Developing a global operational seasonal hydro-meteorological forecasting system: GloFAS-Seasonal v1.0

Rebecca Emerton[1,2], Ervin Zsoter[2,1], Louise Arnal[1,2], Hannah L. Cloke[1,3], Davide Muraro[6], Christel Prudhomme[2,4,5], Elisabeth M. Stephens[1], Peter Salamon[7] and Florian Pappenberger[2]

[1]University of Reading, Reading, UK
[2]European Centre for Medium Range Weather Forecasts (ECMWF), Reading, UK
[3]Uppsala University, Uppsala, Sweden
[4]Centre for Ecology and Hydrology (CEH), Wallingford, UK
[5]University of Loughborough, Loughborough, UK
[6]Image Recognition Integrated Systems (IRIS) Italia, Ispra, Italy
[7]European Commission, Joint Research Centre (JRC), Ispra, Italy

Correspondence to: Rebecca Emerton (r.e.emerton@pgr.reading.ac.uk)

Abstract. Global overviews of upcoming flood and drought events are key for many applications, including disaster risk reduction initiatives. Seasonal forecasts are designed to provide early indications of such events weeks, or even months, in advance, but seasonal forecasts for hydrological variables at large or global scales are few and far between. Here, we present the first operational global scale seasonal hydro-meteorological forecasting system: GloFAS-Seasonal. Developed as an extension of the Global Flood Awareness System (GloFAS), GloFAS-Seasonal couples seasonal meteorological forecasts from ECMWF with a hydrological model, to provide openly available probabilistic forecasts of river flow out to 4 months ahead for the global river network. This system has potential benefits not only for disaster risk reduction through early awareness of floods and droughts, but also for water-related sectors such as agriculture and water resources management, in particular for regions where no other forecasting system exists. We describe the key hydro-meteorological components and computational framework of GloFAS-Seasonal, alongside the forecast products available, before discussing initial evaluation results and next steps.

1 Introduction

Seasonal meteorological forecasts simulate the evolution of the atmosphere over the coming months. They are designed to provide an early indication of the likelihood that a given variable, for example precipitation or temperature, will differ from normal conditions, weeks or months ahead. Will a particular region be warmer or cooler than normal during the next summer? Or will a river have higher or lower flow than normal next winter? Seasonal forecasts of river flow have the potential to benefit many water-related sectors, from agriculture and water resources management, to disaster risk reduction and humanitarian aid through earlier indications of floods or droughts.

Many operational forecasting centres produce long-range (seasonal) global forecasts of meteorological variables, such as precipitation (Weisheimer and Palmer, 2014). However, at present, operational seasonal forecasts of hydrological variables, particularly for large or global scales, are few and far between. A number of continental scale seasonal hydro-meteorological forecasting systems have begun to emerge around the globe over the past decade (Yuan et al., 2015a), using seasonal meteorological forecasts as input to hydrological models to produce forecasts of hydrological variables. These include the European Flood Awareness System (EFAS; Arnal et al., 2018; Cloke et al., 2013), the European Service for Water Indicators in Climate Change Adaptation (SWICCA; Copernicus, n.d.), the Australian Government Bureau of Meteorology Seasonal Streamflow Forecasts (Bennett et al., 2017; BoM, 2018) and the USA's National Hydrologic Ensemble Forecast Service (HEFS; Demargne et al., 2014; Emerton et al., 2016). There are also various ongoing research efforts using seasonal hydro-meteorological forecasting systems for forecast applications and research purposes at regional (Bell et al., 2017; Bennett et al., 2016; Crochemore et al., 2016; Meißner et al., 2017; Mo et al., 2014; Prudhomme et al., 2017; Wood et al., 2002, 2005; Yuan et al., 2013) and global (Candogan Yossef et al., 2017; Yuan et al., 2015b) scales. In addition to the ongoing research into improved seasonal hydro-meteorological forecasts at the global scale, an operational system providing consistent global scale seasonal forecasts of hydrological variables could be of great benefit in regions where no other forecasting system exists, and to organisations operating at the global scale (Coughlan De Perez et al., 2017).

Often, in the absence of hydrological forecasts, seasonal precipitation forecasts are used as a proxy for flooding. It has been shown that forecasts of seasonal total rainfall, the most oft-used seasonal precipitation forecasts, are not necessarily a good indicator of seasonal floodiness (Stephens et al., 2015), and other measures of rainfall patterns, or seasonal hydrological forecasts, would be better indicators of potential flood hazard (Coughlan De Perez et al., 2017).

While it seems a natural next step to produce global scale seasonal hydro-meteorological forecasts, this is not a simple task, not only due to the complexities of geographical variations in rainfall-runoff processes and river regimes across the globe, but also due to the computing resources required and huge volumes of data that must be efficiently processed and stored, and the challenge of effectively communicating forecasts for the entire globe. Indeed, global scale forecasting for medium-range timescales has only become possible in recent years due to the integration of meteorological and hydrological modelling capabilities, improvements in data, satellite observations and land-surface hydrology modelling, and increased resources and computer power (Emerton et al., 2016). In addition to continued improvements in computing capabilities, the recent move towards the development of coupled atmosphere-ocean-land models means that it is now becoming possible to produce seasonal hydro-meteorological forecasts for the global river network.

Despite the chaotic nature of the atmosphere (Lorenz, 1963), which introduces a limit of predictability (generally accepted to be ~2 weeks), seasonal predictions are possible as they rely on components that vary on longer timescales and are themselves

predictable to an extent. This "second type predictability" (Lorenz, 1993) for seasonal river flow forecasts comes from the initial conditions, and large-scale modes of climate variability. The most prominent pattern of climate variability is the El Niño Southern Oscillation (ENSO; McPhaden et al., 2006), which is known to affect river flow and flooding across the globe (Chiew and McMahon, 2002; Emerton et al., 2017; Guimarães Nobre et al., 2017; Ward et al., 2014a, 2014b, 2016). Other

5 teleconnections also influence river flow in various regions of the globe, such as the North Atlantic Oscillation (NAO), Southern Oscillation (SOI), Indian Ocean Dipole (IOD) and Pacific Decadal Oscillation (PDO), and contribute to the seasonal predictability of hydrologic variables (Yuan et al., 2015a). Coupled atmosphere-ocean-land models are key in representing these large-scale modes of variability in order to produce seasonal hydro-meteorological forecasts.

10 This motivates the development of an operational global scale seasonal hydro-meteorological forecasting system as an extension of the Global Flood Awareness System (GloFAS; Alfieri et al., 2013), with openly available forecast products. GloFAS is developed by the European Centre for Medium-Range Weather Forecasts (ECMWF) and the European Commission Joint Research Centre (JRC), and has been producing probabilistic flood forecasts out to 30 days for the entire globe since 2012. In 2016, work began, in collaboration with the University of Reading, to implement a seasonal outlook in GloFAS,

15 aiming to provide forecasts of both high and low river flow for the global river network, up to several months in advance. On 10[th] November 2017, the first GloFAS seasonal river flow forecast was released. This paper introduces the modelling system, its implementation and the available forecast products, and provides an initial evaluation of the potential usefulness and reliability of the forecasts.

2 Implementation

20 The GloFAS seasonal outlooks are produced by driving a hydrological river routing model with meteorological forecasts from ECMWF. The forecasts are run operationally on the ECMWF computing facilities. This section provides an overview of the computing facilities, introduces the key hydro-meteorological components of the modelling platform (the meteorological forecast input, hydrological model and reference climatology), and describes the computational framework of GloFAS-Seasonal.

25 **2.1 ECMWF High Performance Computing Facility**

ECMWF's current High Performance Computing Facility (HPCF) has been in operation since June 2016, and is used for both forecast production and research activities. The HPCF comprises two identical Cray XC40 supercomputers, each of which is self-sufficient with their own storage, and each with equal access to the storage of the other. Each Cray XC40 consists of 20 cabinets of compute notes and 13 storage nodes. One compute node has 2 Intel Broadwell processors, each with 18 cores,

30 giving 192 nodes (6912 cores) per cabinet. The Cray Aries interconnect is used to connect the processing power. The majority

of the nodes of the HPCF are run using the high performance Cray Linux Environment, a stripped-down version of Linux, as reducing the number of operating system tasks is critical for providing a highly scalable environment.

In terms of storage, each Cray XC40 has ~10PB of storage, and the Data Handling System (DHS) also comprises two main applications; the Meteorological Archive and Retrieval System (MARS), which stores and provides access to meteorological data collected or produced by ECMWF, and ECFS, which stores data that is not suitable for storing on MARS. The DHS holds over 210PB of primary data, and the archive increases by ~233TB per day. The reader is referred to the ECMWF website, *www.ecmwf.int*, for further information on the HPCF and DHS.

In addition to the Cray XC40s, the ECMWF computing facility also includes 4 Linux clusters consisting of 60 servers and 1PB of storage. The Linux clusters are currently used to run the river routing model used in GloFAS and to produce the forecast products, while the meteorological forcing and ERA5 reanalysis are produced on the HPCF. All data related to GloFAS-Seasonal are stored on the MARS and ECFS archives.

2.2 Hydro-Meteorological Components

2.2.1 Meteorological Forcing

The first model component of the seasonal outlook is the meteorological forecast input from the ECMWF Integrated Forecast System (IFS, cycle 43r1 (ECMWF, 2018b)). GloFAS-Seasonal makes use of SEAS5, which is the latest version of ECMWF's long-range ensemble forecasting system, made operational in November 2017 (ECMWF, 2017a; Stockdale et al., 2018). SEAS5 consists of 51 ensemble members (50 perturbed members and one unperturbed control member) and has a horizontal resolution of ~36km ($T_{CO}319$). The system, which comprises a data assimilation system and a global circulation model, is run once a month, producing forecasts out to 7 months ahead. Initial pre-implementation testing of SEAS5 has suggested that in comparison to the previous version (System 4), SEAS5 better simulates sea surface temperatures (SSTs) in the Pacific Ocean, leading to improved forecasts of El Niño Southern Oscillation (ENSO, (Stockdale et al., 2018)), which is closely linked to river flow across the globe and can provide added predictability.

SEAS5 is a configuration of the ECMWF IFS (cycle 43r1), including atmosphere-ocean coupling to the NEMO ocean model. SEAS5 is run operationally on the HPCF. Each ensemble member is a complex, HPC-intensive massively parallel code, written in Fortran (version F90). In addition, further complex scripting systems are required to control, prepare, run, post-process and archive all IFS forecasts. The data assimilation systems used to prepare the initial conditions for the forecasts also make use of Fortran and run on the HPCF. For further information, the reader is referred to the IFS documentation (ECMWF, 2018b).

2.2.2 Land Surface Component

Within the IFS, which includes SEAS5, the Hydrology Tiled ECMWF Scheme of Surface Exchanges over Land, HTESSEL (Balsamo et al., 2011), is used to compute the land surface response to atmospheric forcing. HTESSEL simulates the evolution of soil temperature, moisture content and snowpack conditions through the forecast horizon, to produce a corresponding forecast of surface and subsurface runoff. This component allows for each grid box to be divided into tiles, with up to 6 tiles per grid box (bare ground, low and high vegetation, intercepted water and shaded and exposed snow), describing the land surface. For a given precipitation, the scheme distributes the water as surface runoff and drainage, with dependencies on orography and soil texture. An interception layer accumulates precipitation until saturation is reached, with the remaining precipitation partitioned between surface runoff and infiltration. HTESSEL also accounts for frozen soil, redirecting the rainfall and snowmelt to surface runoff when the uppermost soil layer is frozen, and incorporates a snow scheme. Four soil layers are used to describe the vertical transfer of water and energy, with subsurface water fluxes determined by Darcy's law, and each layer has a sink to account for root extraction in vegetated areas. A detailed description of the hydrology of HTESSEL is provided by Balsamo et al., (2011).

HTESSEL comprises a Fortran library of ~20,000 lines of code, using both F77 and F90 Fortran versions, and is implemented modularly. While HTESSEL can be run on diverse architectures from a workstation PC to the HPCF, operationally, it is run on the HPCF.

2.2.3 River Routing Model

As HTESSEL does not simulate water fluxes through the river network, Lisflood (Van Der Knijff et al., 2010), driven by the surface and sub-surface runoff output from HTESSEL interpolated to the 0.1º (~10km) spatial resolution of Lisflood, is used to simulate the groundwater (subsurface water storage and transport) processes and routing of the water through the river network. The initial conditions, used to start the Lisflood model, are taken from the ERA5-R river flow reanalysis (see Sect. 2.2.4).

Lisflood is a spatially distributed hydrological model, including a 1-D channel routing model. Groundwater processes are modelled using two linear reservoirs, the upper zone representing a quick runoff component, including subsurface flow through soil macropores and fast groundwater, and the lower zone representing a slow groundwater component fed by percolation from the upper zone. The routing of surface runoff to the outlet of each grid cell, and the routing of runoff produced by every grid cell from surface, upper and lower groundwater zones through the river network, is done using a four-point implicit finite-difference solution of the kinematic wave equations (Chow et al., 2010). The river network used is that of HydroSHEDS (Lehner et al., 2008), again interpolated to a 0.1º spatial resolution, using the approach of Fekete et al. (2001). For a detailed account of the Lisflood model set-up within GloFAS, the reader is referred to Alfieri et al. (2013).

