objectively and repeatably. When two rain gauges with the same site ID were active simultaneously, priority was given to the gauge with the lowest gauge number.

Tabulated FEH99 daily annual maxima were not used in the concatenation process or to fill gaps in abstracted annual maxima series (see Section 3.4).

## 4.2 Hourly rainfall data

Hourly rainfall data for this project consisted of: data previously available for the FEH13 project; Met Office hourly rainfall data for the whole UK for the period 2006-2013; and Environment Agency rainfall data for the Northwest Region, for the period 2005-2016:

When more than one data source was available to supply rainfall for a specific combination of site, date and time, the highest priority was given to Met Office data, followed by data used in the FEH13 study. This order was informed by the perceived higher quality of the Met Office data and the more comprehensive quality control applied during the FEH13 study in comparison to the current study.

Concatenation was performed on the hourly data, for the same reasons as the daily data. In three cases, the Met Office had assigned different site IDs to gauges less than 300 metres apart, meaning that the Met Office had already decided that it was not appropriate to consider both sites as equal. Hence, these instances were not concatenated.

The tabulated FEH99 annual maxima for hourly durations were not used in the concatenation process, but were used to fill gaps in abstracted annual maxima series (see Section **3.4**).

## 4.3 Quality control procedures

Extensive quality checking was applied to all daily rainfall records used in the FEH13 project, first by the Met Office before supplying the data, according to its own procedures (Met Office, 2001), and then again in CEH, using procedures detailed in Svensson *et al.* (2009). Briefly, these included a comprehensive spatial consistency check (comparing daily totals with nearby gauges with the same aspect) and removal of mislabelled monthly totals.

The hourly rainfall data supplied by the MO, EA and SEPA for use in the FEH13 project were thoroughly quality controlled as part of that project, following procedures detailed in Appendix A of Svensson *et al.* (2009). Therefore, these data were not re-checked during the current project. The hourly data supplied by the EA Northwest region were checked for potential daily accumulations (one large hourly rainfall in one day, with no other rainfall in the same day). Additionally, all single hourly totals greater than 30 mm were checked, flagging 28 suspect sequences, all of which were confirmed to result from equipment faults.

As discussed in the Reservoir Safety report, hourly data may originate from a variety of different recording devices, such as tipping bucket gauges, tilting siphon gauges and optical gauges. The first two of these under-record any rainfall that occurs while the gauge is resetting, and the amount of rainfall lost depends both on the properties of the gauge (e.g. bucket volume) and the rainfall event (e.g. intensity). No attempt was made to correct potential under-recording in hourly data in either the FEH13 or current project, as the type or design of the gauge used at any particular site is not clear from the available data, and not every site has a daily storage gauge against which 09:00-09:00 totals can be compared.

## 4.4 Abstraction of annual maxima

Following concatenation of rainfall records, maximum accumulations at each gauge for each calendar year were extracted from hourly records for 1-, 2-, 4-, 6-, 12-, 18- and 24-hour durations, and from daily records for 1-, 2-, 4- and 8-day durations. Records with fewer than nine valid annual maxima (hereafter ‘short’ records) were not used for the frequency analysis part of this study (Sections **5** and **6**). Records with at least six valid annual maxima were, however, used in the creation of new *RMED* grids (Section **4**).

Ignoring short records, Table 1 and Table 2 detail the quantity of 1-hour and 1-day accumulations available to this study in the context of previous UK-wide rainfall studies: the FSR, original FEH(99) and FEH13. It is shown that the current study adds over 1000 daily gauges to the quantity used in the FEH13 analysis – the first significant increase in number of daily gauges between studies. The number of available gauge-years at daily resolution is increased by almost 20% over the FEH13, now reaching more than double the number available for the FSR. The increase in the number of newly-available hourly gauges and gauge-years is more modest, especially as the quantity of hourly gauges and gauge-years approximately doubled from the FSR to FEH99, and from the FEH99 to FEH13.

Table 1 Number of gauges available to this study, FEH13, FEH99 and FSR (approx.) for 1-hour and 1-day accumulations

| Duration | No. of gauges | | | |
|---|---|---|---|---|
| | This study | FEH13 | FEH99 | FSR (approx.) |
| 1 day | 7,651 | 6,504 | 6,106 | 6,600 |
| 1 hour | 1,036 | 969 | 375 | 200 |

Table 2 Number of gauge-years available to this study, FEH13, FEH99 and FSR (approx.) for 1-hour and 1-day accumulations

| Duration | No. of gauge-years | | | |
|---|---|---|---|---|
| | This study | FEH13 | FEH99 | FSR (approx.) |
| 1 day | 203,116 | 171,904 | 150,245 | 96,000 |
| 1 hour | 19,668 | 17,010 | 7,389 | 2,300 |

The number of valid daily and hourly gauges and gauge-years varies slightly by accumulation period, affecting both the spatial coverage of the UK and uncertainty in individual gauged records. However, the effects are slight.

In addition to the gauges quantified in Table 1 and Table 2, there are approximately 450 further daily gauges and 270 further hourly gauges with six, seven or eight valid annual maxima (depending on duration). These are used in the creation of new grids of median annual rainfall (*RMED*), described in Section **4**.

Figure 1 plots the locations of the daily gauges used in this project and in the Reservoir Safety project, cropped to a rectangle surrounding Cumbria. Figure 2 plots the locations of the hourly gauges used in this project and the reservoir safety project, focused on Cumbria, while Figure 3 plots the same for the whole UK. Yellow × symbols on Figure 2 and Figure 3 show the locations of all gauges supplied by the Environment Agency northwest region, not all of which have nine or more valid annual maxima. Gauges marked by a yellow × symbol without a corresponding red + symbol have less than nine years of valid annual maxima and are used in the creation of *RMED* grids only. Met Office gauges and gauges used in the FEH13 project with less than nine years of valid annual maxima are not plotted, except where these gauges now have nine or more valid annual maxima as a result of the additional data collection occurring between the FEH13 project and this project.

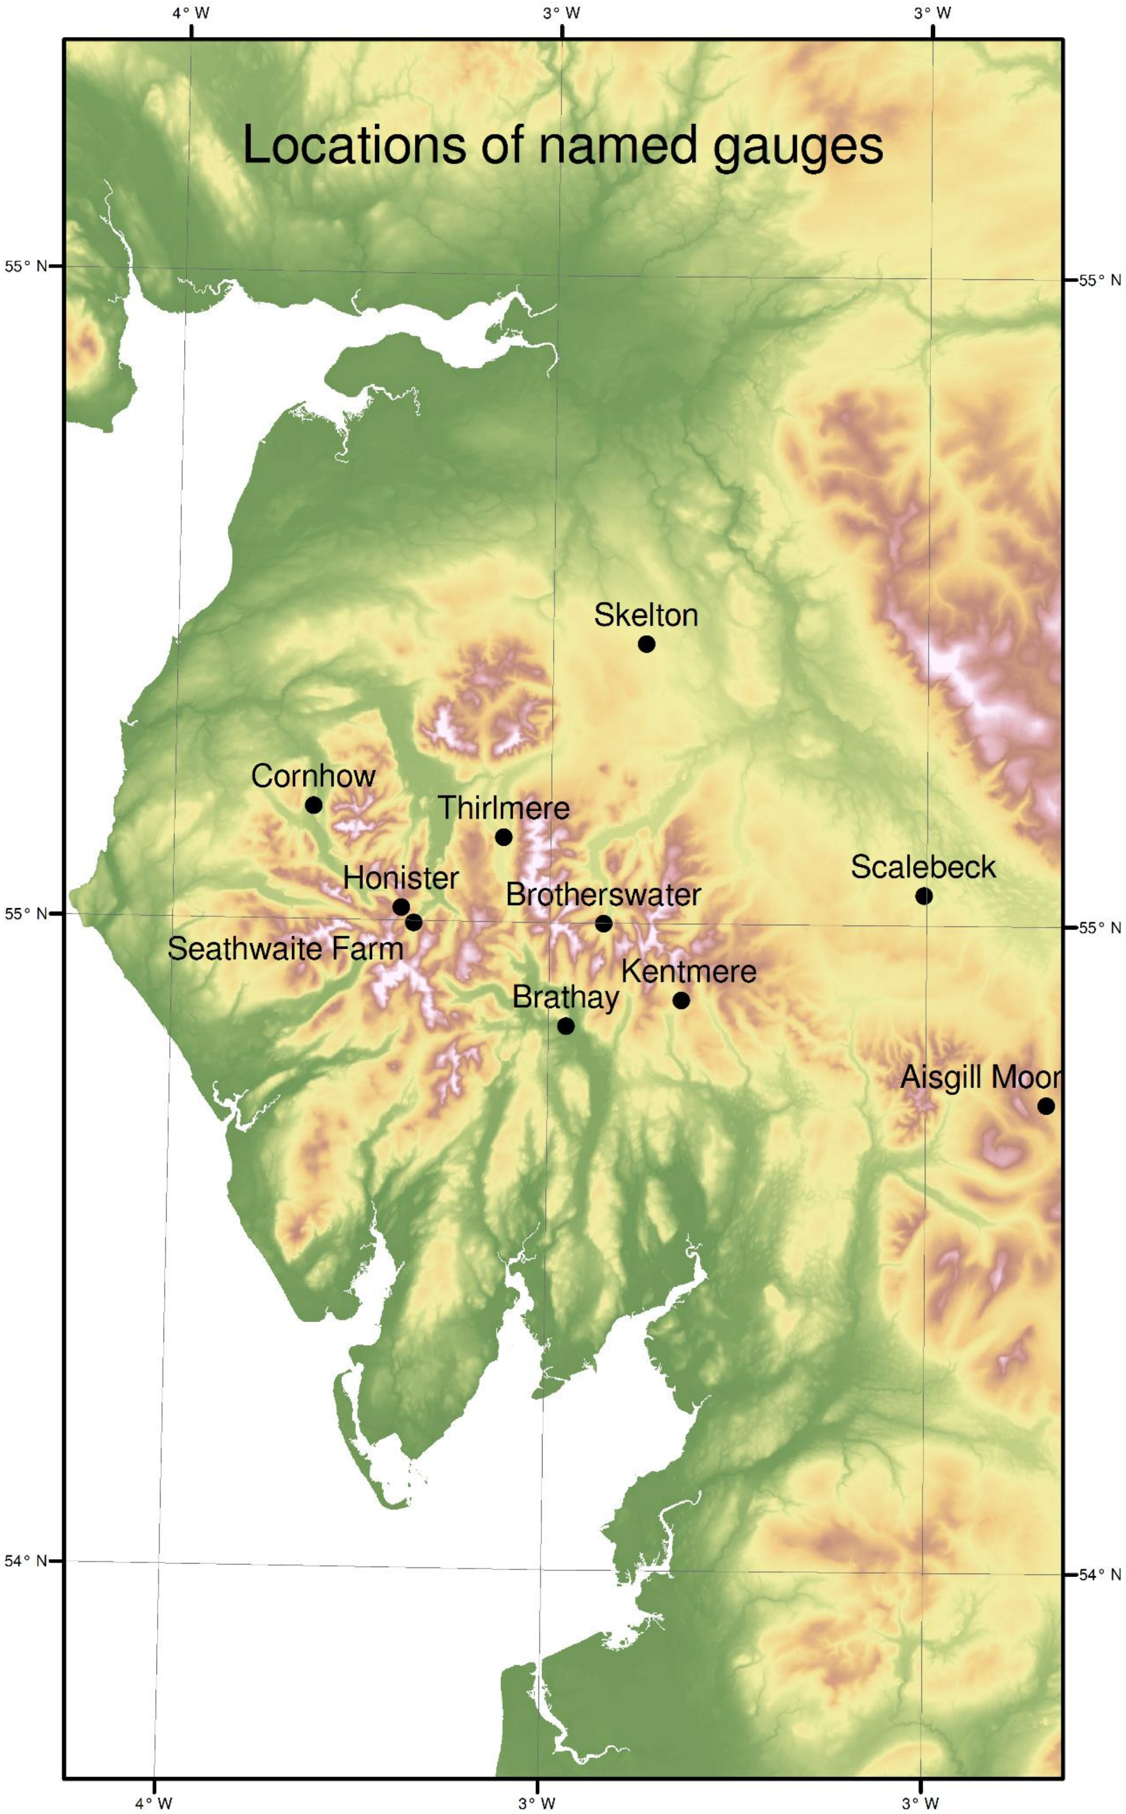


Figure 1 – locations of 10 named gauges

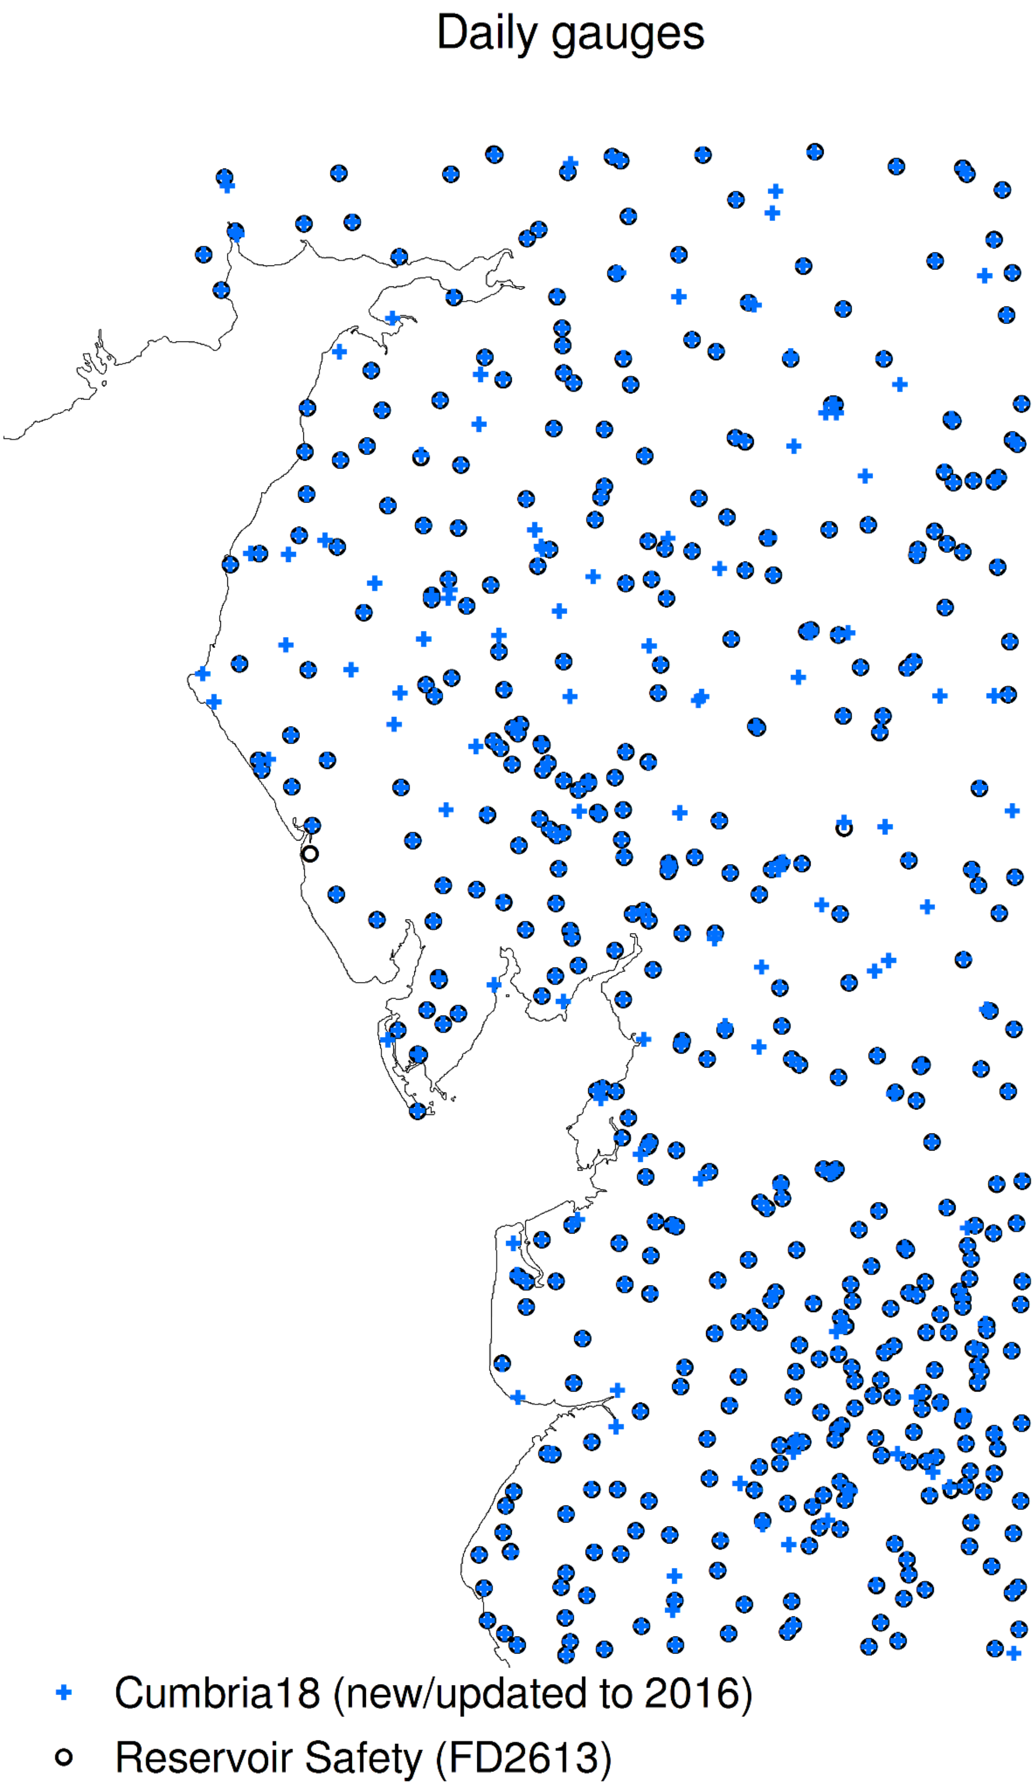


Figure 2 – locations of daily gauges providing data

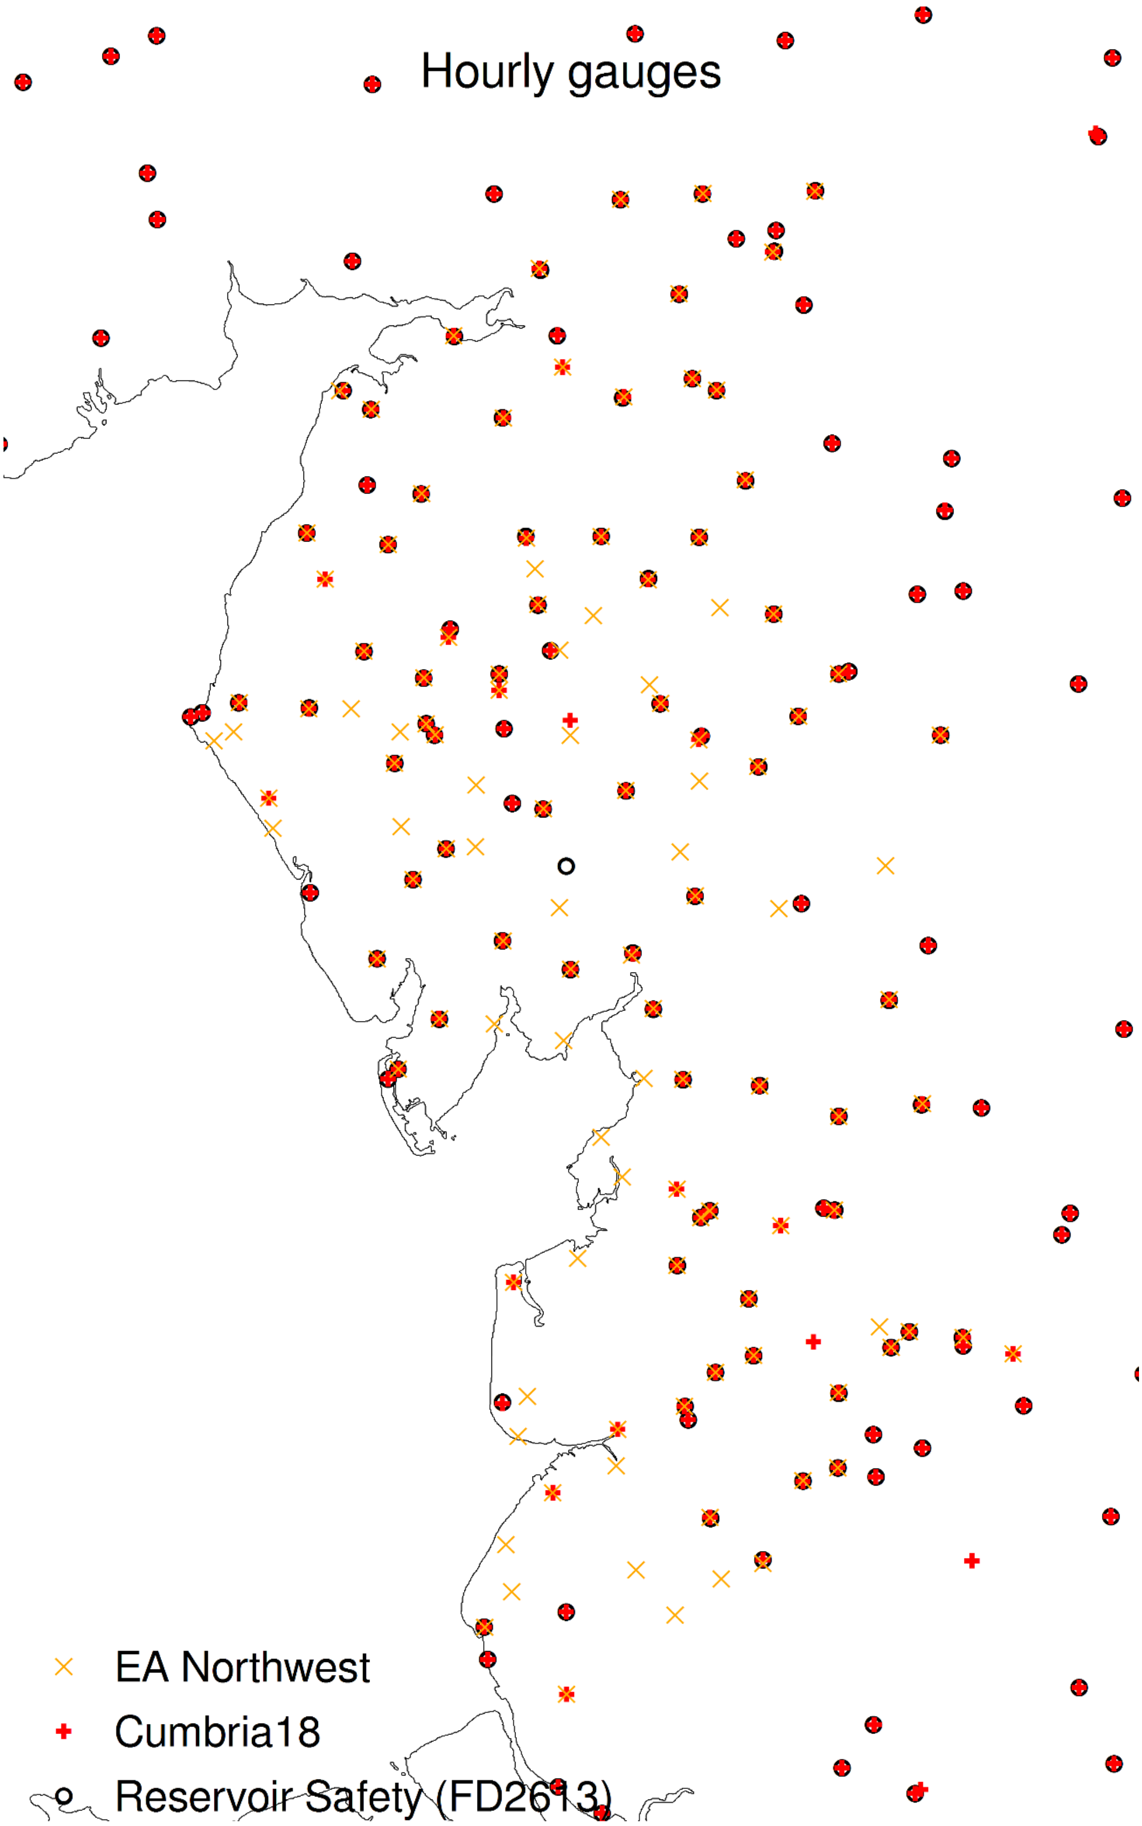


Figure 3 – locations of hourly gauges providing data (zoomed to Cumbria)

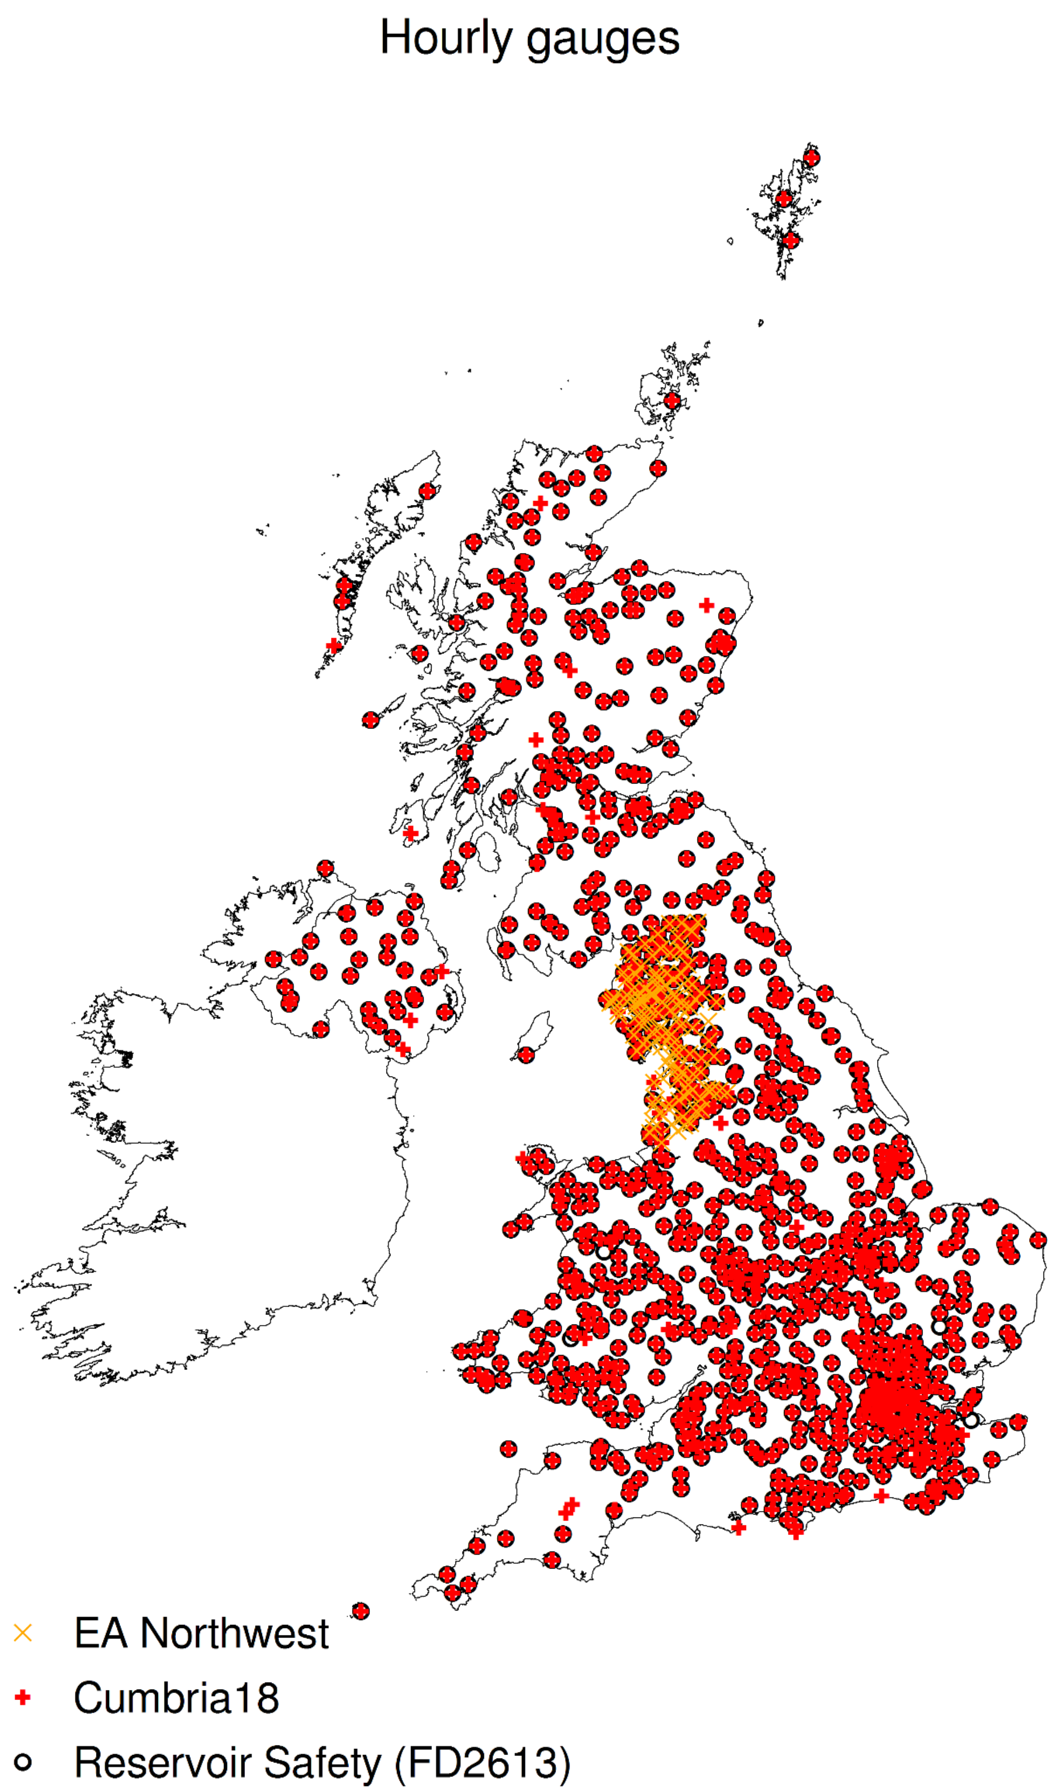


Figure 4 – locations of hourly gauges providing data

# 5 Derivation of new *RMED* grids

As in the original FEH99 model, FEH13 is based on the 'index-flood' method, where a local estimate of an index variable (typically the mean or median annual maximum value) is multiplied by a dimensionless growth curve to obtain a frequency estimate. The method makes the assumption that, following standardisation, the statistical distributions of rainfall at different sites are the same. The FEH13 model modified FEH99's simple standardisation by at-site *RMED*, the median annual maximum rainfall of a specified duration, replacing it with a procedure that uses *RMED* together with a scaling factor that varies from site to site (see Section 5). However, *RMED* remains a key component of the FEH13 standardisation procedure and therefore it was re-estimated on a 1km grid of the UK for each of the key durations using the updated data described in Section 3.

The derivation of the *RMED* grids followed the procedure used in the development of FEH13, which in turn broadly followed the method used in FEH99 (Faulkner, 1999) with some modifications. The FEH13 *RMED* interpolation is composed of two parts: a broad-scale regression equation based on climatic and topographic explanatory variables, which is used to predict gauged *RMED*, and an interpolated grid of errors between the gauged *RMED* values and the regression grid. The main differences between this approach and the FEH99 approach are the choice of external explanatory variables and the combined use of regression and georegression across durations.

*RMED* was re-mapped for the entire UK as annual maxima from a radius of up to 200 km can contribute to rainfall growth curves at any location.

## 5.1 Broad-scale regression equation

The FEH13 broad-scale equation is based on *SAAR*, easting, northing and 'elliptical' distance from a point near Lille, France (an index of continentality). The same form was used in this study, with the coefficients refitted to best match the updated gauged values of *RMED*. This form is presented below.

$$\ln(RMED) = \max(S_1, S_2) + N + E + L \qquad (4.1)$$

where $S_1$ and $S_2$ are *SAAR*-based predictors for high and low *SAAR* respectively, $N$ is a duration-dependent adjustment based on northing, $E$ is an adjustment based on easting and $L$ is an adjustment based on elliptical distance from a point near Lille (750 km east and 80 km north of the origin of the British national grid). $L$ is referred to as an elliptical distance as the true easting offset between the reference and target points is doubled.

The formulation of *RMED* as a regression model and correction grid is used in preference to triangular interpolation between gauged values of *RMED* as *SAAR* and L do not vary linearly and equally in either east-west or north-south directions. Additionally, use of duration in hours (h) as an additional variable within $S_1$, $S_2$ and $N$ ensures that *RMED* increases monotonically and smoothly with duration.

Both 24-hour and 1-day medians contribute to the model fitting for $h$ = 24. Although 24-hour maxima at any site are by definition always equal to or larger than 1-day maxima, differences between them in this dataset were resolved by considering fully-sliding duration rainfall depths throughout. These are the rainfall depths that would be expected in a recording system with no temporal discretisation – while daily rainfall totals are always measured from 0900 UTC to 0900 UTC on the next day, and hourly rainfall totals are always measured from and to zero seconds past an hour, the measuring period for a fully-sliding $x$-hour rainfall total is always that which maximises the rainfall collected over $x$ hours. Approximate conversion from gauged to fully-sliding depths was achieved by multiplying every gauged value of *RMED* by a discretisation conversion factor (DCF), a constant linked to the duration of *RMED* relative to the temporal discretisation of the recording (Table 3). The DCF values used were proposed by Stewart et al. (2013: *Table J.5*) and are those applied in the derivation of the FEH13 results available via the FEH Web Service.

Table 3 Discretisation conversion factors recommended by Stewart *et al.* (2013)

| Duration | 1h | 2h | 4h | 6h | 12h | 18h | 24h |
|---|---|---|---|---|---|---|---|
| DCF | 1.155 | 1.070 | 1.035 | 1.017 | 1.008 | 1.005 | 1.004 |
| Duration | 1d | 2d | 4d | 8d | | | |
| DCF | 1.131 | 1.068 | 1.042 | 1.024 | | | |

Model fitting statistics for the regression model are presented in Table 4, including comparisons to the FEH99 model. The FEH99 model has been used in its original formulation i.e. it has not been re-fitted to the longer and more numerous rain gauge records available to this study, although it has been tested on them. RMSE and $R^2$ are in both cases based on ln(*RMED*).

Table 4 Regression model fitting statistics

| Duration | Number of sites | RMSE (FEH99) | RMSE (this study) | $R^2$ (FEH99) | $R^2$ (this study) |
|---|---|---|---|---|---|
| 1h | 1309 | 0.14984 | 0.14321 | 0.1601 | 0.2329 |
| 2h | 1267 | 0.13657 | 0.12395 | 0.2516 | 0.3835 |
| 4h | 1218 | | 0.11170 | | 0.6047 |
| 6h | 1244 | 0.15990 | 0.11100 | 0.3477 | 0.6857 |
| 12h | 1217 | 0.15775 | 0.11171 | 0.5187 | 0.7586 |
| 18h | 1213 | | 0.11534 | | 0.7810 |
| 24h | 1222 | 0.16645 | 0.11748 | 0.5751 | 0.7883 |
| 1d | 8102 | 0.16428 | 0.10184 | 0.4227 | 0.7781 |
| 2d | 8102 | 0.15819 | 0.09949 | 0.5570 | 0.8248 |
| 4d | 8106 | 0.16120 | 0.09426 | 0.6515 | 0.8809 |
| 8d | 8103 | 0.17196 | 0.09210 | 0.6807 | 0.9084 |

Although failing to re-fit the FEH99 model slightly favours the new model in this comparison, the performance statistics of the new model are clearly better than those of the FEH99 model at every duration: more than 60% of the variance in ln(*RMED*) is explained for durations of 4 hours or more, more than 75% of the variance in ln(*RMED*) is explained for durations of 12 hours or more and almost 91% of the variance in ln(*RMED*) is explained for 8-day rainfalls. In comparison, the FEH99 model never explains more than about 68% of the variance in ln(*RMED*).

## 5.2 Correction layer

Despite the strong performance of the new broad-scale model, particularly at longer durations, it is not and cannot be expected to be perfect, due to potential errors in the *SAAR* data set, sampling variability inherent in short rainfall records or the effect of properties only partially measured by *SAAR* and location, or only useful in certain locations or for certain durations, for example altitude or direction of slope. The purpose of the correction layer is to act as a multiplicative factor applied to the broad-scale model so that the corrected *RMED* grids more accurately match the 'true' *RMED* values. Consequently, the correction layer is not intended to force exact equivalence between gauged and gridded *RMED* values at the locations of rain gauges, as gauged *RMED* values can only ever be estimates of the true values. However, the difference between the gauged and true *RMED* values decreases as gauged record length increases, so the correction grid tends to position 'true' *RMED* closer to the gauged values at sites with longer records.

## 5.3 Final *RMED* grids

The final *RMED* grids are composed of the broad-scale regression layer multiplied by the correction layer. Figure 5 and Figure 6 map fully-sliding 1-hour and 1-day *RMED* respectively, while Figure 7 and Figure 8 compare these to the final *RMED* values mapped during development of the FEH13 rainfall estimates. Note that the grids produced for Scotland are used in the small region where the Northern Irish grid overlaps with Kintyre.

Similarly to previous maps, the wettest parts of the UK remain the mountainous areas along Great Britain's west coast: the Northwest Highlands, Snowdonia and the Lake District.

Comparison of the most recent (14 May 2018) and FEH13 (20 January 2013) *RMED* grids shows very little change, with more than 70% (1-hour) and more than 60% (1-day) of cells exhibiting a change in *RMED* of less than 1%, and no changes outside the range -7.2% to +12.9%. The small size of these changes corresponds to the insensitivity of median values to outliers. The main area where *RMED* has increased is Cumbria, where many previously short records have been augmented with data from recent, wet years, bringing the medians of the longer records closer to the upper ends of the previous, shorter records. However, *RMED* in the other wettest regions of the UK has remained stable or slightly decreased, particularly in Snowdonia.

Unusually, 1-hour *RMED* is slightly decreased in some individual grid points inside Cumbria. As the hourly gauge network is less dense than the daily gauge network, and short-duration extremes are smaller spatially than longer-duration extremes, it is plausible that the small apparent decreases in 1-hour *RMED* in Cumbria are due to a spatial 'mismatch' between short-duration extreme events and the gauge network i.e. some extreme 1-hour events were undetected because they missed a raingauge.

