Project name	Land-Surface Initialization in High-resolution seasonal Prediction (LSIHP)
Research field	PE10_3 Climatology and climate change

Project leader

Title	Dr
Last name	Guemas
First name	Virginie
Organisation name*	Barcelona Supercomputing Center (BSC)
Department*	Earth Sciences Department
Group*	Climate Prediction Group
Country	Spain

1. **Describe your research project. Include discussion of the scientific questions that you are planning to address and the overall scientific goals of the project. It is important that you describe the novelty, impact and timeliness of the proposal. (This section must be no longer than 2 pages).**

The Earth's climate undergoes natural variability at seasonal-to-decadal timescales. Informing public sectors that are vulnerable to its variations is a key societal and economical challenge. Prediction of climate variability tackles this challenge to support stakeholders in sectors such as agricultural production (Challinor et al., 2005), energy production (García-Morales et al., 2007) and human health (Thomson et al., 2006). Current climate forecast systems can provide accurate predictions of the tropical Sea Surface Temperatures (SST) anomalies associated with El Niño Southern Oscillation (ENSO) with forecast times of several months (Saha et al., 2006; Stockdale et al., 2011), although the spread among different forecast systems is substantial, sometimes even differing in the sign of the tropical Pacific SST anomaly (Alessandri et al., 2010). Over the extra-tropics, the skill of current seasonal forecast systems is very limited (Rodwell and Doblas-Reyes, 2006). Recent results suggest however that initializing the land surface from observed soil moisture conditions could increase substantially the forecast quality over Europe for surface temperature and precipitation, in particular during heat waves (Prodhomme et al, 2015) and that an increase of resolution in climate forecast systems is one of the necessary factors to reach a useful level of skill over Europe (Scaife et al., 2014).

In an early study, Schär et al. (1999) had shown the existence of a soil-precipitation feedback over Europe. Later on, soil has been shown to influence precipitations, mean temperature and extreme temperature over Europe (e.g. Douville 2010; Seneviratne et al. 2010, 2013; Bellprat et al. 2013). For instance, Seneviratne et al. (2010) described the soil moisture-temperature coupling feedback loop in which, when an anticyclonic anomaly is present over Europe the soil moisture content will either amplify or moderate the surface temperature response. If the soil is moist (energy limited regime), the available surface energy will preferentially dissipate into latent heat fluxes and dampen surface heating. Conversely, when the soil is dry (soil moisture limited regime) more energy is available for sensible heating, inducing an increase of near-surface air temperature (Seneviratne et al. 2010; Hirschi et al. 2011). As soil moisture partly controls the occurrence of warm events over Europe, a correct initialization of soil moisture content might be essential to correctly forecast summer extreme temperatures. This problem was studied by the global land-atmosphere coupling experiment (GLACE) intercomparison project (http://gmao.gsfc.nasa.gov/research/GLACE). Its second phase (GLACE-2) focused on forecast quality, and assessed the impact of accurate soil-moisture initialization on forecast skill using a multimodel approach (Koster et al. 2011). The multimodel mean in GLACE-2 indicates a significant soil-moisture contribution to surface temperature forecast skill in summer with forecast times of up to two months over North and South America (Koster et al. 2010, 2011). While Europe was not then found as a main region of improvement when soil moisture is initialized during the GLACE project, numerous other studies using coupled ocean-atmosphere models instead have found an impact of soil moisture initialization in Europe (e.g. Douville 2010; Prodhomme et al, 2015).

The EC-Earth model (Hazeleger et al. 2012, https://dev.ec-earth.org) contributed to the projection and prediction protocols of the Coupled Model Intercomparison Project Fifth Phase (CMIP5, Taylor et al 2012) and is regularly used to perform seasonal predictions, using resolutions of ~0.7° and ~1° in the atmosphere and ocean, respectively. While such resolution compares favourably with other CMIP5 models, it is poor in terms of the resolution required for an accurate simulation of important modes of climate and weather variability. Jung et al. (2012) and Kinter et al. (2013) found significant improvements in the simulation of many atmospheric features such as tropical precipitation and the frequency/intensity of both tropical and mid-latitude cyclones in IFS-only simulations (the atmospheric component of EC-Earth) using higher resolutions. Recently, several studies reported the use of coupled climate models in the resolution range ~50 km (atmosphere) and ~10-30 km (ocean), identifying large improvements in a number of systematic errors such as an improved structure of the North Atlantic Ocean circulation (Delworth et al., 2012). LISHP proposes the use of groundbreaking resolution to leverage the benefits from the key source of climate predictability on seasonal timescales that are the land surface conditions.