Lisflood is implemented using a combination of PCRaster GIS and Python, and is currently run operationally on the Linux cluster at ECMWF.

2.2.4 Generation of Reforecasts and Reference Climatology

5   In order to generate a reference climatology for GloFAS-Seasonal, the latest of ECMWF's reanalysis products, ERA5, was used. Reanalysis datasets combine historical observations of the atmosphere, ocean and land surface with a data assimilation system; using global models to "fill in the gaps" and produce consistent global best estimates of the atmosphere, ocean and land state. ERA5 represents the current state of the art in terms of reanalysis datasets, providing a much higher spatial and temporal resolution (30km, hourly) compared to ERA-Interim (79km, 3-hourly), and better representations of precipitation,
10   evaporation and soil moisture (ECMWF, 2017b). In order to produce a river flow reanalysis (ERA5-R) for the global river network, the ERA5 surface and subsurface runoff variables were interpolated to 0.1º (~10km) resolution and used as input to the Lisflood model (see Sect. 2.2.3). ERA5 is currently still in production, and while it will cover the period from 1950 to present when completed, the full dataset will not be available until 2019. ERA5 is being produced in three "streams" in parallel; at the time of producing the ERA5-R reanalysis, 18 years of ERA5 data were available across the three streams (1990-1992,
15   2000-2007 & 2010-2016). In addition to the historical climatology, ERA5 is also produced in near-real-time, with a delay of just ~3 days, allowing its use as initial conditions for the river routing component of the GloFAS-Seasonal forecasts. The ERA5-R reanalysis is thus updated every month prior to producing the forecast. Figure 2 provides an overview of all datasets used in and produced for the development of GloFAS-Seasonal.

20   Once the ERA5-R reanalysis was obtained, a set of GloFAS-Seasonal reforecasts was produced. From the 25-ensemble-member SEAS5 reforecasts produced by ECMWF, the surface and subsurface runoff variables were used to drive the Lisflood model, with initial conditions from ERA5-R. This generated 18 years of seasonal river flow reforecasts (one forecast per month out to 4 months lead time, with 25 ensemble members at 0.1º resolution). It is the weekly-averaged river flow from this reforecast dataset which is used as a reference climatology, including to calculate the high and low flow thresholds used in the
25   real-time forecasts (described in Sect. 2.2.4).

2.3 GloFAS-Seasonal Computational Framework

The GloFAS-Seasonal real-time forecasts are implemented and run operationally on the ECMWF computing facilities using ecFlow (Bahra, 2011; ECMWF, 2012), an ECMWF work package used to run large numbers of programs with dependencies on each other and on time. An ecFlow suite is a collection of tasks and scheduling instructions, with a user interface allowing
30   interaction and monitoring of the suite, the code behind it, and the output. The GloFAS-Seasonal suite is run once per month,

and is used to retrieve the raw SEAS5 forecast data, run this through Lisflood and produce the final forecast products and visualisations using the newly developed GloFAS-Seasonal postprocessing code.

The GloFAS-Seasonal suite performs tasks (detailed below) such as retrieving data, running Lisflood, computing weekly averages and forecast probabilities from the raw Lisflood river flow forecast data, and producing maps and hydrographs for the interface. It is primarily written in Python (version 2.7), with some elements written in R (version 3.1) and shell scripts incorporating Climate Data Operators (CDO). The code was developed and tested on OpenSUSE Leap 42 systems.

When a new SEAS5 forecast becomes available (typically on the 5th of the month at 00:00UTC), the GloFAS-Seasonal ecFlow suite is automatically deployed. The structure of, and tasks within, the ecFlow suite are shown in Fig. 3. Each 'task' represents one script from the GloFAS-Seasonal code. The suite first retrieves the latest raw SEAS5 forecast surface and sub-surface variables for all 51 ensemble members (*stagefc* and *getfc* tasks), alongside the river flow reference climatology (see Sect. 2.2.4) for the corresponding month of the forecast (*copywb* task). The Lisflood river routing model (described in Sect. 2.2.3) is then run for each of the 51 ensemble members (*lisflood* task). Lisflood is initialised using the ERA5-R river flow reanalysis (see Sect. 2.2.4), and driven with the SEAS5 surface and sub-surface runoff forecast, to produce the 4-month ensemble river flow forecast at a daily time step, from which the weekly-averaged ensemble river flow forecast is obtained (*average* task). The weekly averages are computed for every Monday-Sunday, starting from the first Monday of each month, so that the weekly averages correspond from one forecast to the next. While SEAS5 provides forecasts out to 7 months ahead, the first version of GloFAS-Seasonal uses only the first 4 months. This is in order to reduce the data volumes required, and to allow assessment of the forecast skill out to 4 months ahead, before possible extension of the forecasts out to 7 months ahead in the future.

Once the weekly averaging is complete, the *'forecast product'* section of the suite is deployed, which post-processes the raw forecast output to produce the final forecast products displayed on the web interface. The code behind the *'forecast product'* section is provided in the supplementary material. For a full description of the forecast products, including examples, see Sect. 3. The suite computes the full forecast distribution (*distribution* task), followed by the probability of exceedance for each week of the forecast and for every grid point (*probability* task), based on the number of ensemble members exceeding the high flow threshold or falling below the low flow threshold. The high and low flow thresholds are defined as the 80th and 20th percentiles of the reference climatology, for the week of the year corresponding to the forecast week, so as to use thresholds based on time of year of the forecast. From these weekly exceedance probabilities, the maximum probability of exceedance across the 4-month forecast horizon is calculated for each grid point (*maxprob* task). Basin-averaged maximum probabilities are also produced (*basinprob* task), by calculating the mean maximum probability of exceedance across every grid point at which the upstream area exceeds 1500km^2 in each of the 306 major world river basins used in GloFAS-Seasonal (see Sect. 3.1). A minimum upstream area of 1500km^2 is chosen as the current resolution of the global model is such that reliable forecasts for very small rivers are not feasible.

These probabilities are used to produce the forecast visualisation for the web interface (Sect. 3). Firstly, the *map* task produces colour-coded maps of both the river network, again for grid points at which the upstream area exceeds 1500km^2, and the major world river basins. The *reppoint* task then produces an ensemble hydrograph and persistence diagrams for a subset of grid points (the 'reporting points') across the globe. Further details on the location of reporting points are given in Sect. 3.3. Finally, the *web* task collates and subsequently transfers all data required for the web interface.

This process, from the time a new SEAS5 forecast becomes available, takes ~4 hours on average to complete, with up to 10 tasks running in parallel (for example, running Lisflood for 10 ensemble members at the same time). It is possible to speed up this process by running more ensemble members in parallel, however, the speed is sufficient that it is not necessary to use further resources to produce the forecast more quickly. GloFAS-Seasonal forecast products are typically produced by the 5th of the month at 05:00UTC and made available via the web interface on the 10th of the month at 01:00UTC. This is the earliest that the GloFAS-Seasonal forecasts can be provided publicly, under the Copernicus license agreement. Data is automatically archived at ECMWF as the suite runs in real-time; ~285GB of data from each SEAS5 forecast are used as input for GloFAS-Seasonal. Each GloFAS-Seasonal forecast run produces an additional ~1.8TB of data, and makes use of the ~18TB reference climatology.

2.4 GloFAS Web Interface

The GloFAS website is based on a User-Centred Design (UCD), meaning that user needs are core to the design principles (ISO13407). The website uses Web 2.0 concepts such as simplicity, joy of use and usability, that are synonymous with engaging users. It is a Rich Internet Application (RIA), aiming to provide the same level of interactivity and responsiveness as desktop applications. The website is designed for those engaged in flood forecasting and water resources, as users can browse various aspects of the current forecast or past forecasts in a simple and intuitive way, with spatially distributed information. Map layers containing different information, e.g. flood probabilities for different flood severities, precipitation forecasts, seasonal outlooks, etc. can be activated, and the user can also choose to overlay other information such as land use, urban areas or flood hazard maps. The interface consists of three principal modules: MapServer, GloFAS Web Map Service Time and the Forecast Viewer. These are outlined below.

2.4.1 MapServer

MapServer (Open Source Geospatial Foundation, 2016) is an open source development environment for building spatially-enabled internet applications, developed by the University of Minnesota. MapServer has built-in functionality to support industry standard data formats and spatial databases, which is significant to this project, and the support of popular Open Geospatial Consortium (OGC) standards including WMS. In order to exploit the potential of asynchronous data transfer

between server and client, the GloFAS raster data has to be divided into a grid of adequate dimensions and an optimal scale sequence.

2.4.2 GloFAS Web Map Service Time

The OpenGIS Web Map Service (WMS) is a standard protocol for serving geo-referenced map images over the internet. A Web Map Service Time (WMS-T) is a web service that produces maps in several raster formats or in vector format that may come simultaneously from multiple remote and heterogeneous sources. A WMS server can provide support to temporal requests (WMS-T), by providing a TIME parameter with a time value in the request.

The WMS Specification (OGC, 2015) describes three HTTP requests; *GetCapabilities*, *GetMap* and *GetFeatureInfo*. *GetCapabilities* returns an XML document describing the map layers available and the server's capabilities (i.e. the image formats, projections, and geographic bounds of the server). *GetMap* returns a raster map image. The request arguments, such as the layer id and image format should match those listed as available in the GetCapabilities return document. *GetFeatureInfo* is optional, and is designed to provide WMS clients with more information about features in the map images that were returned by earlier GetMap requests. The response should contain data relating to the features nearest to an image coordinate specified in the GetFeatureInfo request. The structure of the data returned is not defined in the specification and is left up to the WMS server implementation. The GloFAS WMS-T (GloFAS, 2018b) can be freely used, allowing access to the GloFAS layers in any GIS environment, such as QGIS (QGIS Development Team, 2017) or ArcMAP (Environmental Systems Research Institute, 2018). The user manual for the GloFAS WMS-T is available via the GloFAS website (GloFAS, 2018a).

2.4.3 Forecast Viewer

The GloFAS forecast viewer is based on the Model View Controller (MVC) architectural pattern used in software engineering. The pattern isolates "domain logic" (the application logic for the user) from input and presentation (User Interface, UI), permitting independent development, testing and maintenance of each. A fundamental part of this is the AJAX (Asynchronous JavaScript and XML) technology used to enhance user-friendly interfaces for web mapping applications. AJAX technologies have a number of benefits; the essential one is removing the need to reload and refresh the whole page after every event. Careful application design and component selection results in a measurably smaller web server load in geodata rendering and publishing, as there is no need to link and send the whole html document, just the relevant part that needs to be changed.

GloFAS uses OpenLayers (OpenLayers, 2018) as a WMS client. OpenLayers is a JavaScript-based web mapping toolkit designed to make it easy to put a dynamic map on any web page. It doesn't depend on the server technology and can display a set of vector data, such as points, with aerial photographs as backdrop maps from different sources. Closely coupled to the

map widget is a layer manager that controls which layers are displayed with facilities for adding, removing and modifying layers. The new layers associated with GloFAS-Seasonal are described in the following section.

3 Forecast Products

The GloFAS seasonal outlook is provided as three new forecast layers in the GloFAS forecast viewer: the basin overview, river network and reporting point layers. Each of the three layers represents a different forecast product, described in the following sections. Information on each of the layers is also provided for end users of the forecasts under the dedicated 'Seasonal Outlook' page of the GloFAS website.

3.1 Basin Overview Layer

The first GloFAS seasonal outlook product is designed to provide a quick global overview of areas that are likely to experience unusually high or low river flow over the coming 4 months. The "Basin Overview" layer displays a map of 306 major world river basins, colour-coded according to the maximum probability of exceeding the high (blue) or low (orange) flow thresholds (the 80^{th} and 20^{th} percentiles of the reference climatology, respectively) during the 4-month forecast horizon. This value is calculated for each river basin by taking the average of the maximum exceedance probabilities at each grid cell within the basin (using only river pixels with an upstream area >1500km^2). The three different shades of orange / blue indicate the probability: dark (>90%), medium (75-90%) and light (50-75%). Basins that remain white are those where the probability of unusually high or low flow does not exceed 50% during the 4-month forecast horizon. An example is shown in Fig. 4.

As mentioned in Sect. 2.2.3, the Lisflood river network is based on HydroSHEDS (Lehner et al., 2008). In order to generate the river basins used in GloFAS-Seasonal, the corresponding HydroBASINS (Lehner and Grill, 2013) data were used. HydroBASINS consists of a suite of polygon layers depicting watershed boundaries at the global scale. These watersheds were manually merged using QGIS (QGIS Development Team, 2017) to create a global polygon layer of major river basins based on the river network used in the model.

3.2 River Network Layer

The second map layer provides similar information at the sub-basin scale, by colour-coding the entire model river network according to the maximum exceedance probability during the 4-month forecast horizon. This allows the user to zoom in to their region of interest and view the forecast maximum exceedance probabilities in more detail. Again, only river pixels with an upstream area >1500km^2 are shown. The same colour scheme is used for both the basin overview and river network layers, with blue indicating high flow (exceeding the 80^{th} percentile) and orange low flow (falling below the 20^{th} percentile) and darker colours indicating higher probabilities. In the river network layer, additional colours also represent areas where the forecast does not exceed 50% probability of exceeding either the high or low flow threshold (light grey), and where the river pixel lies

in a climatologically arid area and the forecast probability cannot be defined (darker grey-brown). Examples of the river network layer can be seen in both Fig. 4 (globally) and Fig. 5 (zoomed in).