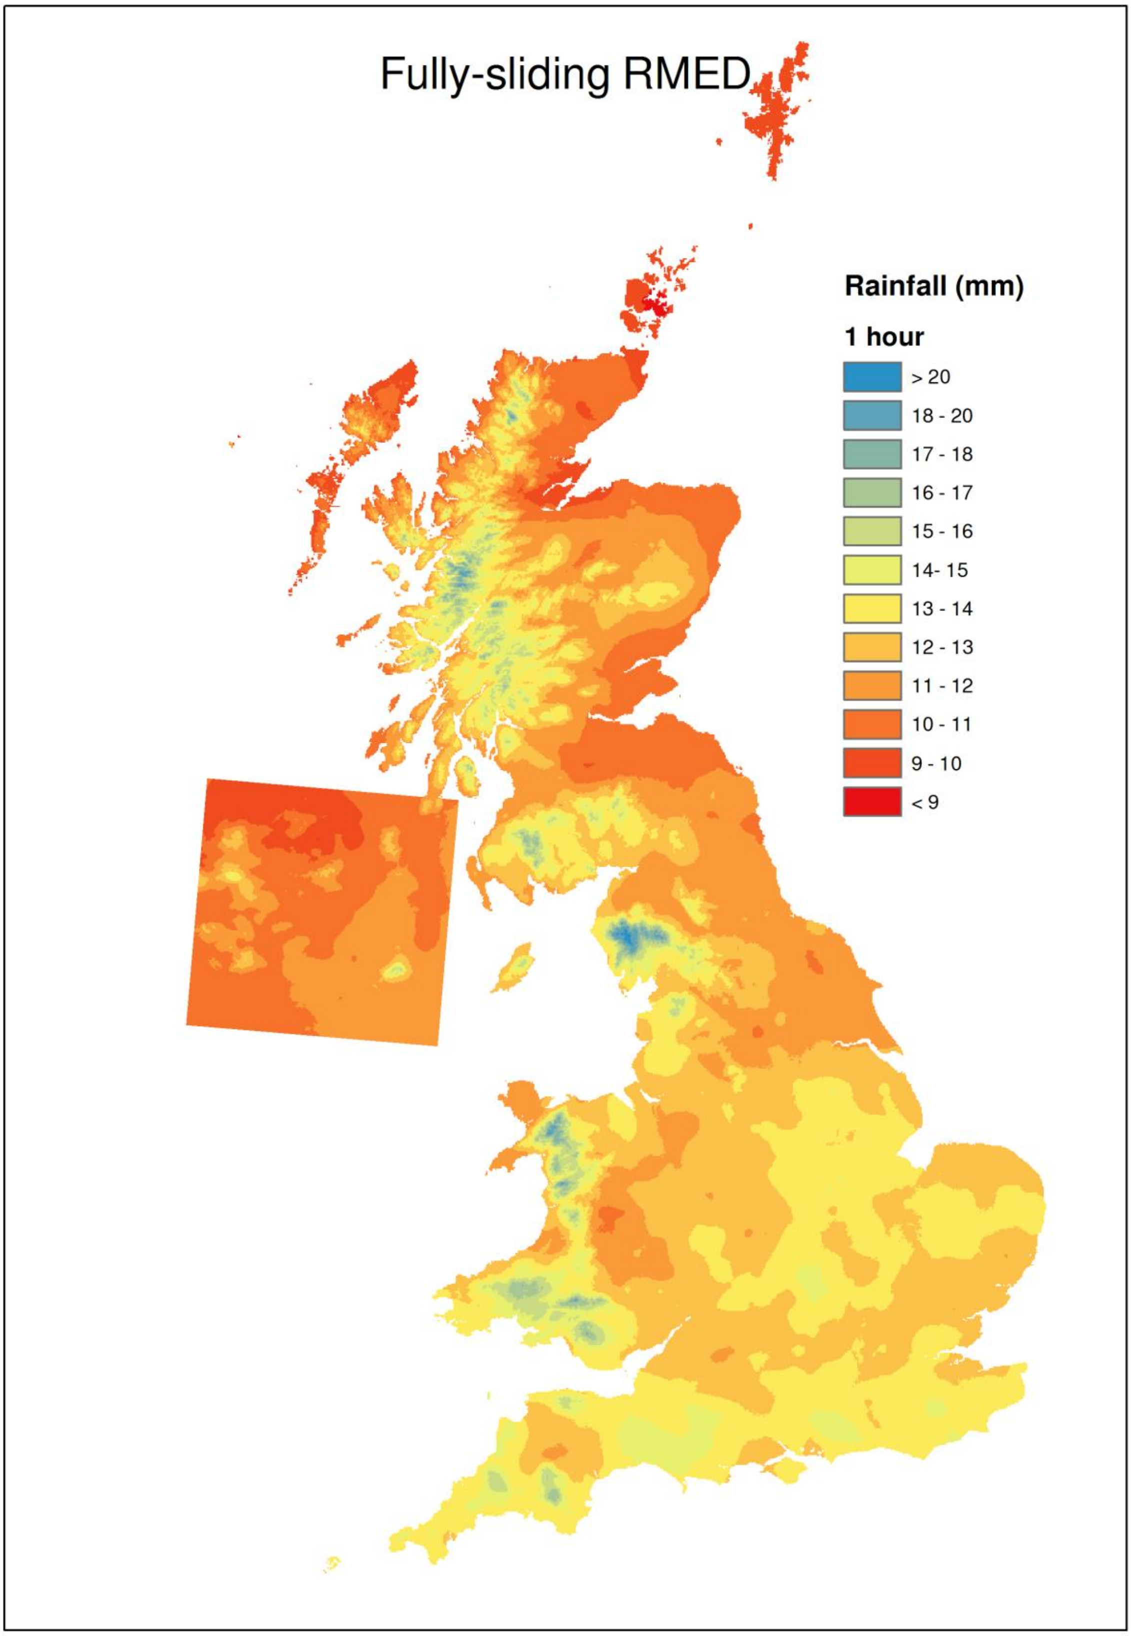


Figure 5 – fully-sliding *RMED* (1-hour)

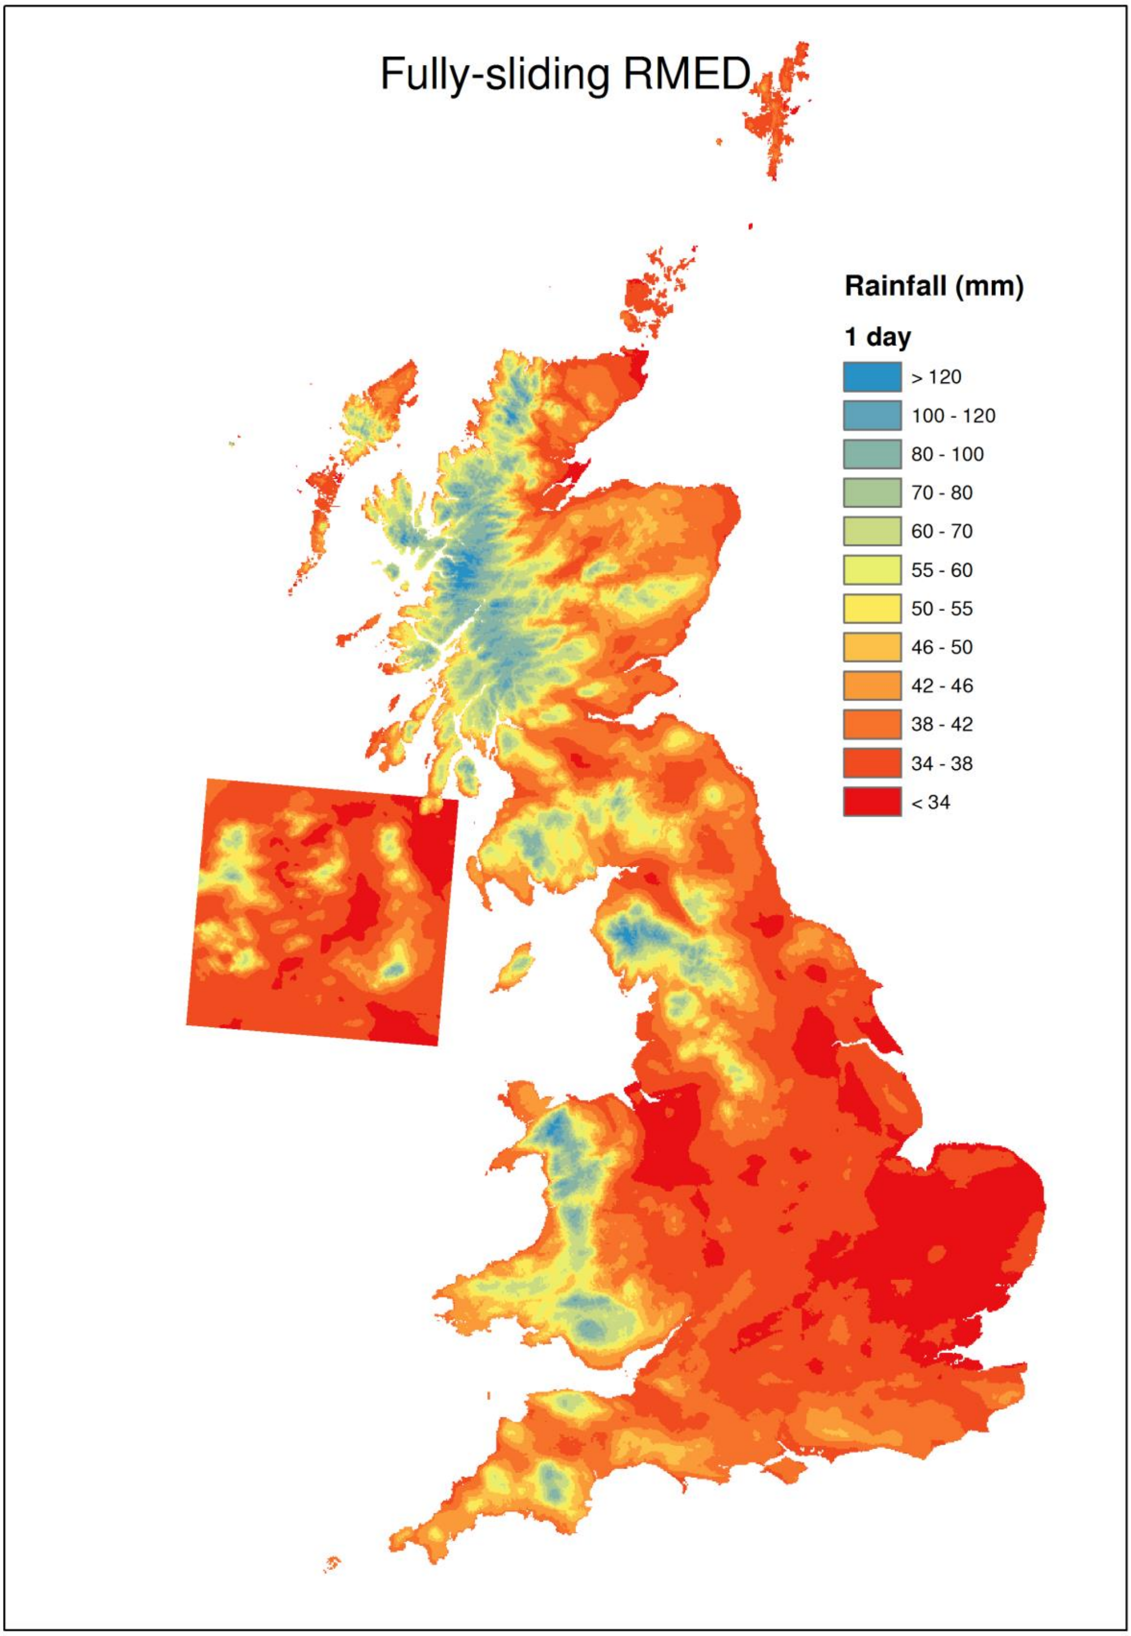


Figure 6 – fully-sliding *RMED* (1-day)

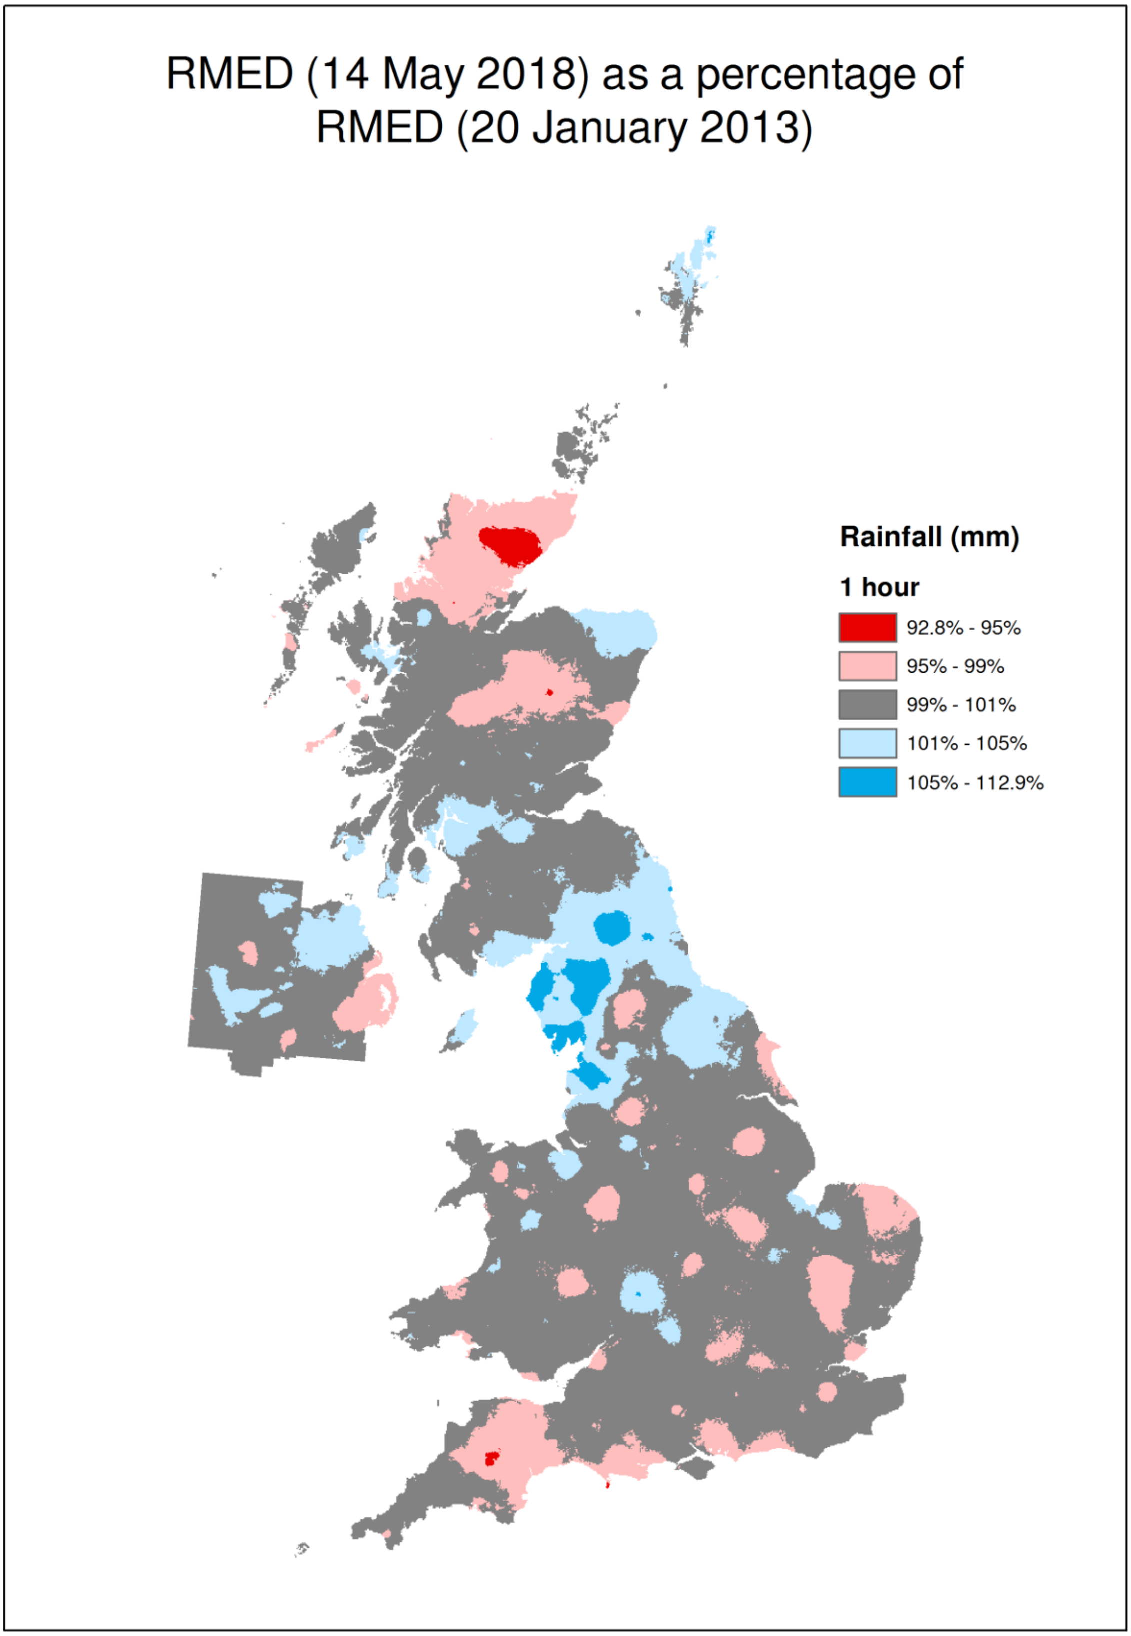


Figure 7 – New fully-sliding *RMED* as a percentage of FEH13 *RMED* (1-hour)

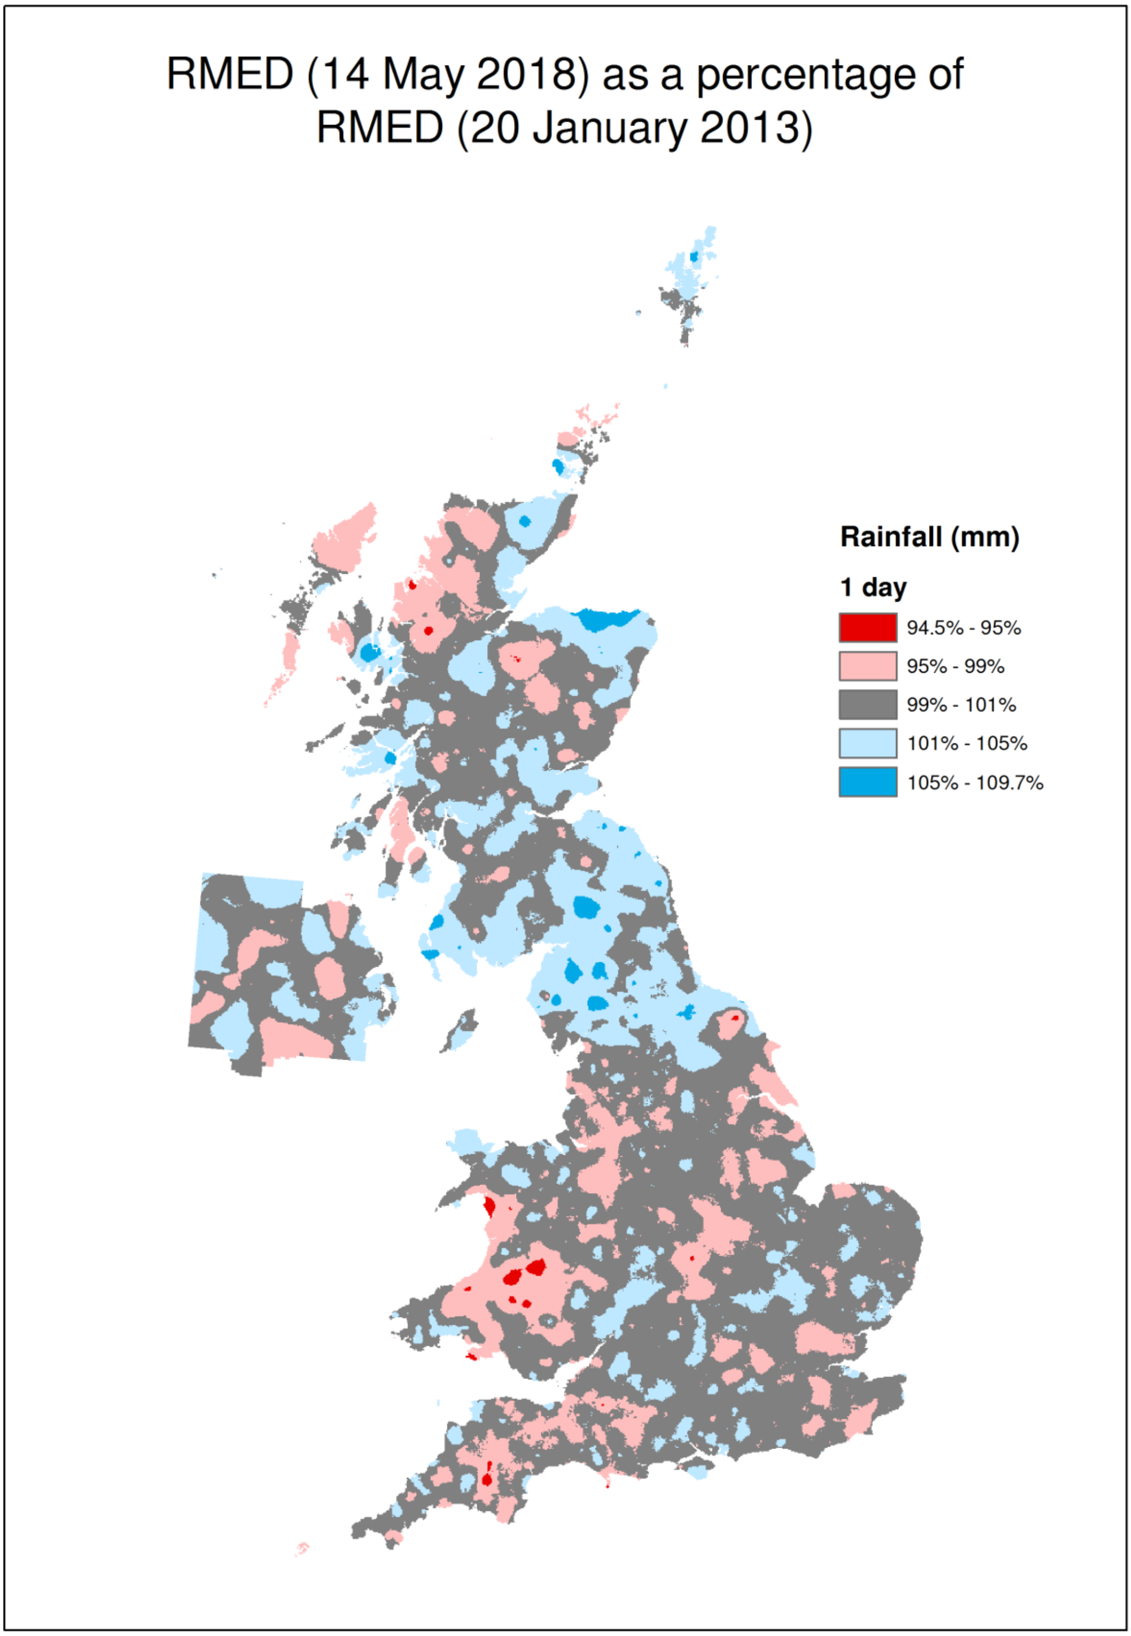


Figure 8 – New fully-sliding *RMED* as a percentage of FEH13 *RMED* (1-day)

# 6 Revised FORGEX procedure

## 6.1 Introduction

FORGEX (FOcused Rainfall Growth EXtension) is the procedure developed during the FEH analysis to form a rainfall growth curve describing the relationship between rainfall depths of different probabilities (return periods) to the index variable, *RMED*. The method is applied individually for each duration of interest and can be centred on any location in the UK (a so-called 'focal point'). Details of FORGEX are given by Faulkner (1999). FORGEX was revised during the FD2613 study and further minor modifications were then made during the development of the FEH13 model but the basic structure and philosophy of the method remain the same.

FORGEX is an empirical method that fits a curve to points plotted on a graph of standardised rainfall against return period. Data points are combined from a hierarchy of expanding circular networks centred on the focal point. The key features of the FEH FORGEX method used in the development of the FEH99 model are:

- the median of at-site annual maxima, *RMED*, is used as the index variable;
- individual durations are treated separately in the construction of rainfall growth curves;
- growth curves are focused on the site of interest rather than applying to pre-defined regions;
- annual maxima are pooled from a network of gauges that expands with return period, giving precedence to the use of local data;
- shifted network maximum rainfalls account for inter-site dependence in rainfall extremes;
- the growth curve is seamlessly extended to long return periods;
- the growth curve is made up of linear segments on a Gumbel scale, avoiding a distributional assumption.

## 6.2 Revised FORGEX (FD2613, 2010)

The FD2613 study made amendments to two key elements of the FEH FORGEX procedure, the standardisation and the spatial dependence model. In addition, the empirical methodology itself was revised to improve both the selection of data points and the method of curve fitting. Details of these revisions and subsequent practical modifications made during the development of FEH13 are discussed in sequence below.

### 6.2.1 Revised standardisation

Stewart *et al.* (2013) introduced a new standardisation step to the Revised FORGEX procedure of the following form:

$$R_{revised} = 1 + \frac{R - RMED}{f \times RMED} = 1 + \frac{1}{f}\left(R_{standardised} - 1\right) \qquad (5.1)$$

where $R_{revised}$ is the revised standardised rainfall, $R$ is the annual maximum value for a given year, $RMED$ is the median of the annual maxima for the site of interest and $f$ is a scaling factor that varies from site to site. The revised standardisation is based on the premise that, although the simple FEH standardisation is effective in bringing the distributions of rainfall together so that they have the same location parameter as defined by the median, other differences still exist for which adjustments can also be made by applying the site-specific factor, *f*. This factor reflects the differences in the spread of the distributions of the standardised values between sites. Since the final method for estimating rainfall needs to be applied at ungauged locations, the scaling factor needs to be specified in terms of readily available geographical variables, rather than being derived directly from rainfall data. It takes the form:

$$f = a + b \times (1000 / SAAR)\ c \times ngy \qquad (5.2)$$

where *SAAR* is standard-period average annual rainfall in mm, *ngy* is northing on the British national grid where 1 = 1000 km, and *a*, *b*, and *c* are coefficients as shown in Table 8. Further details are provided by Stewart *et al.* (2013: *Section 5*).

The formula for the revised standardisation can readily be reversed to allow the determination of a design rainfall from the corresponding value of the standardised value:

$$R = RMED \times \{1 + f(R_{revised} - 1)\} \qquad (5.3)$$

Table 5 Coefficients for revised standardisation of annual maxima

| Duration | *a* | *b* | *c* |
|---|---|---|---|
| **1 hour** | 1.285 | 0.363 | 0 |
| **2 hours** | 0.863 | 0.535 | 0 |
| **4 hours** | 0.646 | 0.530 | 0 |
| **6 hours** | 0.601 | 0.506 | 0 |
| **12 hours** | 0.640 | 0.433 | 0 |
| **18 hours** | 0.706 | 0.395 | 0 |
| **24 hours** | 0.771 | 0.339 | 0 |
| **1 day** | 0.707 | 0.402 | 0.091 |
| **2 days** | 0.608 | 0.374 | 0.236 |
| **4 days** | 0.434 | 0.379 | 0.305 |
| **8 days** | 0.412 | 0.339 | 0.260 |

### 6.2.2 Revised model of spatial dependence

The second key component of the FEH FORGEX method that has been revised is the form of the spatial dependence model that underpins the plotting of network maximum points in the construction of rainfall growth curves. FORGEX originally used the model of spatial dependence in rainfall extremes developed by Dales and Reed (1989) to derive an effective number of independent rain gauges in a network, which is used to define the plotting positions of the network maxima for a given set of gauges. The FD2613 study replaced the Dales and Reed model with a more complex spatial dependence model that allows the degree of dependence to reduce as return period increases. The move away from the Dales and Reed ‘constant shift’ model was made following a graphical analysis of gauge networks throughout the UK, which showed a tendency towards independence in the highest network maxima. The model has two parameters, $\gamma_1$ and $\gamma_2$, and allows the relationship between annual maxima at a single site and network maxima to be described in terms of the network area, the number of gauges and the dominant type of rainfall, as indexed by an average *SAAR* value. The method of assigning plotting positions to network maxima was also revised to treat all plotting positions jointly using a modified maximum likelihood approach.

Subsequently, during the FEH13 development, the Revised FORGEX method was slightly amended again following inspection of the estimated return periods assigned to some of the extreme events used in FD2613 as a ‘reality check’. In some (but not all) cases the estimated return periods were much higher than the FSR values for events that were smaller than PMP. For this reason, in the fitting of FEH13 it was decided to revert to the concept of constant spatial dependence with return period, and the two model parameters were set to equal values. This formulation of the revised spatial dependence model is similar to that used in the original FORGEX method.

### 6.2.3 Revisions to empirical FORGEX methodology

The FD2613 study made a number of other modifications to FORGEX. The most important of these was that the standardised annual maxima from individual rain gauges (known as ‘pooled points’ in FEH) were no longer used and instead the growth curve was fitted to network maxima only. This change was introduced to improve the smoothness of the growth curves at locations close to where the most extreme events have been recorded at numerous gauges and was judged to have little effect at other locations (Stewart *et al.*, 2013: *Section 7.7.1*).

An important consequence of this is that each year can only be represented by one event. As an example, while many rain gauges in Cumbria recorded new highest-ever maxima during December 2015, only the single largest one after standardisation can ever represent 2015 in a network maximum series, hence only one of these events per network can be included in the FORGEX fitting procedure. The included event may differ for smaller and larger networks centred at the same focal point, but no more than one will ever be included in any network.

As standardised rainfalls are used, the network maximum event with the largest standardised growth factor in any year may not originate from the gauge with the greatest rainfall depth. Table 6 compares mm depths and standardised depths for maximum 2-day accumulations at some high-recording gauges in 2009 and 2015.

Table 6 Comparison of rainfall depth (mm) and standardised depth (dimensionless) for selected extreme 2-day accumulations

| Year | Gauge | Rainfall depth (mm) | Standardised depth |
|---|---|---|---|
| 2009 | Seathwaite Farm | 396 | 2.820 |
| 2009 | High Snab Farm | 338 | 3.286 |
| 2015 | Thirlmere | 405 | 4.274 |
| 2015 | Honister Pass | 382 | 2.441 |

This indicates that it is likely that the 2-day rainfall experienced at High Snab Farm in 2009 was rarer than that contributing to the record-breaking 3-day event at Seathwaite Farm, despite being 15% smaller; this was identified previously by Stewart *et al.* (2012). However, the 2015 record-breaking event at Thirlmere was genuinely very rare, having a standardised growth factor of 4.27, compared to 3.29 for the 2009 event at High Snab Farm. The 2015 event as logged at Honister pass was only 5% smaller than at Thirlmere. However, as larger rainfalls are typically expected at Honister Pass, its growth factor of 2.44 (vs 4.27 at Thirlmere) means that it cannot be considered a very extreme event.

The other amendments were:

- the introduction of new rules for the definition of network radii;
- the introduction of new rules for the selection of network maxima in the segment fitting procedure;
- the weighting of network maxima;
- the use of additional networks up to a maximum radius of 300 km.

During the subsequent development of the FEH13 model, practical concerns led to the detail of some of these amendments being reviewed and simplified, for example the rules for defining network radii and the method of weighting. The final amendment listed above was found to lead to estimated return periods that were considered to be implausibly high and therefore the maximum radius of 200 km was adopted in FEH13 except in cases where this returned insufficient data to fit the DDF model up to the required return period (mostly on some islands and some coastal extremities).

## 6.3 Summary of Revised FORGEX procedure

FORGEX is a procedure by which growth curves at a focal point are built by defining a series of progressively larger networks centred at the focal point, identifying each network's annual maximum rainfall series, assigning return periods to these annual maxima, based on the quantity and inter-dependence of gauges contributing to each year in the network, plotting certain annual maxima from each network, where validity is based on growth factor and return period, and fitting least-squares straight-line segments to the plotted points.

The updated rainfall data, described in Section 3, form the basis of the rainfall growth curves to which the DDF model is fitted to produce rainfall estimates. The rainfall growth curves are developed from network maxima via the FORGEX procedure, the most recent version of which was described in the report *Reservoir Safety – Long Return Period Rainfall* (Stewart et al., 2013).

The FORGEX procedure as described by Stewart *et al.* (2013) was used in this project almost unchanged. It is implemented through three custom programs: one that identifies and standardises network annual maximum rainfalls from a series of progressively larger networks centred on a focal point, one that fits a series of connected straight lines as closely as possible to these network annual maximum rainfalls, and one that reviews and (if applicable) adjusts the connected straight lines according to consistency rules described in the step-by-step summary.

In this project, the FORGEX fitting programs were run on 1-km intervals, corresponding to the intersections of exact kilometres on the British national grid, over a rectangle defined by bottom-left and top-right corners of (260,440) and

(399,599) respectively. A buffer zone around Cumbria was included in order to allow production of spatially-smoothed results (see Section 6.3). Points over the sea were skipped.

## 6.4 Changes since FD2613

The implementation of FORGEX here differs somewhat from that reported in the Reservoir Safety report, and slightly from that implemented to produce data for the FEH Web Service.

A detailed step-by-step breakdown of FORGEX in its current form is presented, with supporting notes, in Appendix A.

### *6.4.1 Differences between Reservoir Safety report and FEH Web Service implementations*

- Exclusion of Secondary network: in the method as published in the Reservoir Safety report, a secondary radius was set at 300 km. In order to reduce ridge effects at radii of 200 km from major storms, tapering weights were applied to networks with radius between 200 km and 300 km. Secondary radii were not included in derivation of FEH Web Service data as network reduction weights were instead set for networks with radius between (200 / 6) km and ≤ 200 km.
- Network definition: Network 1 was originally defined to contain 40 effective gauge-years, while the three networks immediately below the Primary were defined to contain 87.36% of the number of gauges of the next largest network. Intermediate gauges were defined to contain 66.67% as many gauges as the next largest network. Network 2 contained 1.5 times as many gauges as Network 1, with no other lower limit, such as the 15 used in the FEH Web Service derivation.
- Spatial dependence variable with return period: The Reservoir Safety report describes a spatial dependence model with a dependency on return period, such that rarer events exhibit lower spatial dependence. While plausible, the dependence on return period was dropped from the model prior to production of FEH Web Service results.
- Segment definition: The Reservoir Safety report states that the range of reduced variates between 0.3665 and the third-rarest event in a network is discretised into segments of width less than one. In both the current and FEH Web Service methods, segment widths greater than one are allowed and common (although widths greater than about 1.2 are very rare). As described in Section 5.4.2, the discretisation rule was changed after the production of FEH Web Service data, resulting in less variable segment widths.
- Weighting of network maxima: The method published in the Reservoir Safety report caused the FORGEX fitting procedure to prioritise fitting to some events more than others. This was achieved through the use of event weighting factors, relating to: uncertainty in the return period estimate of the event; the width of the network containing the event, in terms of reduced variate; and the presence of larger but less rare events in larger networks (for this rule, separate weighting equations applied to secondary networks). None of these weights was applied during development of the FEH Web Service data. However, different weights were (and are still) used, based on: network radius in km; and re-occurrence of the same event in higher-numbered networks.

### *6.4.2 Differences between FEH Web Service implementation and current project*

One further, minor, change was made to the procedure between production of the FEH Web Service results and the current project. A very minor error in the existing segment-fitting program was discovered during checking of the FEH Web Service implementation of FORGEX, which related to the method used to discretise the range of reduced variates from 0.3665 to the third-highest value into segments. The program as used to generate FEH Web Service data selected the number of segments for this range by rounding away from the nearest integer, such that a range with a width of 8.01 would be divided into 9 segments while a range with a width of 7.99 would be divided into 7 segments. This project changed the rule so that all rounding is now towards the nearest integer, as was the probable intention of the code. Figure 9 shows the difference between the FORGEX segments fitted to the current data at Honister Pass with both the previous and current rounding rules. The widths for discretisation of the 6- and 12-hour rainfalls are 8.074 and 7.994 respectively. Under the previous discretisation rules, these would divide into 9 and 7 segments respectively, while under the current rules, both divide into 8 segments. However, it is clear from Figure 9 that the effect of the rule change on the shape of the fitted FORGEX line is minimal, at least when the comparison is between lines with a similar number of segments.

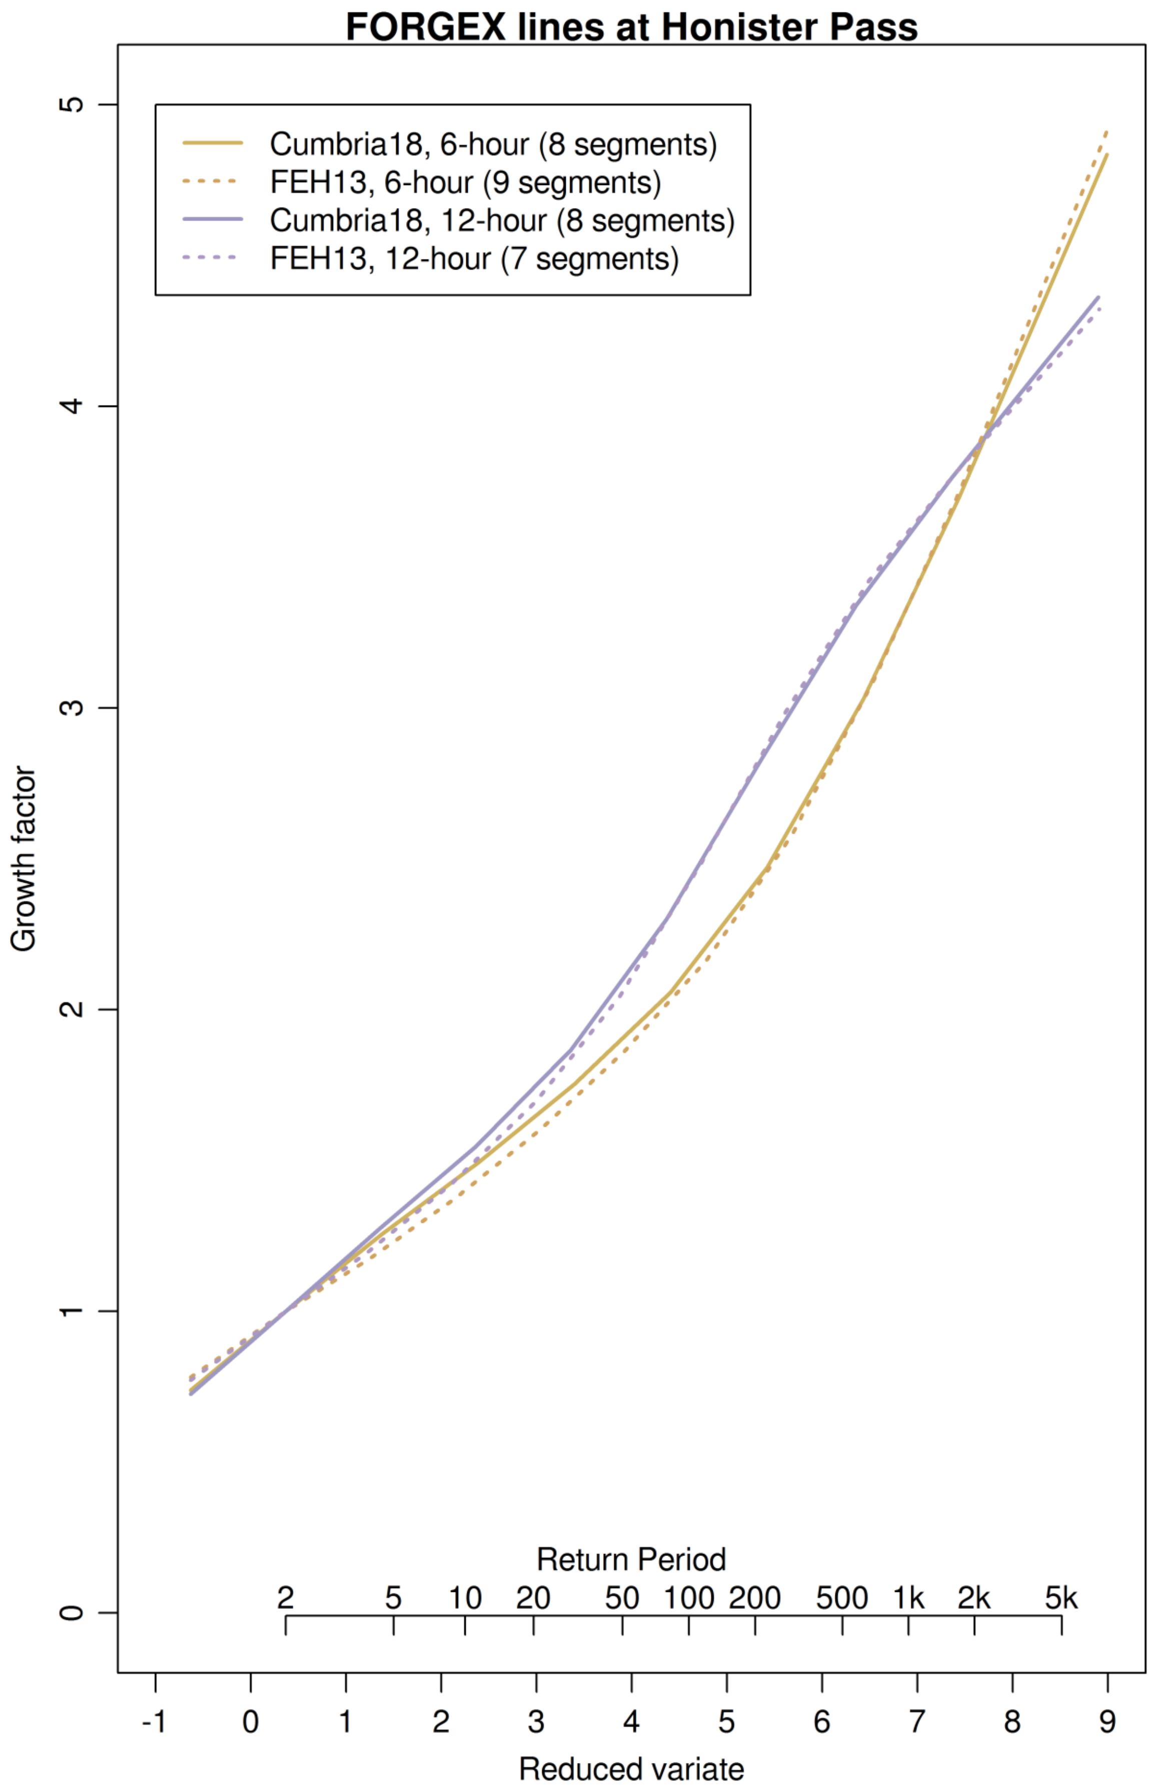


Figure 9 – comparison of FORGEX outputs produced with old and new line discretisation rules

## 6.5 Fitting results

Figure 10 to Figure 13 show the results of the FORGEX fitting procedure at ten locations in Cumbria, for durations of 1 hour, 24 hours, 1 day and 4 days respectively. These present the following general trends:

- Rainfall growth curves are steeper for short durations, reaching higher growth factors at long return periods. All of them show the typical ‘arc-shape’ common to FEH flood frequency curves.
- At 1-day duration, half of the growth curves (at Cornhow, Honister Pass, Seathwaite Farm, Thirlmere and Brotherswater) become more ‘S-shaped’, demonstrating an increase and subsequent decrease in growth rate. All S-shaped growth curves reach their crest between 100 and 1000 years.
- The S-shaped pattern reaches maximum strength at a duration of 4 days, before weakening again at 8 days (not shown).

The prevalence of steeper growth curves for shorter durations is a well-known phenomenon and has been observed during previous FEH analyses (e.g. Figure 2.1 of Faulkner, 1999). It is therefore reassuring to see that the current analysis re-confirms this.

The presence of some S-shaped growth curves for longer durations can be explained by the occurrence of the recent record-breaking events at Honister Pass, Thirlmere and Seathwaite Farm. All of these events have high growth factors for durations of 1-4 days, and by centring a FORGEX analysis at one of these points, it is ensured that the recent record-breaking event from that location will be included in the smallest network, where it will be assigned a relatively low return period as it is the largest event from a small pool. This will cause the FORGEX growth curve to accelerate away from a standardised value of 1 rapidly. It is noted that these record-breaking events may not have the highest standardised growth factors of all those available with a 200-km radius of the focal point, so they may be replaced in larger networks. However, the overall size and presence of the record-breaking events in small networks means that, if or when they are replaced in larger networks, the replacement event is unlikely to have a much higher standardised growth factor but will be assigned a much larger return period. This will cause the growth curve to flatten off as the return period advances through the hundreds to the thousands of years. It is noted that all five stations at which S-shaped growth curves are more prevalent are those located nearest the sites of recent, record-breaking, multi-day rainfalls.

This is explored in more detail in Figure 14 and Figure 16, which show the results of the FORGEX fitting procedure for 4-day rainfall as applied to two locations, and Figure 15 and Figure 17, which show the results of applying the same program to the same locations, but only using the rain gauge data and *RMED* grids available during the Reservoir Safety project. On all four figures, each fitted segment is coloured separately, with the individual events that contribute to each segment’s fitting plotted in the same colour. The colour intensity of each plotted event indicates its weighting in the fitting procedure: points with the maximum weighting are plotted in exactly the same colour as the fitted line; this fades to white as the weighting reduces to zero. These figures demonstrate an S-shaped growth curve at Honister Pass, produced by the recent, record-breaking, multi-day rainfalls, and a more arc-shaped growth curve at Aisgill Moor, which is more distant from Honister Pass, Thirlmere and Seathwaite Farm.

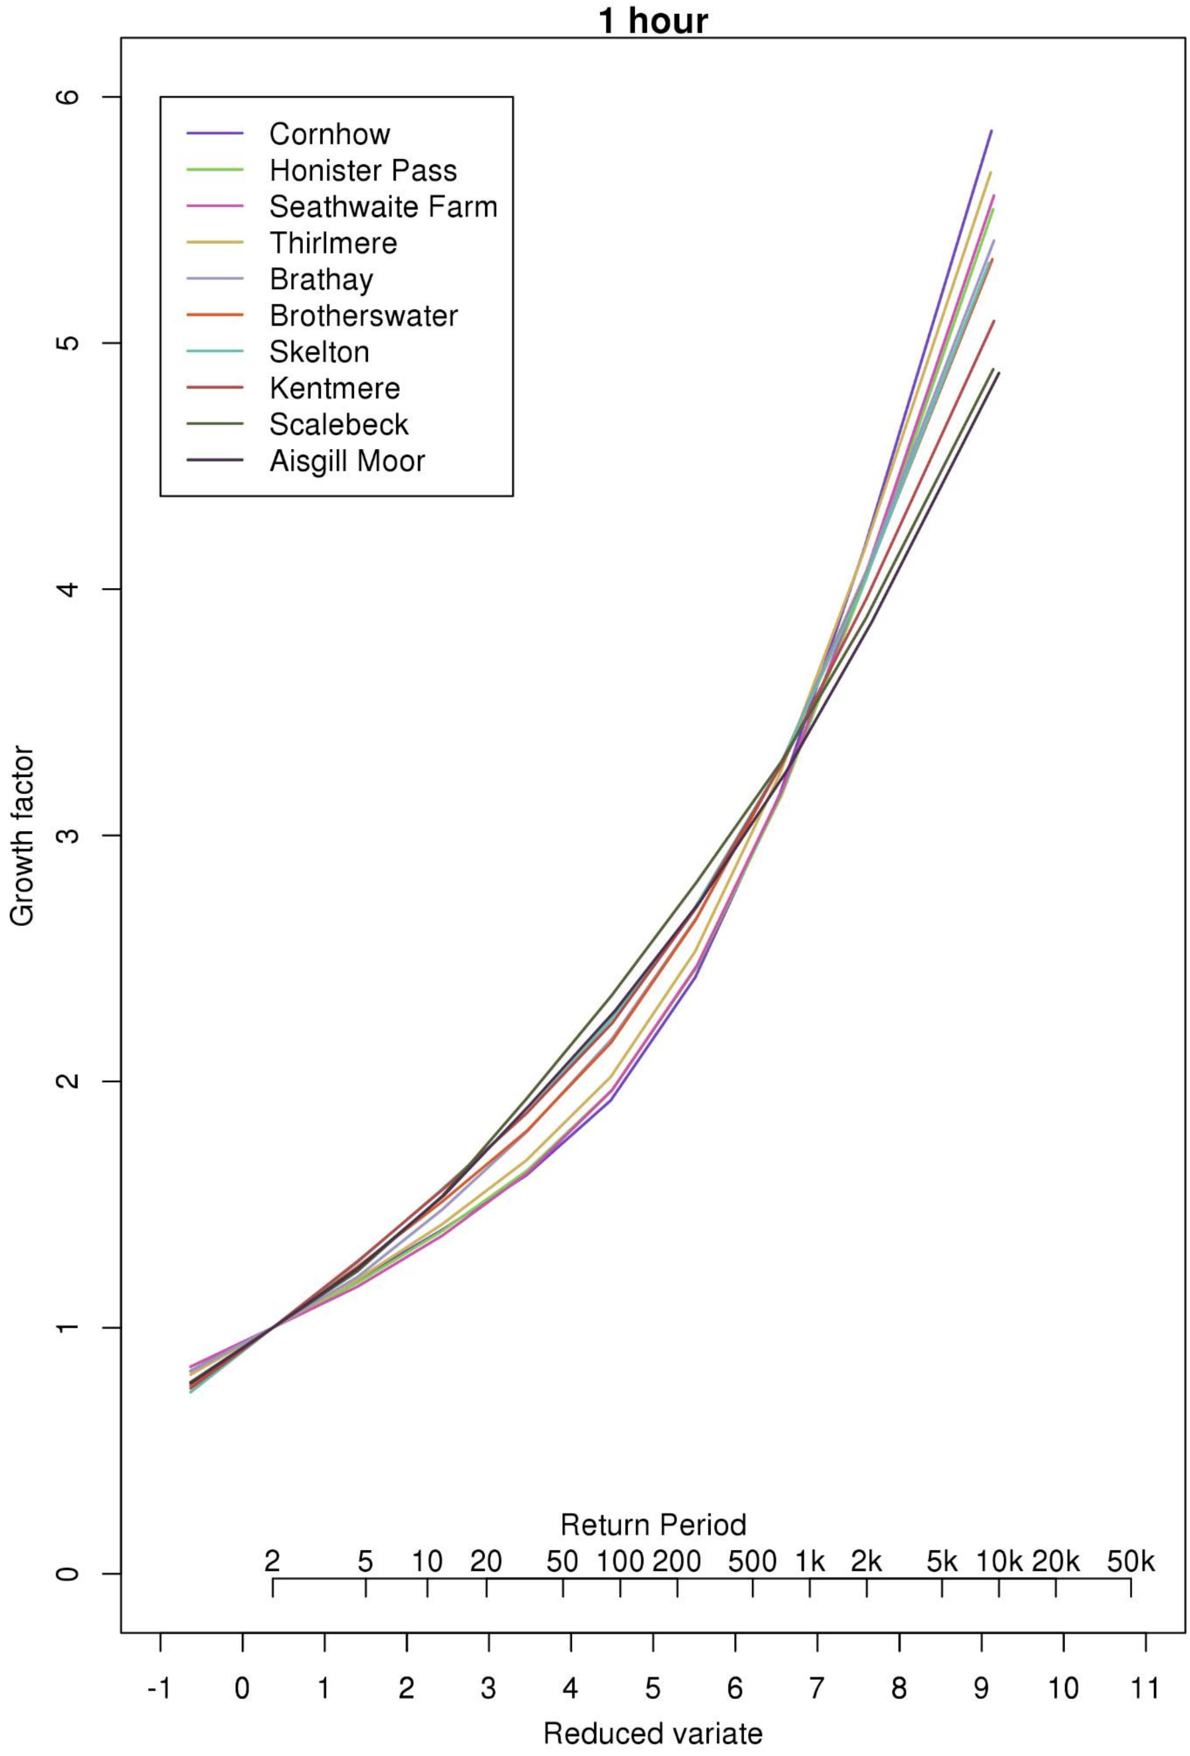


Figure 10 – comparison of FORGEX lines produced for case study sites (1 hour)

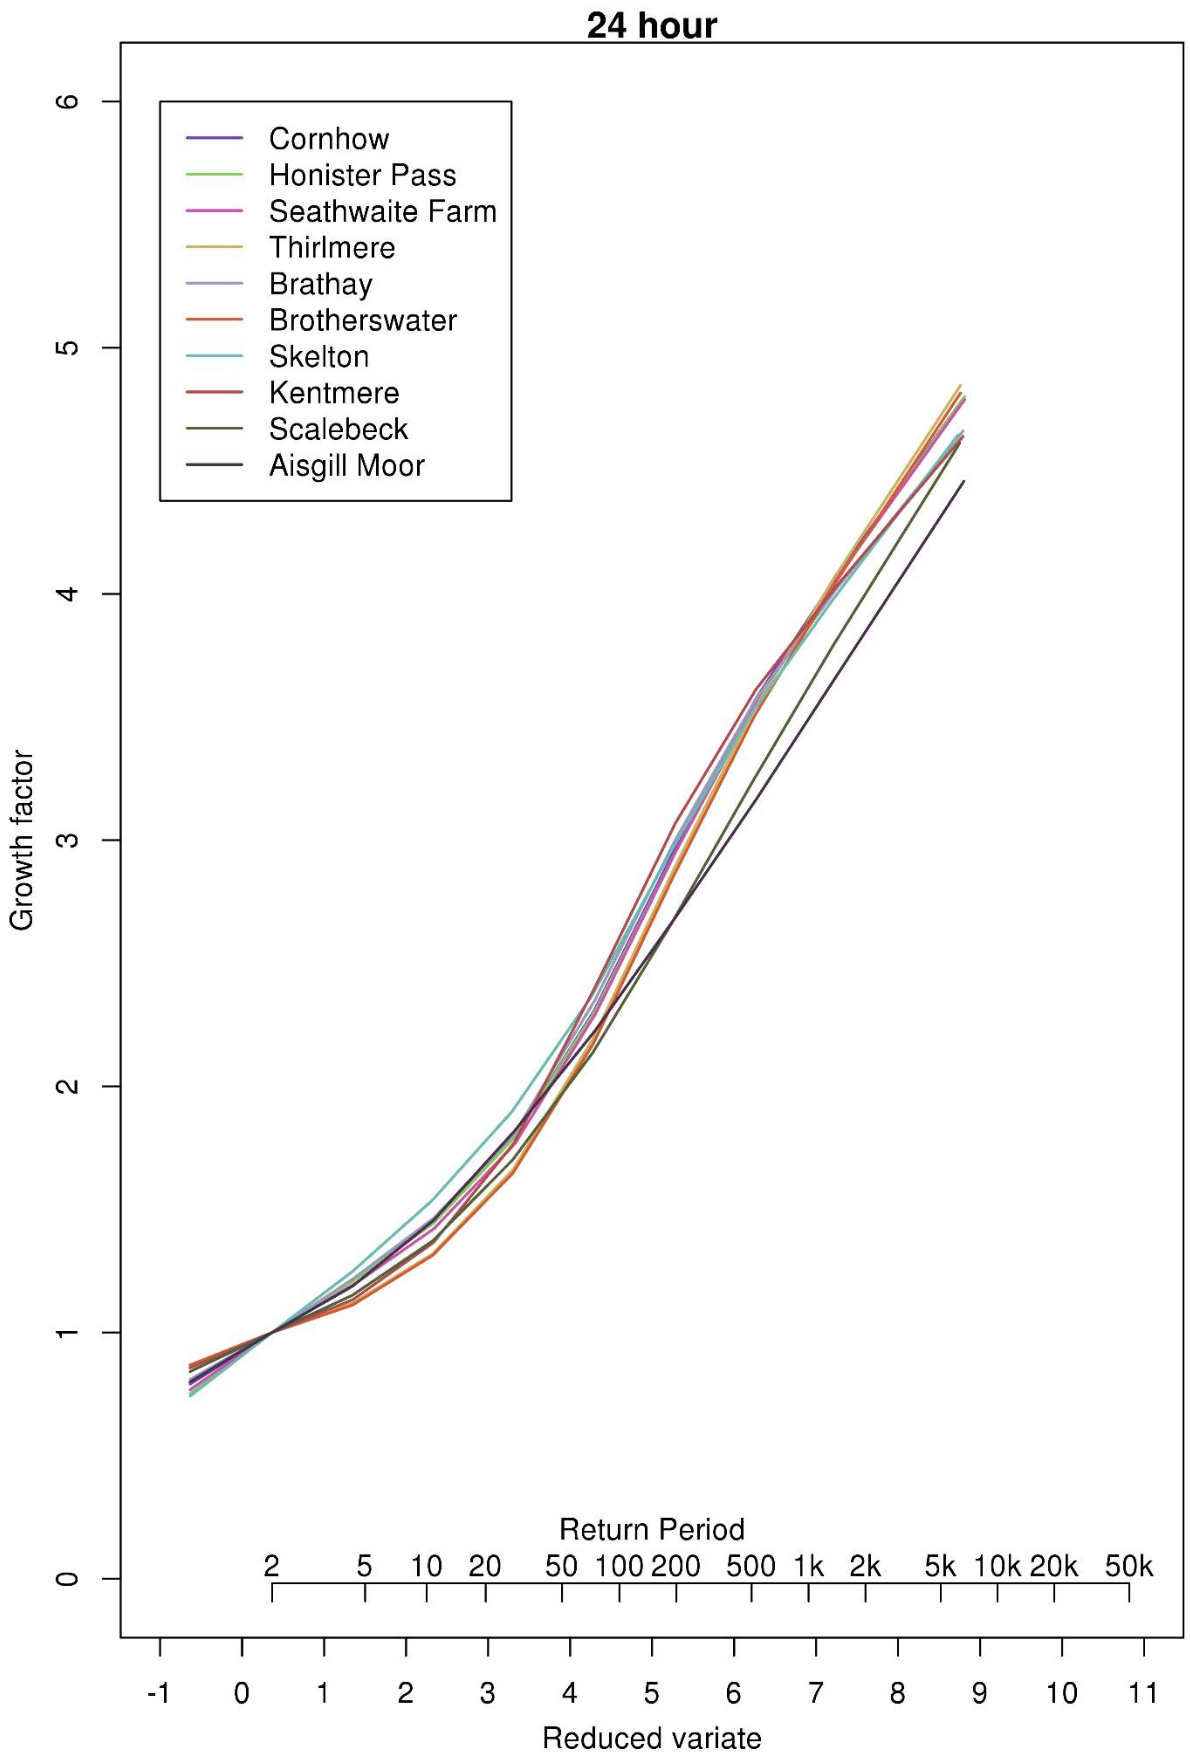


Figure 11 – comparison of FORGEX lines produced for case study sites (24 hour)

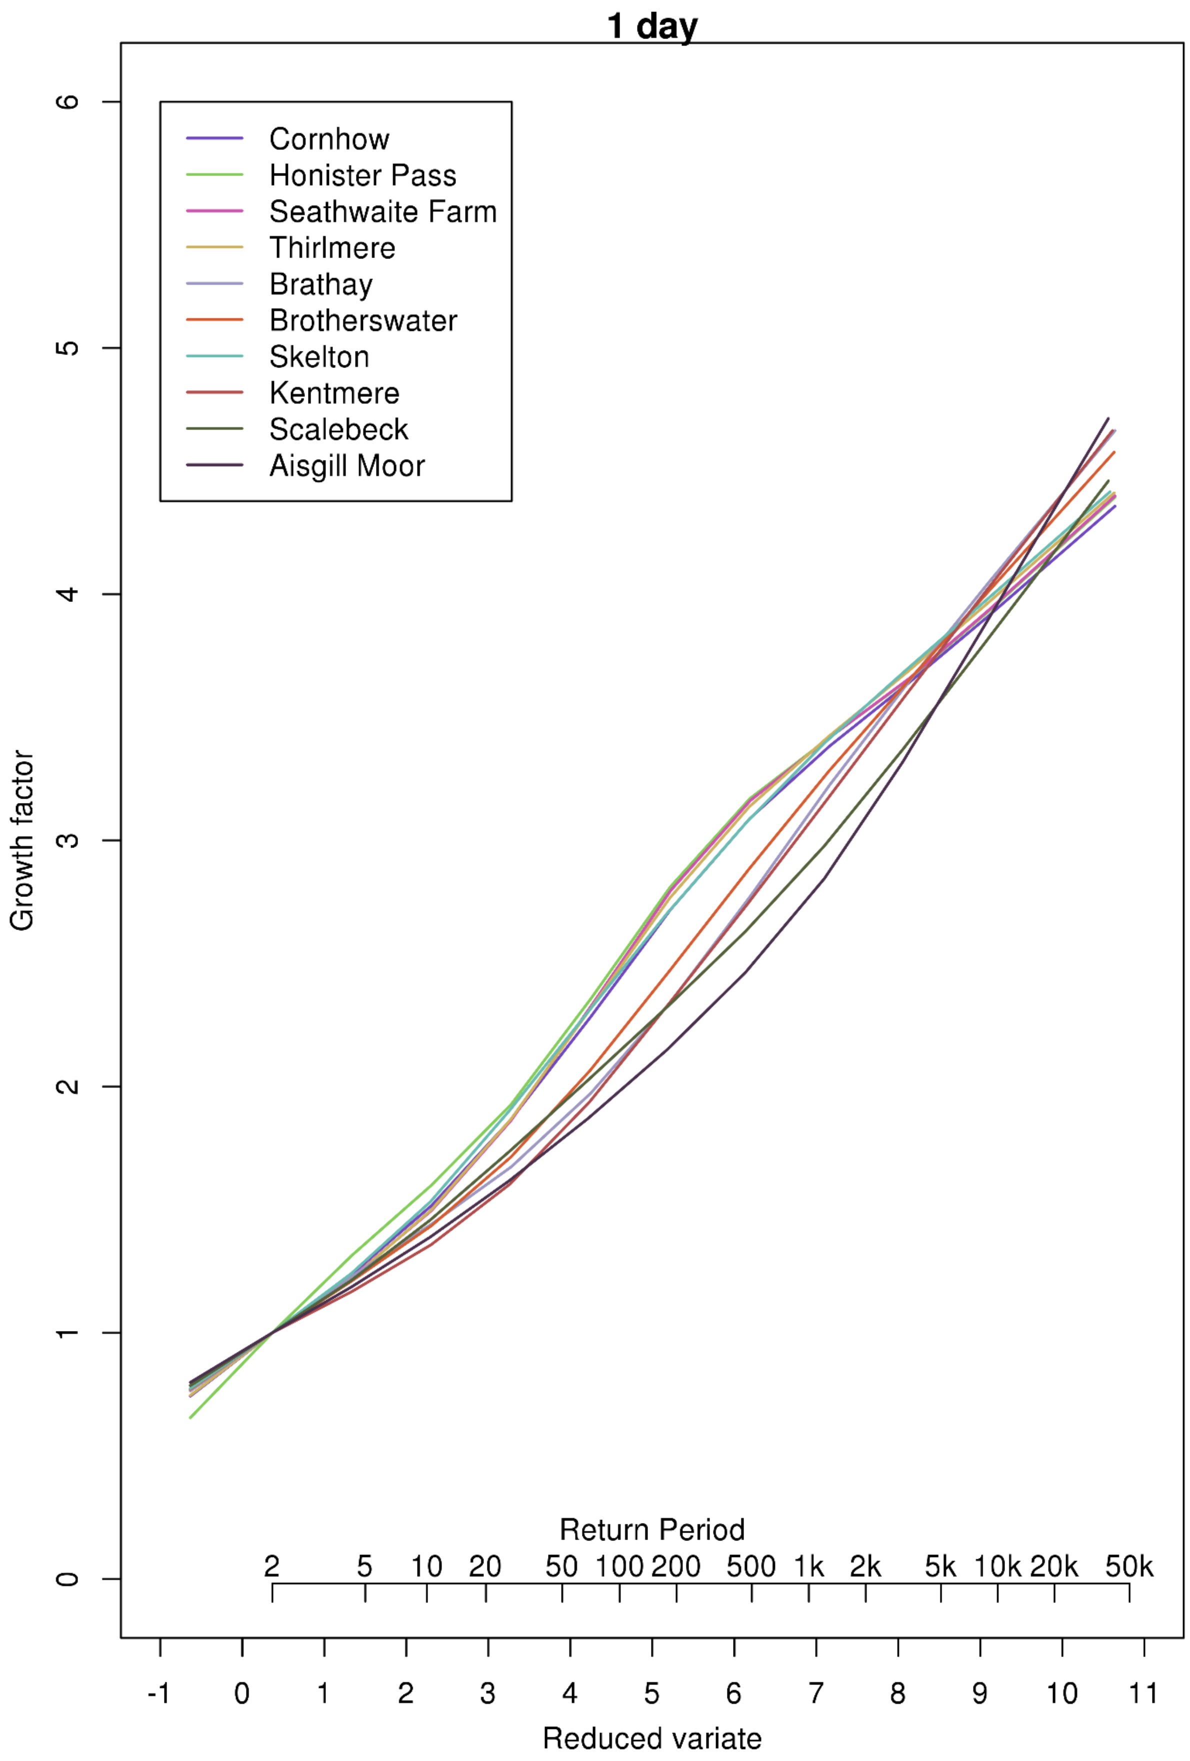


Figure 12 – comparison of FORGEX lines produced for case study sites (1 day)

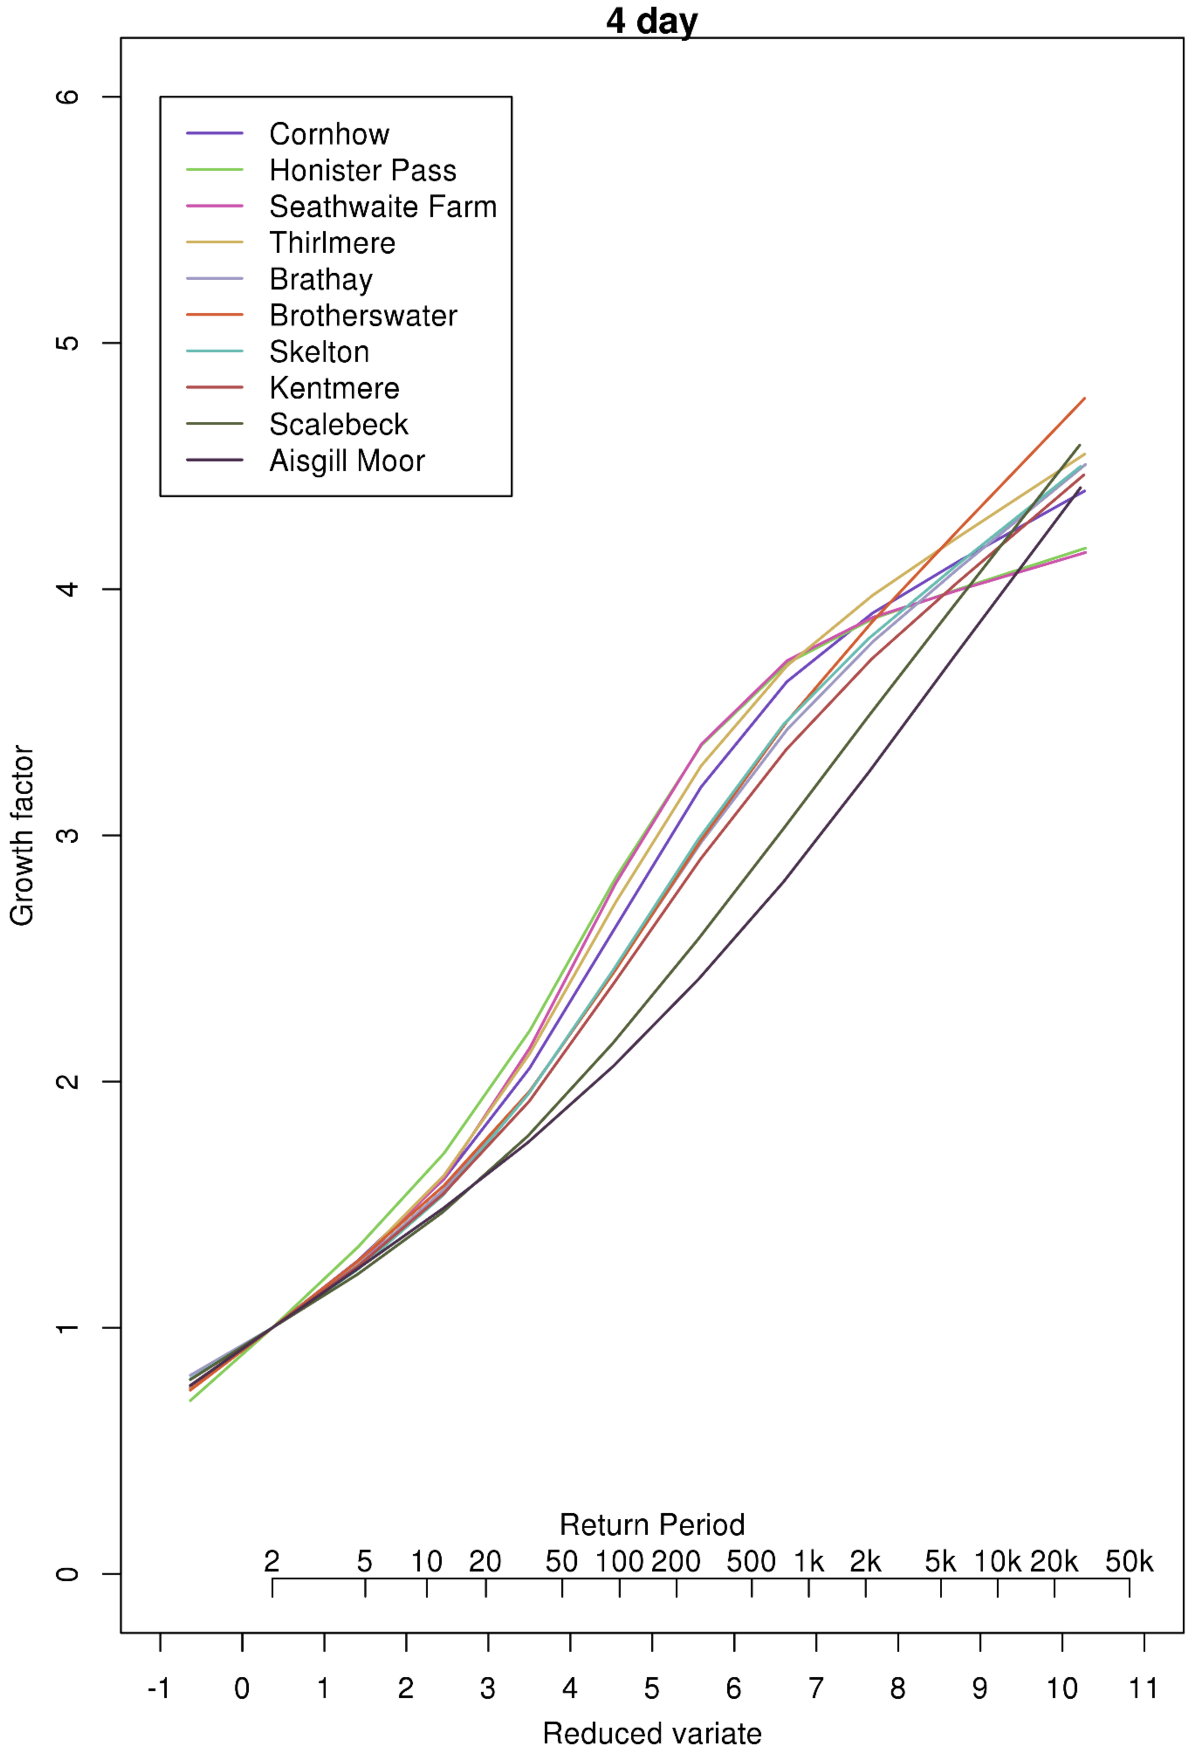


Figure 13 – comparison of FORGEX lines produced for case study sites (4 day)