3.3 Reporting Points Layer

In addition to the two summary map layers, reporting points are provided at both static and dynamic locations throughout the
5   global river network, providing additional forecast information; an ensemble hydrograph and a persistence diagram.

Static points originally consisted of a selection of gauged river stations included in the Global Runoff Data Centre (GRDC; BfG, 2017); this set of points has since been expanded to further include points at locations of particular interest to GloFAS partners. There now exist ~2200 static reporting points in the GloFAS interface.

Dynamic points are generated to provide the additional forecast information throughout the global river network, including river reaches where there are no static points. These points are obtained for every new forecast based on a set of selection criteria, adapted from the GloFAS flood forecast dynamic point selection criteria (Alfieri et al., 2013):

15   -   The maximum probability of high [low] river flow (exceeding [falling below] the 80[th] [20[th]] percentile of the reference climatology) during the 4-month forecast horizon must be ≥50% for at least 5 contiguous pixels of the river network.

     -    The upstream area of the selected point must be ≥4000km^2.

20   -   Dynamic reporting points are generated starting from the most downstream river pixel complying with the previous two selection criteria. A new reporting point is then generated every 300km upstream along the river network, unless a static reporting point already exists within a short distance of the new dynamic point, or the forecasts further upstream no longer comply with the previous two criteria.

25   Reporting points are displayed as black circles in the "reporting points" seasonal outlook layer. An example is shown in Fig. 5. Clicking on a reporting point brings up a new window, containing a hydrograph and persistence diagram alongside some basic information about the location, such as the latitude and longitude, and the upstream area of the point in the model river network. The number of dynamic reporting points can vary from one forecast to the next due to the criteria applied; for example, the March 2018 forecast included ~1600 dynamic points in addition to the static points, thus ~3800 reporting points
30   were available globally.

The ensemble hydrographs (also shown in Fig. 5) display a fan plot of the ensemble forecast of weekly-averaged river flow out to 4 months, indicating the spread of the forecast and associated probabilities. Also shown are thresholds based on the reference climatology; the median, and the 80th and 20th percentiles. These thresholds are displayed as a three-week moving average of the weekly-averaged river flow for the given threshold, for the same months of the climatology as that of the forecast (i.e. a forecast for J-F-M-A also displays thresholds based on the reference climatology for J-F-M-A). This allows comparison of the forecast to typical and extreme conditions for the time of year.

Persistence diagrams (see Fig. 5) show the weekly probability of exceeding the high and low flow thresholds, for the current forecast (bottom row) and previous three forecasts, colour-coded to match the probabilities indicated in the map layers. These diagrams are provided in order to highlight the evolution of the forecast, which can indicate whether the forecast is progressing consistently, or whether behaviour is variable from month to month.

4 Forecast Evaluation

In this section, the GloFAS-Seasonal reforecasts are evaluated using historical river flow observations. Benchmarking a forecasting system is important to evaluate and understand the value of the system, and in order to communicate the skill of the forecasts to end users (Pappenberger et al., 2015). This evaluation is designed to measure the ability of the forecasts to predict the correct category of an 'event', i.e. the ability of the forecast to predict that weekly-averaged river flow will fall in the upper 80th or lower 20th percentile of climatology, using a climatology of historical observations as a benchmark. This can be referred to as the potential usefulness of the forecasts, and is of particular importance for decision-making purposes (Arnal et al., 2018). Another key aspect of probabilistic forecasts to consider is their reliability, which indicates the agreement between forecast probabilities and the observed frequency of events.

The potential usefulness is assessed using the relative operating characteristic (ROC) curve, which is based on ratios of the proportion of events (the probability of detection, POD) and non-events (the false alarm rate, FAR) for which warnings were provided (Mason and Graham, 1999), where in this case warnings are treated as forecasts of river flow exceeding the 80th or falling below the 20th percentile of the reference climatology (see Sect. 2.2.4). These ratios allow for estimation of the probability that an event will be predicted.

For each week of the forecast (out to 16 weeks, corresponding to the forecasts provided via the interface, for example the hydrograph shown in Fig. 5), the POD (eq. 1) and FAR (eq. 2) are calculated for both the 80th and 20th percentile events at each observation station:

$$POD = \frac{hits}{hits+misses} \qquad (1)$$

$$FAR = \frac{\text{false alarms}}{\text{hits+false alarms}} \tag{2}$$

where a *hit* is defined when the forecast correctly exceeded [fell below] the 80th [20th] percentile of the reference climatology during the same week that the observed river flow exceeded [fell below] the 80th [20th] percentile of the observations at that station. It follows that a *miss* is defined when an event was observed but the forecast did not exceed the threshold, and a *false alarm* when the forecast exceeded the threshold but no event was observed. From these, the area under the ROC curve (AROC) is calculated, again for both the 80th and 20th percentile events. The AROC ($0 \leq AROC \leq 1$, where 1 is perfect) indicates the skill of the forecasts compared to the long-term average climatology (which has an AROC of 0.5) and is used here to evaluate the potential usefulness of the forecasts. The maximum lead time at which forecasts are more skilful than climatology (AROC > 0.5) is identified; a forecast with an AROC < 0.5 would be less skilful than climatology, and thus not useful.

The reliability of the forecasts is assessed using attributes diagrams, which show the relationship between the forecast probability and the observed frequency of the events. While the ROC measures the ability of a forecasting system to predict the correct category of an event, the reliability assesses how closely the forecast probabilities correspond to the actual chance of observing the event. As such, these evaluation metrics are useful to consider together. As with the ROC calculations, the reliability is assessed for each week of the forecast (out to 16 weeks), and for both the 80th and 20th percentile events. The range of forecast probabilities is divided into 10 bins (0-10%, 10-20%, etc.), and the forecast probability is plotted against the frequency at which an event was observed for forecasts in each probability bin. Perfect reliability is exhibited when the forecast probability and the observed frequency are equal, for example if a forecast predicts that an event will occur with a probability of 60%, then the event should occur on 60% of the occasions that this forecast was made. Attributes diagrams can also be used to assess the sharpness and resolution of the forecasts. Forecasts that do not discriminate between events and non-events are said to have no resolution (a forecast of climatology would have no resolution), and forecasts which are capable of predicting events with probabilities that differ from the observed frequency, such as forecasts of high or 0 probability, are said to have sharpness.

25 The GloFAS-Seasonal reforecasts (of which there are 216, covering 18 years, as described in Sect. 2.2.4 and Fig. 2) are compared to river flow observations that have been made available to GloFAS, covering 17 years of the study period up to the end of 2015, when the data were collated (see Fig. 2). To ensure a large enough sample size for this analysis, alongside the best possible spatial coverage, the following criteria are applied to the data:

30 - The weekly river flow data record available for each station must contain no more than 53% (9 years) missing data. The high and low flow thresholds (the 80th and 20th percentile, respectively) are calculated using the observations for each station, and for each week, across the 17 years of data, so a sample size of 17 is the maximum possible. A threshold of (up to) 53% missing data allows for a minimum sample size of 8. Selecting a smaller threshold reduced

the number of stations, and the spatial coverage across the globe, significantly. The percentage of missing data is calculated at each station and for each week of the dataset independently, and as such the number of stations used can vary slightly with time.

5   -   The upstream area of the corresponding grid point in the model river network must be at least 1500km^2.

These criteria allow for the use of 1140±14 stations globally. While the dataset contains 6122 stations, just 1664 of these contain data during the 17-year period, and none have the full 17 years of data available. Data from human-influenced rivers have not been removed, as in this study we are interested in identifying the ability of the forecasting system in its current state
10   to predict observed events, rather than the ability of the hydrological model to represent natural flow.

4.1 Potential Usefulness

In order to gain an overview of the potential usefulness of the GloFAS-Seasonal forecasts across the globe, we map the maximum lead time at which the forecasts are more skilful than climatology (i.e. AROC > 0.5), at each observation station, averaged across all forecast months. These results are shown in Fig. 6, and it is clear that forecasts of both high and low flow
15   events are more skilful than climatology across much of the globe, with potentially useful forecasts at many stations out to 4 months ahead. However, there are regions where the forecasts are (on average, across all forecast months) not useful (i.e. AROC < 0.5), such as the western USA and Canada (excluding coastlines), much of Africa, and additionally across parts of Europe for low flow events. As forecasts with an AROC larger than but close to 0.5 could be deemed as only marginally more skilful than climatology, we apply a skill buffer, setting the threshold to AROC > 0.6 for a forecast to be deemed as potentially
20   useful. These results are mapped in Fig. 7, and clearly indicate the reduction in the lead time at which forecasts are potentially useful (for both high and low flow events) at many stations, implying that in some locations, forecasts beyond the first 1-2 months are only marginally more skilful than climatology. There are, however, stations in some rivers with an AROC > 0.6 out to 4 months lead time, and many locations across the globe that still indicate that forecasts are potentially useful 1-2 months ahead for both high and low flow events.

These results can be further broken down by season, indicating whether the forecasts are more potentially useful at certain times of the year. Maps showing the maximum lead time at which AROC > 0.6 for each season (for forecasts started during the season; e.g. DJF indicates the average results for forecasts produced on 1st December, 1st January and 1st February) are provided for high and low flow events in the supplementary material, Fig. S1 and S2, respectively.

The following paragraphs provide an overview of these results for each continent; for further detail please refer to the maps.

South America: For high flow events, forecasts for the Amazon basin in DJF and MAM are potentially useful out to longer lead times (up to 3-4 months) and at more stations than in JJA and SON, with similar results in MAM for low flow events. In contrast, further south, forecasts are most potentially useful JJA and SON, up to 4 months ahead. In the more mountainous regions of western South America, forecasts in JJA and SON are generally less skilful than climatology for high and low flow events. In the northwest, however, for some stations, forecasts started in DJF and MAM are potentially useful up to 3 months ahead.

North America: In eastern North America, JJA and SON forecasts are most potentially useful, with more stations indicating an AROC > 0.6 out to 2-3 months ahead. However, during all seasons there are several stations in the east showing skill out to varying lead times. Much of the western half of the continent (excluding coastal areas) sees forecasts that are less skilful than climatology during all seasons, although some stations do indicate skill up to 4 months ahead for high flow, for forecasts started in MAM and JJA, and for low flow in MAM. At many coastal stations in the west, forecasts of high flow events started in DJF, MAM and JJA do indicate skill out to 3-4 months, and out to ~6 weeks in SON.

Europe: Forecasts for European rivers generally perform best for high flow events in SON and DJF, with the exception of some larger rivers in eastern Europe, for which the forecasts are more potentially useful in JJA and SON. In MAM and JJA, the number of stations indicating no skill is generally higher. In contrast, forecasts for low flow events are less skilful than climatology across much of Europe. Particularly in northeast Europe and Scandinavia, forecasts produced in the summer months of JJA have an AROC < 0.6 at all stations, with only a few stations indicating any skill in other seasons, whereas in central and southeast Europe forecasts of low flow events are most skilful in JJA and SON, out to 3-4 months ahead in the larger rivers. These results are similar to those of Arnal et al. (2018) for the potential usefulness of the EFAS seasonal outlook.

Asia: Although the number of available stations is very limited, the few stations available in southeast Asia indicate that the forecasts are potentially useful out to 3-4 months ahead, particularly for forecasts started in DJF and MAM, preceding the start of the wet season. For low flow events, this skill extends into JJA, whereas forecasts made in SON, towards the end of the wet season, tend to be less skilful than climatology.

Australia & New Zealand: Forecasts are most skilful out to longer lead times in the Murray-Darling river basin in the southeast, in particular for forecasts started in JJA and SON during the southern hemisphere winter and spring. In northern Australia, forecasts started in DJF and MAM for high flow events, and MAM and JJA for low flow events, are potentially useful out to 3-4 months ahead. This corresponds with the assessment of the skill of the Bayesian joint probability modelling approach for sub-seasonal to seasonal streamflow forecasting in Australia by Zhao et al. (2016), who found that forecasts in northern Australian catchments tend to be more skilful for the dry season (May to October) than the wet season (December to March). At the 3 stations in New Zealand, forecasts are only skilful for high flow events during the first month of lead time, in DJF

and MAM; however, for low flow events forecasts made in SON for the southern stations are potentially useful out to 4 months ahead.

Africa: While the spatial distribution of stations is limited, for high flow events forecasts are seen to be potentially useful at some of the stations in eastern Africa, particularly in SON and to a lesser extent in DJF. In southern Africa, there is skill in DJF and MAM, although the maximum lead time varies significantly from station to station. For low flow, there is little variation between the seasons; forecasts are generally less skilful than climatology across the continent, with some stations in DJF in southern and western Africa indicating skill in the first 1-2 months only.