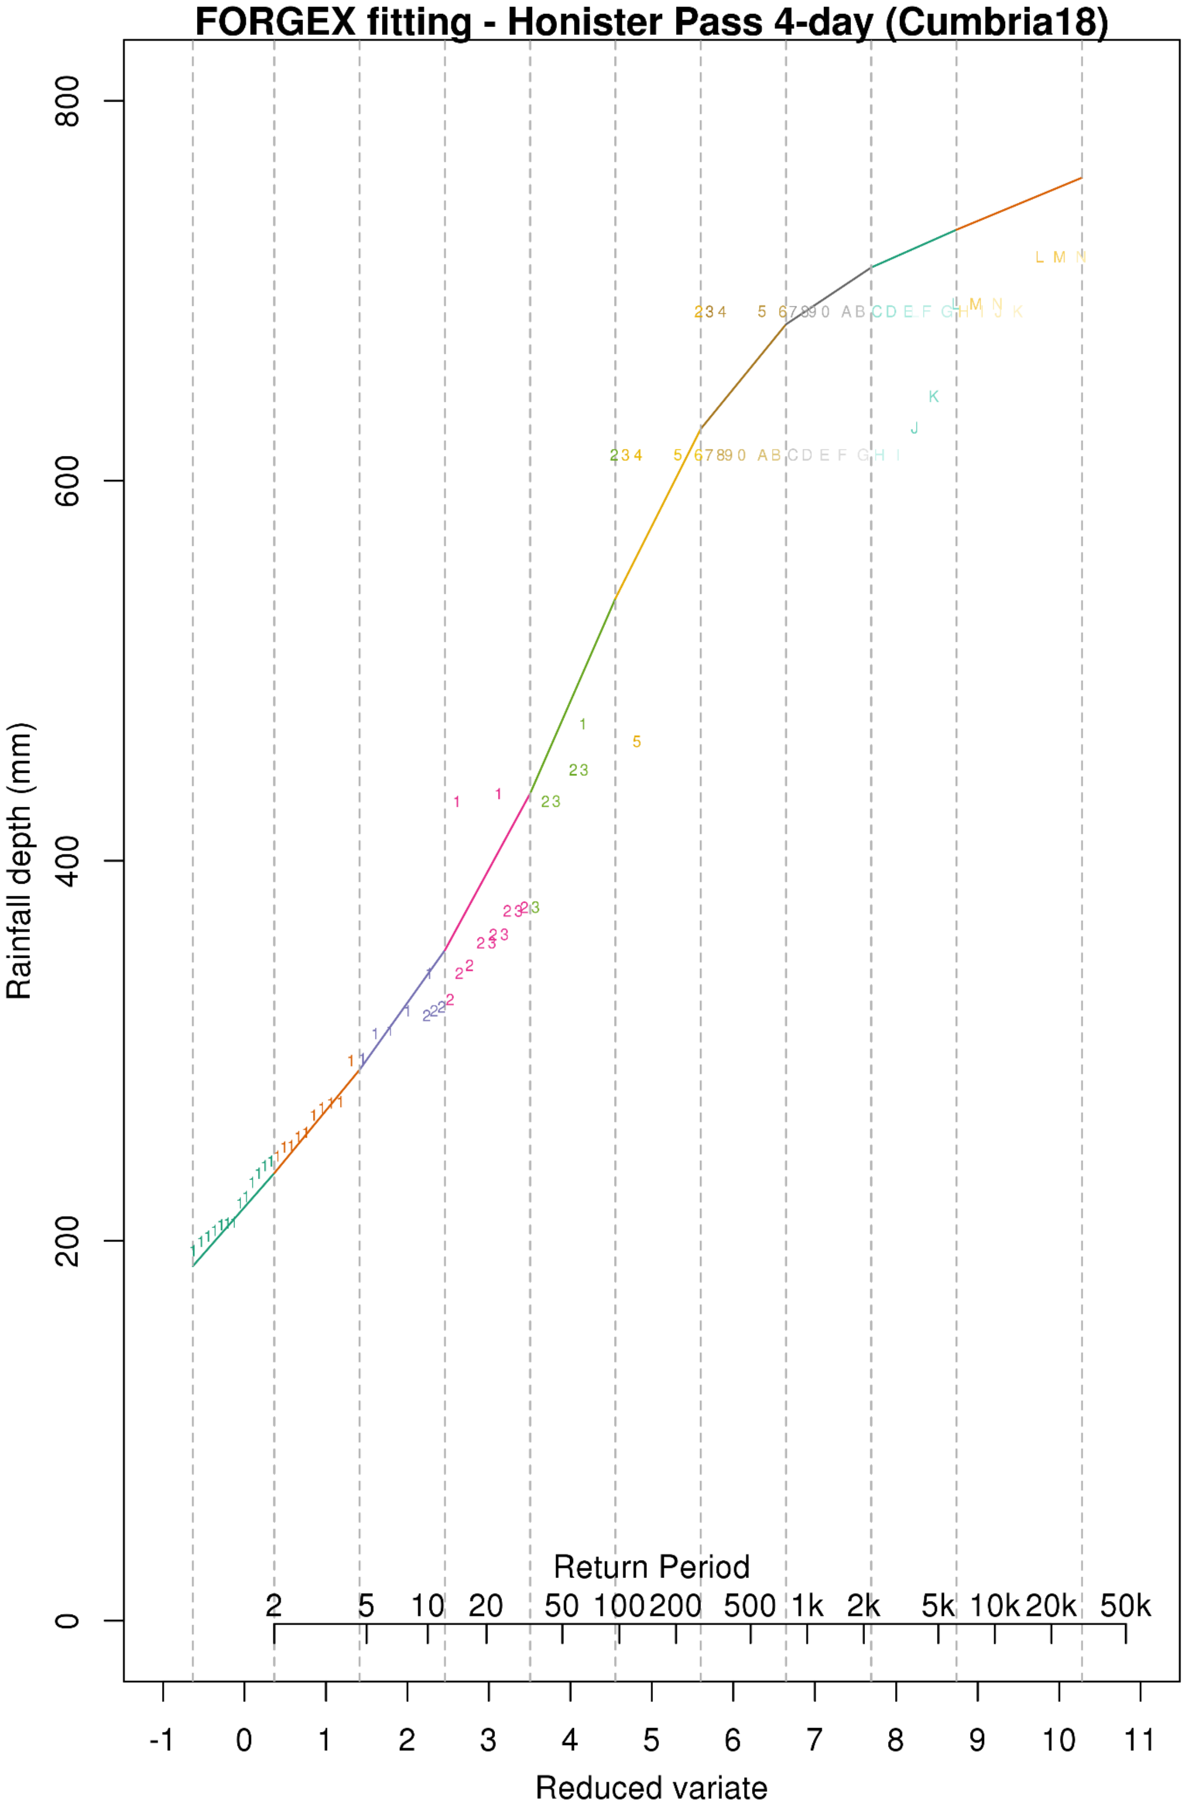


Figure 14 – discretised FORGEX line for Honister Pass (4 day) using new data

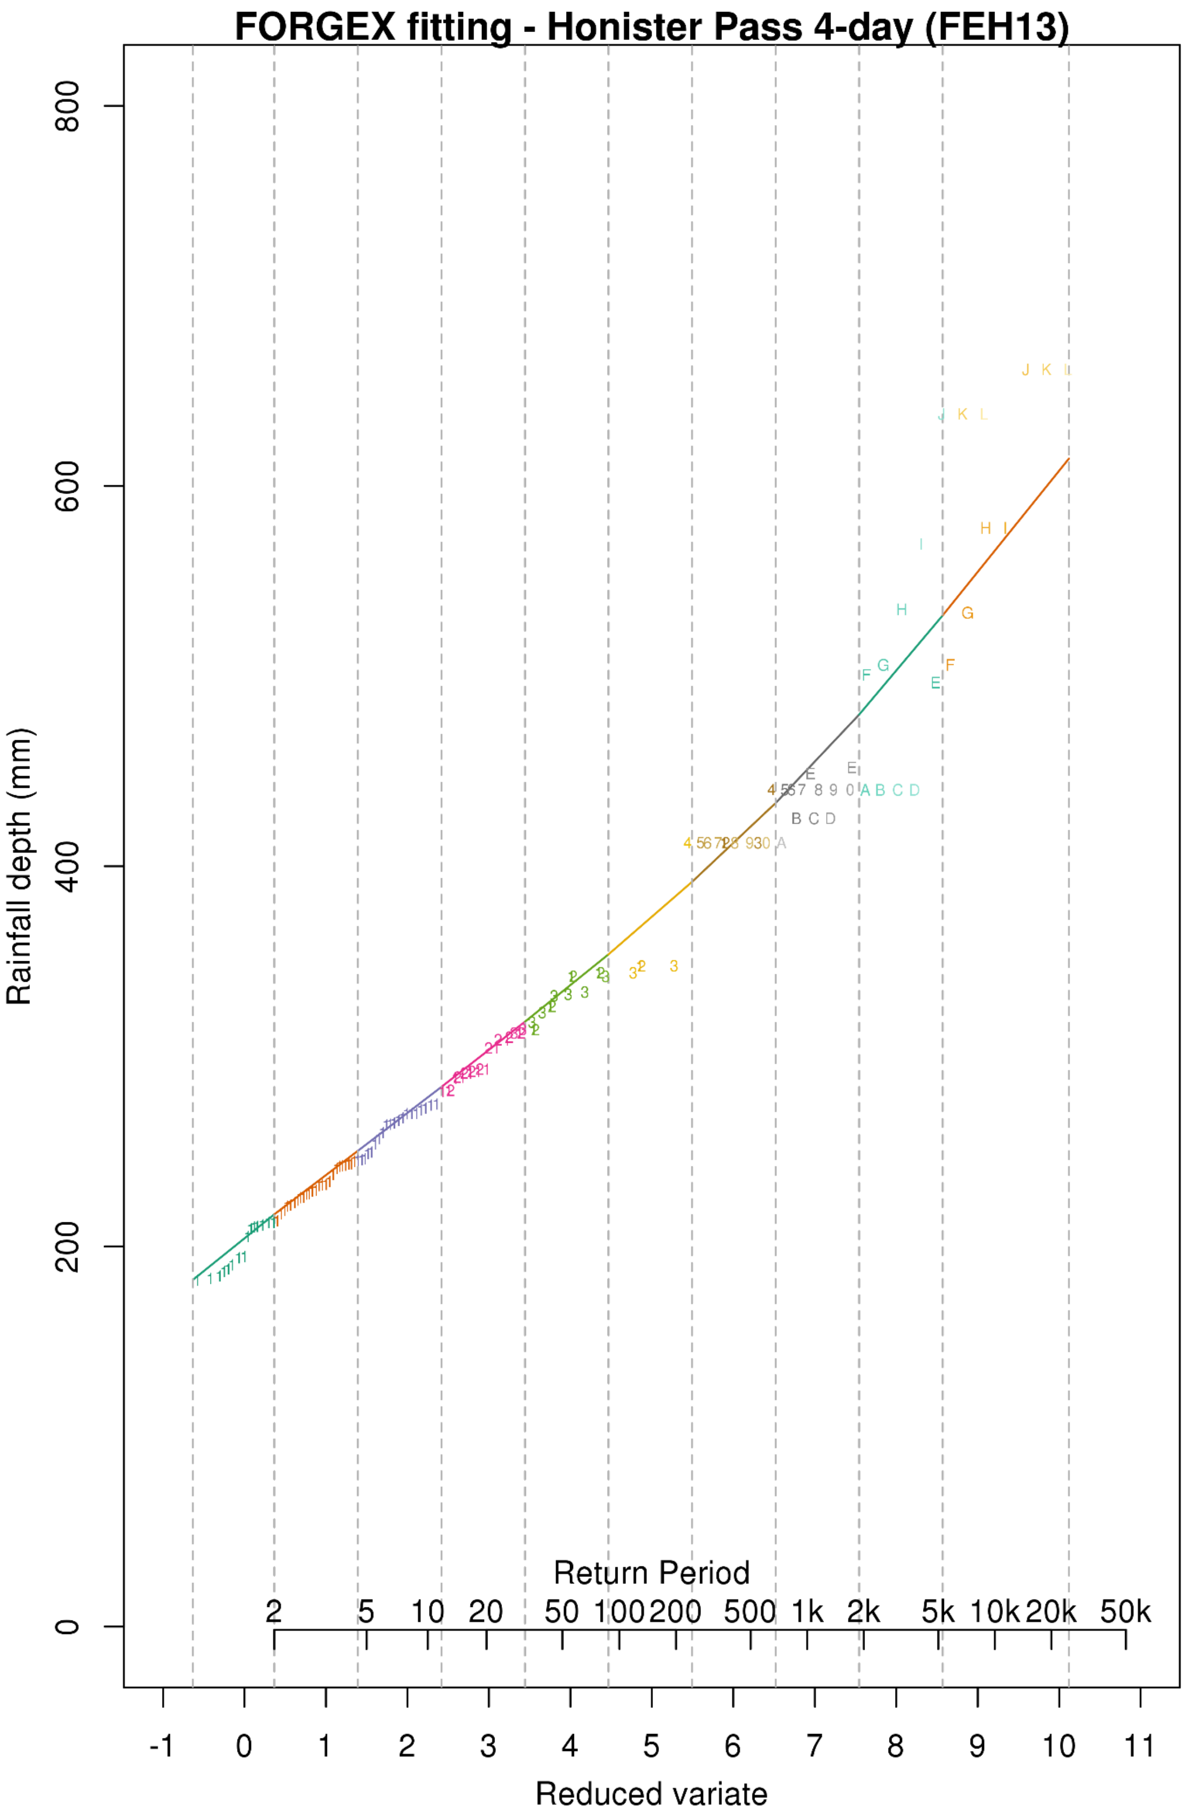


Figure 15 – discretised FORGEX line for Honister Pass (4 day) using FEH13 data

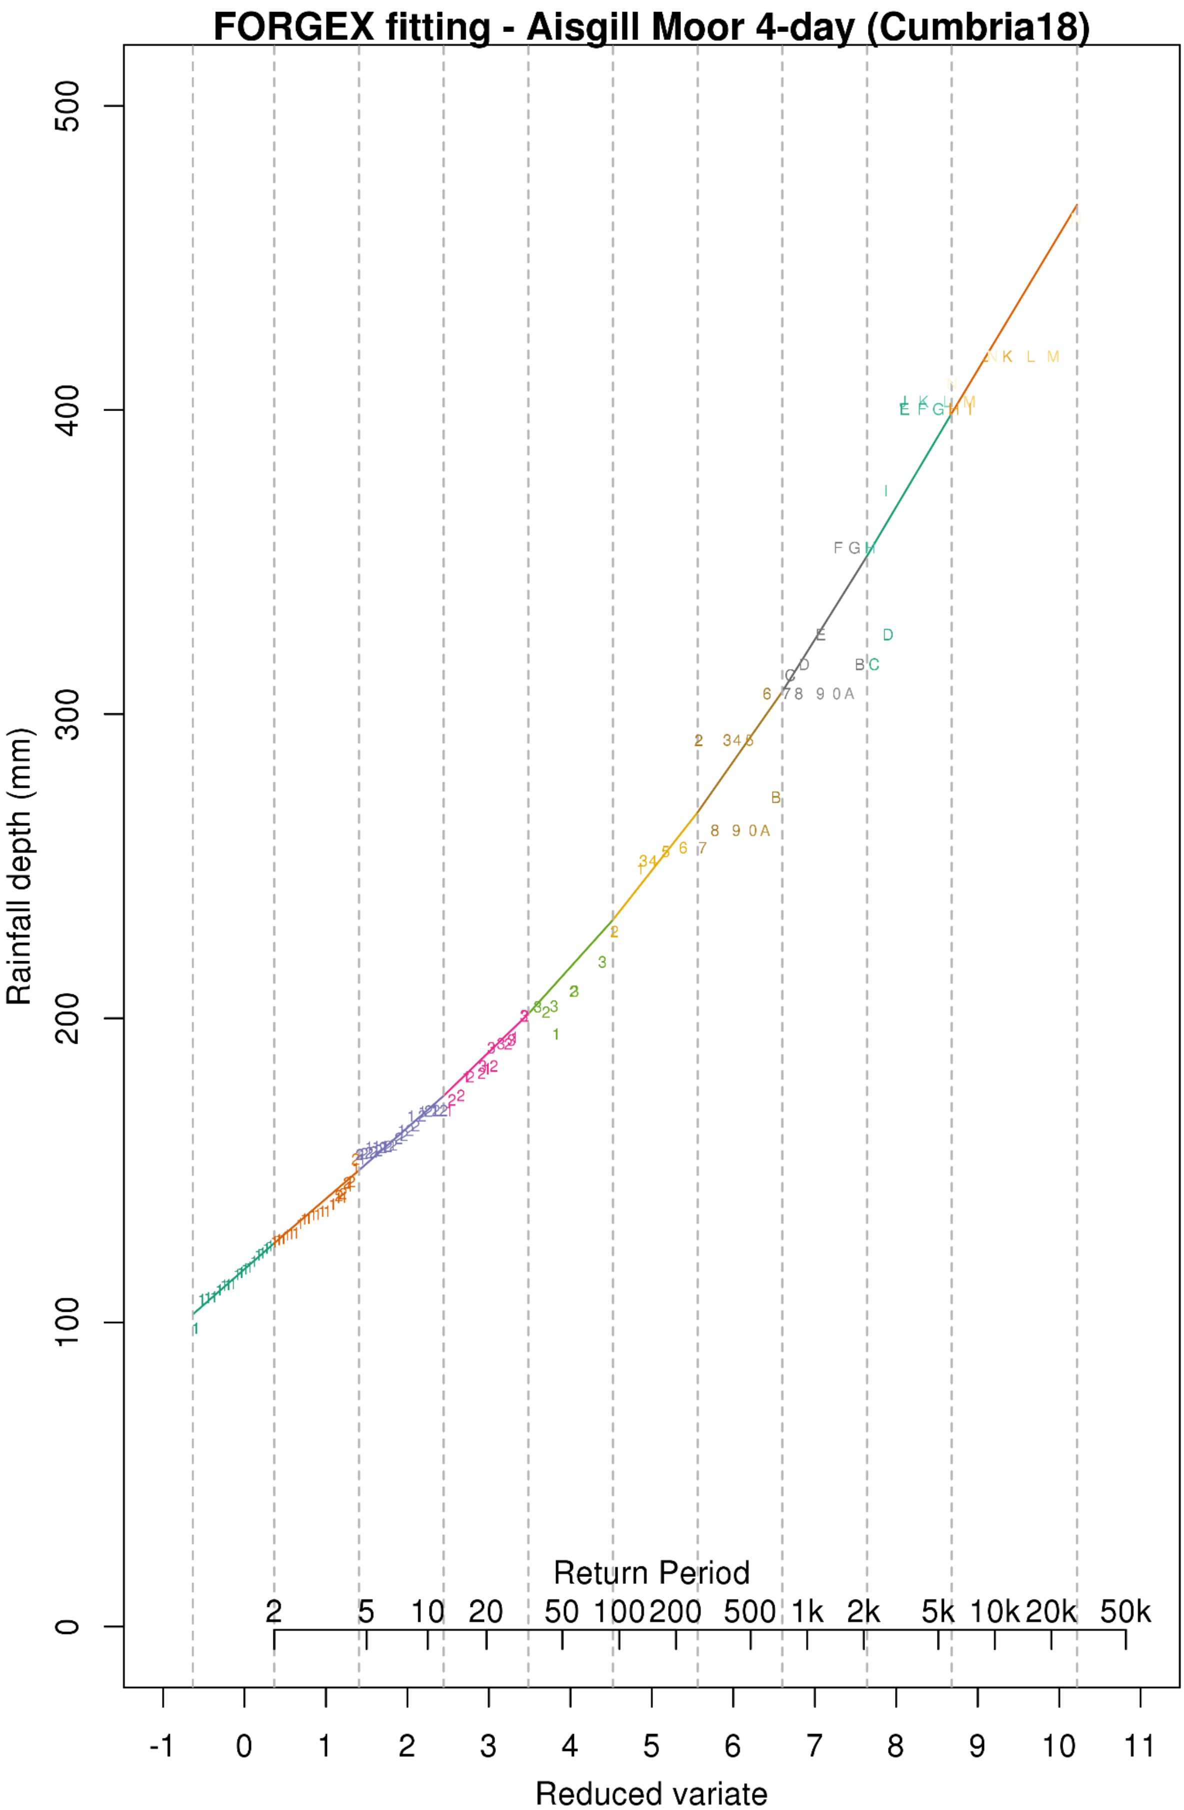


Figure 16 – discretised FORGEX line for Aisgill Moor (4 day) using new data

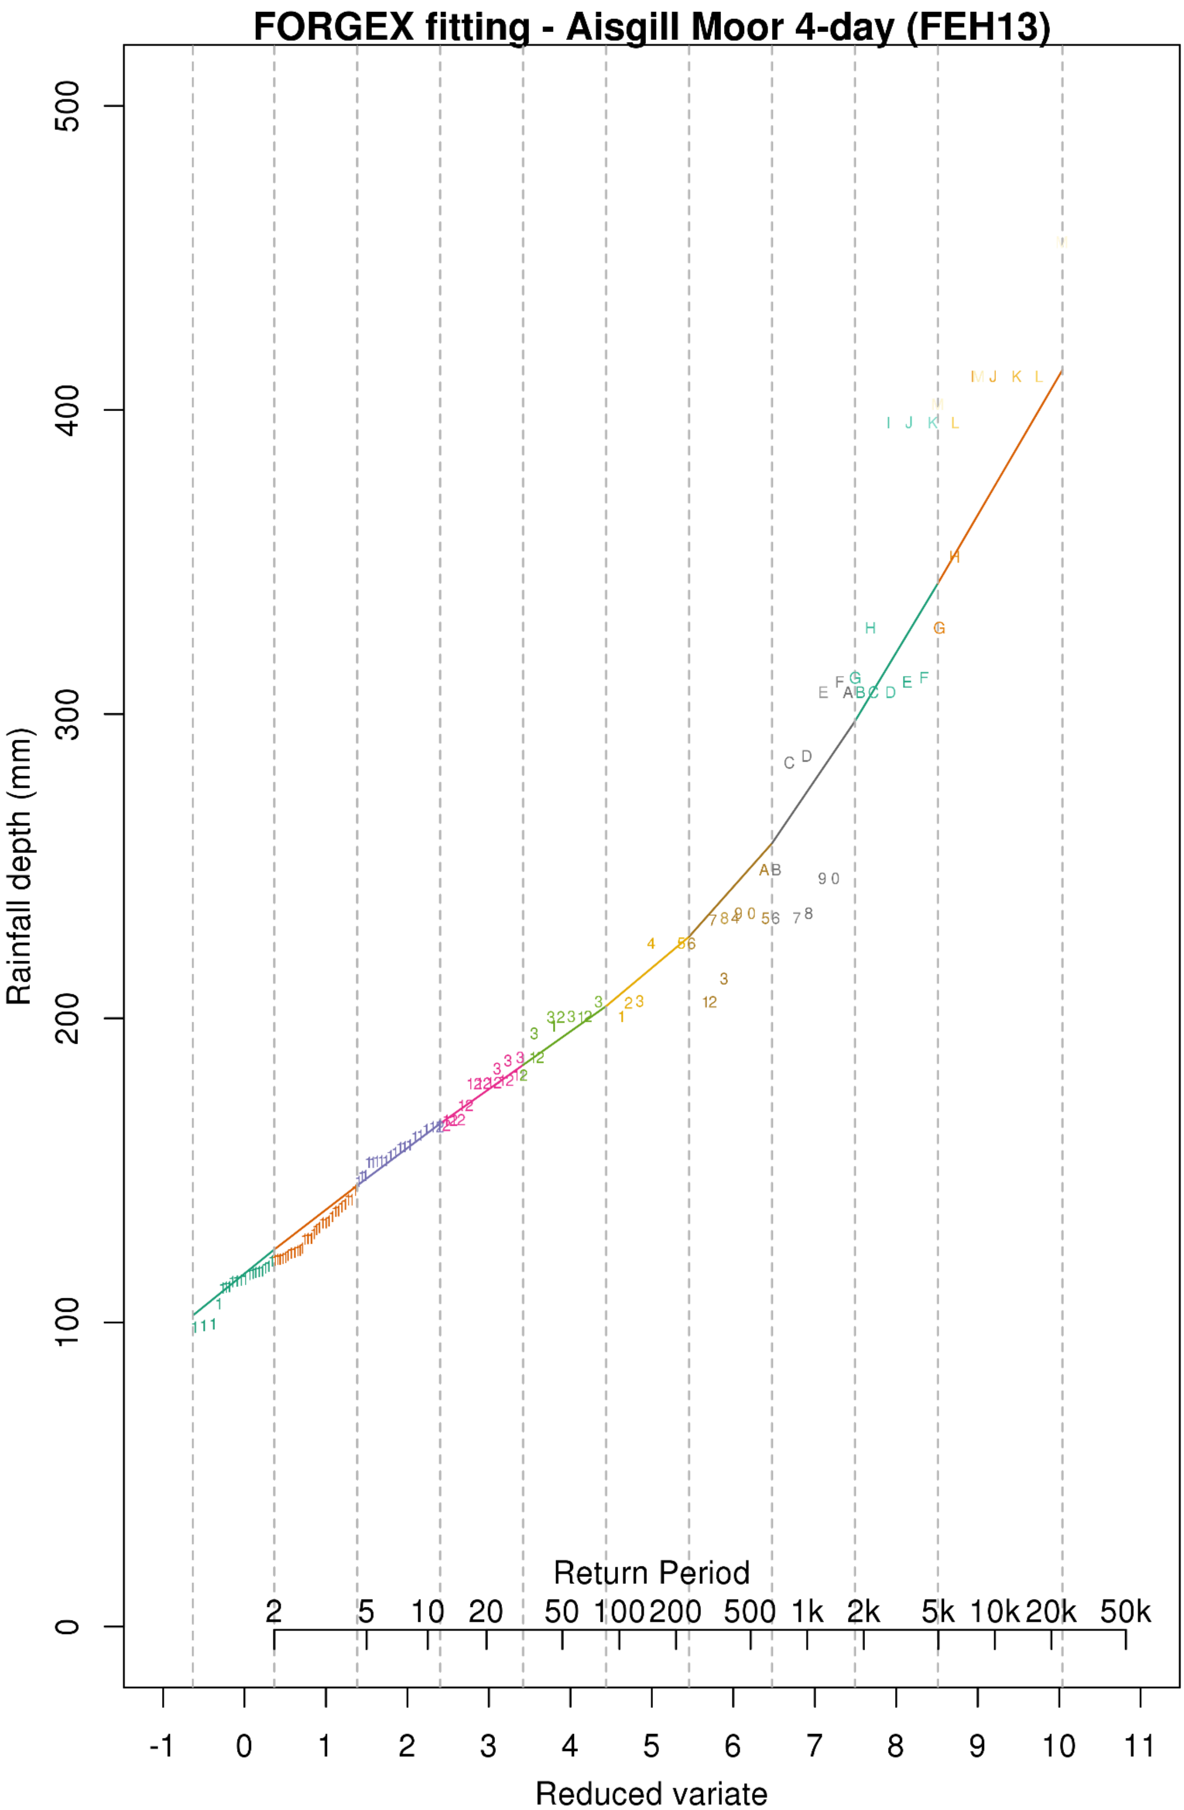


Figure 17 – discretised FORGEX line for Aisgill Moor (4 day) using FEH13 data

It was seen from Figure 10 to Figure 17 that growth curves for shorter durations are generally arc-shaped, whereas some could be S-shaped for longer durations of 2-4 days. This appears to be related to the intensity properties of the storms causing the S-shaped growth curves – as long-duration, frontal events, there were no significant peaks that could have had the same S-shaping effect on the FORGEX plots for shorter durations at the same sites. However, Figure 12 clearly shows some S-shaped growth curves for pooled daily data, while Figure 11 shows only arc-shaped curves for pooled 24-hour accumulated hourly data at the same sites. The difference in pooled growth curve shape between 24-hour and 1-day duration events at certain locations is probably more related to the relative densities of hourly-recording and daily-recording rain gauges: there are more than seven times as many of the latter, giving them greater opportunity to capture the spatial centres of extreme events.

# 7 FEH13 DDF model

Following the creation of rainfall growth curves for specified rainfall durations in FORGEX, a depth-duration-frequency (DDF) model is fitted. This serves several purposes:

- It ensures that rainfall depths for any duration always increase with increasing return period, and that rainfall depths for any return period always increase with increasing duration,
- It allows interpolation, so that rainfall depths can be estimated at durations in between those for which rainfall maxima were extracted and rainfall growth curves produced,
- It can allow extrapolation to longer and shorter durations than plotted in FORGEX, and extrapolation to longer return periods than those of the rainfall network maxima data. It must be noted that extrapolations should only ever be treated as indicative.

## 7.1 DDF model structure

The FEH13 DDF model is based on a weighted total of two duration-dependent Gamma distributions raised to a power. This takes the form:

$$F(z; D) = (p_1\Gamma(z; \alpha_1(D), \beta_1(D)) + (1 - p_1)\Gamma(z; \alpha_2(D), \beta_2(D)))^v \quad (6.1)$$

$$\alpha(D) = \lambda_0 + \lambda_1 D \quad (6.2)$$

$$\beta(D) = \gamma_1 D + \gamma_2(1 - 1/(1 + \gamma_3 D)) \quad (6.3)$$

$$z = 100x/RMED_{24h} \quad (6.4)$$

Where $\Gamma(z; \alpha(D), \beta(D))$ are two independently-parameterised Gamma distributions with scale and shape parameters $\alpha$ and $\beta$ respectively, $p_1$ is a weighting factor in the range 0 to 1, and $v$ is an exponent, included to allow $F(z; D)$ to take forms similar to a Generalised Extreme Value (GEV) distribution. This is considered an important requirement of the DDF model, given that block maxima from single series will tend towards the GEV distribution as the series increases in length.

$z$ is a partly-standardised variable based on rainfall depth, $x$, whose purpose is to give the 24-hour *RMED* a nominal value of 100 at every point in the UK. The DDF model has 12 parameters, however only 11 are ‘free’ due to the use of $z$ rather than $x$. Specifically, $v$ is derived from the 11 other parameters according to:

$$v = \ln(0.5) / \ln(p_1\Gamma(100; \alpha_1(24), \beta_1(24)) + (1 - p_1)\Gamma(100; \alpha_2(24), \beta_2(24))) \quad (6.5)$$

## 7.2 DDF model fitting

The input data to the DDF fitting program consist of the fully-sliding rainfall depths for all 11 durations plotted in FORGEX, discretised from the FORGEX lines at intervals of 0.2 on the Gumbel reduced variate scale, starting at -0.6 and continuing up to, but not beyond, the Gumbel reduced variate corresponding to the return period of the rarest event contributing to each line. An extra point, corresponding to a depth of fully-sliding *RMED* at a Gumbel reduced variate of 0.3665 (2 years) is also included – this is the only point that the DDF model is obliged to match exactly.

Fitting to the points is performed via a Nelder-Mead simplex routine, based on a version of Algorithm AS 47 (O'Neill, 1971) incorporating subsequent remarks AS R11 (Chambers & Ertel, 1974), AS R15 (Benyon, 1976) and AS R28 (Hill, 1978), and further modified to more closely match the Polytope algorithm (Gill *et al.*, 1981). Both 24-hour and 1-day rainfall data contribute to the 24-hour duration in the DDF model.

The DDF model outputs two data types: fitted parameter values and rainfall depths rounded to the nearest 0.1 mm for all 312 combinations of 13 durations and 24 return periods. The 13 durations are comprised of the 10 durations for which FORGEX lines were produced (with 24-hour and 1-day lines combined), plus 0.4, 0.5 and 240 hours. Extrapolation to shorter and longer durations is possible because rainfall depths at each site are defined purely by the DDF model and its parameters, some of which are duration-dependent. The 24 return periods are set at values from 1.3 to 500,000 years on the annual maximum scale and, over the range from 10 to 100,000 years, are approximately geometrically spaced, so that each value is on average 77.8% larger than the last, equal to a ratio of 10:1 between any value and the value four places before or after it when they are ordered. Extrapolated rainfall depths for return periods beyond the end of the DDF input data (i.e. adjusted FORGEX lines) are produced by fitting the DDF model to the available data, then extending the fitted DDF lines without changing gradient on a plot of rainfall depth vs Gumbel reduced variate. This ensures that rainfall estimates do not grow exponentially outside of the calibrated zone, as they did in the FEH99 model.

The DDF model was fitted three times at each point where FORGEX outputs were produced, first with generic parameter starting values, then twice more starting with the parameter values that resulted from the previous model fitting.

## 7.3 DDF model smoothing

The output rainfall depths from the third DDF fitting are smoothed to promote spatial consistency between nearby points. All smoothing is performed on re-standardised rainfalls to avoid over-trimming extreme observed rainfall depths in the final results. The smoothing procedure for a single point is as follows:

1. Identify all other points within a 10 km radius of that single point that have associated rainfall data. Assuming a regular 1-km grid, 316 points will be identified, excluding the central one.
2. Assign each point a weight according to $1 / (1 + d)^2$, where $d$ is straight-line distance to the centre of the radius.
3. Re-standardise all mm rainfall depths into growth factors, according to *RMED*, *SAAR* and, for daily data, northing.
4. Discard the top 3% of growth factors, corresponding to the largest 10 growth factors.
5. Calculate one-over-the the largest remaining growth factor and discard all growth factors smaller than this, up to a maximum of 30% of the original total number of points (95 points). If more than 30% of growth factors are smaller than this value, only discard the bottom 30% of growth factors (i.e. 95 values).

6. Calculate a weighted average of the non-discarded growth factors.
7. Unstandardise the weighted average.
8. Check that the fractional rise between consecutive smoothed rainfall depths at 24-hour duration is no more than five times smaller or larger than before smoothing. Constrain any that are.
9. Check the same, working down durations from 18 hours to 1 hour. Simultaneously check that the fractional rise between consecutive durations for the same return period is no more than three times smaller or larger than before smoothing, and constrain any that are.
10. Check as before, working up durations from 48 to 192 hours.
11. Set the 0.4- and 0.5-hour rainfalls to be equal to their unsmoothed values multiplied by the ratio between smoothed and unsmoothed 1-hour rainfall at each return period, subject to the same constraints applied to all other sub-daily durations.
12. Set the 240-hour rainfalls to be equal to their unsmoothed values multiplied by the ratio between smoothed and unsmoothed 192-hour rainfall at each return period, subject to the same constraints applied to other multi-day durations, with the following exception: rainfall depths between consecutive return periods are permitted to grow (but not shrink) by more than five times as much as in the unsmoothed data.

Regarding point 11, it is noted that the FEH Web Service does not derive smoothed sub-hourly durations in this way. Instead, these are produced by multiplying the smoothed 1-hour rainfalls by the ratio between sub-hourly and 1-hour FSR rainfall estimates for each return period. Furthermore, FSR-based smoothed grids were produced for durations of 5, 15, 30 and 45 minutes, not 24 and 30 minutes as here.

It is also noted that no special rules apply to the point at the centre of the radius: if the point to be smoothed is within either the discarded top or bottom outliers in terms of growth factor, it is discarded as if it were any other point and the smoothing continues as normal without it.

As the smoothing requires a 10-km radius of points around the point of interest, a 10-km ‘frame’ around the edge of the FORGEX/DDF results cannot be smoothed. This reduces the area for which final, smoothed DDF model results are produced to a rectangle defined by bottom-left and top-right corners of (270,450) and (389,589) respectively.

Although some steps of the smoothing procedure seem as if they could significantly alter the fitted DDF results, potentially to a point of divergence from the input FORGEX lines (e.g. discarding the at-site DDF results if they are outliers), there is generally minimal difference between smoothed rainfall depths and unsmoothed depths produced by the third DDF model run. This is shown in Figure 18, which expresses the percentage change due to smoothing for every combination of duration, return period and location in the rectangle over which smoothing was performed (3,763,620 values). While outlier combinations of location, duration and frequency do exist, at which smoothed rainfall depths may be as much as 118% or as little as 93% of the unsmoothed value, 86% of adjustments to rainfall depth are within the range ±1% and only 2097 out of 3,763,620 are outside the range ±5%. Just over one-eighth (12.6%) of depths are identical before and after smoothing.

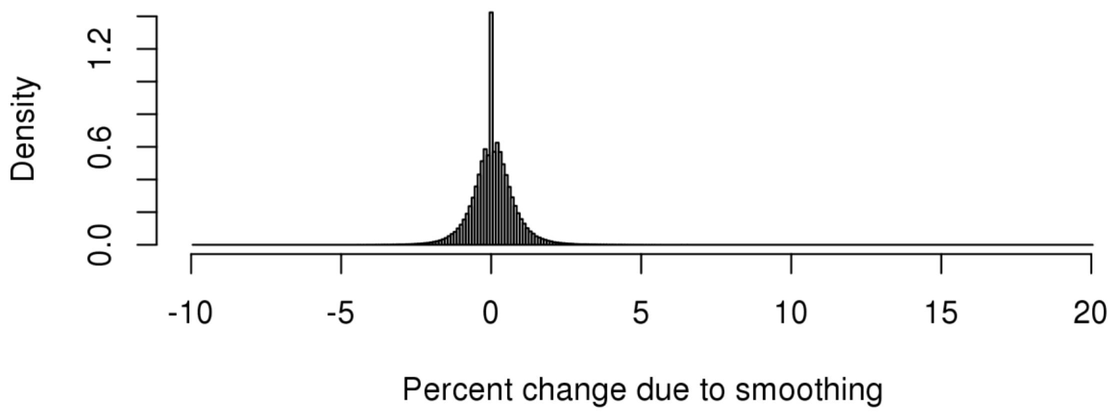


Figure 18 – percent change to fitted DDF model results introduced by smoothing

## 7.4 DDF plots

Figure 19 and Figure 20 compare fitted, unsmoothed DDF curves against adjusted FORGEX lines converted to fully-sliding depths, at Aisgill Moor and Honister Pass respectively. These two sites are chosen for comparison due to their contrasting arc- and S-shaped daily FORGEX lines. Colour-coding indicates which adjusted FORGEX lines contribute to which DDF curves: note that both the 24-hour and 1-day FORGEX lines contribute to the 24-hour DDF curve, but that no FORGEX lines contribute to the 0.4-, 0.5- or 240-hour DDF curve, as these are extrapolated.

In Figure 19, the DDF model is shown to fit well to hourly FORGEX lines for durations of 1 and 12 hours, and quite well for intermediate durations. The 18-hour FORGEX line, however, shows faster growth than the DDF curve for return periods longer than about ten years, and is very close to the 24-hour DDF curve for return periods above 50 years. This is because the 24-hour DDF curve must compromise between the 24-hour and 1-day FORGEX lines; it lies between these for return periods from about 100 to 10,000 years. Although the 18-hour line is higher than the 1-day line, the fact that event duration in hours features in the parameterisation of both parameters for both Gamma distributions on which the DDF model is based, to enforce consistency between durations, means that the 18-hour curve must move down to provide smooth changes in depth across the full range of durations.

Considering daily durations, it also becomes apparent that the model fits a very gentle S-shape where none is obvious in the input lines. This could be a compromise to take account for the parameter formulation (which includes duration) being unable to model the small difference between the 2-day and 4-day lines simultaneously with the relatively larger differences found between other lines. This is suggested by considering the relative spacing between the 1-, 2-, 4- and 8-day FORGEX lines and the same-duration DDF curves.

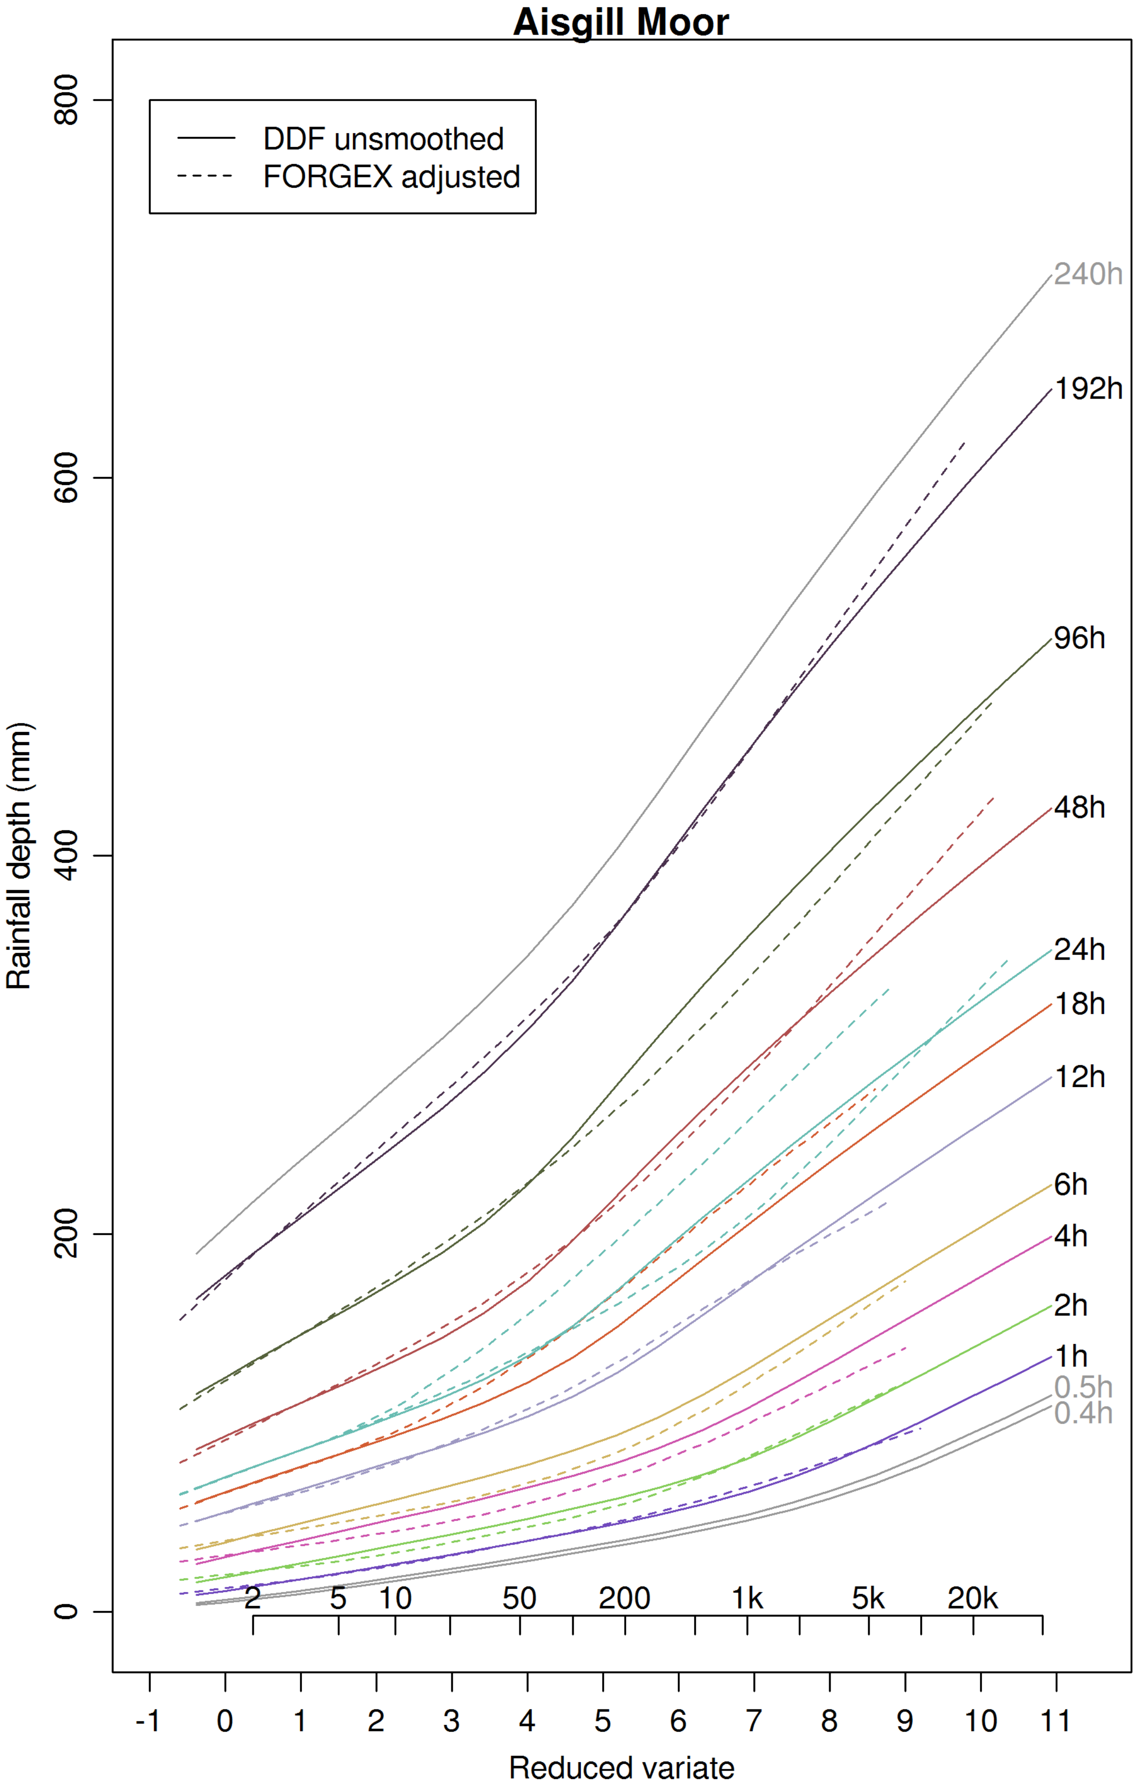


Figure 19 – comparison of unsmoothed DDF curves and FORGEX lines for Aisgill Moor

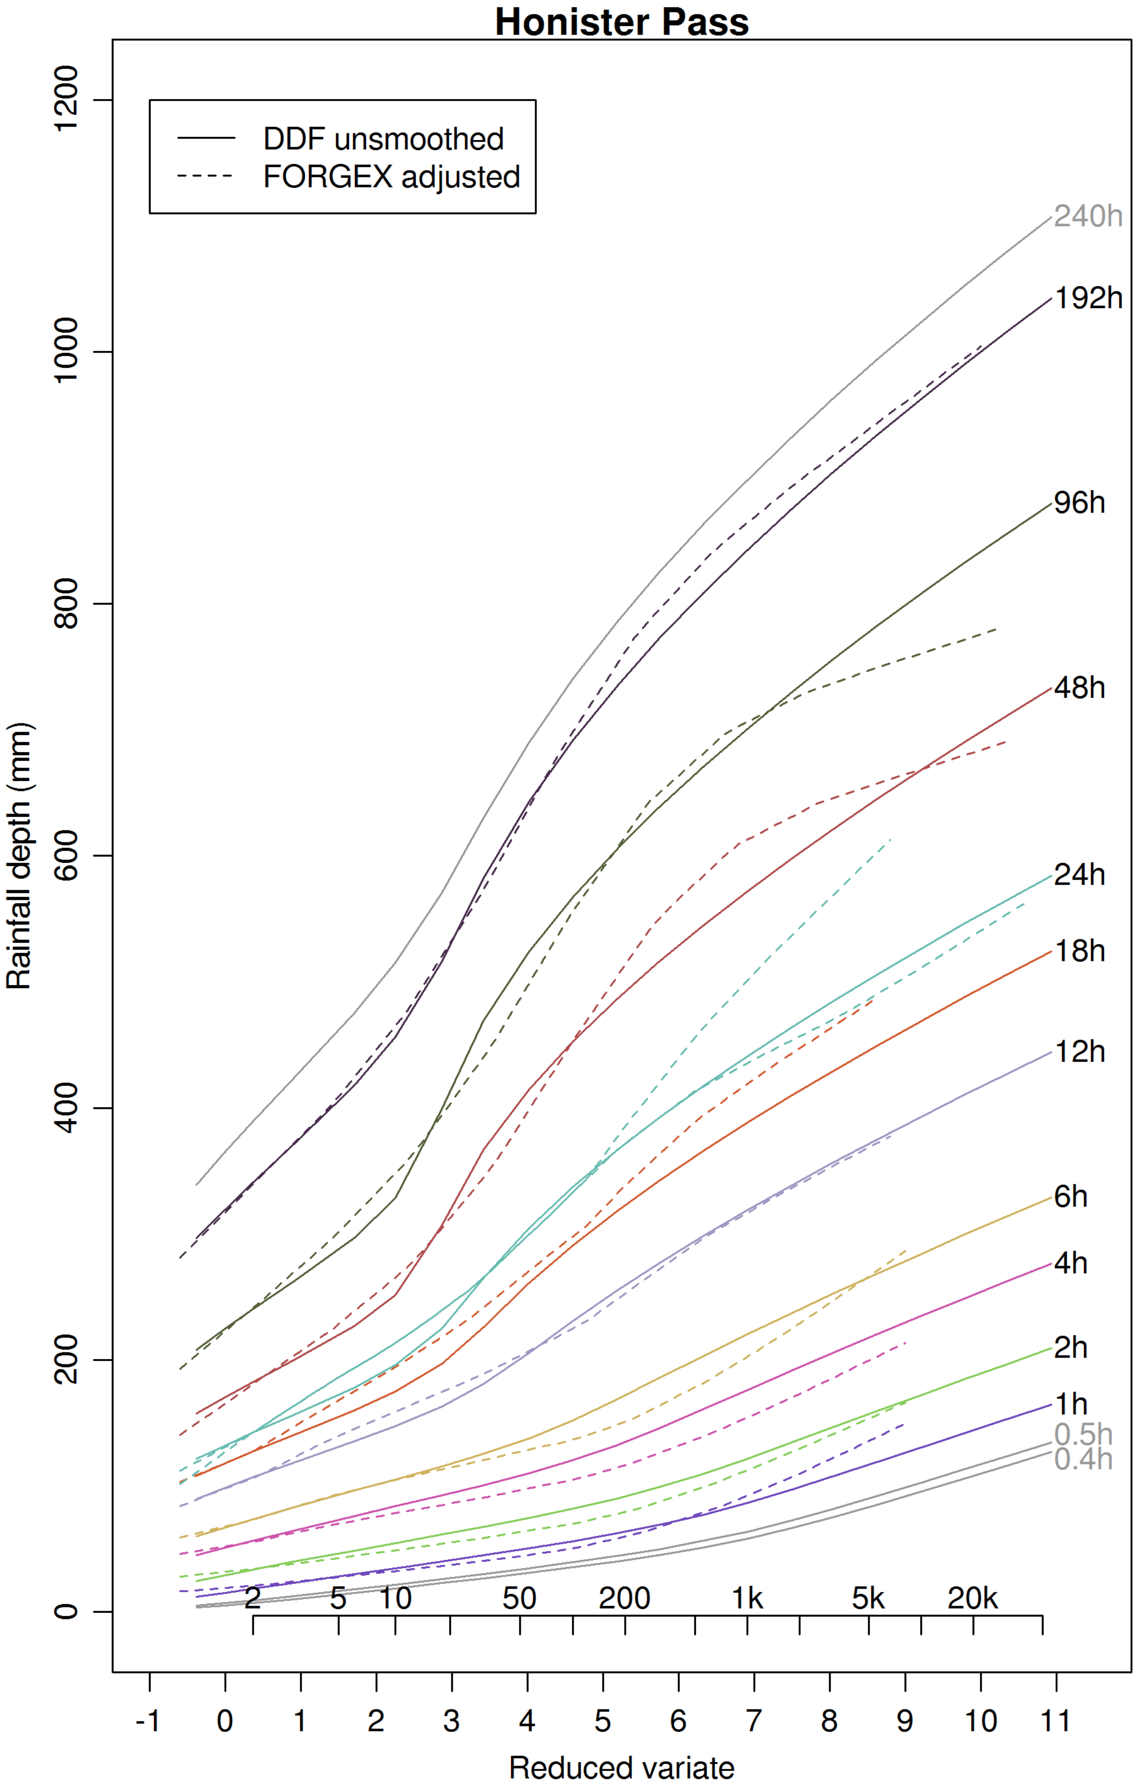


Figure 20 – comparison of unsmoothed DDF curves and FORGEX lines for Honister Pass

Considering Figure 20, there is again a close fit between DDF curves and FORGEX lines for hourly durations; there may even be a closer fit for both 18- and 24-hour durations than was observed in Figure 19. For daily durations, the very obvious S-shape most visible in the 2- and 4-day FORGEX lines is partly represented in the fitted DDF curves. The reason for this only partial representation is due to the use of a single model to unify and optimise to both arc-shaped and S-shaped input data. The fitted DDF curves become more S-shaped and less arc-shaped as duration increases; the potential to model this, where necessary, has been a design requirement of the model since the Reservoir Safety report, where the change in shape between shorter and longer durations was referred to as 're-curving behaviour'. However, both the 2- and 4-day lines are more S-shaped than the longer duration 8-day lines, hence the incomplete fitting to them. In fact, this incomplete fitting solves a potential problem that could have resulted from the 2- and 4-day lines flattening off at long return periods. The plot also illustrates the continuation in linear rainfall growth against reduced variate that is how the model extrapolates beyond the end of the FORGEX lines. At shorter return periods, more of an S-shape is imposed on the 1-day DDF curve than can be justified, resulting in this curve falling below the 1-day, 24-hour and, briefly, 18-hour FORGEX lines. This is again due to the requirement to unify both arc-shaped and S-shaped lines, for the S-shaping of the DDF curves to always increase with duration, and for the spacing between adjacent curves to be consistent with the change in duration between them. This gradual transition between S-shaped and arc-shaped curves persists into sub-daily durations.

Comparing Figure 19 and Figure 20, both figures show the DDF estimating larger 4- and 6-hour rainfalls than suggested by FORGEX, as well as smaller 18-hour rainfalls and a larger gap between 2- and 4-day rainfalls. The spacing between DDF curves for different durations is controlled within the DDF model by the duration-dependent components of the scale and shape parameters of the two Gamma distributions on which the model is based. This parameter structure enforces consistency between durations along the full scale produced by the model and prevents sudden changes in the DDF relationships. This may be a defence against unrepresentative periods of record with respect to specific durations. However, if an apparently inconsistent relationship between durations, such as the relative closeness of the 2-day and 4-day rainfall estimates for given return periods, is in fact a feature of the rainfalls and not due to sampling error, then the form of the model, or of the model parameters, may need to be reinvestigated in the future to allow more flexibility.

Figure 21 and Figure 22 demonstrate the effect of the smoothing procedure by comparing DDF model results before and after smoothing for Aisgill Moor and Honister Pass, respectively.

Throughout previous sections of this report, Aisgill Moor has been shown to be a fairly 'standard' site, in that its FORGEX lines follow a standard arc shape and the DDF model form is well suited to fitting the FORGEX data closely. This is because, unlike Honister Pass, for example, there has not been a recent series of large events to cause sudden growth at return periods in the hundreds of years. As a result of the 'standardness' of the rainfall characteristics at Aisgill Moor, smoothing has had almost no effect on the DDF model outputs: the smoothed and unsmoothed estimates are almost exactly concurrent for all plotted durations and return periods.

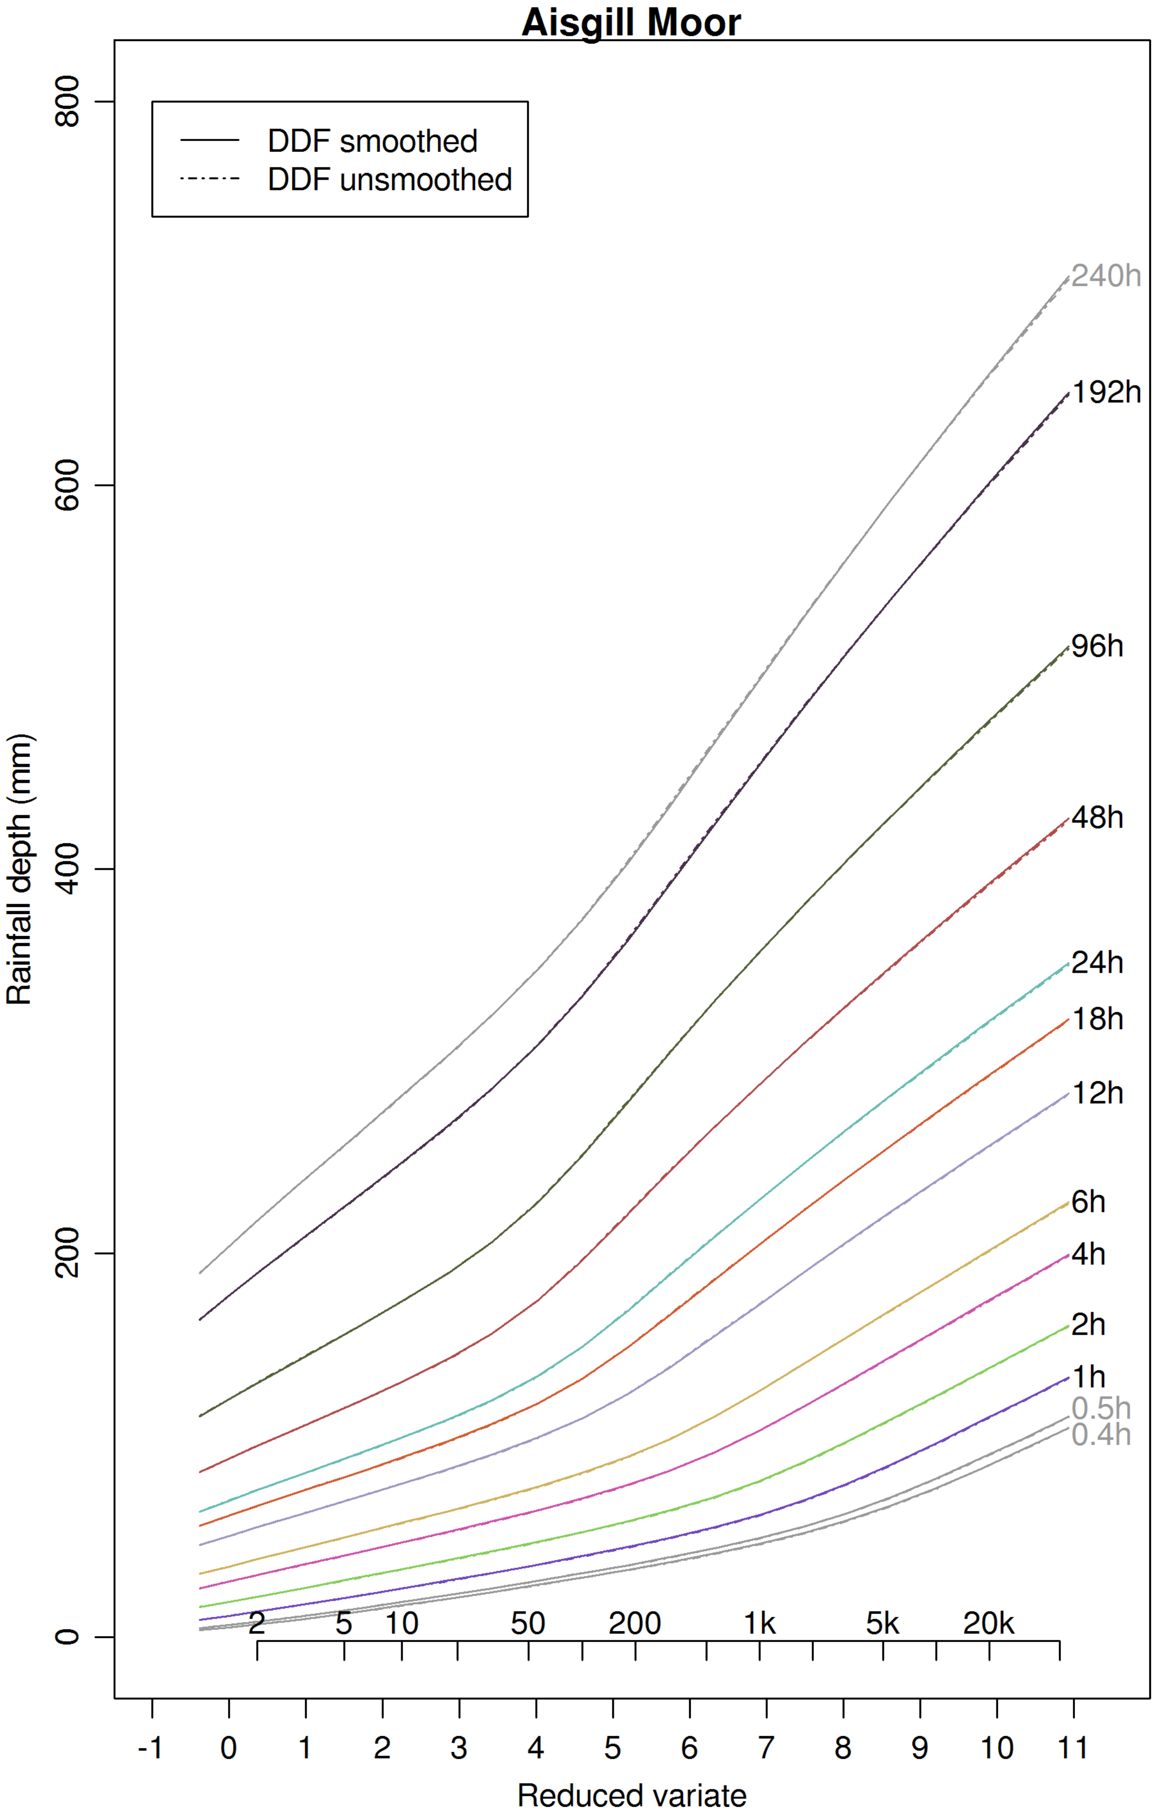


Figure 21 – Comparison of smoothed and unsmoothed DDF curves for Aisgill Moor

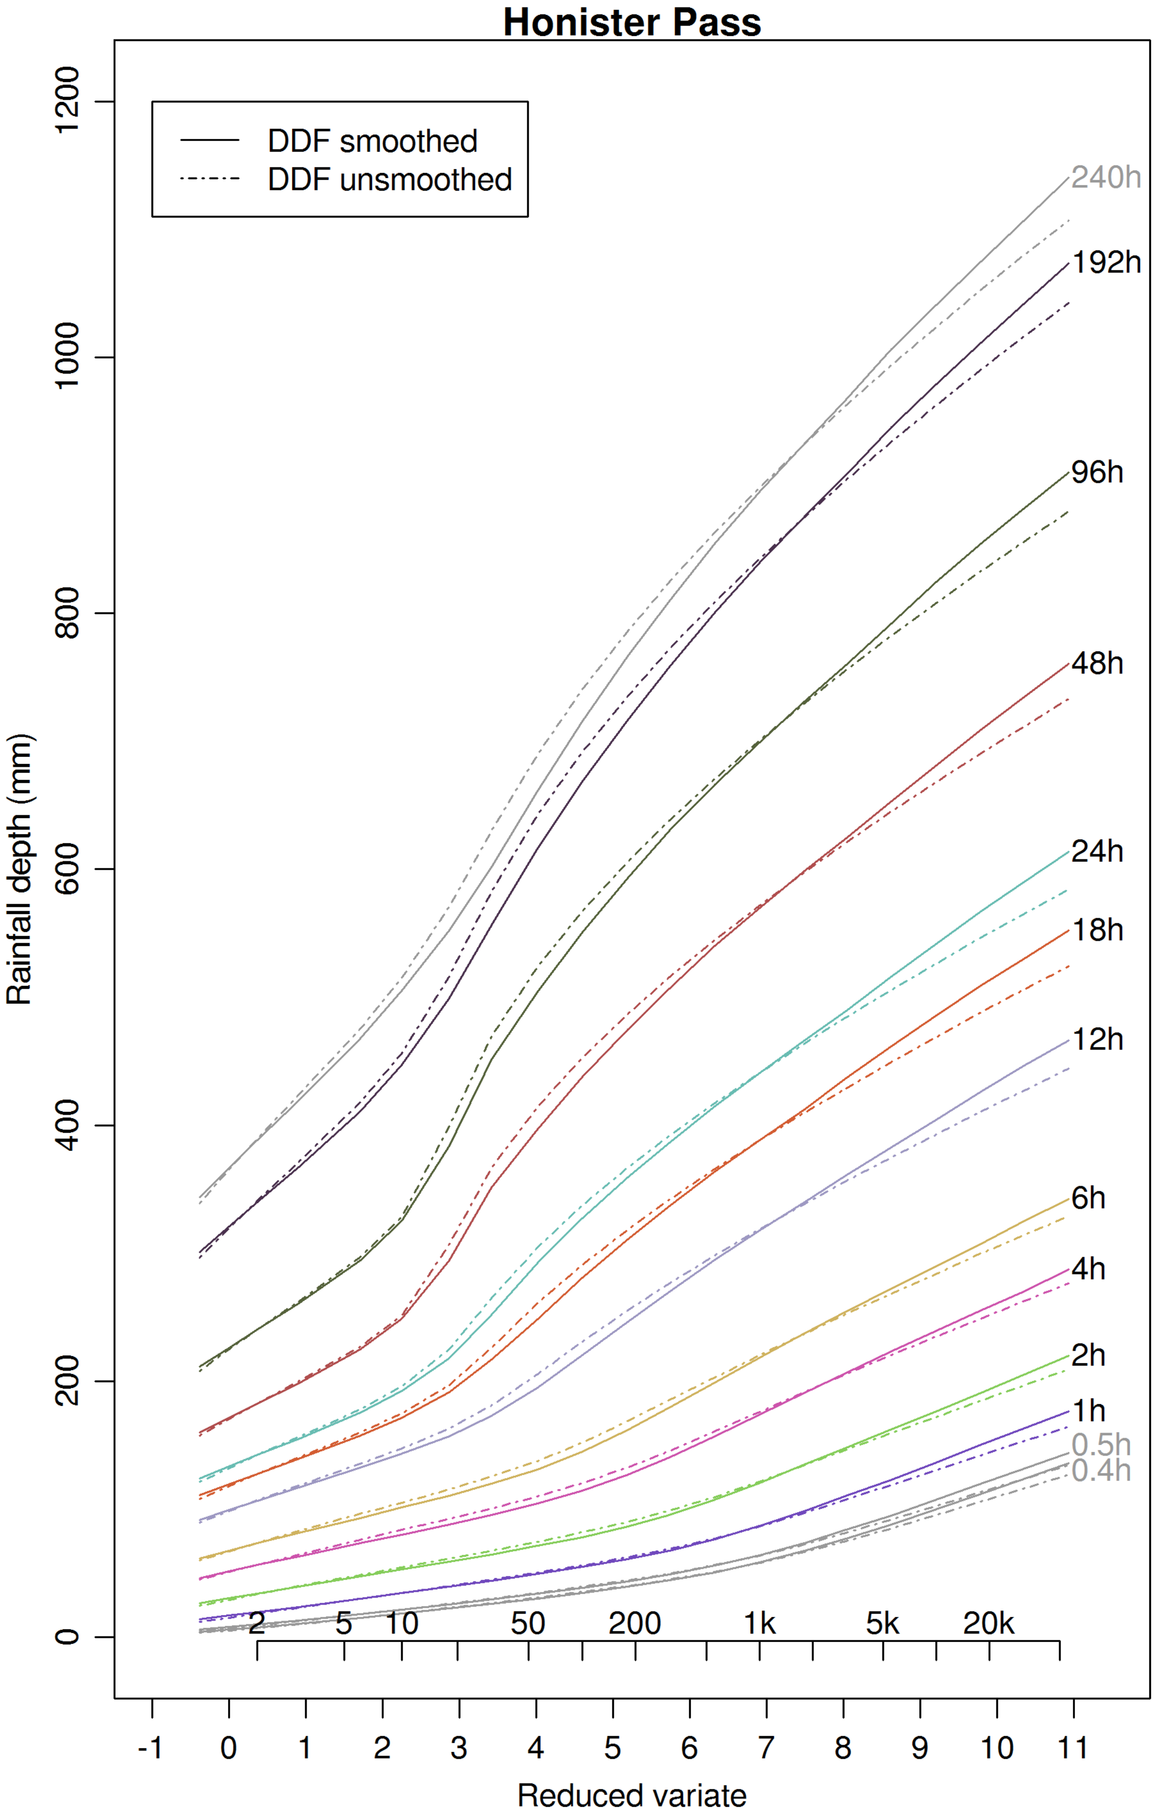


Figure 22 – Comparison of smoothed and unsmoothed DDF curves for Honister Pass

Conversely, Honister Pass is an example of a site where annual maxima data suggest S-shaped relationships between rainfall depth and frequency for some durations, resulting in the DDF model experiencing some difficulty in unifying both S-shaped and arc-shaped relationships into a single distribution. As a result, the DDF model attenuates two distinct shape profiles to give one average. The smoothing procedure continues this to a small extent by reducing the depths of multi-day rainfalls with return periods up to around 1000 years and increasing the depths of multi-day rainfall associated with longer return periods. The net effect of this is to further reduce the rapid growth and subsequent flattening found in the FORGEX lines, producing slightly straighter DDF curves. A similar effect is found for sub-daily rainfall durations, albeit reduced due to the less pronounced S-shaping. It is noted that the smoothing procedure results in increases of 17% and 11% to the 1-hour (and consequently 0.5- and 0.4-hour) rainfall depth, for return periods of 1.30 and 1.58 years respectively. While these are some of the largest proportional increases due to smoothing in Cumbria as a whole, the practical effect on rainfall frequency estimation is negligible for extreme events. Although the changes due to smoothing may be of interest for 1-in-12 month rates on the smallest of sites, where the critical duration is 1 hour or less, the actual quantities of rainfall are so small (15.3 vs 17.0 mm for 1-hour duration) that designing a project to cope with the higher, smoothed rainfall depth is unlikely to involve significant extra expenditure. Furthermore, the difference due to smoothing is reduced to 1% at 4 hours.

Despite the more significant changes brought by smoothing at Honister Pass in comparison to Aisgill Moor, it is important to note that the smoothed DDF curves at each site are highly dissimilar in shape, meaning that differences in the input data (i.e. network pooled annual maximum rainfalls) are preserved.

# 8 Results

## 8.1 Historical events

Figure 23 shows the return period of the maximum 36-hour total rainfall during the November 2009 event estimated by both the FEH13 and Cumbria 2018 models. The FEH13 estimates are similar, but not identical, to those shown in Figure 9 of Stewart *et al*. (2012), which was based on an earlier version of the FEH13 model than was initially made available on the FEH Web Service.

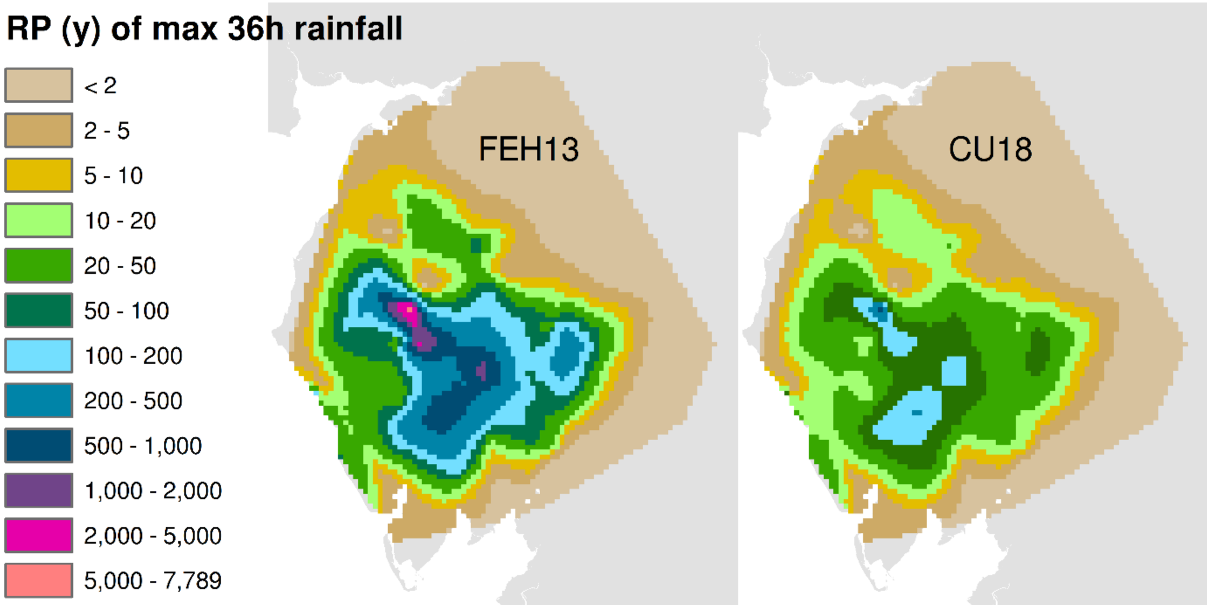


Figure 23 – Return period of maximum 36-hour total rainfall during November 2009 event estimated by FEH13 model (left) and Cumbria 2018 model (right)

The overall spatial patterns of return periods are very similar for both models, with the longest return period occurring in the same square kilometre in both cases. However, the return period of the 36-hour rainfall in that square kilometre is reduced from 7,789 to 502 years. While both values are extreme, one is almost 16 times lower than the other. The return periods around the periphery of the event are similar for both models; they only diverge for return periods greater than about 10 years. The differences in return period estimated by the FEH13 and Cumbria 2018 models are due almost entirely to the calibration data input to each model, as the DDF model structure was not modified, FORGEX was modified only very slightly and neither model assumes any trend in the rainfall data.

Table 7 compares return periods estimated by the CU18, FEH13 and FEH99 models for recent events in Cumbria. In all cases, the closest 1-km grid point is used, rather than the exact gauge location. Estimated discretisation factors of 1.002 and 1.052 are used for the 38-clock hour rainfall at Thirlmere and 3-clock day rainfall at Seathwaite Farm respectively.

Table 7 New (Cumbria 2018), FEH13 and FEH99 estimated return periods for extreme events occurring in Cumbria in 2009 and 2015

| Location | Date | Depth (mm) | Duration | Return period | | |
|---|---|---|---|---|---|---|
| | | | | Cumbria 2018 | FEH13 | FEH99 |
| **Honister Pass** | Dec 2015 | 341.4 | 24 clock hours | 131 | 988 | 1118 |
| **Thirlmere** | Dec 2015 | 405.0 | 38 clock hours | 8293 | >100k | 4017 |
| **Thirlmere** | Dec 2015 | 405.0 | 2 clock days | 7020 | >100k | 4751 |
| **Seathwaite Farm** | Nov 2009 | 316.4 | 24 clock hours | 150 | 980 | 160 |
| **Seathwaite Farm** | Nov 2009 | 392.6 | 36 clock hours | 192 | 2604 | 172 |
| **Seathwaite Farm** | Nov 2009 | 456.4 | 3 clock days | 132 | 3224 | 133 |
| **Seathwaite Farm** | Nov 2009 | 495.0 | 4 clock days | 113 | 2847 | 109 |

Table 7 shows that the probabilities of the 2009 and 2015 events in Cumbria are now far higher than they were using the FEH13 model. In particular, the return period of the record-breaking rainfall at Thirlmere is no longer ‘off the scale’ but a few thousand years. The return periods given by the new model and the FEH99 model show very close correspondence for the extreme event at Seathwaite Farm, and good correspondence for the event at Thirlmere. At Honister Pass, however, the FEH13 and FEH99 models give similar estimates for the 2009 event.

It is important to note that the return periods generated by this study are consistent with the standardised growth factors of the events: The Thirlmere event is estimated to be considerably rarer than either the Seathwaite or Honister Pass event, and its growth factor is over four. The Seathwaite and Honister Pass events both have growth factors of roughly two-and-a-half, and therefore have similar estimated return periods. It is also worth noting that the FEH99 estimated return period for the Thirlmere event is higher for the longer time period (two clock days) than the shorter time period (38 clock hours), even though the depth is the same in both cases. This is due to the high discretisation conversion factor used in the FEH99 model, which gives a fully-sliding depth of 450 mm over two clock days. The new model gives a fully-sliding depth of 426 mm over two clock days, which results in a slightly shorter return period for the longer duration than the shorter one.

Given that Cumbria has experienced two record-breaking rainfalls in the six year period from 2009 to 2015, it is definitely possible that the new return period estimates, which are more similar to those from the FEH99, are more plausible than the FEH13 estimates, which start around 1,000 years and extend beyond 100,000 years.

## 8.2 Fixed duration and return period events

Figure 24 to Figure 43 display final outputs for the 1-hour, 6-hour, 12-hour, 24-hour and 48-hour, 100- and 1000-year rainfall depths, and compare them to the FEH13 model outputs. For all durations and return periods, the greatest rainfall depths occur over the highest-altitude land. As expected, the smallest rainfall depths are strongly associated with low-lying land, more so as the rainfall duration increases. For 1-hour, events of both 100- and 1000-year return period, the smallest depths are associated with an arc from Cockermouth to Carlisle, although these are still large when compared to the UK as a whole. As duration increases, rainfall depths in this arc increase more rapidly than in areas where they are always high, leaving only the Eden Valley as an area with (relatively) low rainfall depths.

For the shortest duration (1 hour), depths associated with extreme rainfalls have decreased over considerable parts of the study area. This can be explained by the lack of notable short-duration events, like thunderstorms, gauged in and around Cumbria over the past decade. As a result, the main effect of the decade of new rainfall data is to increase the relative rareness of unremarkable events, as they are now the largest among a greater total number. Nevertheless, the 1-hour, 100-year rainfall depth is still increased by more than 5% over a region centred on Windermere, and the greatest reductions in 1-hour extreme events in this study region occur outside Cumbria: despite an area around Dumfries having relatively high 1-hour rainfall depths, these are frequently at least 15% smaller than estimated by the FEH13 model.

For 6-hour events, the 100- and 1000-year rainfall depths in Cumbria are largely unchanged but may be up to 10% higher. Figure 19 and Figure 20 show that the 6-hour DDF curve is fitted somewhat above the 6-hour FORGEX line at both sites. Hence, the increase in modelled 6-hour rainfall depths is (at least partly) not a result of equal increases in the observed 6-hour rainfall depths. As the DDF model has to unify rainfall maxima across durations from 1 to 192 hours, it could be suggested that the model's assumed relationship between 6-hour duration rainfalls and other-duration rainfalls is incorrect, and that there should be more flexibility to allow reduction in e.g. the ratio of 6-hour to 12-hour rainfalls for given return periods.

For 12-, 24- and 48-hour events, considerable increases in the 100- and 1000-year rainfalls are estimated for almost the whole of Cumbria. In all cases, the greatest proportional increases over the FEH13 model are centred on three locations: Honister Pass/Seathwaite Farm, Mosedale and Wet Sleddale. The reason that Thirlmere is not one of these locations, despite experiencing Cumbria's most extreme recorded event ever in terms of standardised growth factor, is due mainly to the relationship between pre-2007 and post-2007 rainfall depths recorded there. Similarly, Mosedale is not among the wettest places in Cumbria according to these results but does experience one of the greatest proportional increases over FEH13.

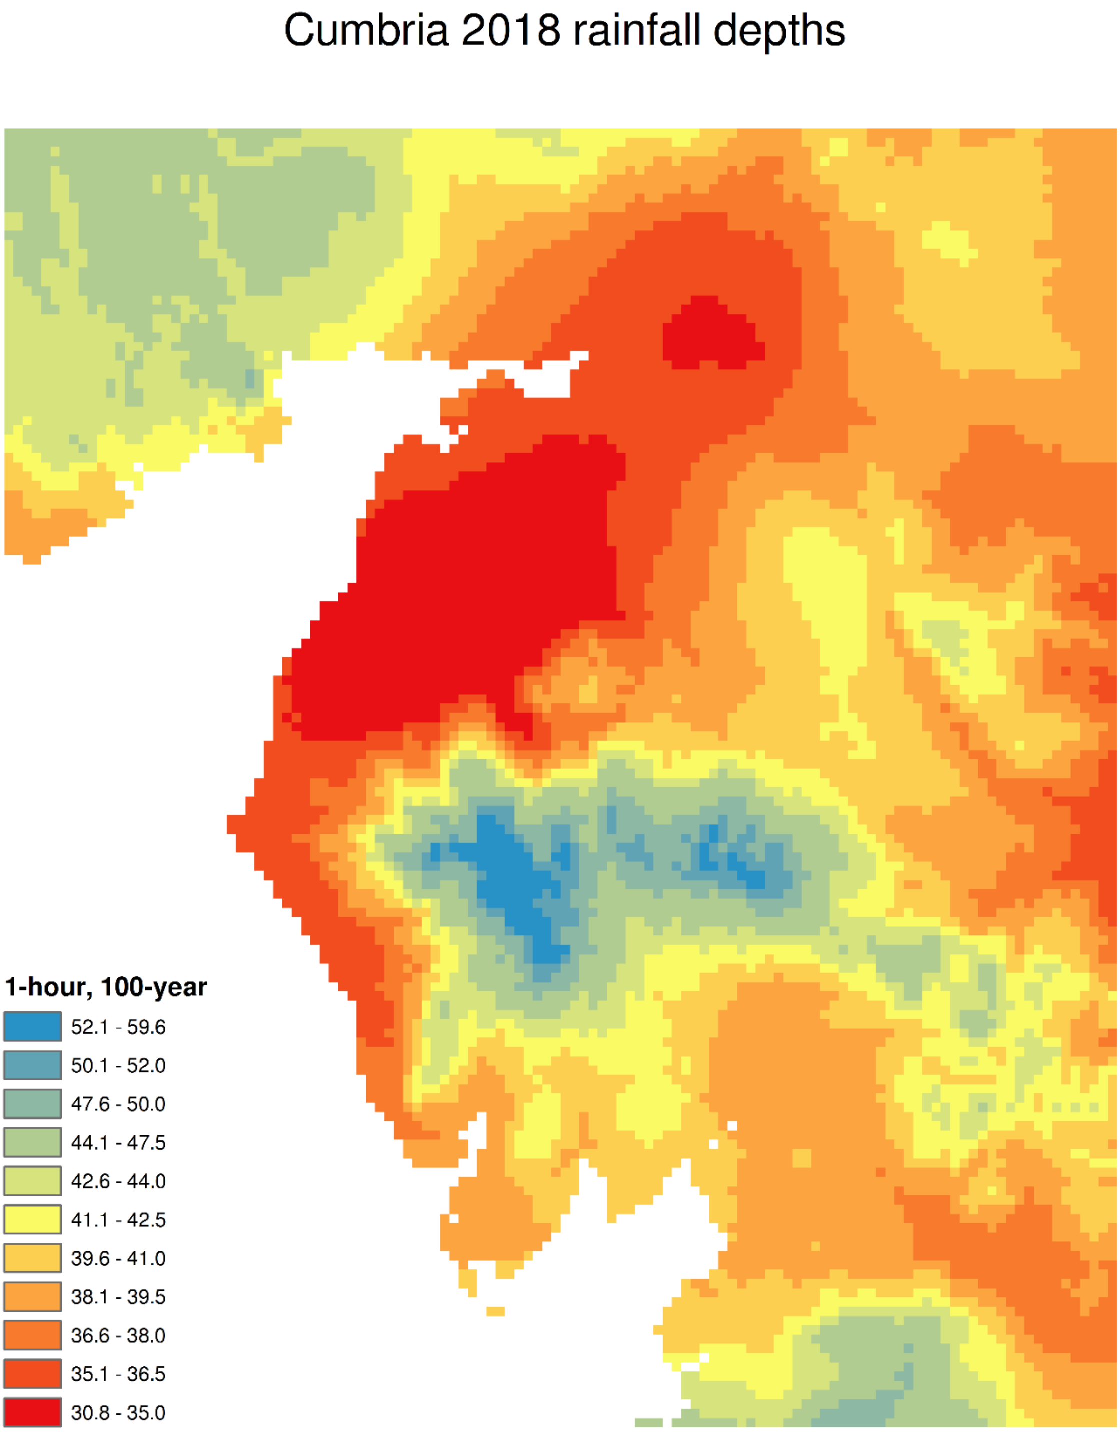


Figure 24 – New model 1-hour, 100-year rainfall depth (mm)

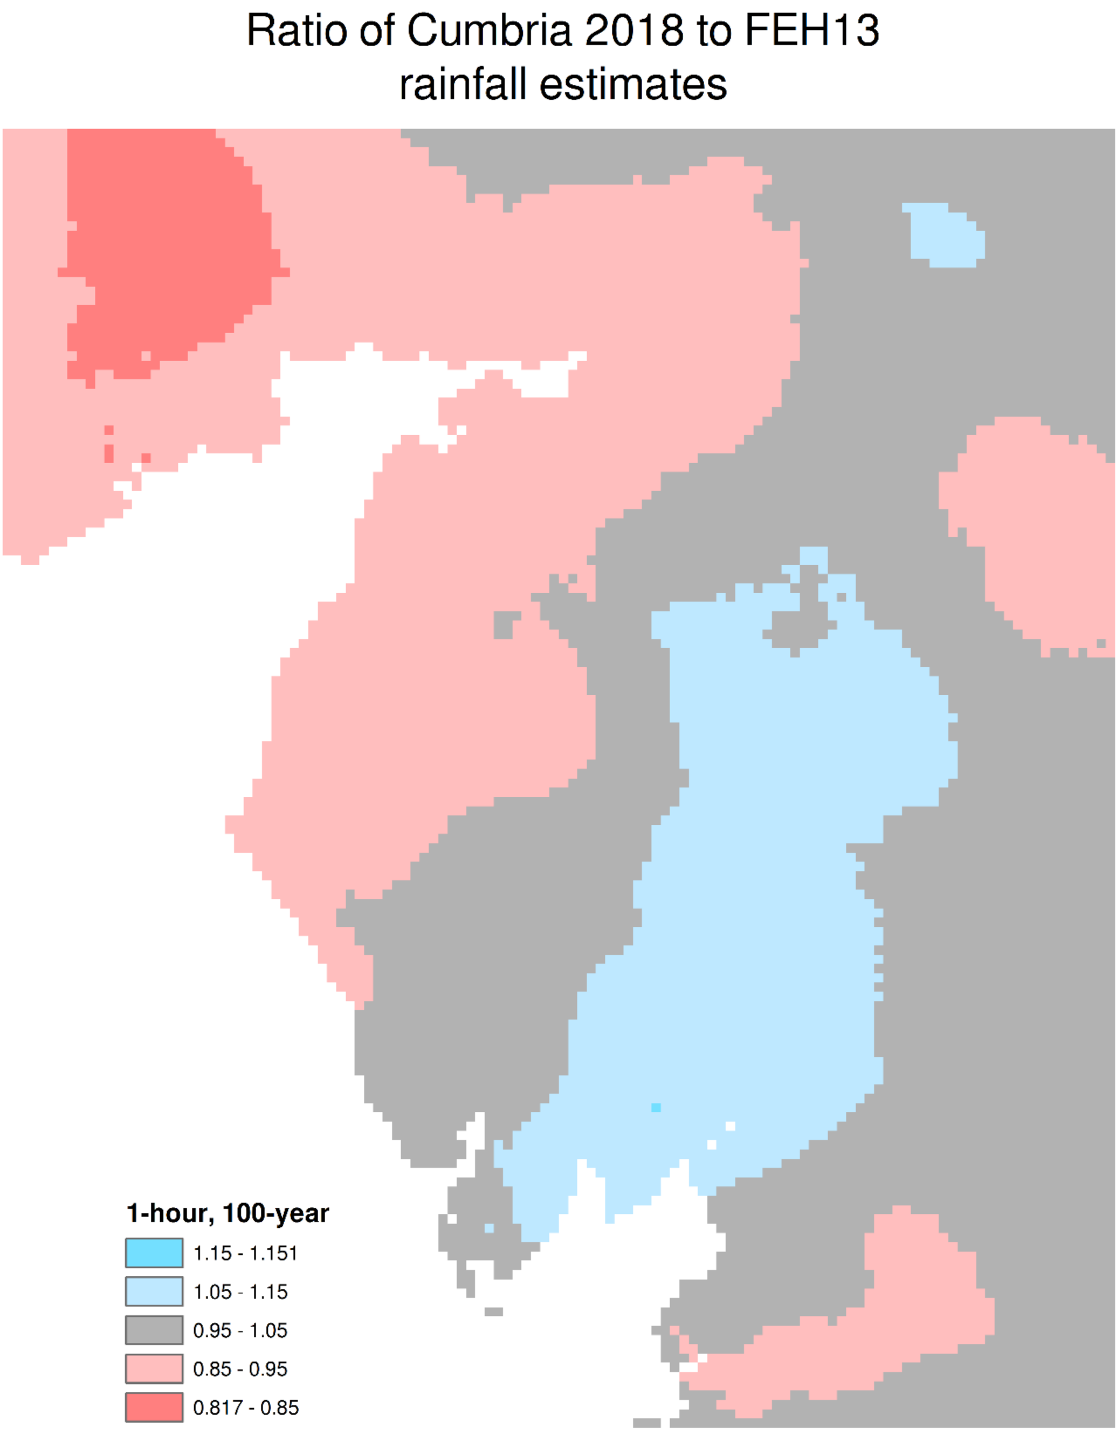


Figure 25 – Ratio of new model 1-hour, 100-year rainfall depth to FEH13 depth

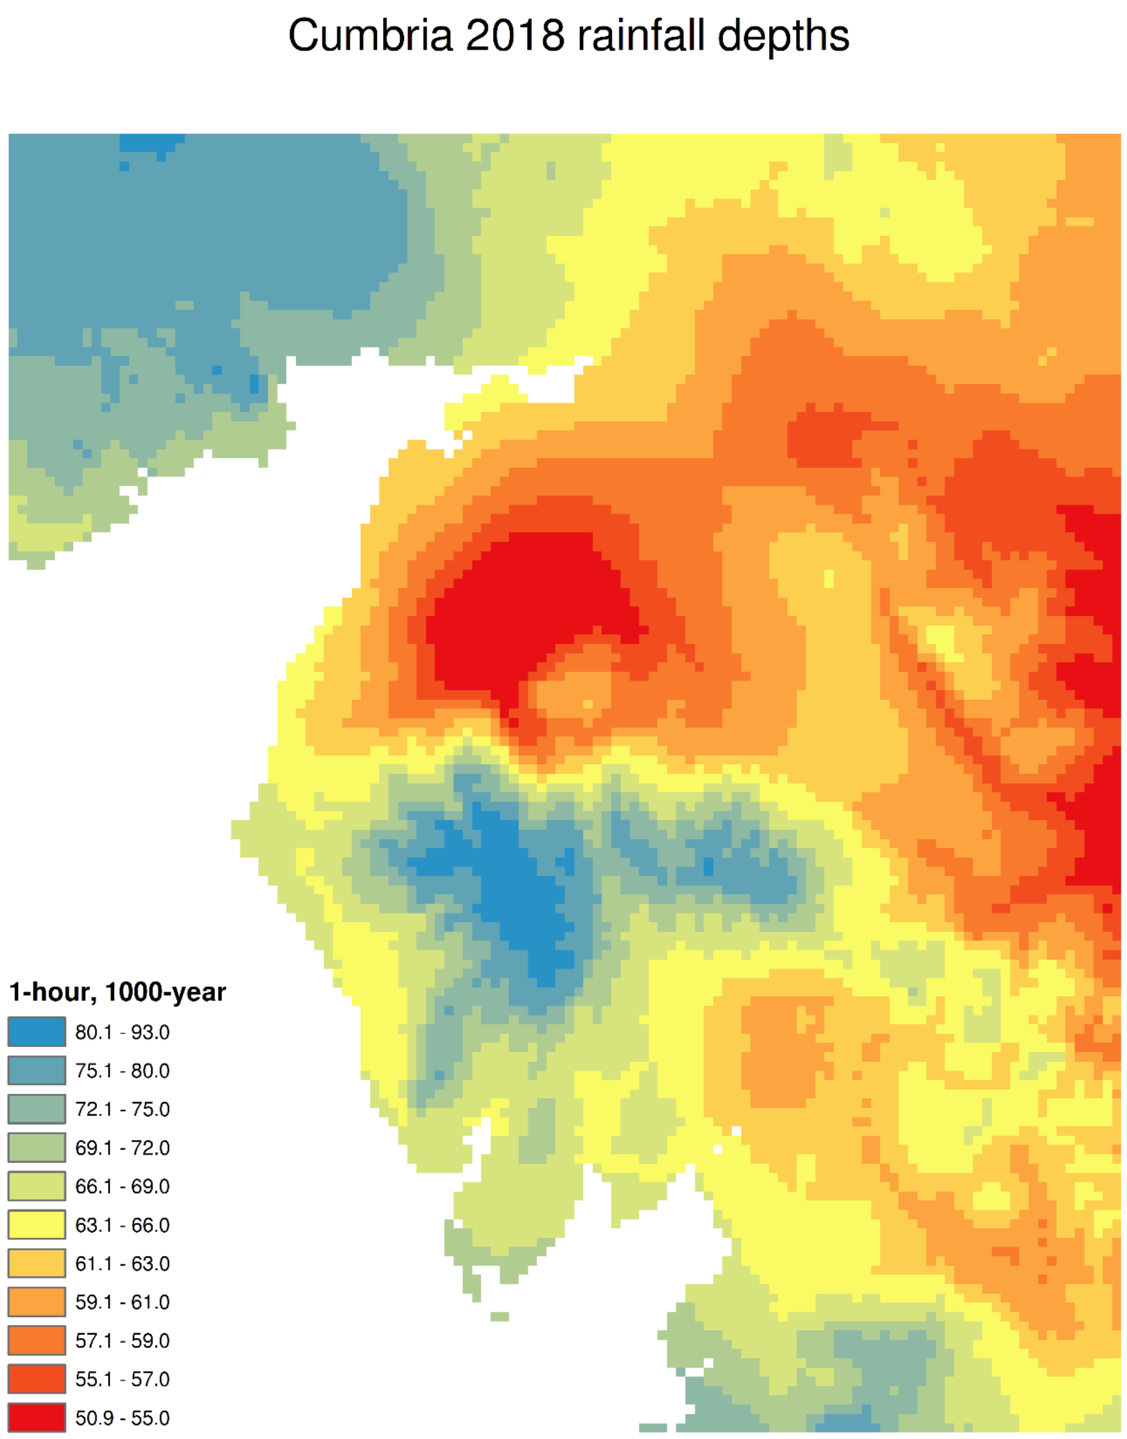


Figure 26 – New model 1-hour, 1000-year rainfall depth (mm)