4.2 Reliability

To provide an overall picture of the reliability of the GloFAS-Seasonal forecasts, attributes diagrams are produced for forecasts aggregated across all observation stations globally, for both the 80[th] and 20[th] percentile events. In order to assess geographical differences in forecast reliability, attributes diagrams are also produced for forecasts aggregated across the stations within each of the major river basins used in the GloFAS-Seasonal forecast products (see Sect. 3.1). Many of these river basins do not contain a large enough number of stations to produce useful attributes diagrams, and as such, results in this section are presented for one river basin per continent for this initial evaluation. The river basin chosen for each continent is that which contains the largest number of observation stations.

The globally aggregated results (Fig. 8) indicate that, in general, the forecasts have more reliability than a forecast of climatology, though the reliability is less than perfect. It is important to note that the globally aggregated results shown in Fig. 8 mask any variability between river basins. Overall, the reliability appears to be slightly better for forecasts of high flow events than low flow events, and for lower probabilities, indicated by the steeper positive slope showing that as the forecast probability increases, so does the verified chance of the event. The forecasts for both high and low flow events exhibit sharpness, although more so for high flow events, meaning that they have the ability to forecast probabilities that differ from the climatological average. This is indicated by the histograms inset within the attributes diagrams in Fig. 8; a forecast with sharpness will show a range of forecast probabilities differing from the climatological average (20%), and a forecast with perfect sharpness will show peaks in the forecast frequency at 0% and 100%. Forecasts with no, or low, sharpness will show a peak in the forecast frequency near to the climatological average. A forecast can have sharpness but still be unreliable. Fig. 8 also suggests that in general, GloFAS-Seasonal forecasts have a tendency to over-predict the likelihood of an event occurring.

The following paragraphs summarise the forecast reliability for one river basin per continent; for a map of the location of these river basins, please refer to Fig. S3. The attributes diagrams for these river basins for both the 80[th] and 20[th] percentile events, and for each season, are provided in Fig. S4 – S8. Each attributes diagram displays the results for forecast weeks 4,8,12 and

16, representing the reliability out to 1,2,3 and 4 months ahead. There are no river basins in Asia containing enough stations to produce an attributes diagram.

South America, Tocantins River (Fig. S4). For high flow events, forecasts for the Tocantins River indicate good reliability in
5 all seasons, particularly up to 50% probability. Forecasts in the higher probability bins tend to over-predict, and this over-prediction worsens with lead time. In MAM and JJA, the forecasts tend to slightly under-predict in the lower probability bins. The forecasts have sharpness, but it is clear that the sample size of high probability forecasts is limited. There is a tendency to over-predict the likelihood of low flow events in all seasons, but the forecasts show good reliability for the lower probability bins, particularly in SON and DJF. In JJA, the resolution of the forecasts is low.

North America, Lower Mississippi River (Fig. S5). For high flow events, the sample size of high probability forecasts is small, and as such it is difficult to evaluate the reliability of these forecasts. The forecasts at lower probabilities have good reliability, particularly out to 2 months ahead in MAM and JJA. In SON and DJF, forecasts are more reliable at longer lead times. There is a tendency to under-predict at low probabilities and over-predict at high probabilities. For low flow events, the forecasts
15 have a tendency to over-predict in all seasons, and the resolution of the forecasts is lower than for high flow events. At higher probabilities, forecasts of low flow events are more reliable than climatology, but the resolution is particularly low for probabilities up to 50-60%. The forecasts for both high and low flow events have sharpness.

Europe, River Rhone (Fig. S6). For the River Rhone, the reliability is better than climatology at all lead times for high flow
20 events, although there is a lack of forecasts of higher probabilities, particularly in MAM and JJA, as may be expected in the summer months. In SON, the reliability of forecasts up to 60-70% is good at all lead times, and in DJF the forecasts are more reliable in the first 2 months of lead time for most probability bins. The reliability is less good for low flow events, but is generally better than climatology, particularly in summer (JJA). In winter (DJF), the resolution and reliability of the forecasts is poor. For all seasons and lead times, and for both events, the forecasts have sharpness.

Australia, River Murray (Fig. S7). The attributes diagrams for both high and low flow events indicate that forecasts are often over-confident in this river basin, with probabilities of 0-10% for low flow events, and 0-30% and 90-100% for high flow events, occurring frequently. As such, the sample size of forecasts in several of the bins is low. For high flow events, forecast tend to over-predict at high probabilities, and under-predict at low probabilities. The reliability is very good up to ~30%, after
30 which the sample size is too small. For low flow events, there is a tendency to under-predict, but based on the forecasts available, the reliability is better than climatology at all lead times. The reliability for low flow events is better in SON and DJF (spring and summer), than MAM and JJA (autumn and winter) and for high flow events there is less differentiation between the seasons.

Africa, Orange River (Fig. S8). For the Orange River, forecasts of high flow events exhibit good reliability for lower probabilities in SON, DJF and MAM (spring through autumn), particularly at longer lead times in SON and DJF, with a tendency to over-predict at higher probabilities. Resolution and reliability are poor for high flow events in JJA (winter), with probabilities of 90-100% predicted too frequently. For low flow events, forecasts of 0-10% are very frequent, and the forecasts under-predict in all seasons, although the reliability is better than climatology at all lead times (based on a limited sample of forecasts for most probability bins). Reliability for low flow events is best in DJF (summer).

4.3 Discussion

The results presented provide an initial evaluation of the potential usefulness and reliability of GloFAS-Seasonal forecasts. For decision-making purposes, it is important to measure the ability of a forecasting system to predict the correct category of an event. As such, an event-based evaluation of the forecasts is used to assess whether the forecasts were able to correctly predict observed high and low river flow events over a 17-year period, and whether it is able to do so with good reliability. The initial results are promising, indicating that the forecasts are, on average, potentially useful up to 1-2 months ahead in many rivers worldwide, and up to 3-4 months ahead in some locations. The GloFAS-Seasonal forecasts have sharpness, i.e. they are able to predict forecasts with probabilities that differ from climatology, and overall have better reliability than a forecast of climatology, but with a tendency to over-predict at higher probabilities. It is also clear that there exists a frequency bias in the reliability results, as often there is a small sample of high probability forecasts. Typically, the reliability is seen to be better when there is a higher forecast frequency on which to base the results. As would be expected, the potential usefulness and reliability of the forecasts vary by region, season and forecast lead time.

Considering the evaluation results by season allows further analysis of the times of year in which the forecasts are potentially useful and/or reliable. For example, in southeast Australia, forecasts are seen to be potentially useful up to 4 months ahead in JJA and SON, but for forecasts produced in DJF the skill only extends to 1 month ahead, and forecasts are less skilful than climatology at several of the stations in MAM. In many rivers across the globe, it is the case that forecasts are potentially useful in some seasons, but not in others, and may be more reliable in certain seasons than others. As such, the maps provided in Fig. S1 and S2 are intended to highlight where and when the forecasts are likely to be useful, information that is key in terms of decision-making.

It is clear that there are regions and seasons where the forecasts are less skilful than climatology and do not have good reliability, and thus in these rivers it would be more useful to use a long-term average climatology than seasonal hydro-meteorological forecasts of river flow. This lack of skill could be due to several factors, such as certain hydrological regimes that may not be well-represented in the hydrological model or may be difficult to forecast at these lead times (for example snow dominated-catchments, or regions where convective storms produce most of the rainfall in some seasons), poor skill of

the meteorological forecast input, poor initial conditions from the ERA5-R reanalysis, extensive management of rivers that cannot be represented by the current model, or the lack of model calibration. While this initial evaluation is designed to provide an overview of whether the forecasts are potentially useful and reliable in predicting high and low flow events, more extensive analysis is required to diagnose the sources of predictability in the forecasts and the potential causes of poor skill. Additionally,

5   it is evident that observations of river flow, particularly covering the reforecast period, are both spatially and temporally limited across large areas of the globe. A more extensive analysis should make use of the globally consistent ERA5-R river flow reanalysis as a benchmark in order to fully assess the forecast skill worldwide, including in regions where no observations are available.

10   The verification metrics used also require that a high or low flow event is predicted with the correct timing, in the same week as that in which it occurred. This is asking a lot of a seasonal forecasting system and for many applications, such as water resources and reservoir management, a forecast of the exact week in which an event is expected at a lead time of several months ahead may not be necessary. That such a system shows real skill despite this being a tough test for the model, and is able to successfully predict observed high or low river flow in a specific week, several weeks or months ahead, provides

15   optimism for the future of global scale seasonal hydro-meteorological forecasting. Further evaluation should aim to assess the skill of the forecasts with a more relaxed constraint on the event timing, and also make use of alternative skill measures to cover different aspects of the forecast skill, such as the spread and bias of the forecasts. It will also be important to assess whether the use of weekly-averaged river flow is the most appropriate way to display the forecasts. While this is commonly used for applications such as drought early awareness and water resources management, there may be other aspects of decision-

20   making, such as flood forecasting, for which other measures may be more appropriate, for example daily averages or floodiness (Stephens et al., 2015).

Future development of GloFAS-Seasonal will aim to address these evaluation results and improve the skill and reliability of the current forecasts, and will also aim to overcome some of the grand challenges in operational hydrological forecasting, such

25   as seamless forecasting and the use of data assimilation. Seamless forecasting will be key in the future development of GloFAS; the use of two different meteorological forecast inputs for the medium-range and seasonal versions of the model means that discrepancies can occur between the two timescales thus providing confusing, inconsistent forecast information to users. Additionally, the use of river flow observations could lead to significant improvements in skill, through calibration of the model using historical observations, and assimilation of real-time data to adjust the forecasts. This remains a grand challenge

30   due to the lack of openly available river flow data, particularly in real time.

5 Conclusions

In this paper, the development and implementation of a global scale operational seasonal hydro-meteorological forecasting system, GloFAS-Seasonal, was presented, and an event-based forecast evaluation was carried out using two different but complementary verification metrics, to assess the capability of the forecasts to predict high and low river flow events.

GloFAS-Seasonal provides forecasts of high or low river flow out to 4 months ahead for the global river network through three new forecast product layers via the openly available GloFAS web interface at *www.globalfloods.eu*. Initial evaluation results are promising, indicating that in many rivers, forecasts are both potentially useful, i.e. more skilful than a long-term average climatology, out to several months ahead in some cases, and overall more reliable than a forecast of climatology.Forecast skill and reliability vary significantly by region and by season.

The initial evaluation however also indicates a tendency of the forecasts to over-predict, in general, and in some regions forecasts are currently less skilful than climatology; future development of the system will aim to improve the forecast skill and reliability with a view to providing potentially useful forecasts across the globe. Development of GloFAS-Seasonal will continue based on results of the forecast evaluation, and on feedback from GloFAS partners and users worldwide, in order to provide a forecast product that remains state-of-the-art in hydro-meteorological forecasting, and caters to the needs of its users. Future versions are likely to address some of the grand challenges in hydro-meteorological forecasting in order to improve forecast skill, such as data assimilation, and will also include more features, such as flexible percentile thresholds and indication of the forecast skill via the interface. A further grand challenge that is important in terms of global scale hydro-meteorological forecasting and indeed for the development of GloFAS, is the need for more observed data (Emerton et al., 2016), which is essential not only for providing initial conditions to force the models, but also for evaluation of the forecasts and continuous improvement of forecast accuracy.

While such a forecasting system requires extensive computing resources, the potential for use in decision-making across a range of water-related sectors, and the promising results of the initial evaluation, suggest that it is a worthwhile use of time and resources to develop such global scale systems. Recent papers have highlighted that seasonal forecasts of precipitation are not necessarily a good indicator of potential floodiness, and called for investment in better forecasts of seasonal flood risk (Coughlan De Perez et al., 2017; Stephens et al., 2015). Coughlan de Perez et al. (2017) state that "ultimately, the most informative forecasts of flood hazard at the seasonal scale could be seasonal streamflow forecasts using hydrological models", and that better seasonal forecasts of flood risk could be hugely beneficial for disaster preparedness.

GloFAS-Seasonal represents a first attempt at overcoming the challenges of producing and providing openly-available seasonal hydro-meteorological forecast products, which are key for organisations working at the global scale, and for regions where no

other forecasting system exists. We provide, for the first time, seasonal forecasts of hydrological variables for the global river network, by driving a hydrological model with seasonal meteorological forecasts. GloFAS-Seasonal forecasts could be used in addition to other forecast products such as seasonal rainfall forecasts and short-range forecasts from national hydro-meteorological centres across the globe, to provide useful added information for many water-related applications, from water

5 resources management and agriculture to disaster risk reduction.

Code Availability

The ECMWF IFS source code is available subject to a license agreement, and as such access is available to the ECMWF member-state weather services and other approved partners. The IFS code is also available for educational and academic purposes as part of the OpenIFS project (ECMWF, 2011, 2018a), with full forecast capabilities and including the HTESSEL

10 land surface scheme, but without modules for data assimilation. Similarly, the GloFAS river routing component source code is not openly available; however, the *'forecast product'* code (prior to implementation in ecFlow) that was newly developed for GloFAS-Seasonal, used for a number of tasks such as computing exceedance probabilities and producing the graphics for the interface, is provided in the supplementary material.