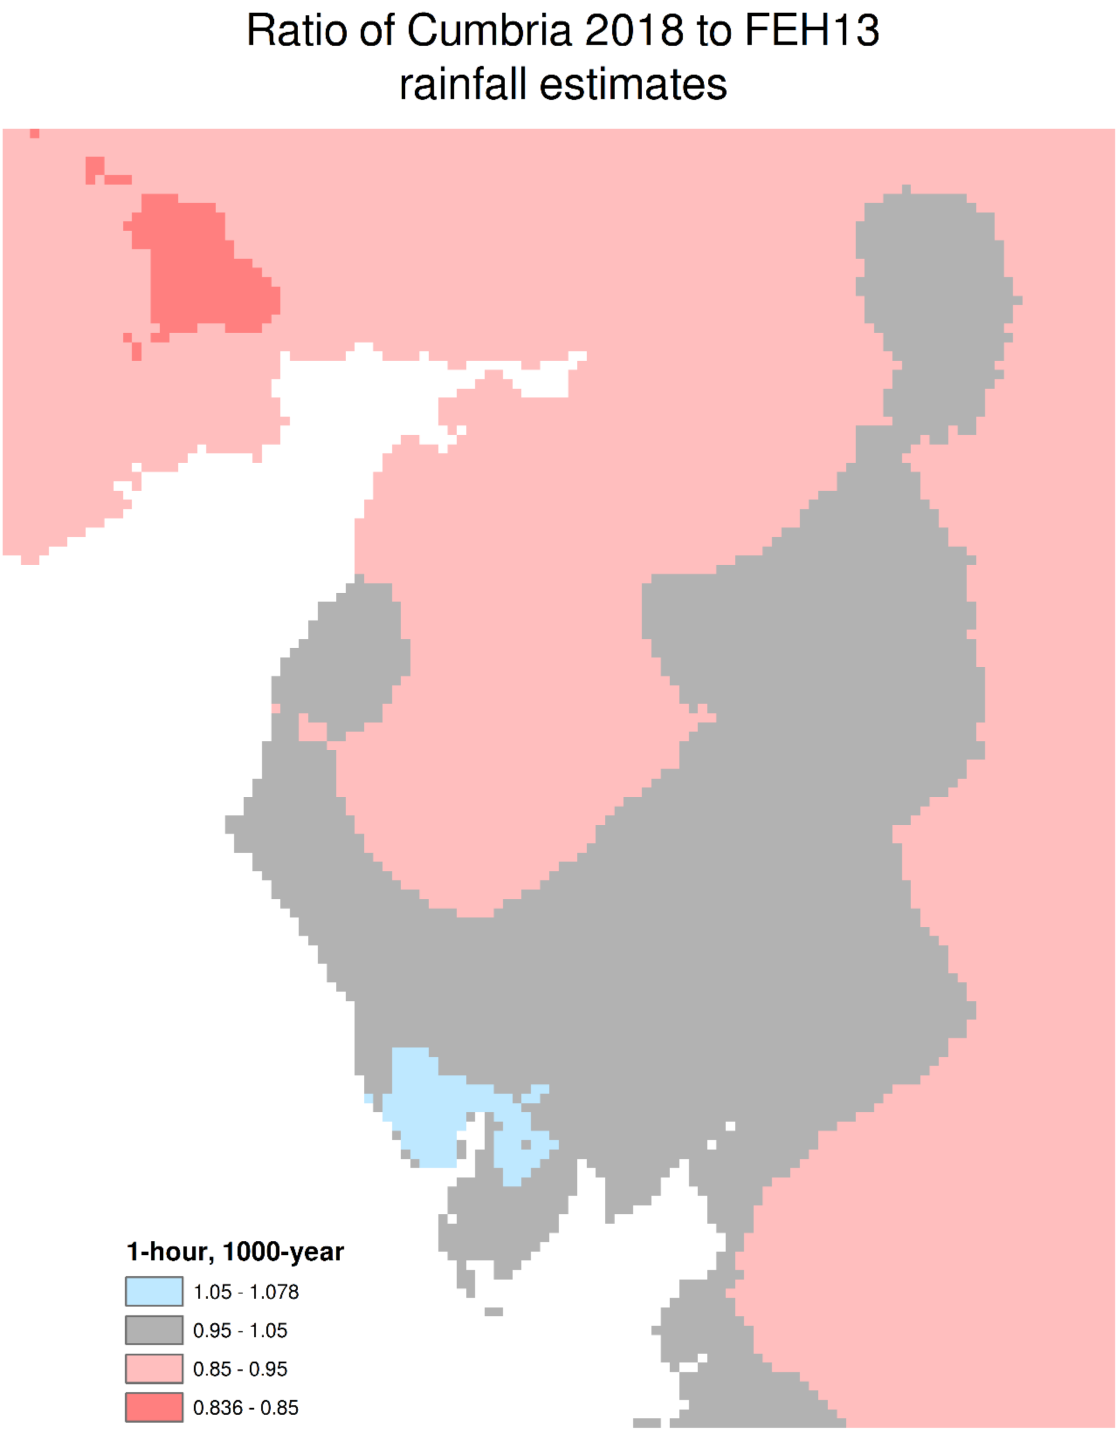


Figure 27 – Ratio of new model 1-hour, 1000-year rainfall depth to FEH13 depth

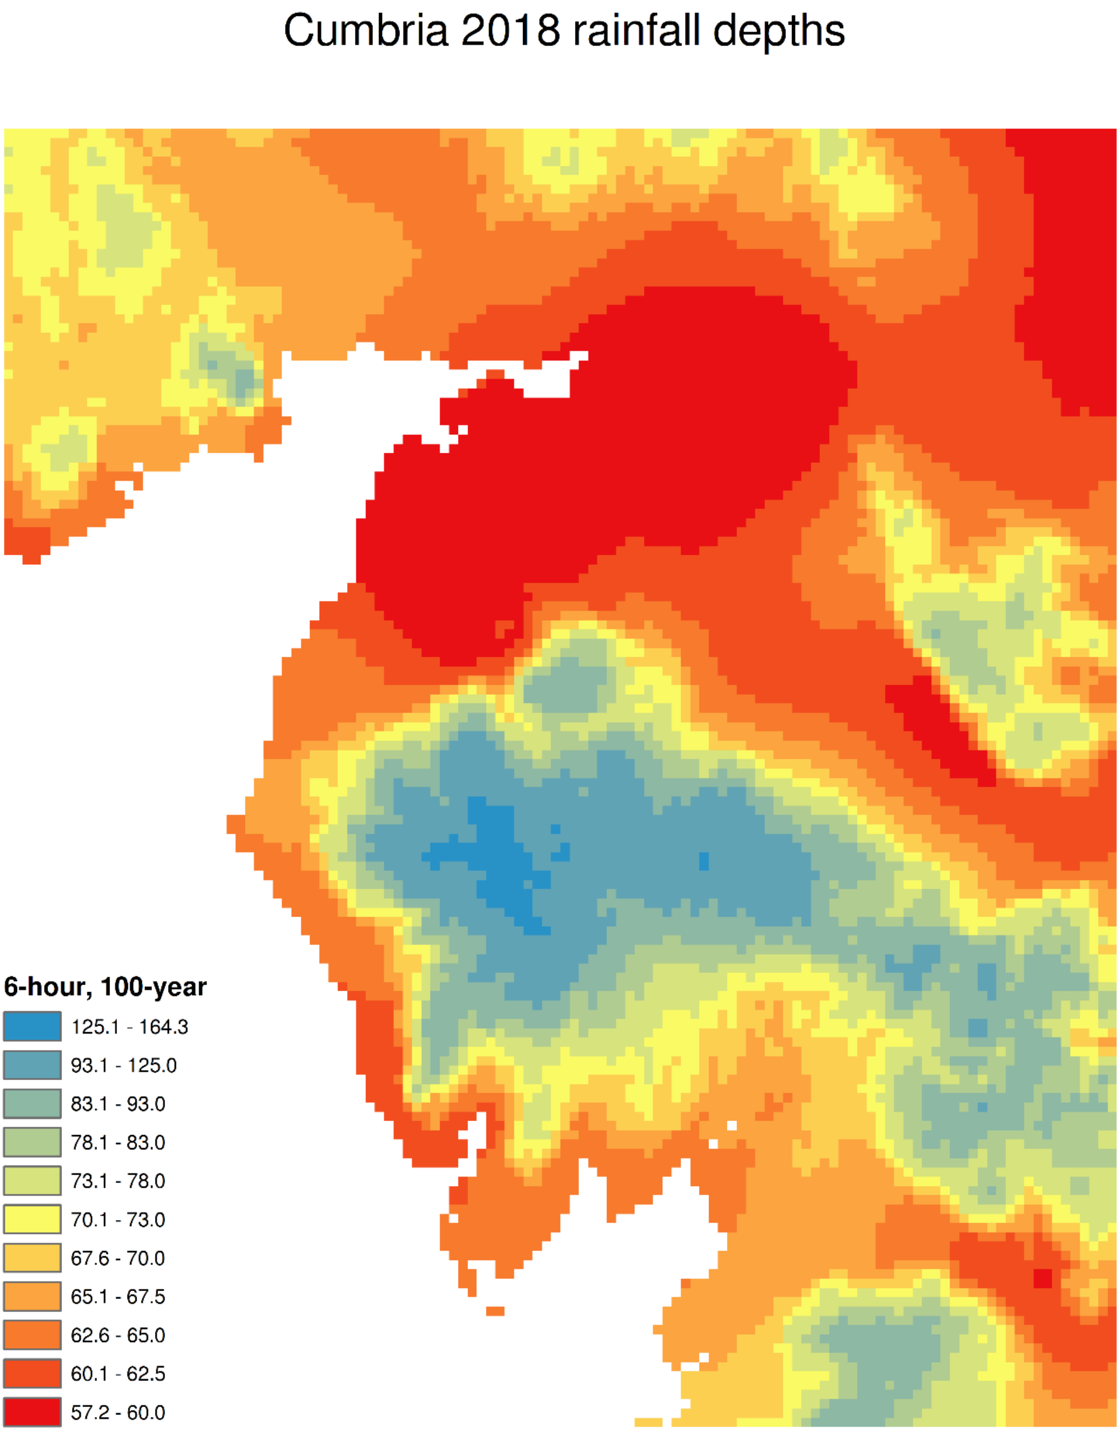


Figure 28 – New model 6-hour, 100-year rainfall depth (mm)

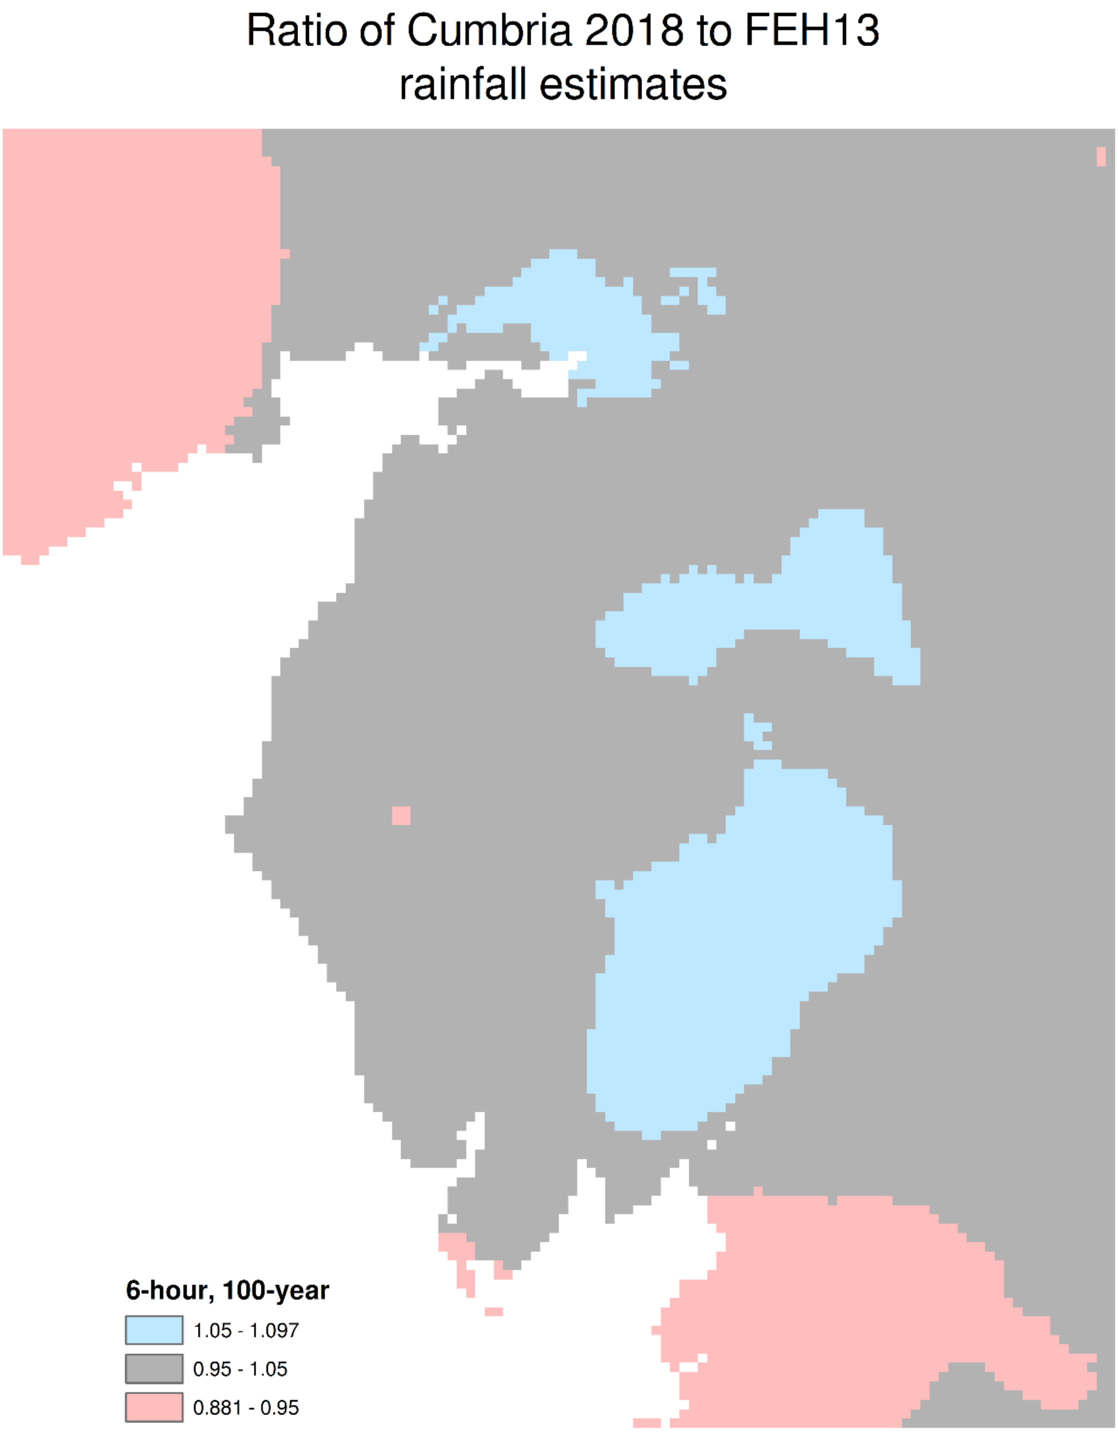


Figure 29 – Ratio of new model 6-hour, 100-year rainfall depth to FEH13 depth

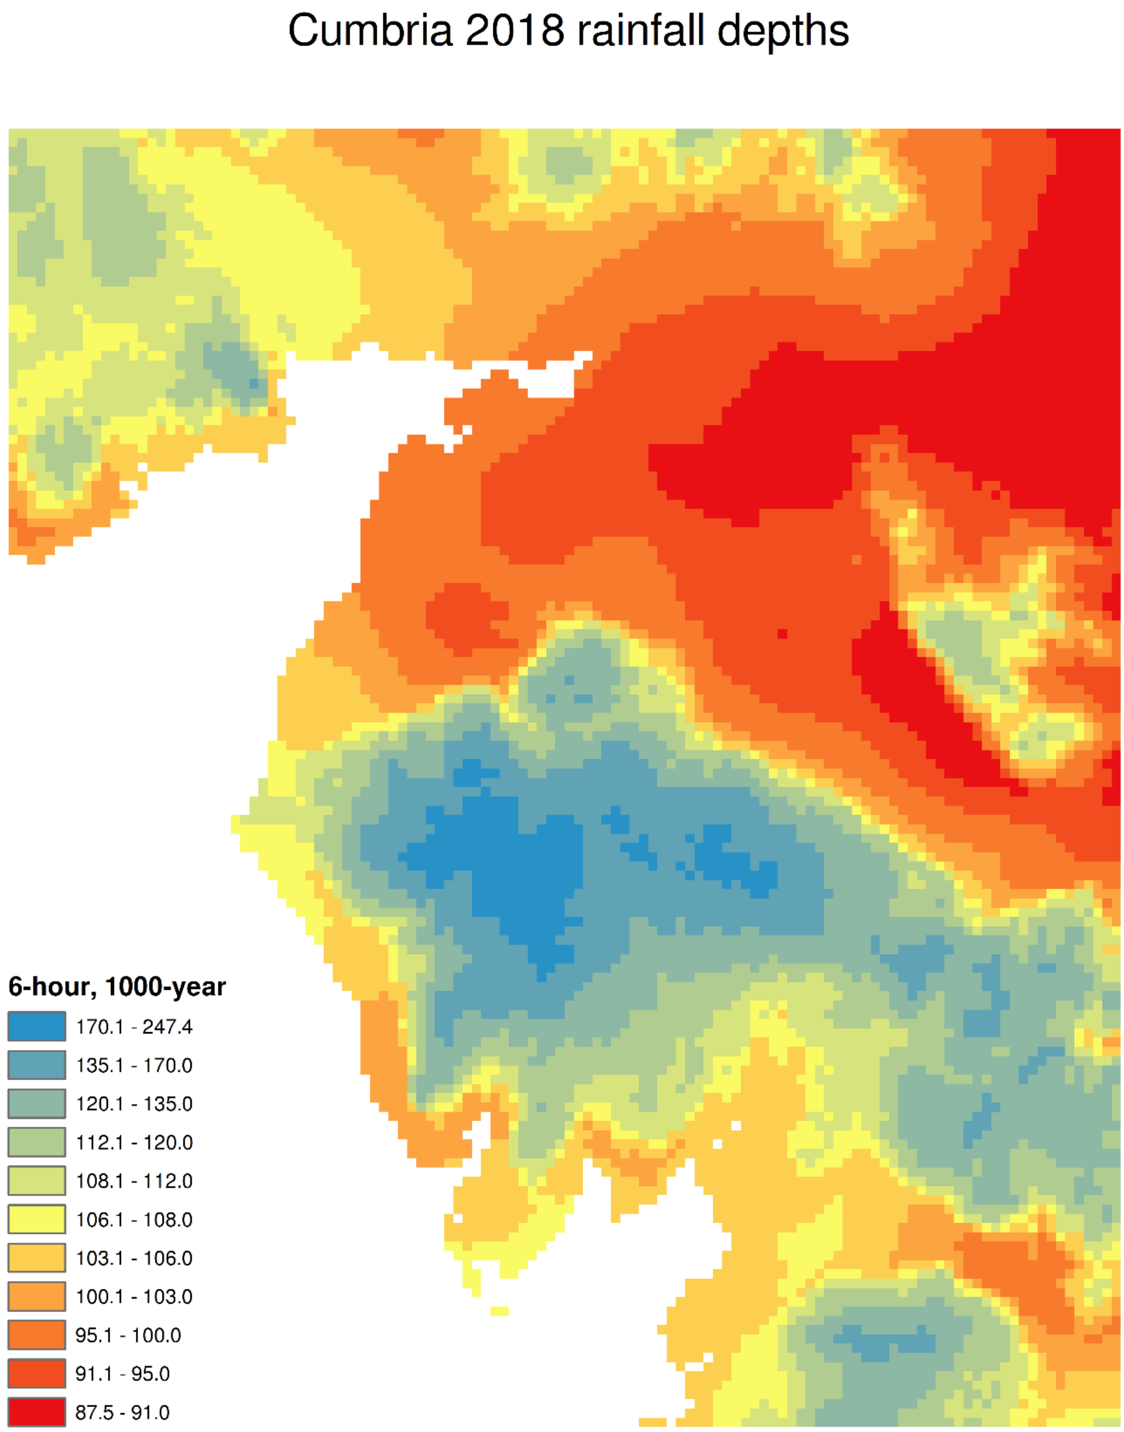


Figure 30 – New model 6-hour, 1000-year rainfall depth (mm)

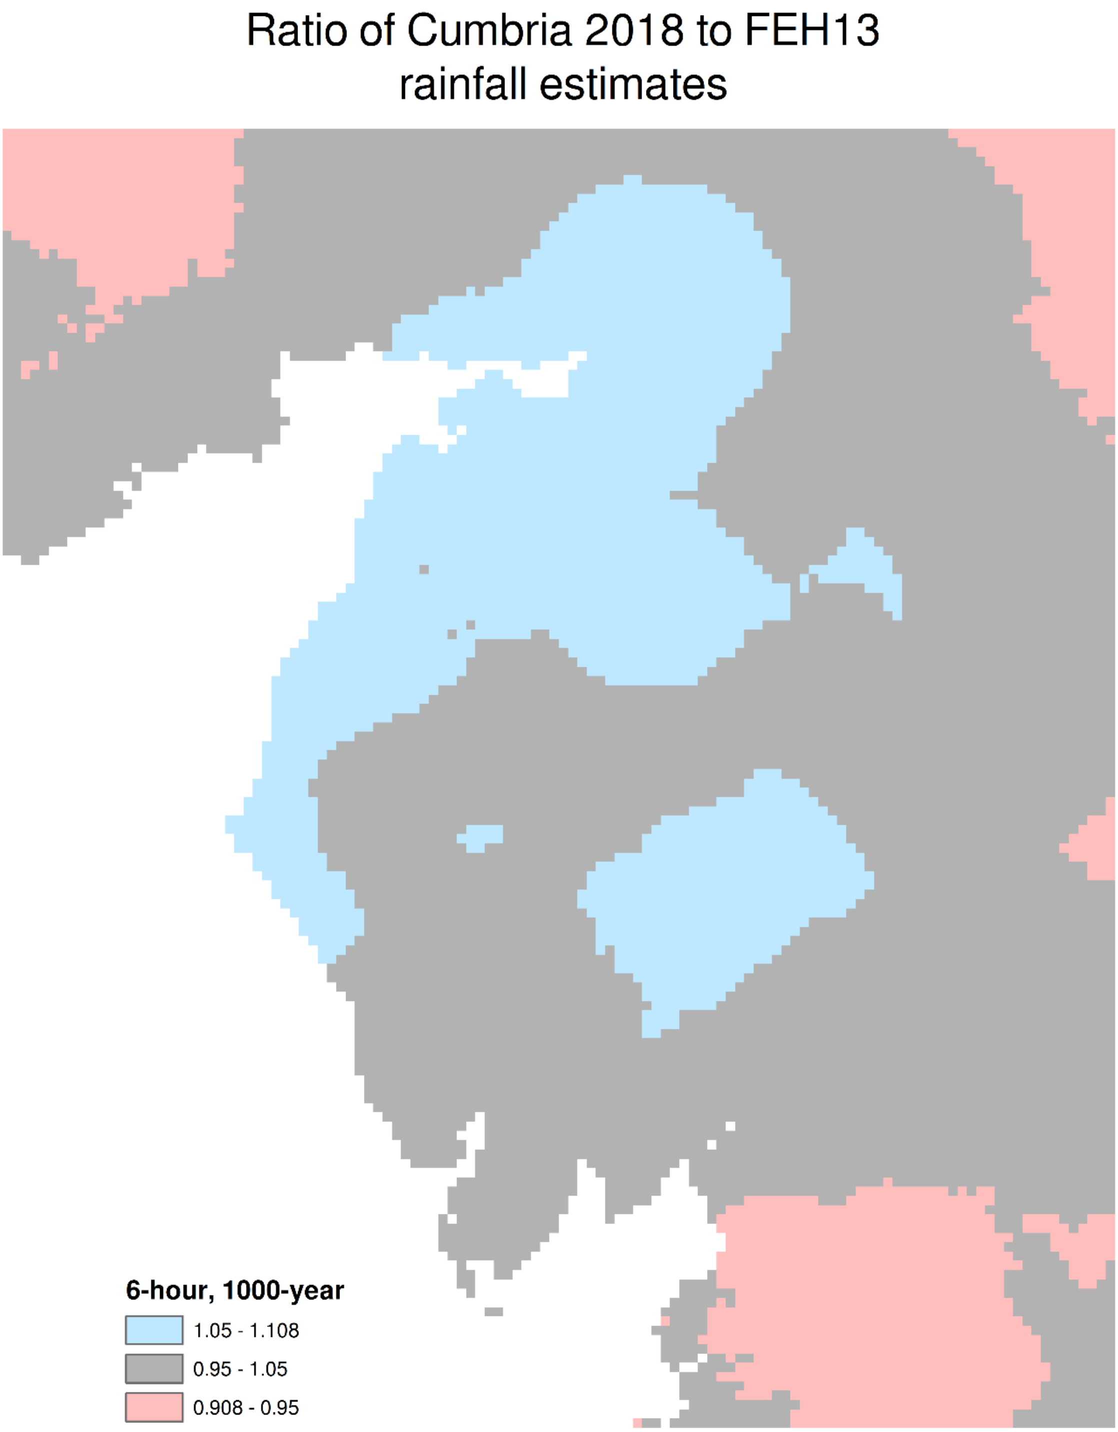


Figure 31 – Ratio of new model 6-hour, 1000-year rainfall depth to FEH13 depth

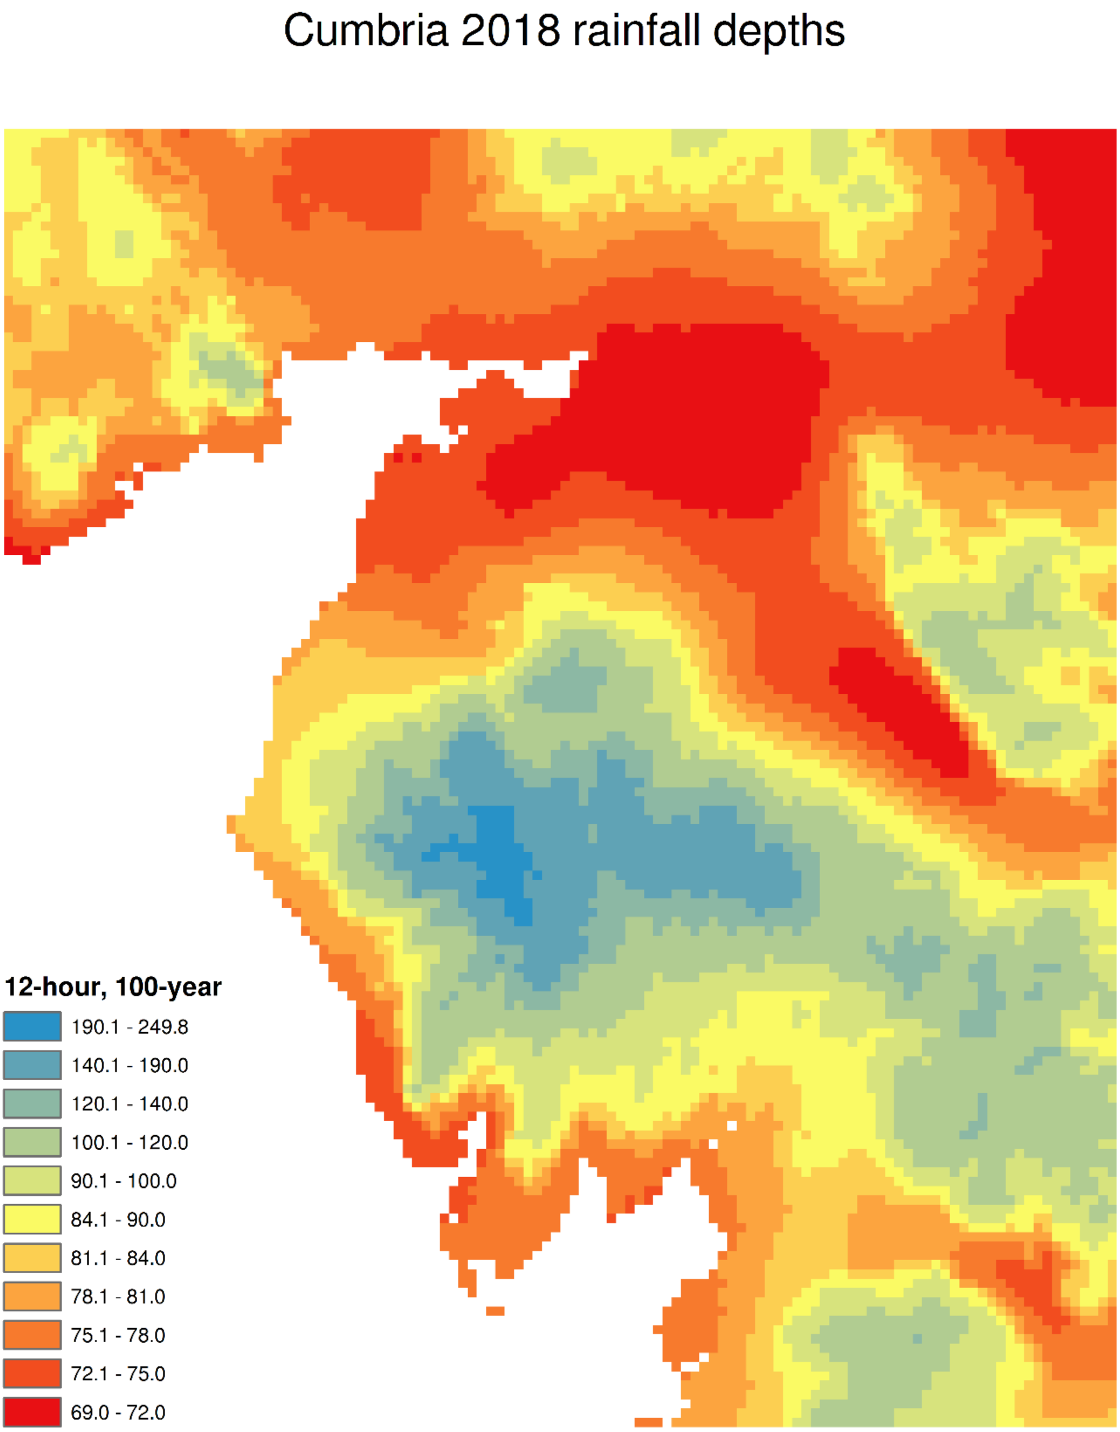


Figure 32 – New model 12-hour, 100-year rainfall depth (mm)

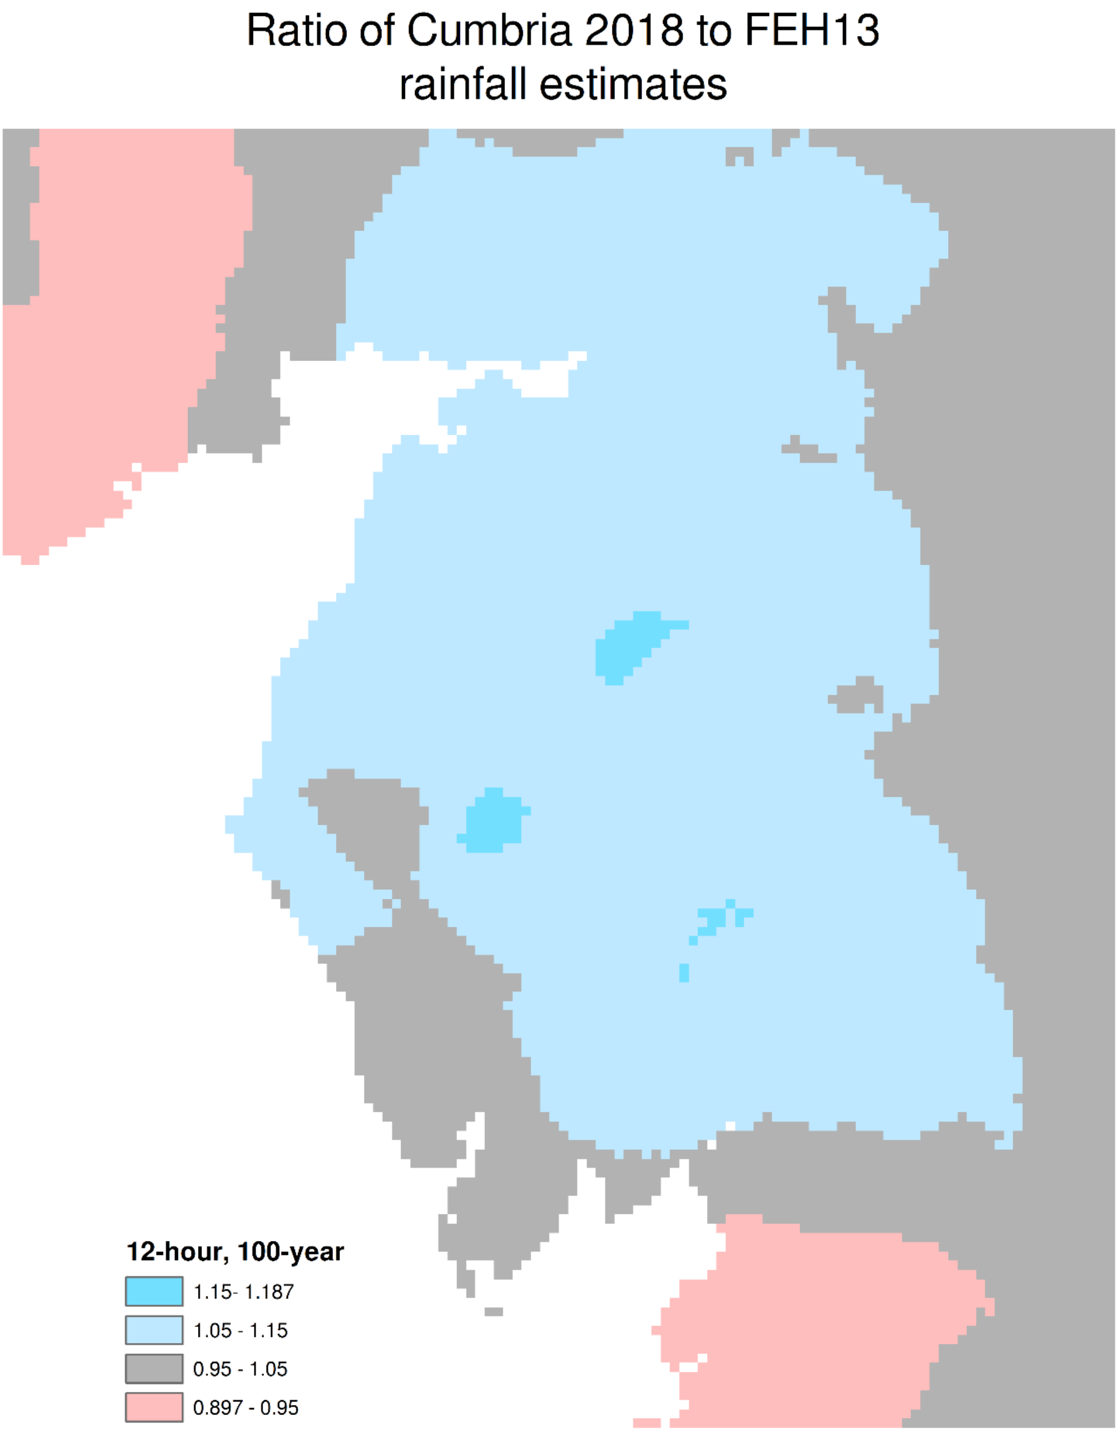


Figure 33 – Ratio of new model 12-hour, 100-year rainfall depth to FEH13 depth

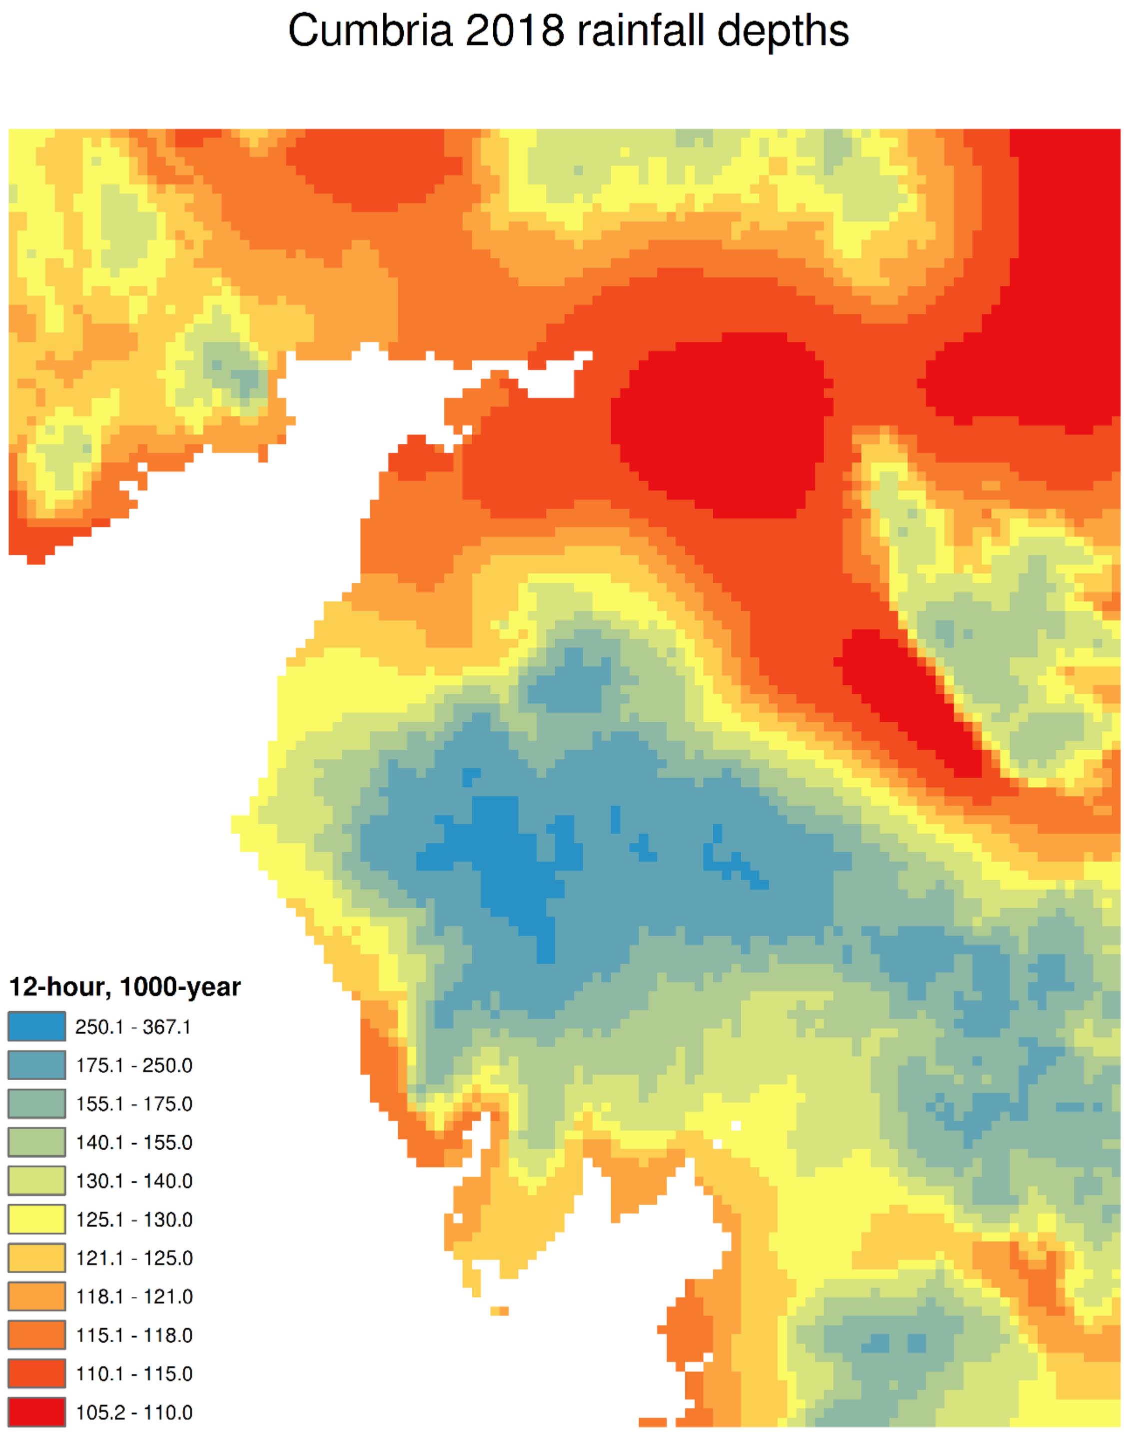


Figure 34 – New model 12-hour, 1000-year rainfall depth (mm)

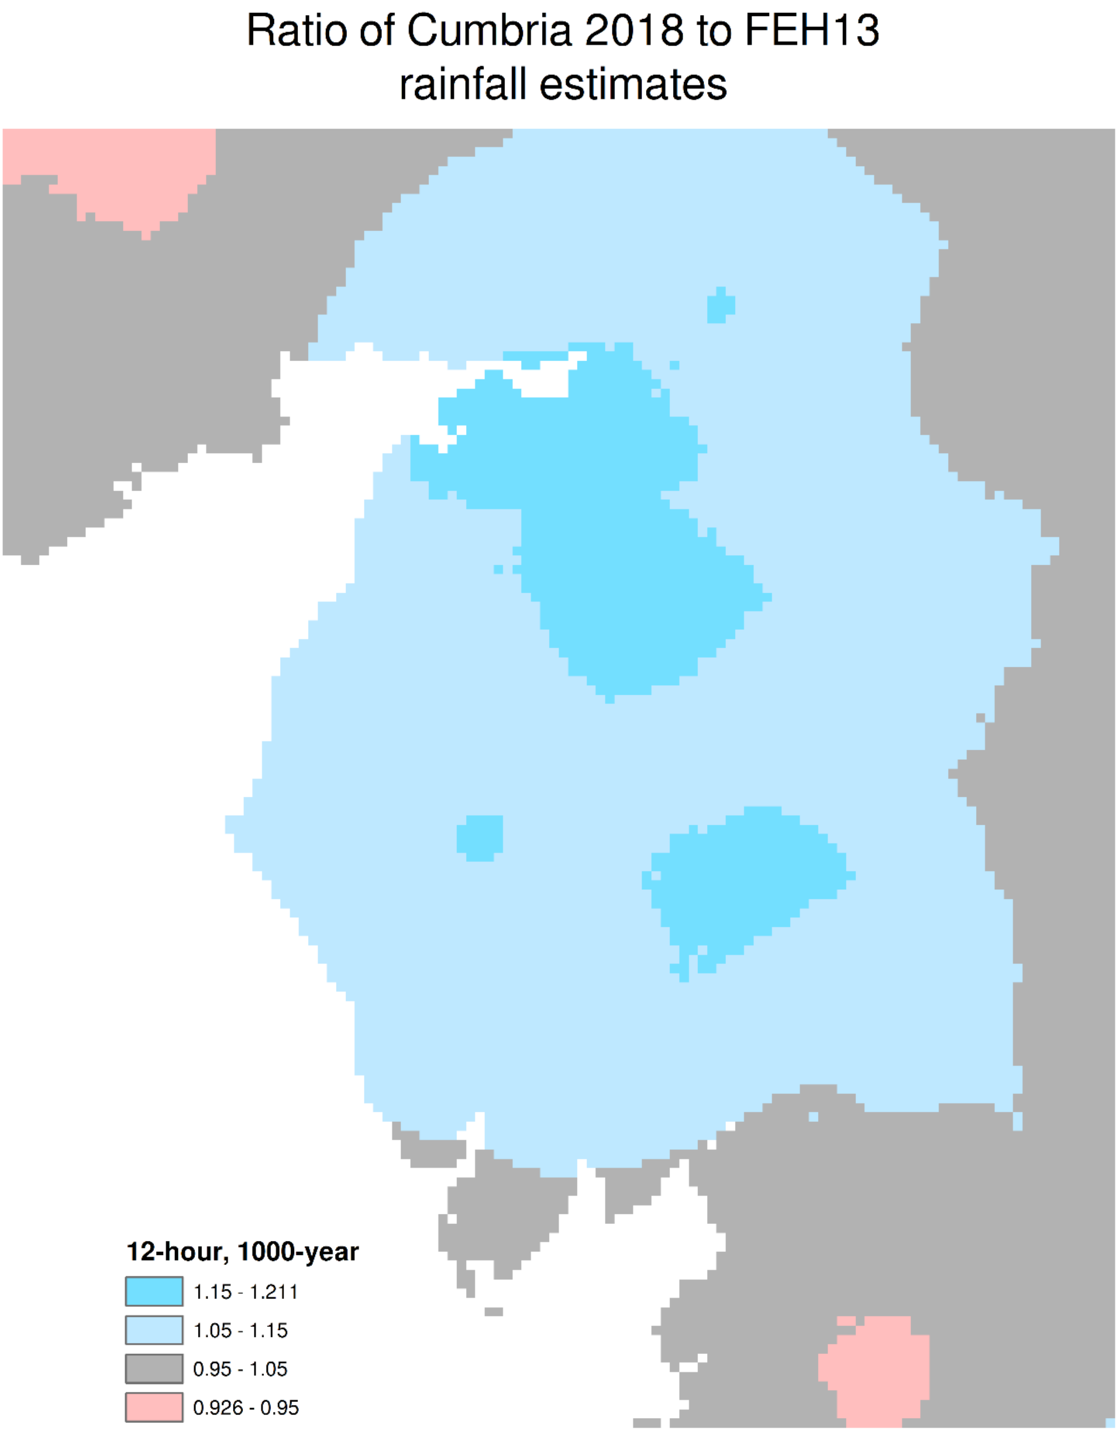


Figure 35 – Ratio of new model 12-hour, 1000-year rainfall depth to FEH13 depth

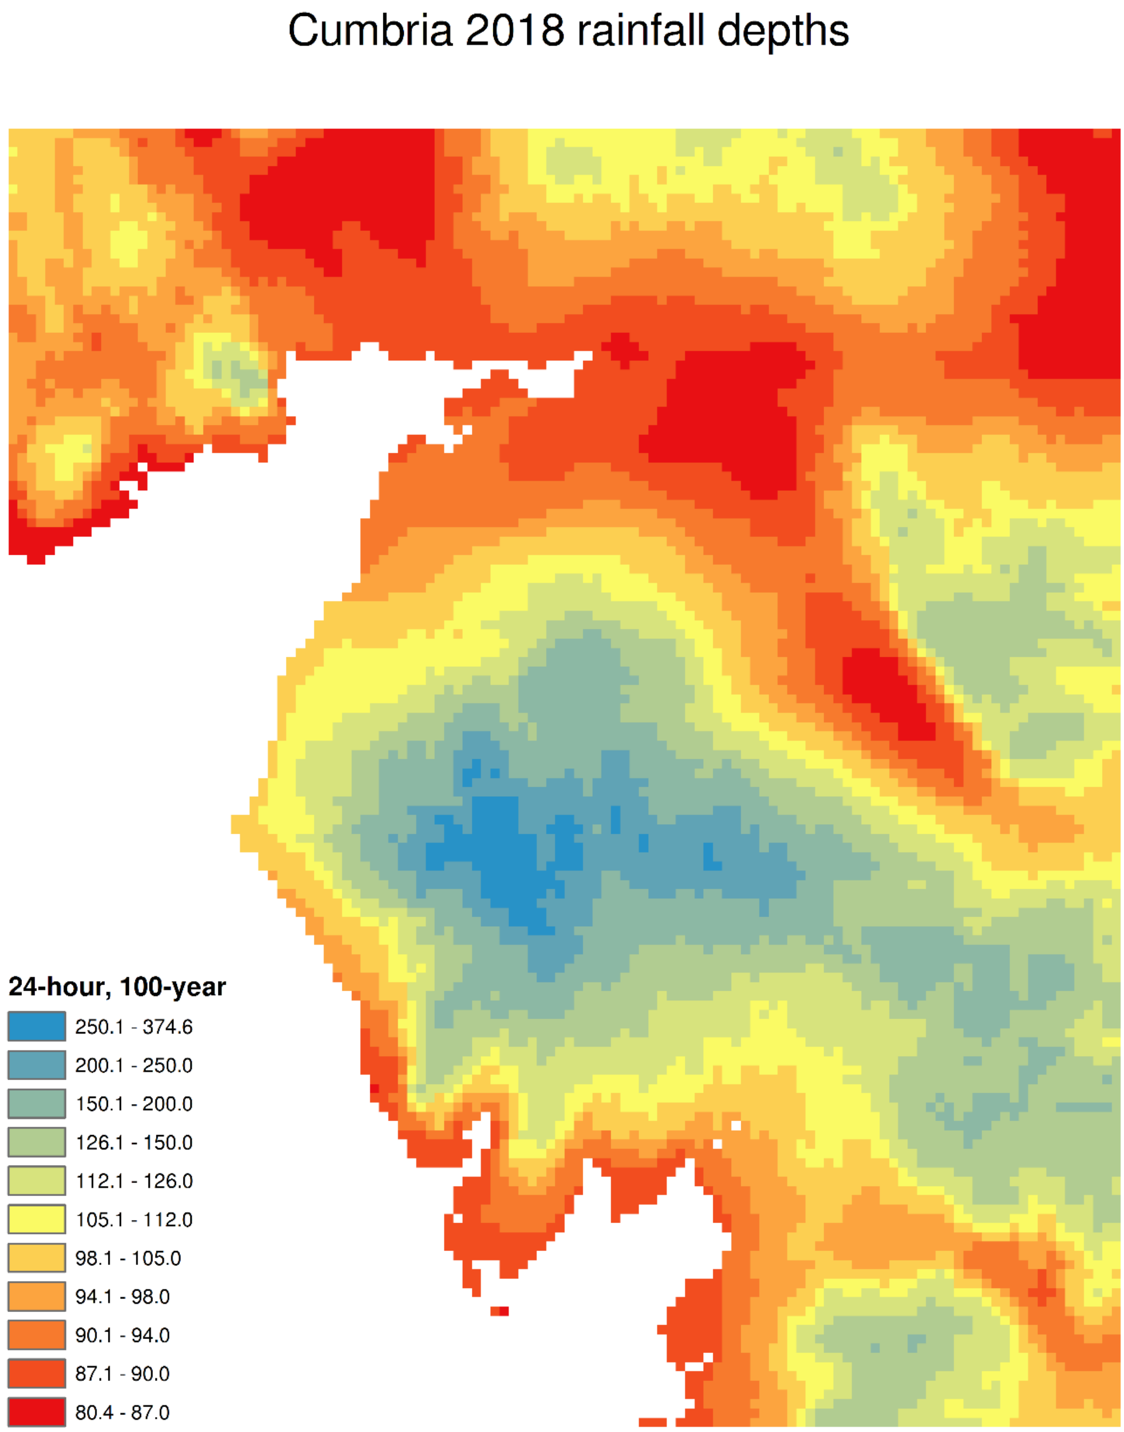


Figure 36 – New model 24-hour, 100-year rainfall depth (mm)

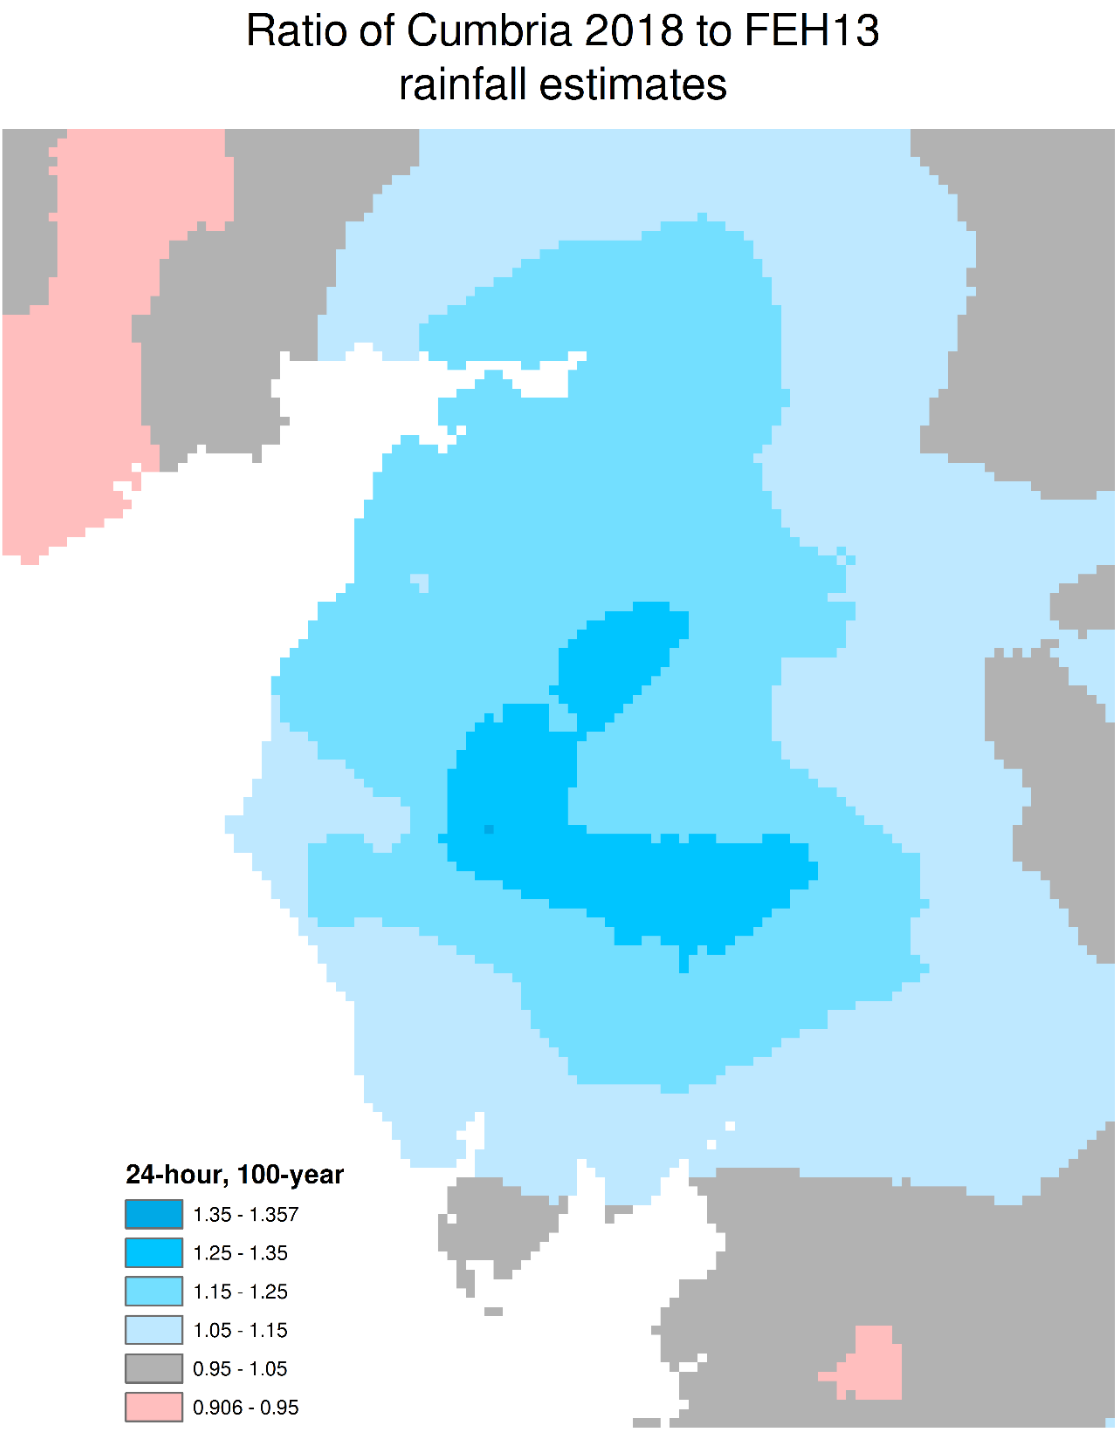


Figure 37 – Ratio of new model 24-hour, 100-year rainfall depth to FEH13 depth

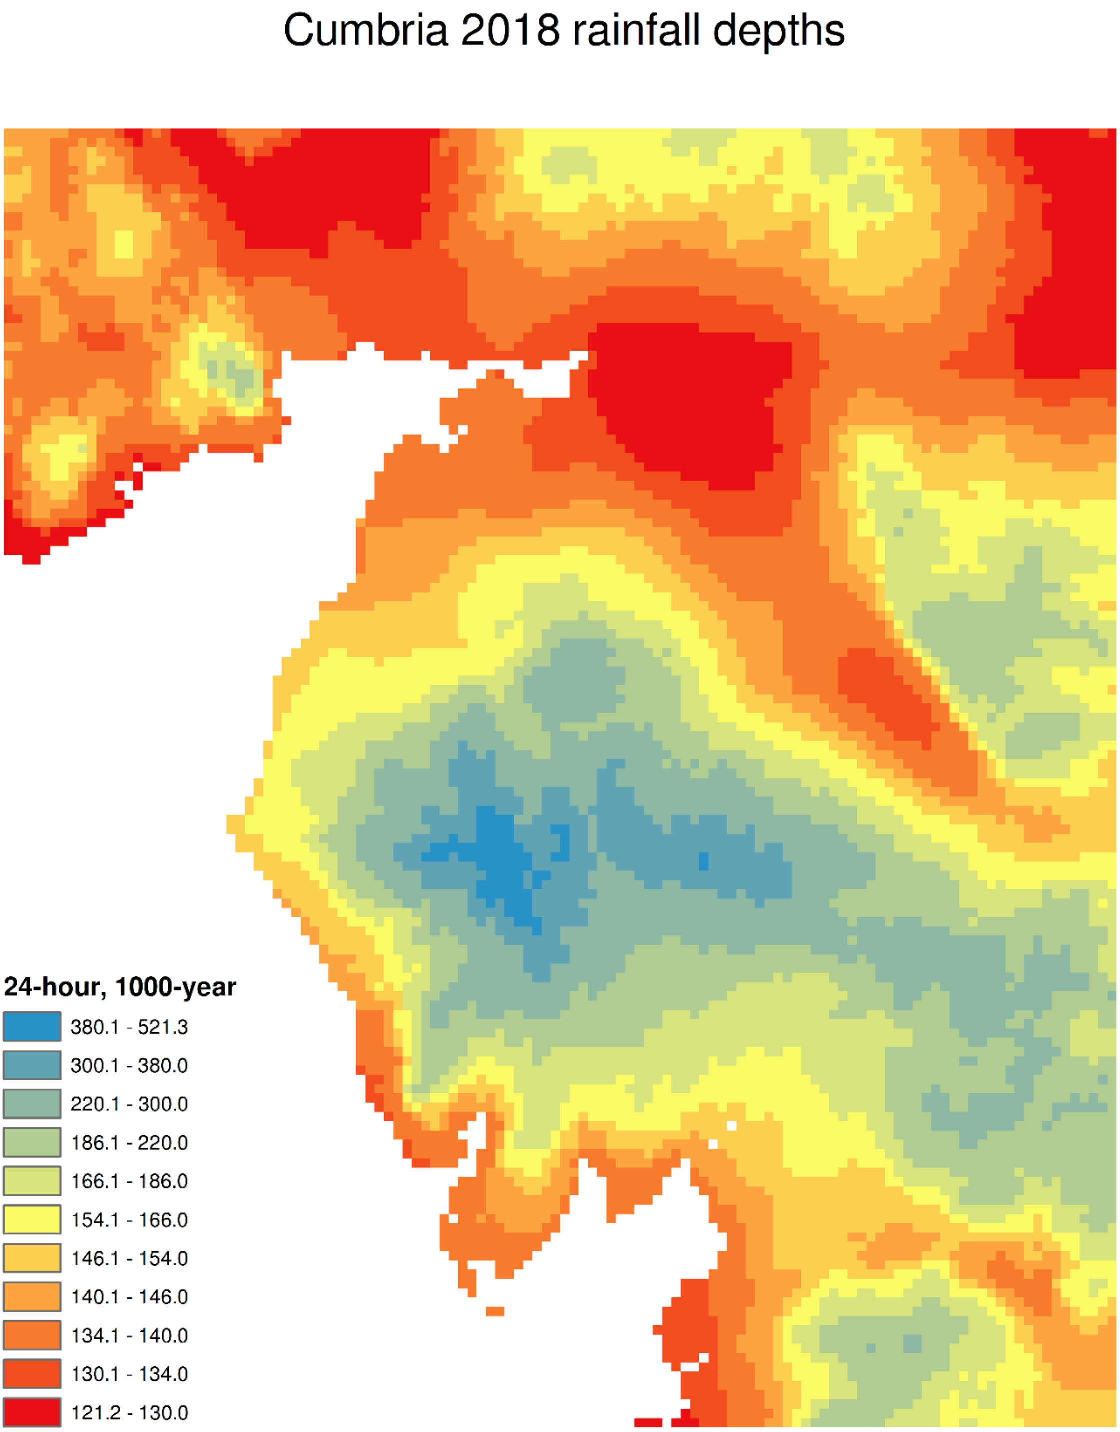


Figure 38 – New model 24-hour, 1000-year rainfall depth (mm)

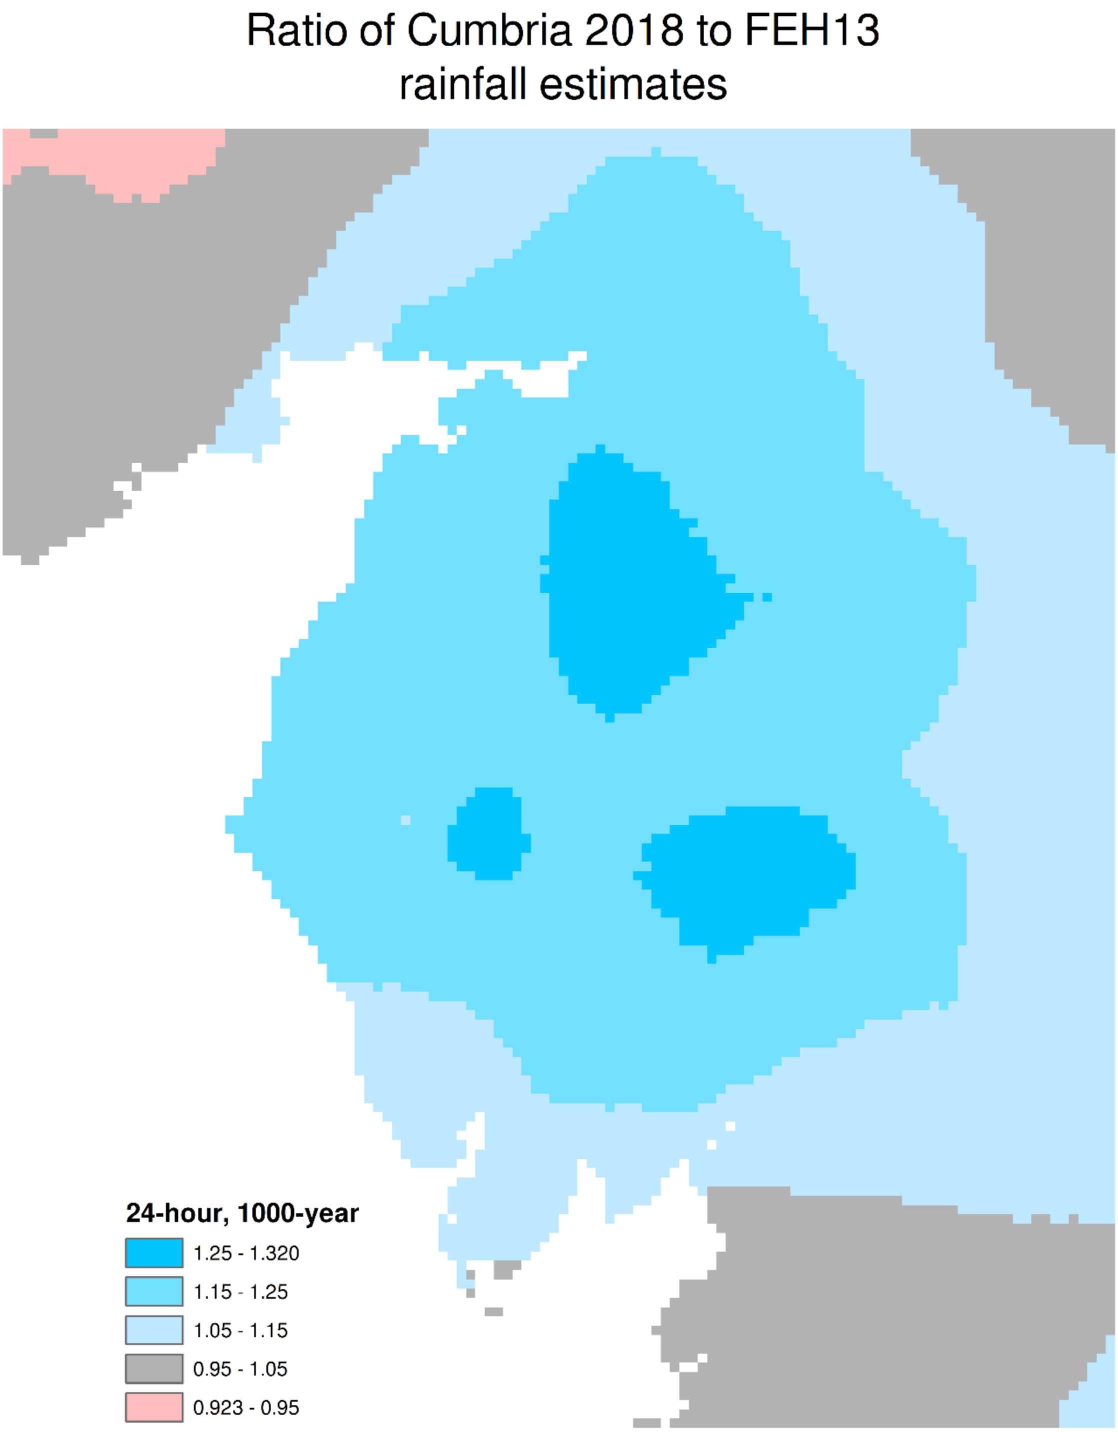


Figure 39 – Ratio of new model 24-hour, 1000-year rainfall depth to  FEH13 depth

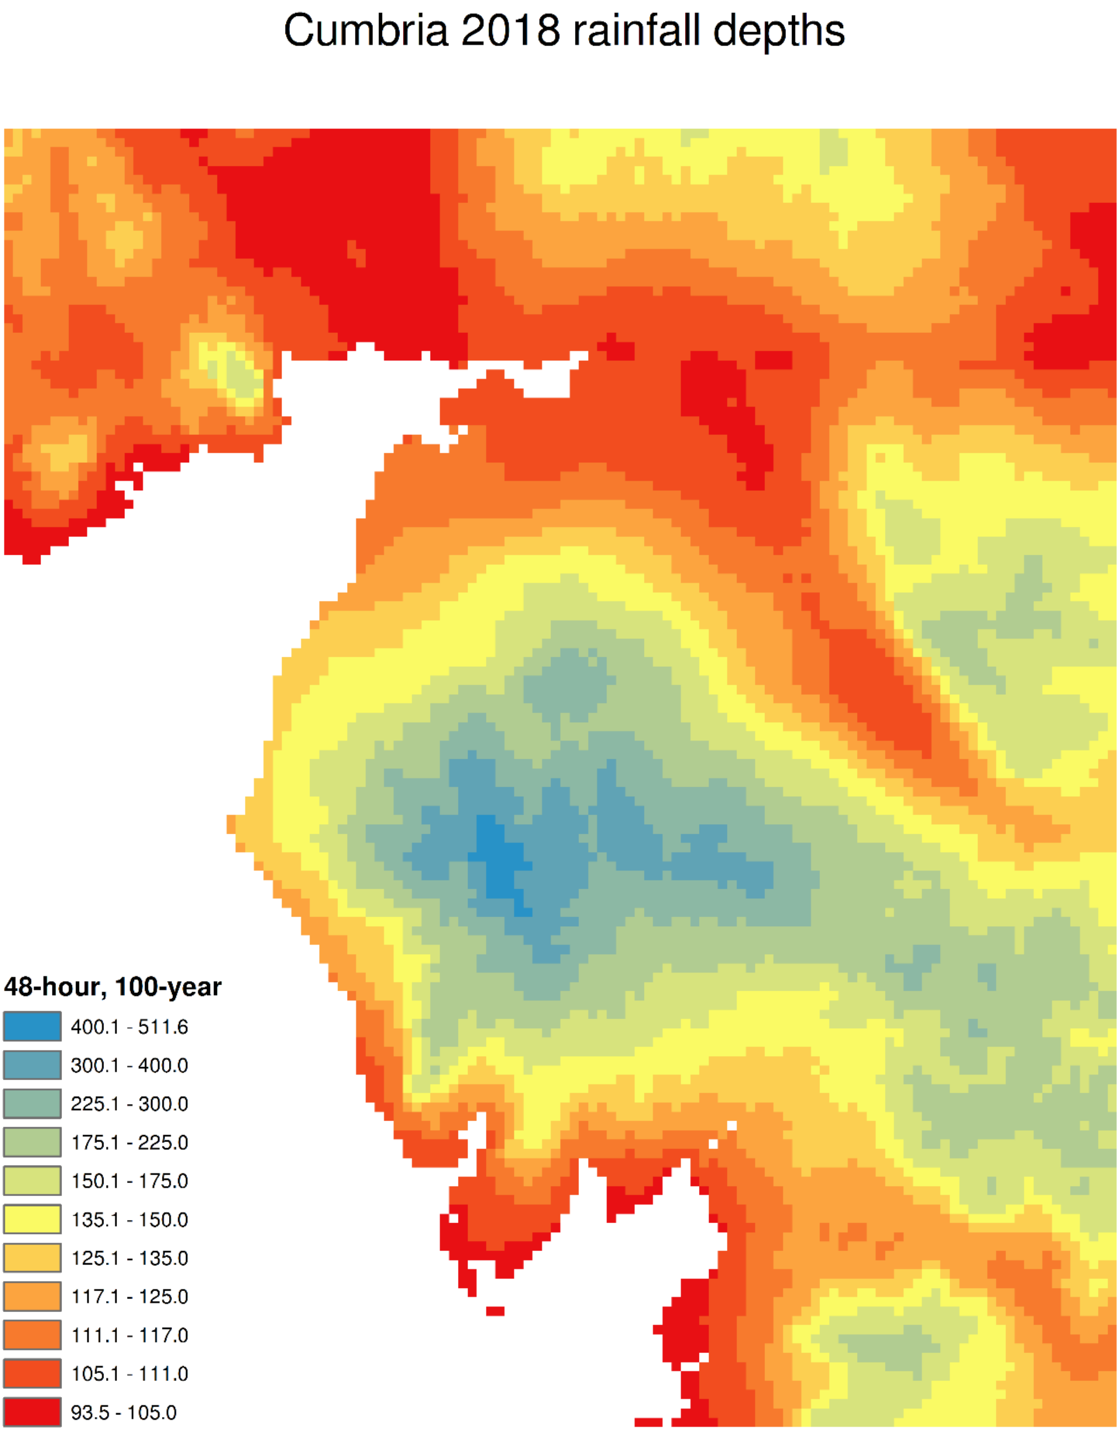


Figure 40 – New model 48-hour, 100-year rainfall depth (mm)

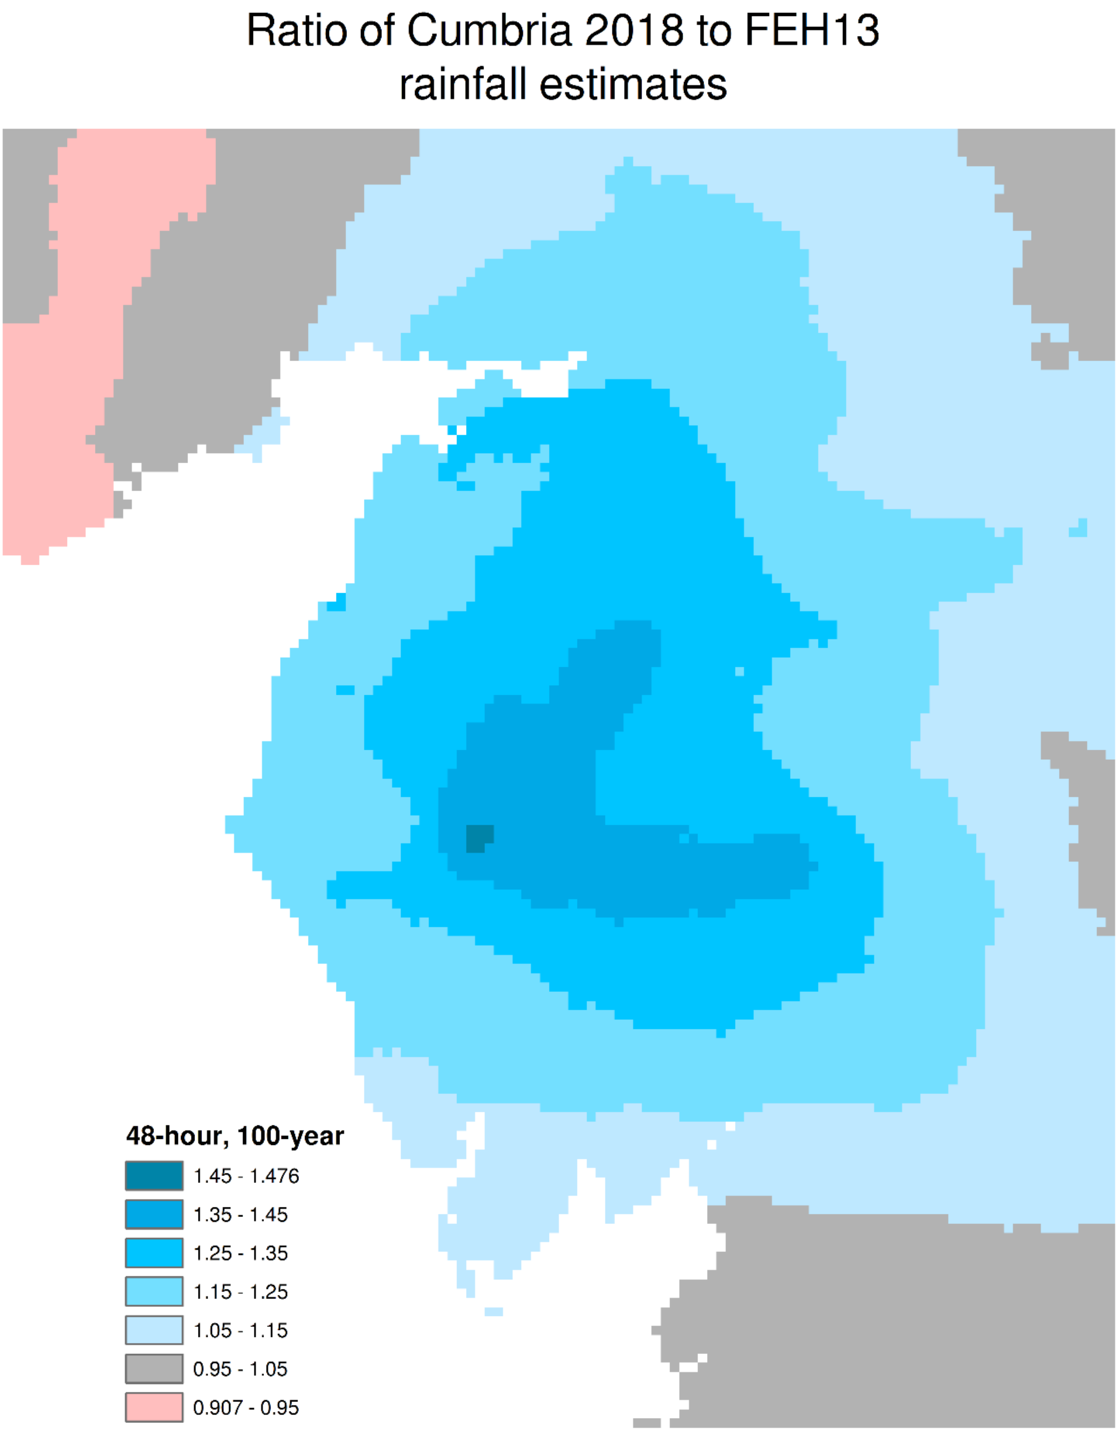


Figure 41 – Ratio of new model 48-hour, 100-year rainfall depth to FEH13 depth

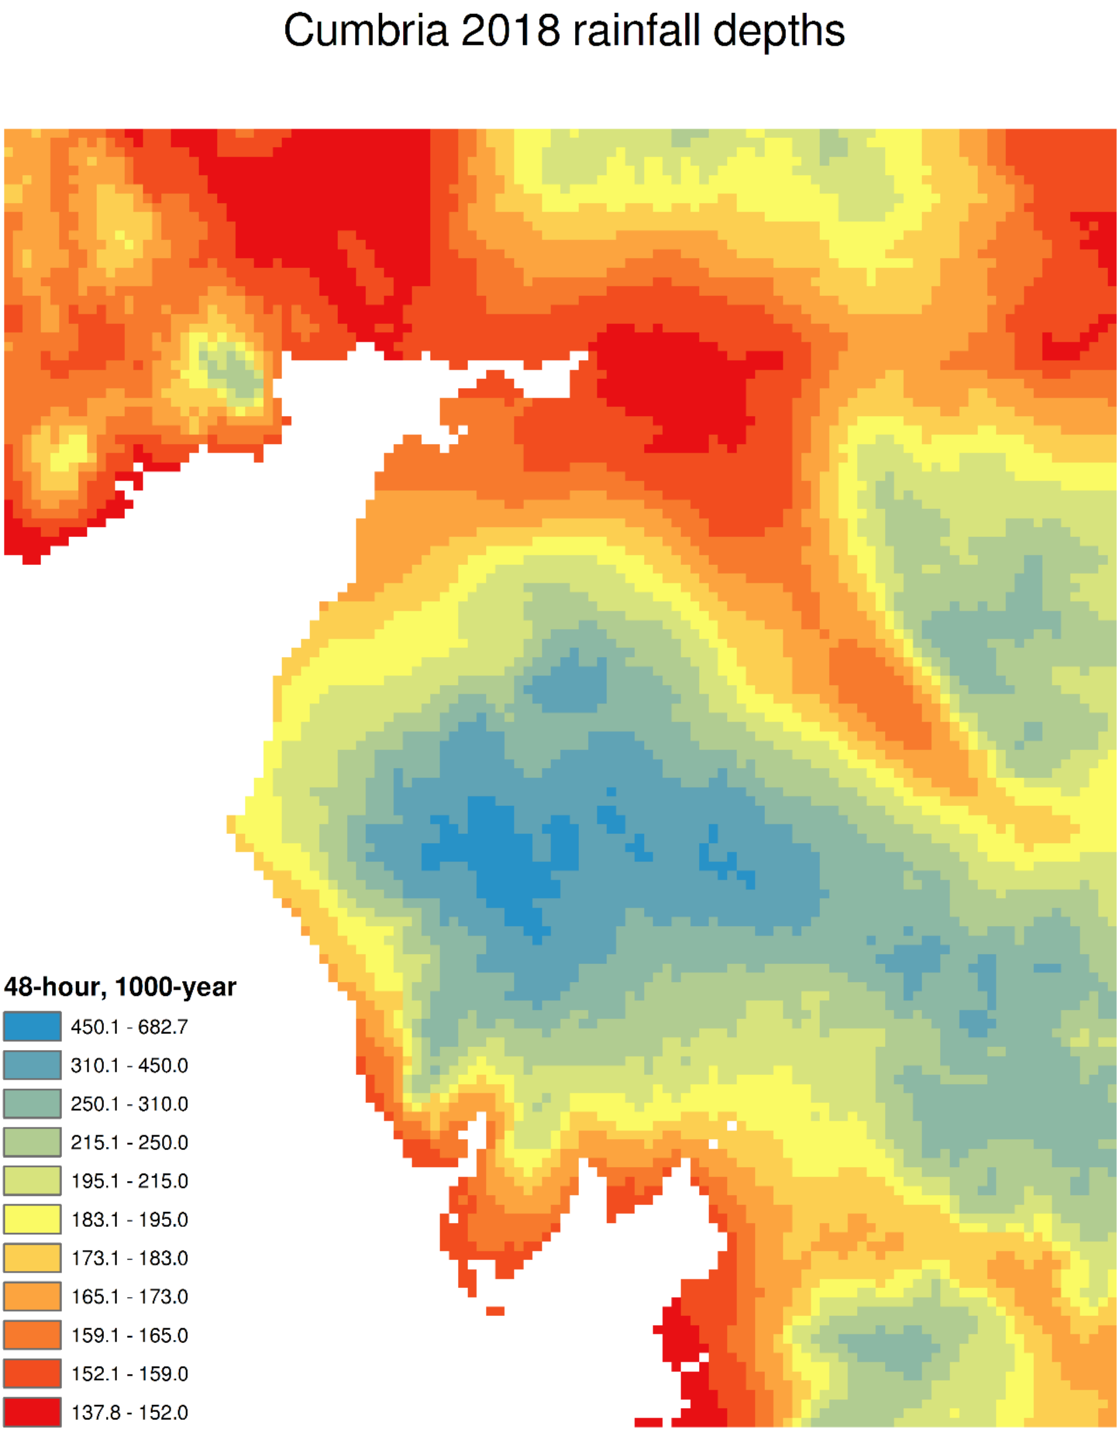


Figure 42 – New model 48-hour, 1000-year rainfall depth (mm)