Data Availability

15 ECMWF's ERA5 reanalysis and SEAS5 reforecasts are available through the Copernicus Climate Data Store (Copernicus, 2018). The ERA5-R river flow reanalysis and the GloFAS-Seasonal reforecasts (daily data) are currently available from the authors on request, and will be made available through ECMWF's data repository in due course. The majority of the observed river flow data was provided by the Global Runoff Data Centre (GRDC; BfG, 2017). This data is freely available from *www.bafg.de/GRDC*. Additional data was provided by the Russian State Hydrological Institute (SHI, 2018), the European

20 Flood Awareness System (EFAS, 2017), Somalia Water and Land Information Management (SWALIM, 2018), South Africa Department for Water and Sanitation (DWA, 2018), Colombia Institute of Hydrology, Meteorology and Environmental Studies (IDEAM, 2014), Nicaragua Institute of Earth Studies (INETER, 2016), Dominican Republic National Institute of Hydraulic Resources (INDRHI, 2017), Brazil National Centre for Monitoring and Forecasting of Natural Hazards (Cemaden, 2017), Environment Canada Water Office (Environment Canada, 2014), Nepal Department of Hydrology and Meteorology (DHM,

25 2017), Red Cross Red Crescent Climate Centre (RCCC, 2018), Chile General Water Directorate (DGA, 2018), Historical Database on Floods (BDHI, 2018).

Author Contributions

FP proposed the development of GloFAS-Seasonal, RE wrote the GloFAS-Seasonal forecast product code, RE and LA designed the forecast products. EZ built the ecFlow suite and produced ERA5-R and the GloFAS-Seasonal reforecasts, and DM provided technical support for the website and operational implementation. RE evaluated the forecasts and wrote the paper, with the exception of Sect. 2.4, written by DM. All authors were involved in discussions throughout development, and all authors commented on the manuscript.

The authors declare that they have no conflict of interest.

Acknowledgements

This work has been funded by the Natural Environment Research Council (NERC) as part of the SCENARIO Doctoral Training Partnership under grant NE/L002566/1. EZ, DM, CP and PS were supported by the Copernicus Emergency Management Service – Early Warning Systems (CEMS-EWS (EFAS)). LA, HLC and FP acknowledge financial support from the Horizon 2020 IMPREX project (grant agreement no. 641811). ES is thankful for support from NERC and the Department for International Development (grant number NE/P000525/1) under the Science for Humanitarian Emergencies and Resilience (SHEAR) research programme (project FATHUM (Forecasts for AnTicipatory HUManitarian action)).

References

Alfieri, L., Burek, P., Dutra, E., Krzeminski, B., Muraro, D., Thielen, J. and Pappenberger, F.: GloFAS – global ensemble streamflow forecasting and flood early warning, Hydrol. Earth Syst. Sci, 17, 1161–1175, doi:10.5194/hess-17-1161-2013, 2013.

Arnal, L., Cloke, H. L., Stephens, E., Wetterhall, F., Prudhomme, C., Neumann, J., Krzeminski, B. and Pappenberger, F.: Skilful seasonal forecasts of streamflow over Europe?, Earth Syst. Sci, 225194, 2057–2072, doi:10.5194/hess-22-2057-2018, 2018.

Bahra, A.: Managing work flows with ecFlow, ECMWF Newsl., 129, 30–32 [online] Available from: https://www.ecmwf.int/sites/default/files/elibrary/2011/14594-newsletter-no129-autumn-2011.pdf (Accessed 18 April 2018), 2011.

Balsamo, G., Pappenberger, F., Dutra, E., Viterbo, P. and van den Hurk, B.: A revised land hydrology in the ECMWF model: a step towards daily water flux prediction in a fully-closed water cycle, Hydrol. Process., 25(7), 1046–1054, doi:10.1002/hyp.7808, 2011.

Balsamo, G., Viterbo, P., Beljaars, A., Hurk, B., Hirschi, M., Betts, A. K. and Scipal, K.: A Revised Hydrology for the ECMWF Model: Verification from Field Site to Terrestrial Water Storage and Impact in the Integrated Forecast System, , doi:10.1175/2008JHM1068.1, n.d.

BDHI: Base de Donnees Historiques sur les Inondations, [online] Available from: http://bdhi.fr/appli/web/welcome (Accessed 23 April 2018), 2018.

Bell, V. A., Davies, H. N., Kay, A. L., Brookshaw, A. and Scaife, A. A.: A national-scale seasonal hydrological forecast system: development and evaluation over Britain, Hydrol. Earth Syst. Sci., 21(9), 4681–4691, doi:10.5194/hess-21-4681-2017, 2017.

Bennett, J. C., Wang, Q. J., Li, M., Robertson, D. E. and Schepen, A.: Reliable long-range ensemble streamflow forecasts: Combining calibrated climate forecasts with a conceptual runoff model and a staged error model, Water Resour. Res., 52(10), 8238–8259, doi:10.1002/2016WR019193, 2016.

Bennett, J. C., Wang, Q. J., Robertson, D. E., Schepen, A., Li, M. and Michael, K.: Assessment of an ensemble seasonal streamflow forecasting system for Australia, Earth Syst. Sci, 215194, 6007–6030, doi:10.5194/hess-21-6007-2017, 2017.

BfG: The GRDC, [online] Available from: http://www.bafg.de/GRDC/EN/Home/homepage_node.html (Accessed 23 April 2018), 2017.

BoM: Seasonal Streamflow Forecasts: Water Information: Bureau of Meteorology, [online] Available from: http://www.bom.gov.au/water/ssf/about.shtml (Accessed 24 April 2018), 2018.

Candogan Yossef, N., van Beek, R., Weerts, A., Winsemius, H., P Bierkens, M. F. and to, C.: Skill of a global forecasting system in seasonal ensemble streamflow prediction, Earth Syst. Sci, 215194, 4103–4114, doi:10.5194/hess-21-4103-2017, 2017.

Cemaden: Cemaden – Centro Nacional de Monitoramento e Alertas de Desastres Naturais, [online] Available from: http://www.cemaden.gov.br/ (Accessed 23 April 2018), 2017.

Chiew, F. H. S. and McMahon, T. A.: Global ENSO-streamflow teleconnection, streamflow forecasting and interannual variability, Hydrol. Sci. J., 47(3), 505–522, doi:10.1080/02626660209492950, 2002.

Chow, V. Te, Maidment, D. R. and Mays, L. W.: Applied hydrology, Tata McGraw-Hill Education. [online] Available from: https://books.google.co.uk/books/about/Applied_Hydrology.html?id=RRwidSsBJrEC&redir_esc=y (Accessed 17 November 2017), 2010.

Cloke, H., Pappenberger, F., Thielen, J. and Thiemig, V.: Operational European Flood Forecasting, in Environmental Modelling, pp. 415–434, John Wiley & Sons, Ltd, Chichester, UK., 2013.

Copernicus: Copernicus Climate Data Store, [online] Available from: https://climate.copernicus.eu/climate-data-store (Accessed 23 April 2018), 2018.

5 Copernicus: SWICCA | Service for Water Indicators in Climate Change Adaptation, [online] Available from: http://swicca.climate.copernicus.eu/ (Accessed 12 January 2018), n.d.

Coughlan De Perez, E., Stephens, E., Bischiniotis, K., Van Aalst, M., Van Den Hurk, B., Mason, S., Nissan, H. and Pappenberger, F.: Should seasonal rainfall forecasts be used for flood preparedness?, Earth Syst. Sci, 215194, 4517–4524, doi:10.5194/hess-21-4517-2017, 2017.

10 Crochemore, L., Ramos, M.-H. and Pappenberger, F.: Bias correcting precipitation forecasts to improve the skill of seasonal streamflow forecasts, Hydrol. Earth Syst. Sci, 20, 3601–3618, doi:10.5194/hess-20-3601-2016, 2016.

Demargne, J., Wu, L., Regonda, S. K., Brown, J. D., Lee, H., He, M., Seo, D.-J., Hartman, R., Herr, H. D., Fresch, M., Schaake, J., Zhu, Y., Demargne, J., Wu, L., Regonda, S. K., Brown, J. D., Lee, H., He, M., Seo, D.-J., Hartman, R., Herr, H. D., Fresch, M., Schaake, J. and Zhu, Y.: The Science of NOAA's Operational Hydrologic Ensemble Forecast Service, Bull. Am. Meteorol.
15 Soc., 95(1), 79–98, doi:10.1175/BAMS-D-12-00081.1, 2014.

DGA: Ministerio de Obras Públicas - Dirección de General de Aguas, [online] Available from: http://www.dga.cl/Paginas/default.aspx (Accessed 23 April 2018), 2018.

DHM: Department of Hydrology and Meteorology, [online] Available from: http://www.dhm.gov.np/ (Accessed 23 April 2018), 2017.

20 DWA: Department: WAter and Sanitation, [online] Available from: http://www.dwa.gov.za/default.aspx (Accessed 23 April 2018), 2018.

ECMWF: OpenIFS, , https://www.ecmwf.int/en/research/projects/openifs, 2011.

ECMWF: ecFlow Documentation, [online] Available from: https://software.ecmwf.int/wiki/display/ECFLOW/Documentation (Accessed 18 April 2018), 2012.

25 ECMWF: SEAS5 user guide. [online] Available from: https://www.ecmwf.int/sites/default/files/medialibrary/2017-10/System5_guide.pdf (Accessed 18 April 2018a), 2017.

ECMWF: What are the changes from ERA-Interim to ERA5? - Copernicus Knowledge Base - ECMWF Confluence Wiki, [online] Available from: https://software.ecmwf.int/wiki/pages/viewpage.action?pageId=74764925 (Accessed 24 April 2018b), 2017.

ECMWF: About OpenIFS, [online] Available from: https://software.ecmwf.int/wiki/display/OIFS/About+OpenIFS (Accessed 26 April 2018a), 2018.

ECMWF: ECMWF IFS Documentation CY43R1, [online] Available from: https://www.ecmwf.int/search/elibrary/IFS?secondary_title=%22IFS Documentation CY43R1%22 (Accessed 18 April 2018b), 2018.

EFAS: European Flood Awareness System (EFAS), [online] Available from: https://www.efas.eu/ (Accessed 23 April 2018), 2017.

Emerton, R., Cloke, H. L., Stephens, E. M., Zsoter, E., Woolnough, S. J. and Pappenberger, F.: Complex picture for likelihood of ENSO-driven flood hazard, Nat. Commun., 8, 14796, doi:10.1038/ncomms14796, 2017.

Emerton, R. E., Stephens, E. M., Pappenberger, F., Pagano, T. C., Weerts, A. H., Wood, A. W., Salamon, P., Brown, J. D., Hjerdt, N., Donnelly, C., Baugh, C. A. and Cloke, H. L.: Continental and global scale flood forecasting systems, Wiley Interdiscip. Rev. Water, 3(3), 391–418, doi:10.1002/wat2.1137, 2016.

Environment Canada: Water Level and Flow - Environment Canada, [online] Available from: https://wateroffice.ec.gc.ca/ (Accessed 23 April 2018), 2014.

Environmental Systems Research Institute: ArcMap | ArcGIS Desktop, [online] Available from: http://desktop.arcgis.com/en/arcmap/ (Accessed 26 April 2018), 2018.

Fekete, B. M., Vörösmarty, C. J. and Lammers, R. B.: Scaling gridded river networks for macroscale hydrology: Development, analysis, and control of error, Water Resour. Res., 37(7), 1955–1967, doi:10.1029/2001WR900024, 2001.

GloFAS: GloFAS Web Map Service Time (WMS-T) User Manual, [online] Available from: http://www.globalfloods.eu/static/downloads/GloFAS-WMS-T_usermanual.pdf (Accessed 26 April 2018a), 2018.

GloFAS: GloFAS WMS-T, [online] Available from: http://globalfloods-ows.ecmwf.int/glofas-ows/?service=WMS&request=GetCapabilities, 2018b.

IDEAM: IDEAM, [online] Available from: http://www.ideam.gov.co/ (Accessed 23 April 2018), 2014.

INDRHI: INDRHI - National Institute of Hydraulic Resources, [online] Available from: http://indrhi.gob.do/ (Accessed 23 April 2018), 2017.

INETER: Ineter | Instituto Nicaragüense de Estudios Territoriales, [online] Available from: http://www.ineter.gob.ni/ (Accessed 23 April 2018), 2016.

5 Van Der Knijff, J. M., Younis, J. and De Roo, A. P. J.: LISFLOOD: a GIS-based distributed model for river basin scale water balance and flood simulation, Int. J. Geogr. Inf. Sci., 24(2), 189–212, doi:10.1080/13658810802549154, 2010.

Lehner, B. and Grill, G.: Global river hydrography and network routing: baseline data and new approaches to study the world's large river systems, Hydrol. Process., 27(15), 2171–2186, doi:10.1002/hyp.9740, 2013.

Lehner, B., Verdin, K. and Jarvis, A.: New Global Hydrography Derived From Spaceborne Elevation Data, Eos, Trans. Am.
10 Geophys. Union, 89(10), 93, doi:10.1029/2008EO100001, 2008.

Lorenz, E. N.: Deterministic Nonperiodic Flow, J. Atmos. Sci., 20(2), 130–141, doi:10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2, 1963.

Lorenz, E. N.: The essence of chaos, University of Washington Press., 1993.

Mason, S. J. and Graham, N. E.: Conditional Probabilities, Relative Operating Characteristics, and Relative Operating Levels,
15 Weather Forecast., 14(5), 713–725, doi:10.1175/1520-0434(1999)014<0713:CPROCA>2.0.CO;2, 1999.

McPhaden, M. J., Zebiak, S. E. and Glantz, M. H.: ENSO as an integrating concept in earth science., Science, 314(5806), 1740–5, doi:10.1126/science.1132588, 2006.

Meißner, D., Klein, B. and Ionita, M.: Development of a monthly to seasonal forecast framework tailored to inland waterway transport in central Europe, Hydrol. Earth Syst. Sci., 21(12), 6401–6423, doi:10.5194/hess-21-6401-2017, 2017.

20 Mo, K. C., Lettenmaier, D. P., Mo, K. C. and Lettenmaier, D. P.: Hydrologic Prediction over the Conterminous United States Using the National Multi-Model Ensemble, J. Hydrometeorol., 15(4), 1457–1472, doi:10.1175/JHM-D-13-0197.1, 2014.

OGC: OGC Web Map Service v1.3.0, , doi:10.3173/air.21.76, 2015.

Open Source Geospatial Foundation: MapServer 7.0.1 documentation, [online] Available from: http://mapserver.org/uk/index.html (Accessed 26 April 2018), 2016.

OpenLayers: OpenLayers, [online] Available from: http://openlayers.org/ (Accessed 18 April 2018), 2018.

Pappenberger, F., Ramos, M. H., Cloke, H. L., Wetterhall, F., Alfieri, L., Bogner, K., Mueller, A. and Salamon, P.: How do I know if my forecasts are better? Using benchmarks in hydrological ensemble prediction, J. Hydrol., 522, 697–713, doi:10.1016/J.JHYDROL.2015.01.024, 2015.

5   Prudhomme, C., Hannaford, J., Harrigan, S., Boorman, D., Knight, J., Bell, V., Jackson, C., Svensson, C., Parry, S., Bachiller-Jareno, N., Davies, H., Davis, R., Mackay, J., McKenzie, A., Rudd, A., Smith, K., Bloomfield, J., Ward, R. and Jenkins, A.: Hydrological Outlook UK: an operational streamflow and groundwater level forecasting system at monthly to seasonal time scales, Hydrol. Sci. J., 62(16), 2753–2768, doi:10.1080/02626667.2017.1395032, 2017.

QGIS Development Team: Quantum GIS Geographical Information System, [online] Available from: https://www.qgis.org/
10   (Accessed 4 December 2017), 2017.

RCCC: Home - Red Cross Red Crescent Climate Centre, [online] Available from: http://www.climatecentre.org/ (Accessed 23 April 2018), 2018.

SHI: "State Hydrological Institute" (SHI) | Russian Federal State Budgetary Organization, [online] Available from: http://www.hydrology.ru/en (Accessed 23 April 2018), 2018.

15   Stephens, E., Day, J. J., Pappenberger, F. and Cloke, H.: Precipitation and floodiness, Geophys. Res. Lett., 42(23), 10,316-10,323, doi:10.1002/2015GL066779, 2015.

Stockdale, T., Johnson, S., Ferranti, L., Balmaseda, M. and Briceag, S.: ECMWF's new long-range forecasting system SEAS5, ECMWF Newsl., 154, 15–20 [online] Available from: www.ecmwf.int/en/about/news-centre/media-resources (Accessed 18 April 2018), 2018.

20   SWALIM: FAO SWALIM: Somalia Water and Land Information Management, [online] Available from: http://www.faoswalim.org/ (Accessed 23 April 2018), 2018.

Ward, P. J., Eisner, S., Flörke, M., Dettinger, M. D. and Kummu, M.: Annual flood sensitivities to El Niño–Southern Oscillation at the global scale, Hydrol. Earth Syst. Sci., 18(1), 47–66, doi:10.5194/hess-18-47-2014, 2014a.

Ward, P. J., Jongman, B., Kummu, M., Dettinger, M. D., Sperna Weiland, F. C. and Winsemius, H. C.: Strong influence of El
25   Niño Southern Oscillation on flood risk around the world., Proc. Natl. Acad. Sci. U. S. A., 111(44), 15659–64, doi:10.1073/pnas.1409822111, 2014b.

Ward, P. J., Kummu, M. and Lall, U.: Flood frequencies and durations and their response to El Niño Southern Oscillation:

Global analysis, J. Hydrol., 539, 358–378, doi:10.1016/J.JHYDROL.2016.05.045, 2016.

Weisheimer, A. and Palmer, T. N.: On the reliability of seasonal climate forecasts., J. R. Soc. Interface, 11(96), 20131162, doi:10.1098/rsif.2013.1162, 2014.

Wood, A. W., Maurer, E. P., Kumar, A. and Lettenmaier, D. P.: Long-range experimental hydrologic forecasting for the eastern
5   United States, J. Geophys. Res., 107(D20), 4429, doi:10.1029/2001JD000659, 2002.

Wood, A. W., Kumar, A. and Lettenmaier, D. P.: A retrospective assessment of National Centers for Environmental Prediction climate model–based ensemble hydrologic forecasting in the western United States, J. Geophys. Res., 110(D4), D04105, doi:10.1029/2004JD004508, 2005.

Yuan, X., Wood, E. F., Chaney, N. W., Sheffield, J., Kam, J., Liang, M., Guan, K., Yuan, X., Wood, E. F., Chaney, N. W.,
10   Sheffield, J., Kam, J., Liang, M. and Guan, K.: Probabilistic Seasonal Forecasting of African Drought by Dynamical Models, J. Hydrometeorol., 14(6), 1706–1720, doi:10.1175/JHM-D-13-054.1, 2013.

Yuan, X., Wood, E. F. and Ma, Z.: A review on climate-model-based seasonal hydrologic forecasting: physical understanding and system development, Wiley Interdiscip. Rev. Water, 2(5), 523–536, doi:10.1002/wat2.1088, 2015a.

Yuan, X., Roundy, J. K., Wood, E. F., Sheffield, J., Yuan, X., Roundy, J. K., Wood, E. F. and Sheffield, J.: Seasonal
15   Forecasting of Global Hydrologic Extremes: System Development and Evaluation over GEWEX Basins, Bull. Am. Meteorol. Soc., 96(11), 1895–1912, doi:10.1175/BAMS-D-14-00003.1, 2015b.