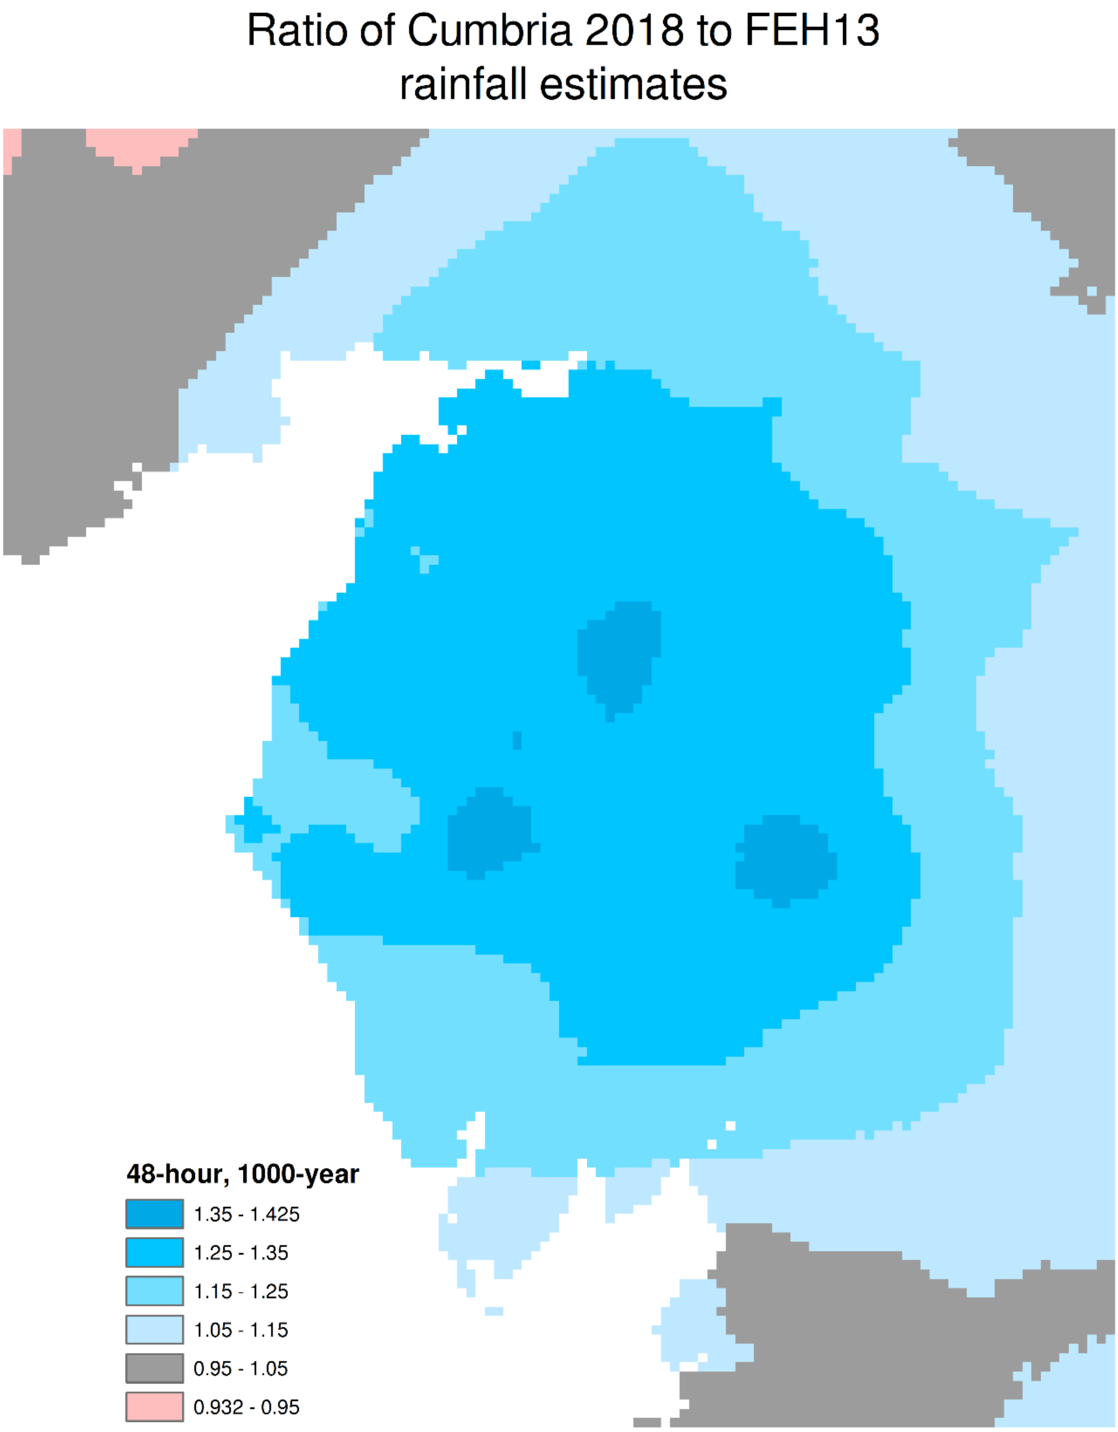


Figure 43 – Ratio of new model 48-hour, 1000-year rainfall depth to FEH13 depth

# 9 Discussion

Cumbria has experienced severe flooding, driven by extreme rainfall, four times over the 11 year-period from January 2005 to December 2015. During this time, new records were set for total rainfall accumulation over a range of durations from one to four days, in some cases twice. Therefore, recalibration of the FEH13 model reduces the estimated return periods associated with extreme events of these durations and increases the estimated rainfall depths associated with defined-probability extreme events of these durations. DDF relationships for return periods up to 10 years are largely unaffected, even for the 36-hour duration at which DDF relationships are strongly affected for longer return periods. Additionally, the depth-duration-frequency relationships for short-duration events are relatively unchanged as none of the record-breaking events was particularly intense over shorter durations.

As a consequence of the recalibration affecting rarer 1-day or 2-day events most significantly, the 150-year, 1000-year and 10,000-year events used for assessing Category D, Category C and Category B reservoirs respectively could be increased by 35% or more for critical durations around 36 hours, but only 5-15% for critical durations of 12 hours and potentially not at all for critical durations of 1 hour. The greatest proportional increases in long-return-period rainfall depths are centred on Honister Pass/Seathwaite Farm, Mosedale and Wet Sleddale, all of which are upstream of places that have been affected by severe flooding in the early 21$^{st}$ century – Cockermouth for the first two and Carlisle for the third. However, the fact that rainfall estimates are most increased on higher ground in catchment headwaters, rather than on lower-lying, more developed land, allows for flow attenuation or natural flood management (NFM) techniques to be implemented before the runoff reaches built-up areas.

As discussed in Section 2, the FEH13 method does not attempt to quantify uncertainty in the model outputs. Uncertainty due to natural variability should be lower in theory than it is in the original FEH13 method, simply because of the additional data collection over the period 2006-2016. However, model accuracy depends on whether the data collected over 2006-2016, including the two occurrences of daily rainfalls in excess of 300 mm, are representative of the true average Cumbrian climate. Unfortunately, in order to know whether these events have return periods of 100 years, 1000 years, or longer, it is necessary to collect data for several times the true return period of the event. Since the return period of the median annual rainfall is 2 years, the additional 11 years of data collection should reduce uncertainty due to natural variability in *RMED*, helping to better characterise the standardisation of annual maximum rainfalls.

Uncertainty in the structure of the FEH13 model is much more difficult to quantify. However, both the original FEH13 method and this recalibration use the same model structure. Hence, neither the record-breaking events of November 2009 and December 2015 nor other large events recorded between 2006 and 2016 were used to inform possible uncertainty-reducing revisions to the model structure at long return periods. Since there are, by the very definition of “extreme”, very few extreme events in the calibration dataset, any revisions to the model structure resulting from the few extreme events captured between 2006 and 2016 could easily be superseded when the next few extreme events are captured and integrated.

Considerable increases in total winter precipitation over the 20th century (Osborn and Maraun, 2008) and considerable increases in the magnitudes of autumn and winter rainfall-driven floods (Blöschl *et al.*, 2019) strongly suggest that climate change has altered the hydro-climate of Cumbria over the last few decades, hence that the climate of Cumbria is non-stationary when observed over this period. Otto *et al.* (2018) compare the current climate of Cumbria to that of a simulated pre-industrial climate and find that industrialisation increased the chance of the December 2015 event occurring by 40%. Faulkner *et al.* (2020) used trend tests and non-stationary analysis to conclude that non-stationary flow estimates were up to 55% higher than stationary flow estimates in north-west England.

All of the above suggests that events of a defined probability are becoming larger and events of a defined depth are becoming more frequent in Cumbria and may continue to become even more frequent in the future. However, this can be resolved without requiring an explicitly non-stationary model – a simple recalibration of the stationary FEH13 model with more recent data reduces the return periods associated with the November 2009 and December 2015 events by at least seven times and in some cases more than 25 times, depending on the duration over which total rainfall is summed. A future recalibration will change these probabilities again according to what rainfall events are observed between now and then. Griffin *et al.* (2019) highlight that just one large observation can greatly alter a flood frequency analysis, and this study shows that the same is clearly applicable to rainfall DDF analyses.

Taking the evidence of a changing climate in Cumbria together with the demonstrable outsized impact of single events, it is imperative that statistical models for rainfall DDF estimation be regularly recalibrated with the latest quality-controlled rainfall data, independently of any other considerations for model improvement.

# 10 Conclusions

This report presents the results of a recalibration of the FEH13 rainfall model in Cumbria, to account for very large rainfalls that occurred after the collation of the data set used to calibrate the initial release of the model (i.e. the model results available on the FEH Web Service). As the model was almost unchanged (fixing only one 'greater than/less than' condition in FORGEX line discretisation), the results can be considered to follow updated stationarity. This study did not attempt to compare the relative advantages and disadvantages of updated stationarity versus non-stationarity for rainfall depth-duration-frequency estimation.

The data set collated for this project includes over 7600 daily gauges, representing an increase of over 1000 compared to the Flood Studies Report (6600), Flood Estimation Handbook (6100) and current FEH13 model (6500). However, if the Cumbria 2018 model were to be applied to the whole UK, additional quality control on the area outside Cumbria may see the number of daily gauges reduced. As in the current FEH13 model, the *RMED* model developed from both daily and hourly annual maxima here uses only *SAAR* and location as explanatory variables, so is simpler than that published in the Flood Estimation Handbook, but also explains a greater proportion of variance for most durations and is unified across both data types. This report includes the first published information on the *RMED* model underpinning the 2-year FEH13 estimates available to practitioners through the FEH Web Service. The updated *RMED* model showed an increase in 1-hour *RMED* of 1-13% across all of Cumbria, but a more nuanced set of changes in 1-day *RMED*, with some parts of Cumbria (e.g. Furness Peninsula) showing no change.

The FORGEX procedure bases depth-return period analysis at a site on 'network maxima', which are the largest events recorded by any raingauge in a series of concentric circles around the site of interest. Study of network maxima suggest that the relationship between rainfall depth in mm and Gumbel reduced variate (approximately the natural logarithm of the return period in years) is best represented by an arc-shaped line for maxima based on hourly accumulations, but that an S-shaped line is more appropriate for maxima based on daily accumulations at sites near to the centres of the 2009 and 2015 storms. This is because they are brought into the analysis for small-radius circles corresponding to shorter return periods, but due to their extreme depths, are rarely exceeded even as the circle is extended to take in potential events from further and further away. The same effect is not seen in hourly data as the most extreme parts of the two record-breaking events were only captured by daily gauges.

As the network maxima are standardised (i.e. 'mapped' onto a defined growth curve with median 1 and second *L*-moment ratio 0.15), it becomes apparent that the largest events in mm depth are not the highest-quantile on the standard growth curve. For example, the estimated return period of the 2009 event is longer at High Snab Farm than either Seathwaite or Honister Pass. This is consistent with the published findings of Stewart *et al.* (2012).

The DDF model is the same as that used in the FEH13 analysis, the outputs of which are on the FEH Web Service. This unifies the 11 FORGEX lines for different hourly and daily durations into a single relationship. In common with previous UK rainfall models, there is a strong correlation between increased altitude and increased rainfall depth for any specific rainfall duration and return period. The recalibration increases the vast majority of rainfall depths of specified duration and frequency, with the greatest increases occurring for durations around one or two days. However, 1-hour rainfall depths of 100- and 1000-year return period are reduced across large parts of the north of the study area. Typical extreme-value plots in hydrology follow either arc-shaped or relatively straight lines. However, S-shaped lines were seen on FORGEX plots centred near to the spatial centres of the November 2009 and December 2015 events, like Honister Pass and Seathwaite Farm. These were most pronounced for 2-day and 4-day durations, while FORGEX lines for shorter, hourly and multi-hourly durations followed a more typical arc-shape, consistent with the descriptions of the November 2009 and December 2015 as sustained, with near-constant intensity and no peaks. There is therefore some evidence that the DDF model structure might need to be more flexible, in order to: fit arc- and S-shaped lines as part of the same unified model structure; and to allow more variation in the vertical distance between the 4- and 6-hour, and 12- and 18-hour curves in some cases.

The recalibrated FEH13 model gives new return periods to the 2009 and 2015 events as gauged at Honister Pass, Seathwaite and Thirlmere. Depending on what rainfall duration is considered, these are at least seven and sometimes more than 25 times shorter than the FEH13 return periods. They are broadly comparable to the FEH99 return periods estimated for accumulations at Seathwaite Farm (~100-200 years) and Thirlmere (~4000-8000 years), but not Honister Pass, where the FEH99 and original FEH13 model give similar return period estimates to each other (~1000 years). Spatial analysis of the November 2009 event over a 36-hour period shows that the FEH13 DDF relationship is unaffected for return periods < 10 years even at one of the durations most affected by recalibration.

The recalibrated return periods are potentially more realistic as more data were used in the calibration. However, a very high level of uncertainty is involved in defining return periods in the hundreds or thousands of years using records with typical lengths of 20-30 years. Hence, while uncertainty in these latest estimates is lower than in equivalent FEH13, FEH99 and FSR estimates, it is still high. Uncertainty can only be reduced first by recalibrating the model with more extreme events recorded over a longer time period and then by using the growing dataset of extreme events to re-evaluate the behaviour of the model at extreme return periods.

# References


Babtie Group, CEH Wallingford and Rodney Bridle Ltd. 2000. *Reservoir Safety – floods and reservoir safety: Clarification on the use of FEH and FSR design rainfalls. Final Report*. Babtie Group, Glasgow.

Benyon P R. 1976. Remark AS R15: Function Minimization Using a Simplex Procedure. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **25.1**, 97.

Blöschl G, Hall J, Viglione J, Perdigão R A P, Parajka J, Merz B, Lun D, Arheimer B, Aronica G T, Bilibashi A, Boháč M, Bonacci O, Borga M, Čanjevac I, Castellarin A, Chirico G B, Claps P, Frolova N, Ganora D, Gorbachova L, Gül A, Hannaford J, Harrigan S, Kireeva M, Kiss A, Kjeldsen T R, Kohnová S, Koskela J J, Ledvinka O, Macdonald N, Mavrova-Guirguinova M, Mediero L, Merz R, Molnar P, Montanari A, Murphy C, Osuch C, Ovcharuk V, Radevski I, Salinas J L, Sauquet E, Šraj M, Szolgay J, Volpi E, Wilson D, Zaimi K and Živković N. 2019. Changing climate both increases and decreases European river floods. *Nature* **573**, 108-111.

Chambers J M and Ertel J E. 1974. Remark AS R11: A Remark on Algorithm AS 47 'Function Minimization using a Simplex Procedure' *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **23.2**, 250-251.

Chiverrell R C, Sear D A, Warburton J, Macdonald N, Schillereff D N, Dearing J A, Croudace I W, Brown J and Bradley J. 2019. Using lake sediment archives to improve understanding of flood magnitude and frequency: Recent extreme flooding in northwest UK. *Earth Surface Processes and Landforms* **44**, 2366-2376.

Cox D R. 2003. Some comments on 10,000 year return period rainfall. Report for Defra.

Cumbria County Council. 2017. *Population Density, Mid-2016*. Cumbria County Council, Carlisle.

Defra. 2004. *Floods and Reservoir Safety - Revised Guidance for Panel Engineers*. Department for Environment, Food and Rural Affairs, London.

Efron B. 1979. Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, **7.1**, 1-26.

Environment Agency. 2006. *Cumbria Floods Technical Report: Factual report on meteorology, hydrology and impacts of January 2005 flooding in Cumbria*. Environment Agency, Bristol.

Faulkner D. 1999. *Rainfall frequency estimation (Flood Estimation Handbook volume 2)*. Institute of Hydrology, Wallingford. ISBN: 978-0-948540-90-7.

Faulkner D, Warren S, Spencer P and Sharkey P. 2020. Can we still predict the future from the past? Implementing non-stationary flood frequency analysis in the UK. *Journal of Flood Risk Management* **13**, e12582.

Gill P E, Murray W and Wright M H. 1981. *Practical Optimization*. Elsevier Academic Press, Amsterdam. ISBN: 978-0-12-283952-8.

Griffin A, Vesuviano G and Stewart E. 2019. Have trends changed over time? A study of UK peak flow data and sensitivity to observation period*. Natural Hazards and Earth System Sciences* **19**, 2157-2167.

Hill I D. 1978. Remark AS R28: A Remark on Algorithm AS 47: Function Minimization Using a Simplex Procedure. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **27.3**, 380-382.

Hosking J R M and Wallis J R. 1997. *Regional Frequency Analysis: An Approach Based on L-Moments*. Cambridge University Press, Cambridge. ISBN: 978-0-521-43045-6.

Lake District National Park Authority. 2018. *Lake District facts and figures*. Lake District National Park Authority, Kendal.

Luke A, Vrugt J A, AghaKouchak A, Matthew R and Sanders B F. 2017. Predicting nonstationary flood frequencies: Evidence supports an updated stationarity thesis in the United States. *Water Resources Research* **53.7**, 5469-5494.

MacDonald D E and Scott C W. 2001. FEH vs FSR rainfall estimates: an explanation for the discrepancies identified for very rare events. *Dams & Reservoirs* **11.2**, 28-31.

Met Office. 2001. *Quality control of rainfall observations 01/0329*. Met Office Observations Supply, Bracknell.

Met Office. 2012. *Heavy rainfall/flooding in the Lake District, Cumbria – November 2009*. Met Office, Exeter.

Met Office. 2019. *Flooding in Cumbria December 2015*. Met Office, Exeter.

NERC. 1975. *Flood Studies Report (5 volumes)*. Natural Environment Research Council, London. ISBN: 978-0-948540-50-9.

O'Neill R. 1971. Algorithm AS 47: Function Minimization Using a Simplex Procedure. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **20.3**, 338-345.

Osborn T and Maraun D. 2008. *Changing intensity of rainfall over Britain (Climatic Research Unit information sheet 15)*. University of East Anglia, Norwich.

Otto F E L, van der Wiel K, van Oldenborgh G J, Philip S, Kew S F, Uhe P and Cullen H. 2018. Climate change increases the probability of heavy rains in northern England/sothern Scotland like those of storm Desmond – a real-time event attribution revisited. *Environmental Research Letters* **13**, 024006.

Serinaldi F, Kilsby C G and Lombardo F. 2018. Untenable nonstationarity: An assessment of the fitness for purpose of trend tests in hydrology. *Advances in Water Resources* **111**, 132-155.

Stewart E J, Jones D A, Svensson C, Morris D G, Dempsey P, Dent J E, Collier C G and Anderson C W. 2013. *Reservoir Safety - Long Return Period Rainfall*. Department for Environment, Food and Rural Affairs, London.

Stewart E J, Morris D G, Jones D A and Gibson H S. 2012. Frequency analysis of extreme rainfall in Cumbria, 16–20 November 2009. *Hydrology Research* **43.5**, 649-662.

Svensson C, Folwell S, Dempsey P, Dent J and Fung C F. 2009. *Report on data consolidation, and appendices. Report No. 4 to Defra, Contract WS 194/2/39. Revised April 2009*.

Watts G, Battarbee R W, Bloomfield J P, Crossman J, Daccache A, Durance I, Elliott J A, Garner G, Hannaford J, Hannah D M, Hess T, Jackson C R, Kay A L, Kernan M, Knox J, Mackay J, Monteith D T, Ormerod S J, Rance J, Stuart M E, Wade A J, Wade S D, Weatherhead K, Whitehead P G and Wilby R L. 2015. Climate change and water in the UK – past changes and future prospects. *Progress in Physical Geography: Earth and Environment* **39.1**, 6-28.

# *Appendix A*

## Step-by-step summary of FORGEX

The step-by-step summary of the FORGEX procedure as implemented in this project is as follows. Differences between this procedure and the FEH13 procedure, both as implemented for the FEH Web Service and as described in Stewart *et al.* (2013), are noted in Section 5.4 (Changes).

1. Calculate the distance between the focal point and every valid rain gauge with within 200 km of the focal point.
2. Standardise all valid annual maxima, using *SAAR* and *RMED* grids, and northing for daily recording gauges.
3. Define the largest network as equivalent to all valid gauges identified in bullet point 1, then define progressively smaller networks, with each omitting the most distant 20% of gauges from the next largest network, until the smallest network has only two gauges. These are the 'pre-calculated' networks.
4. Find the mean easting, northing, *SAAR* and inter-gauge distance of all gauges in each network and use these to calculate the spatial dependence and therefore effective number of independent gauges in each network.
5. Build 'netmax' series for each network: this is equivalent to producing an AMAX series for the network as a whole, where each year is only represented by one gauge.
6. Determine the plotting position of each event in each netmax series.
7. Define Network 1 as the smallest possible network (not necessarily pre-calculated) with 25 + 9.5*d* valid gauge years in a *d*-km radius from the focal point. If this criterion cannot be met, then refer to Section 5.4 for additional criteria.
8. Define Network 2 as the smallest pre-calculated network with at least 15 gauges and at least one more gauge than Network 1. Define Networks 3, 4, etc. as the complete and ascending series of pre-calculated networks larger than Network 2. The largest network is the 'Primary' network.
9. Calculate minimum permitted return period and growth factor for each network. Events exceeding both are known as 'eligible'.
10. Set outer network reduction weight for each network.
11. Discretise the range of reduced variates between 0.3665 and the third-rarest event in the Primary network into segments.
12. Separately discretise the range of reduced variates from the third-rarest to the rarest event in the Primary network into segments.
13. Assign each eligible event to a segment, based on its reduced variate (i.e. return period).
14. Reduce weighting of repeat events i.e. those occurring in more than one network.
15. Generate a final weighting for each event to be used in the fitting procedure.
16. Perform constrained least-squares optimisation, using NAG routine E04NCF, in order to produce rainfall growth curves that: are produced from straight line

segments; are connected end to end; pass through the point (0.3665, 1); penalise large changes in gradient between connected segments.

17. Sample the rainfall growth curves at intervals of 0.2 in reduced variate and convert the sampled growth factors to fully-sliding rainfall depths.
18. Ensure that each combination of duration and return period has a rainfall depth at least 0.1 mm greater than the next smallest return period and 1 mm greater than the next previous duration, for daily- and hourly-durations separately.
19. If the 24-hour line is ever below the 1-day line, calculate the fractional increase (FI) required to bring its depth up to that of the 1-day line.
20. For *h*-hour durations, apply a fractional increase of $(1 + h/24)/2 \times FI_{24h}$, then check that each combination of duration and return period maintains a depth at least 0.1 mm greater than the next smallest return period and 1 mm more than the next previous duration.

## Supporting details for current FORGEX method

Supporting details to the step-by-step procedure above are presented here; each numbered bullet point in this section corresponds to the same-numbered bullet point above.

1. 'Valid' rain gauges are considered as those with nine or more valid annual maxima. If there are fewer than 75 valid hourly-recording gauges within a 200 km radius of the focal point, the radius is automatically extended until 75 valid hourly-recording gauges are included.
2. The standardisation is of the form

   $R_{revised} = 1 + (R - RMED) / (f \times RMED)$

   where *RMED* is the standardised rainfall (expressed as a growth factor), R is the unstandardized rainfall (in mm) and *f* is a site- and duration-specific scaling factor of the form:

   $f = a + b \times (1000 / SAAR)\ c \times y$

   where *SAAR* is the catchment descriptor of the same name, *y* is northing on the British national grid where 1 = 1000 km, and *a*, *b*, and *c* are coefficients derived from an ordinary least-squares fitting on $f = (\lambda_2 / (0.15 \times RMED))$, where $\lambda_2$ is the second *L*-moment (Hosking & Wallis, 1997) of the at-site annual maxima. This scaling factor was introduced in order to reduce the variation in at-site rainfall growth curves, to try to ensure that pooled data come from similar distributions. The coefficients *a*, *b* and *c* are presented in Table 8.

Table 8 Coefficients for revised standardisation of annual maxima

| Duration | Nº. catchments | *a* | *b* | *c* |
|---|---|---|---|---|
| **1 hour** | 1036 | 1.285 | 0.363 | 0 |
| **2 hours** | 1007 | 0.863 | 0.535 | 0 |
| **4 hours** | 956 | 0.646 | 0.530 | 0 |
| **6 hours** | 984 | 0.601 | 0.506 | 0 |
| **12 hours** | 956 | 0.640 | 0.433 | 0 |
| **18 hours** | 951 | 0.706 | 0.395 | 0 |
| **24 hours** | 959 | 0.771 | 0.339 | 0 |
| **1 day** | 7651 | 0.707 | 0.402 | 0.091 |
| **2 days** | 7651 | 0.608 | 0.374 | 0.236 |
| **4 days** | 7678 | 0.434 | 0.379 | 0.305 |
| **8 days** | 7656 | 0.412 | 0.339 | 0.260 |

3. The number of gauges in each progressively smaller network is defined as 80% of that in the next larger network, always rounded down (i.e. floor function). This means that, for example, 204.8 is rounded down to 204 gauges.
4. Effective number of gauges is calculated according to

   $N_{eff} = Ne^{(1 - \gamma \ln(N))}$

   where $N_{eff}$ and $N$ are effective and actual number of gauges, and γ is spatial dependence, where

   $\gamma = a + b \times \ln(2.5 \times distg^2)\ c \times (\ln(N)/(1 + 0.5\ln(N)) + d \times (SAAR/1000)$

   where *distg* is mean inter-gauge distance, *SAAR* is the catchment descriptor of the same name and *a*, *b*, *c* and *d* are fitted coefficients, as detailed in Table 9.

Table 9 Coefficients for calculating effective number of gauges in a network

| Duration | *a* | *b* | *c* | *d* |
|---|---|---|---|---|
| **1 hour** | 0.191 | –0.016 | –0.034 | 0.074 |
| **2 hours** | 0.256 | -0.017 | –0.033 | 0.029 |
| **4 hours** | 0.409 | –0.031 | –0.011 | –0.008 |
| **6 hours** | 0.464 | –0.032 | –0.006 | –0.043 |
| **12 hours** | 0.613 | –0.039 | –0.026 | –0.067 |
| **18 hours** | 0.687 | –0.047 | –0.009 | –0.067 |
| **24 hours** | 0.669 | –0.039 | –0.036 | –0.086 |
| **1 day** | 0.822 | –0.060 | 0.073 | –0.109 |
| **2 days** | 0.842 | –0.063 | 0.089 | –0.083 |
| **4 days** | 0.829 | –0.068 | 0.130 | –0.051 |
| **8 days** | 0.873 | –0.066 | 0.115 | –0.058 |

It should be noted that the effective number of gauges per year can vary, as it is unlikely that every gauge in a network recorded a valid annual maximum during every year represented in the network. Because of this, γ is calculated separately for each year, based on the network-mean *SAAR* and inter-gauge distance of only the gauges active in each year.

5. The 'netmax' series uses standardised rainfalls, so the 'largest' rainfall in any given year, hence the rainfall representing that year, may not be the absolute largest in mm.
6. The plotting position for each event in a network is solved iteratively, but is generally based upon Gringorten plotting positions of ordered events in a series comprising as many independent gauge-years as are found cumulatively across all network-years.
7. If 25 + 9.5*d* independent gauge-years cannot be found within a *d*-km radius of the focal point, then the next attempted definition of Network 1 is the smallest possible network with ≥ 120 effective gauge-years and ≥ 5 gauges, with a maximum permitted radius defined as the smallest of $r_{40}$ + 15 km and $r_{40}^3$ km, where $r_{40}$ is the smallest radius to contain 40 effective gauge-years (if $r_{40}^3$ is less than 15 km, then a radius of 15 km is used).
8. The lower limit of 15 gauges for Network 2 is set so that all points (particularly adjacent ones) have a similar number of gauges in Network 2, 3, 4, etc. regardless of potential differences in the size of Network 1.
9. Within the FORGEX fitting procedure, only events that fall within return periods defined by the beginning and end of a segment are used to fit that segment. These events are permitted to originate from any network, although lower segments will not be fitted to events from higher networks due to the eligibility criteria mentioned in bullet point 9 above.

   The minimum permitted return period for each network is defined in terms of Gumbel reduced variate and is as follows:
   Network 1: –0.6335
   Network 2: Effective gauge-years in previous network (Network 1) / 16
   Network 3: Effective gauge-years in previous network (Network 2) / 8
   Networks 4 to Primary: Effective gauge-years in previous network / 2

   If Network 1 has a radius under 10 km, then the minimum return period for Network 2 is adjusted by a factor 1 + 0.2 × (10 – Network 1 radius)

   There is no maximum permitted return period for any network.

   The minimum permitted standardised growth factor for each network is defined as follows:
   Network 1: –0.6335
   Networks 2 to Primary: That of the rarest event (in terms of reduced variate) whose reduced variate is below the minimum permitted reduced variate for the next largest network.

10. Outer network reduction weight is set as follows:
    All networks with radius ≤ 200/6 km: 1
    All networks with radius > 200/6 km:
    1 – 0.9 × (network radius – 200/6) / (5 × 200/6)
11. Segments between 0.3665 and the reduced variate of the third-rarest event in any network are defined according to the following rules:
    Each segment occupies the same width in terms of reduced variate

There are an integer number of segments, equal to the width from 0.3665 to the third-rarest event rounded to the nearest integer.

11.1 A segment of width 1 is always defined, starting at a reduced variate of −0.6335 and ending at a reduced variate of 0.3665

12. Segments from the third-rarest to the rarest event are defined according to the following rules:
Each segment occupies the same width in terms of reduced variate
There are an integer number of segments, with minimum width 1 and maximum width 2 on the reduced variate scale.

13. Each event is assigned to a segment, according to which segment encompasses the reduced variate of the event.

14. Events that occur in more than one network are assigned a weight of 1 for their first occurrence, $0.9^{0.5}$ for their second occurrence, $0.9^{1}$ for their third occurrence, $0.9^{1.5}$ for their fourth occurrence, etc., counting the first occurrence as belonging to the smallest network to contain that event, second occurrence as belonging to the second-smallest network, etc.

15. The final weighting for each event is calculated as:
outer network reduction weight × repeat event reduction factor

16. Each segment is fitted using only the events that have been assigned to it.

17-20. One key site where rainfalls are not consistent across durations and return periods is Thirlmere. There, only a daily-recording gauge was able to capture the core of the December 2015 event, resulting in the 1-day line briefly exceeding the 24-hour line, even before discretisation conversion factors are applied.

Figure A1 uses the fitted FORGEX lines at Thirlmere to demonstrate how the consistency rules are applied.

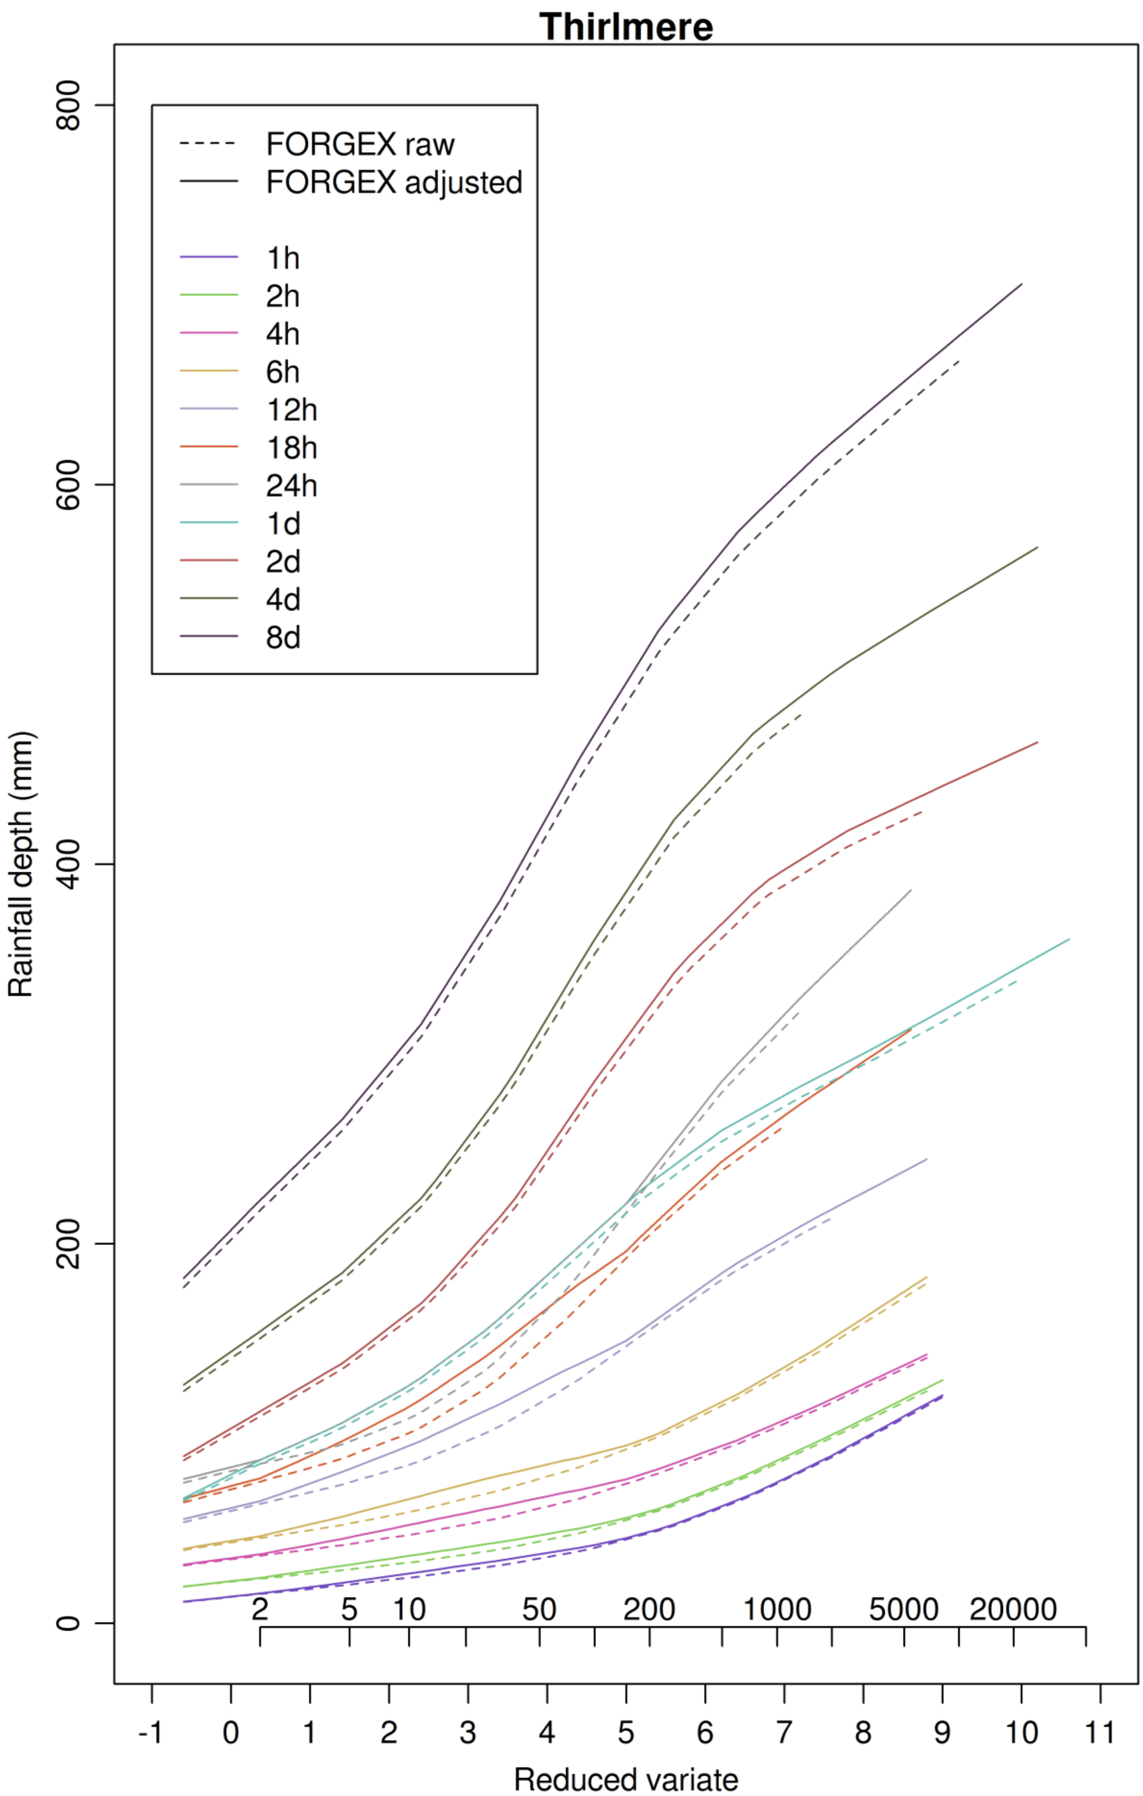


Figure A1 – comparison of FORGEX lines at Thirlmere before and after consistency rules ('adjustments') are applied

**BANGOR**
UK Centre for Ecology & Hydrology
Environment Centre Wales
Deiniol Road
Bangor
Gwynedd
LL57 2UW
United Kingdom
T: +44 (0)1248 374500
F: +44 (0)1248 362133

**LANCASTER**
UK Centre for Ecology & Hydrology
Lancaster Environment Centre
Library Avenue
Bailrigg
Lancaster
LA1 4AP
United Kingdom
T: +44 (0)1524 595800
F: +44 (0)1524 61536

**EDINBURGH**
UK Centre for Ecology & Hydrology
Bush Estate
Penicuik
Midlothian
EH26 0QB
United Kingdom
T: +44 (0)131 4454343
F: +44 (0)131 4453943

**WALLINGFORD (Headquarters)**
UK Centre for Ecology & Hydrology
Maclean Building
Benson Lane
Crowmarsh Gifford
Wallingford
Oxfordshire
OX10 8BB
United Kingdom
T: +44 (0)1491 838800
F: +44 (0)1491 692424

enquiries@ceh.ac.uk

www.ceh.ac.uk