Zhao, T., Schepen, A. and Wang, Q. J.: Ensemble forecasting of sub-seasonal to seasonal streamflow by a Bayesian joint probability modelling approach, J. Hydrol., 541, 839–849, doi:10.1016/J.JHYDROL.2016.07.040, 2016.
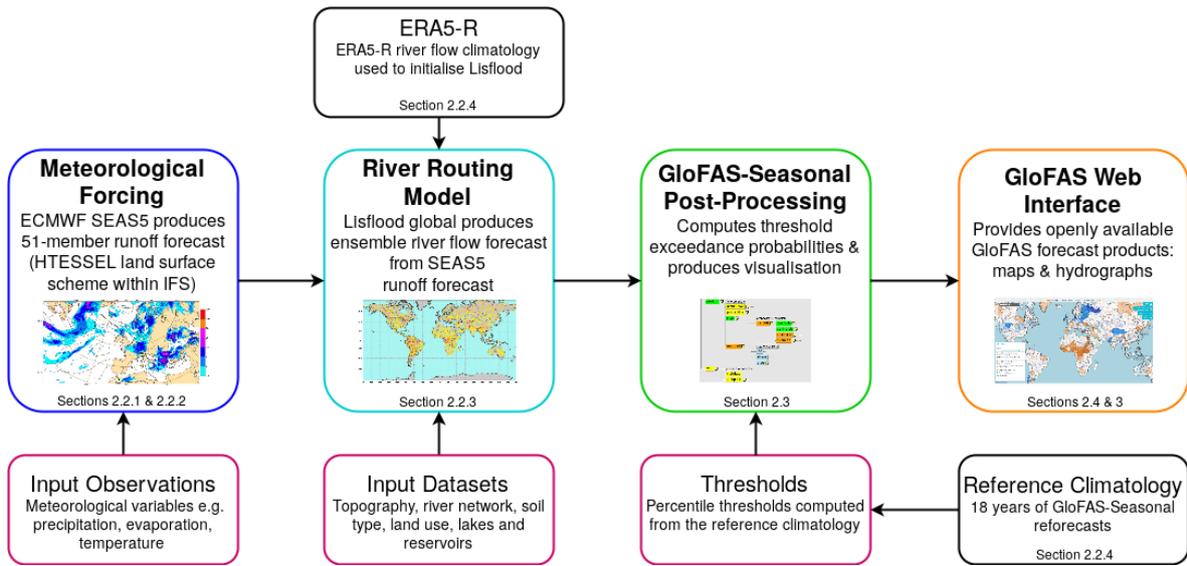
Figure 1: Flowchart depicting the key GloFAS-Seasonal forecasting system components.
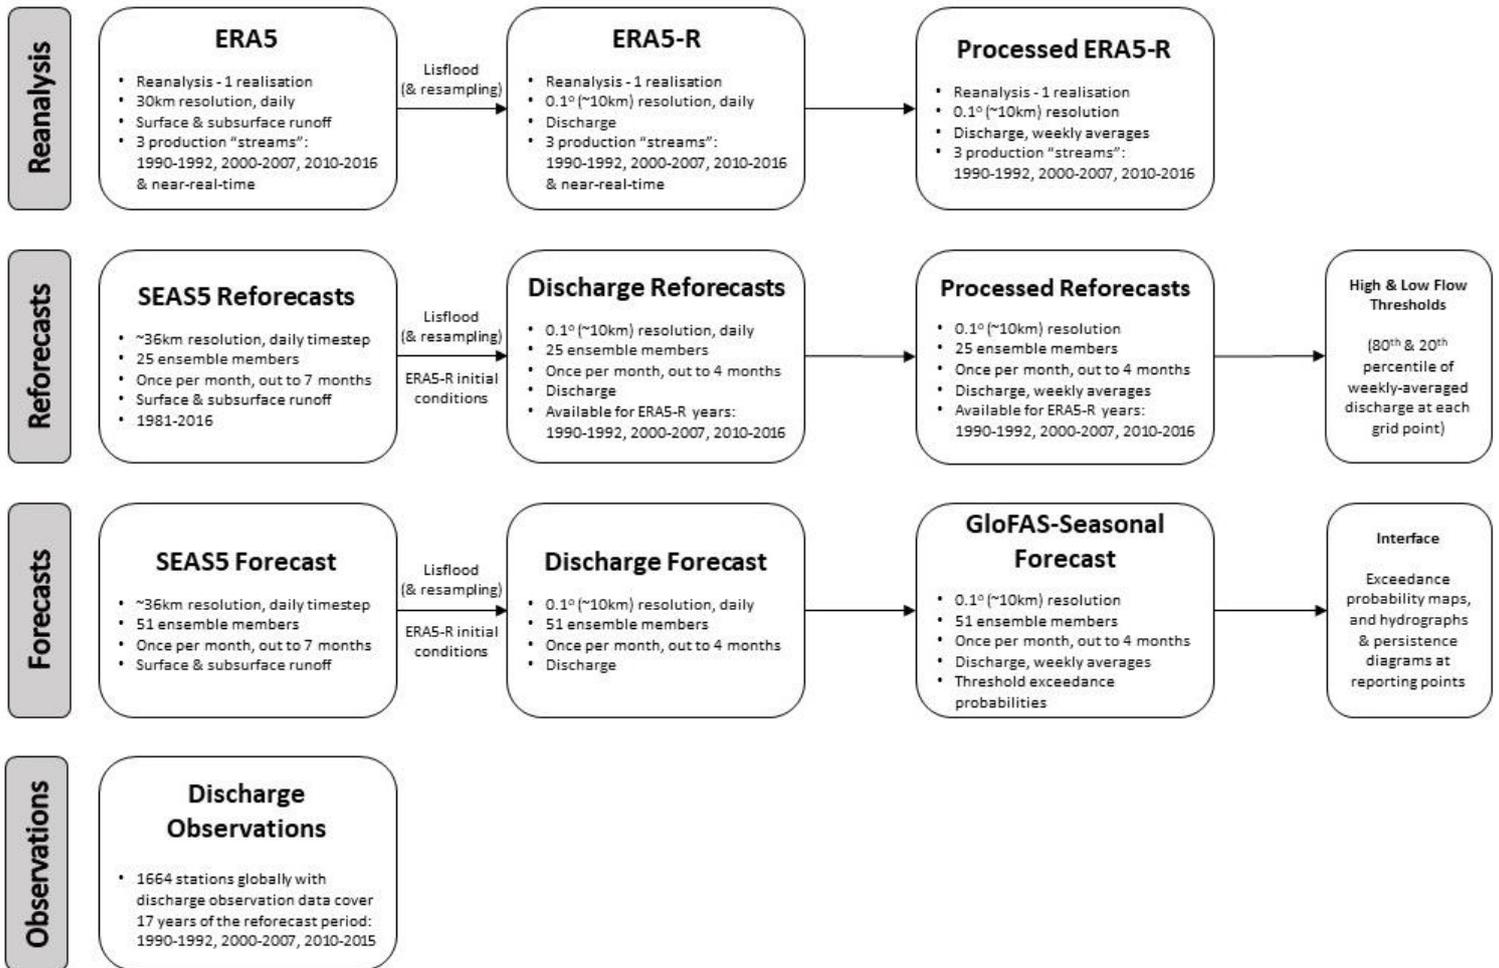
Figure 2: All datasets used and produced for GloFAS-Seasonal, including reanalysis, reforecasts, real-time forecasts and observations.
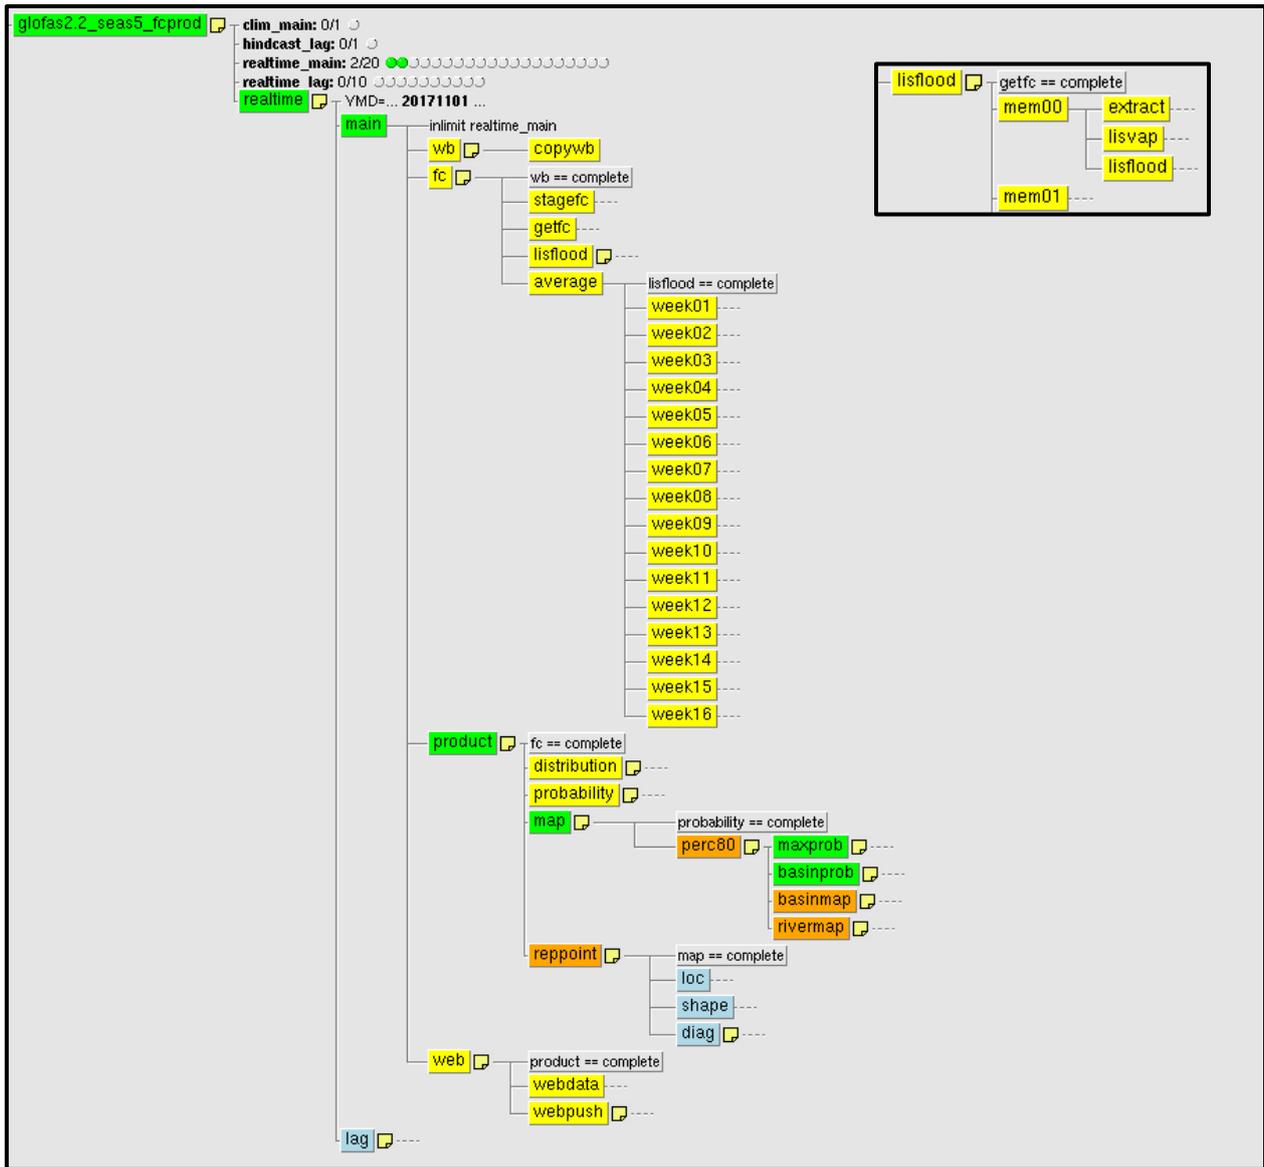
Figure 3: The GloFAS-Seasonal ecFlow suite. The inset image shows the subtasks within the lisflood task, for 1 of the 51 ensemble members. Colours indicate the status of each task, where yellow = complete, green = active, orange = suspended, pale blue = waiting, turquoise (not shown) = queued and red (not shown) = aborted / failed. Grey boxes indicate dependencies, for example "lisflood == complete" indicates that the lisflood task and all lisflood subtasks must have successfully completed in order for the average task to run.
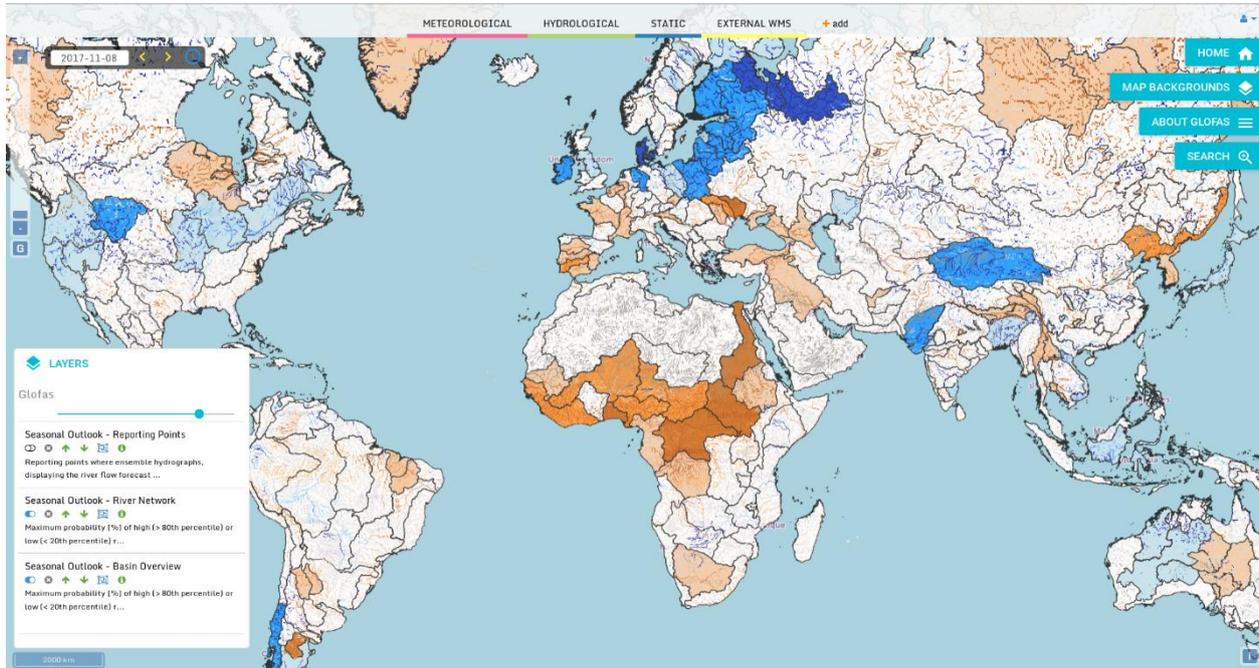
Figure 4: Example screenshot of the seasonal outlook layers in the GloFAS web interface. Shown here are both the "basin overview" layer and "river network" layer, both indicating the maximum probability of unusually high (blue) or low (orange) river flow during the 4-month forecast horizon. The darker the colour, the higher the probability: darkest shading = >90% probability, medium shading = 75-90% probability, light shading = 50-75% probability. A white basin or light grey river pixel indicates that the forecast does not exceed 50% probability of high or low flow during the forecast horizon. Legends providing this information are available for each layer by clicking on the green "i" next to the layer toggle (shown at the bottom left in this example).
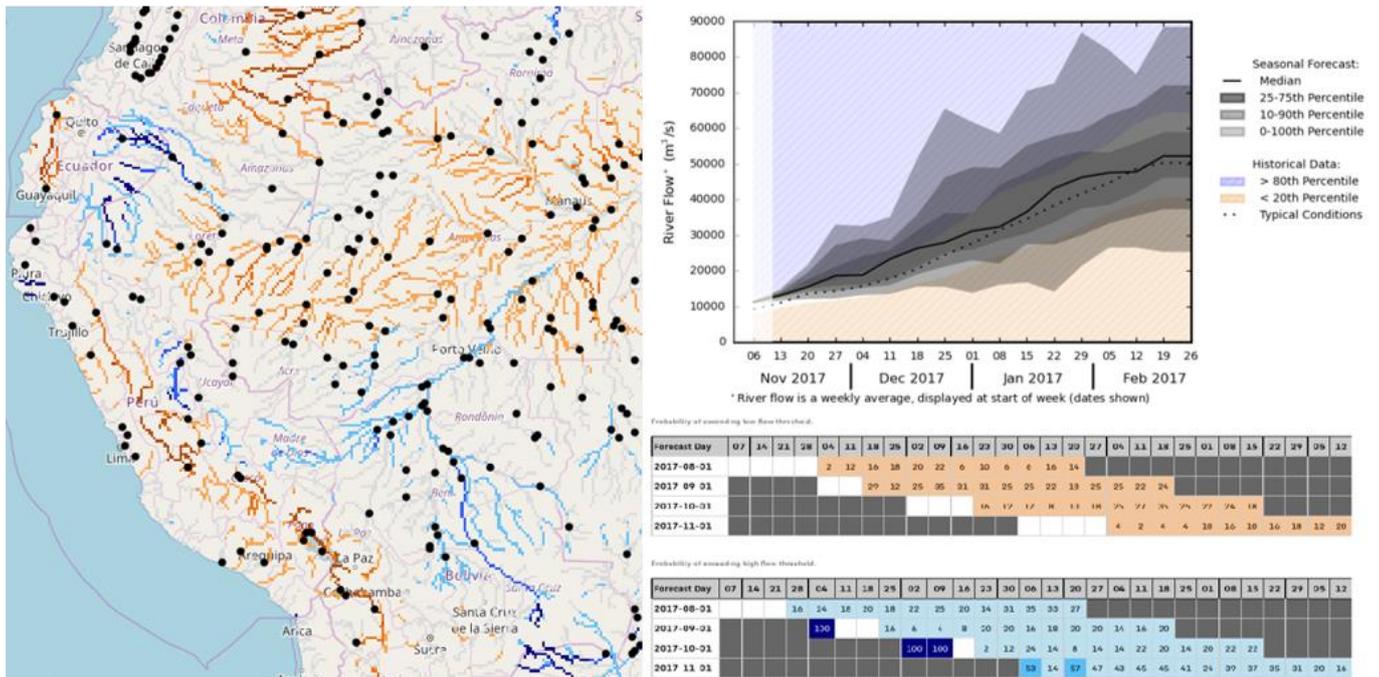
Figure 5: Example of the "reporting points" GloFAS seasonal outlook layer in the web interface. Black circles indicate the reporting points, which provide the ensemble hydrograph (top right) and persistence diagrams for both low flow (centre right) and high flow (bottom right). Also shown is an example section of the "river network" seasonal outlook layer, indicating the maximum probability of high (blue) or low (orange) river flow during the 4-month forecast horizon. The darker the colour, the higher the probability.
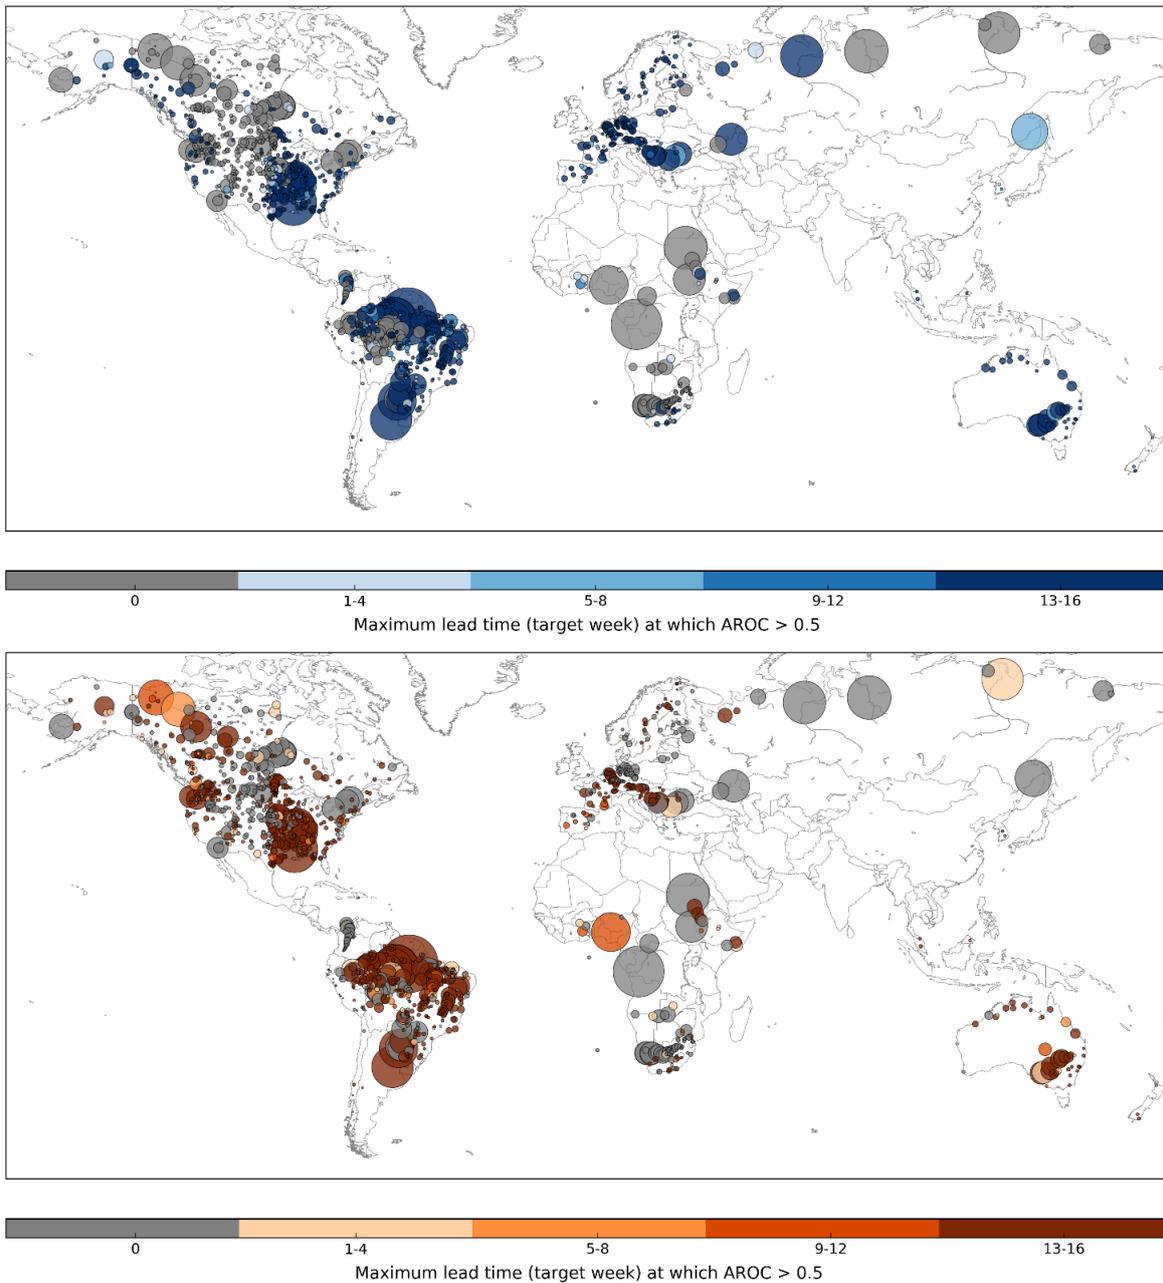
Figure 6: Maximum forecast lead time (target week, averaged across all months) at which the area under the ROC curve (AROC) is greater than 0.5 for high flow events (flow exceeding the 80th percentile of climatology, top panel) and low flow events (flow below the 20th percentile of climatology, bottom panel), at each observation station. This is used to indicate the maximum lead time at which forecasts are more skilful than the long-term average. Dot size corresponds to the upstream area of the location – thus larger dots represent larger rivers and vice versa. Grey dots indicate that (on average, across all months) forecasts are less skilful than climatology at all lead times.
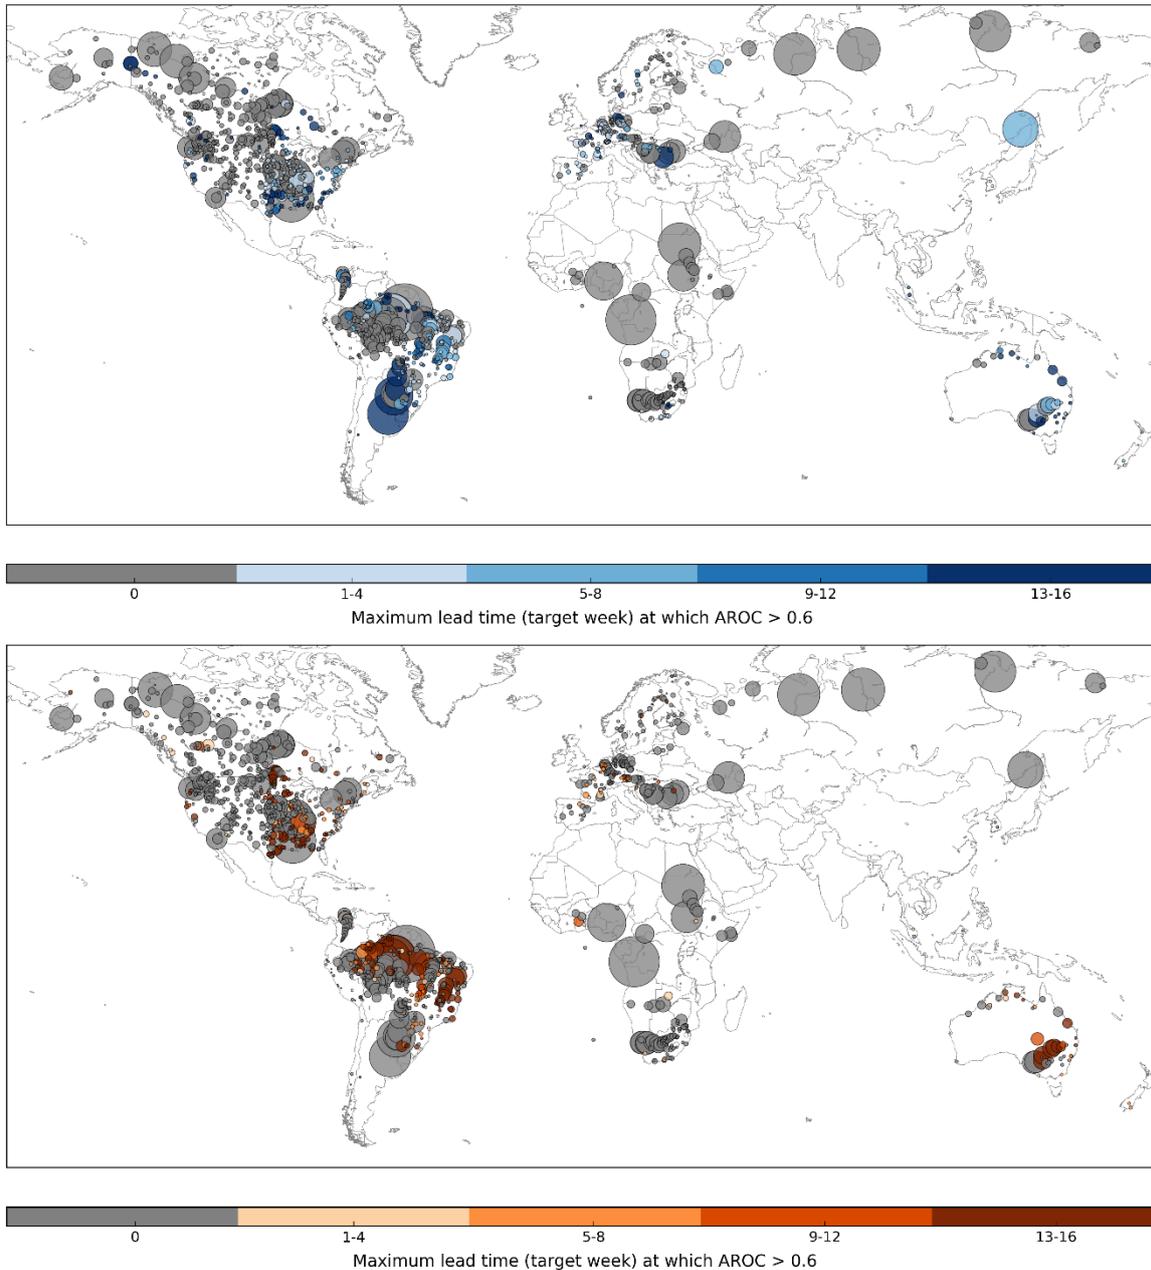
Figure 7: Maximum forecast lead time (target week, averaged across all months) at which the area under the ROC curve (AROC) is greater than 0.6 for high flow events (flow exceeding the 80^{th} percentile of climatology, top panel) and low flow events (flow below the 20^{th} percentile of climatology, bottom panel), at each observation station. This is used to indicate the maximum lead time at which forecasts are deemed skilful. Dot size corresponds to the upstream area of the location – thus larger dots represent larger rivers and vice versa. Grey dots indicate that (on average, across all months) forecasts are less skilful than climatology at all lead times. Maps for each season are provided in the supplementary material.
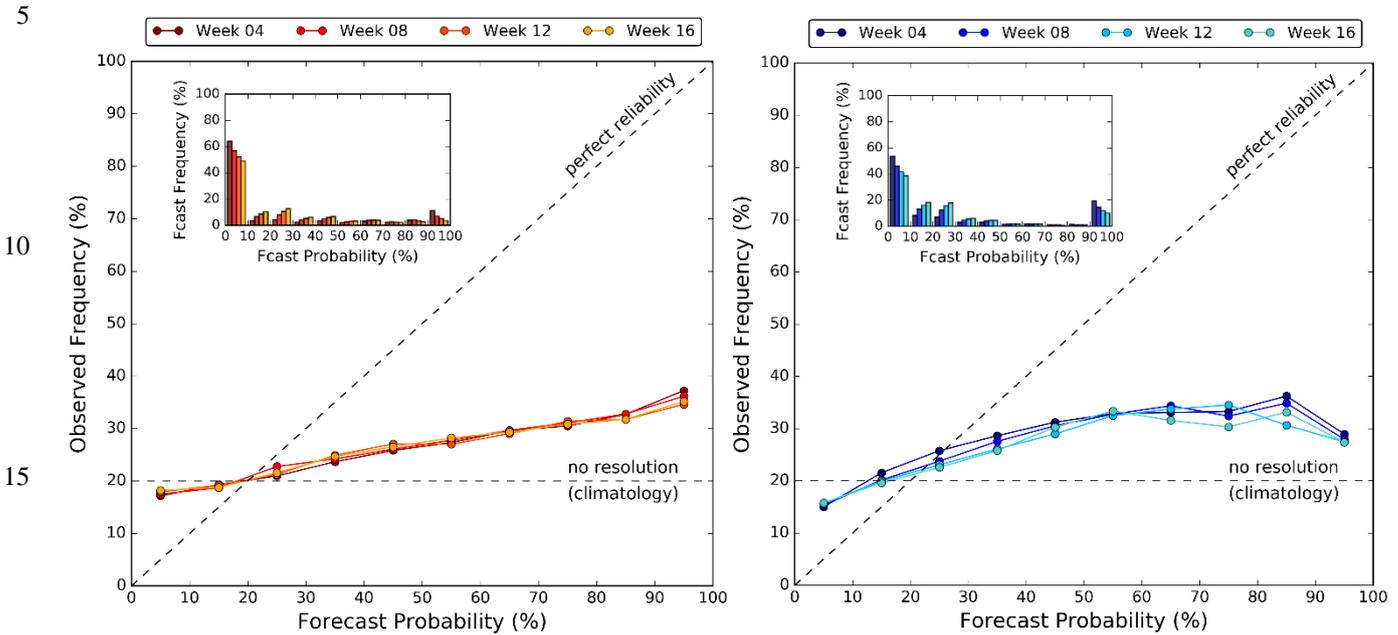
Figure 8: Attributes diagram for forecasts of low flow events (flow below the 20th percentile of climatology, left) and high flow events (flow exceeding the 80th percentile of climatology, right), aggregated across all observation stations globally. Results are shown for lead time weeks 4, 8, 12 and 16, and indicate the reliability of the forecasts. The histograms (inset) show the frequency at which forecasts occur in each probability bin, and are used to indicate forecast sharpness. Attributes diagrams for selected river basins are provided in the supplementary material.