We plan to use EC-Earth3 with a spectral truncation of the atmospheric model (IFS) at T511 (approx. 40 km globally) and 91 vertical levels and a grid resolution of the ocean model (NEMO3.3) of 0.25° globally (approximately 25 km) with 75 vertical levels which thickness increases from 1m below surface up to 500m in the deep ocean. Compared to the resolution of the model used in CMIP5 simulations, the horizontal resolution is increased by a factor of up to 4, with an increased number of vertical levels in both ocean and atmosphere by factors of 1.5 to 2. With this high-resolution version, our objective is to reproduce the second phase of the GLACE-2 project, but with a forecasting period of 17 years instead of 10 years. Two retrospective seasonal forecasting exercises will be conducted over the 1993-2009 period: one exercise using the best possible estimate from observed land surface initial conditions and another one using a simple climatology of these land surface conditions. The rest of the experimental setup will be exactly the same for both exercises. Their comparison will allow for the most robust identification performed up-to-date of the added-value from land-surface initialisation from observations on seasonal forecast quality, i.e. the most robust identification of the role of land sources of predictability, thanks to seasonal forecasts ran at the highest resolution ever used in a seasonal forecasting context, and over a exceptionally long reforecasting period.

The forecast quality will be assessed in terms of amplitude and characteristic timescales of the climate forecast drift (the climate model develops its own biases/systematic error when initialised from observations), as well as in terms skill after bias correction under a deterministic (anomaly correlation and root mean square skill score) and probabilistic framework (Brier score and reliability diagram). Various bias correction techniques will be tested. The forecast quality assessment will be carried out using the most complete framework available at the time of finishing the simulations and will be based on the new verification packages developed by the BSC Earth Sciences Department in the European-funded SPECS project (e.g. s2dverification). The forecast quality assessment will focus, in particular, on climate modes such as the North Atlantic Oscillation, and its potential teleconnections, as well as extreme temperature and precipitation events (through metrics based on daily data such as quantiles). The predictability mechanisms linking the land surface initialisation with the European climate will be identified through water budget and analyses of the large-scale dynamics. A particular emphasis will be put on preconditioning by the soil moisture conditions via regression analyses and classification of start dates and prediction members. Case studies such as the 2003 European heat wave will be investigated in detail.

LISHP will deliver crucial information about seasonal forecast quality and enhance our predictive capability over Europe, in particular for extreme events such as heat waves and droughts. Through close collaboration with the Earth Services Group from the BSC Earth Department, we will ensure that the resulting progress feeds the climate services team and is exploited to generate user-relevant climate information in sectors such as the wind energy or the viticulture. The results of the project will be widely presented in scientific conferences and will feature at the meetings of the Scientific Steering Group (SSG) CLIVAR (Climate and Ocean Variability Predictability and Change) project, of which Virginie Guemas is a member, as well as the meeting of CLIVAR’s WGSIP Working Group on Seasonal-to-Interannual Prediction of which Francisco Doblas-Reyes is a member. These results will contribute to the SPECS, IMPREX and PRIMAVERA projects funded by the European Commission in which the applying team participates, the first one focusing on improving seasonal forecast quality, the second one on a better understanding of hydrological extremes and the third one on the added-value from increasing climate model resolution.

2. Describe how you will manage the resources requested? Use a Gantt chart or equivalent to illustrate this. (1 page).

The hindcast period will be 1993 to 2013 with ten-member ensembles and two start dates per year (first of May and November) that will be run for four forecast months. The ten members are generated by introducing perturbations of the atmospheric initial conditions using singular vectors. This makes a total of 1680 (21 years, 2 start months, 4 forecast months, 10 members) months of simulation for each experiment. In both experiments, the atmosphere will be initialised with the ERA-Interim atmospheric reanalysis (Dee et al., 2011) interpolated to the T511L91 resolution, the ocean with the GLORYS2v1 ocean reanalysis (Ferry et al., 2010) and the sea ice will be initialised following the methodology of Guemas et al (2014) where the sea-ice model is forced by fluxes from reanalyses. In ***LandInit***, the land surface scheme will use data from an offline simulation performed with the ERA-Land system (Balsamo et al., 2014) at the T511 resolution as initial conditions for the retrospective forecasts. In ***LandClim***, climatological conditions will be used for the land-surface initialisation. The climatology is computed for a window of 10 days around the starting month of the predictions. These experiments will follow the schedule indicated in Table 1.

	M1-M2	M3-M4	M5-M6	M7-M8	M9-M10	M11-M12
LandInit	🟦	🟦	🟦			
LandClim				🟦	🟦	🟦

Table 1: Approximate schedule of the experiments to be performed. M5 stands for Month 5.

Table 2 summarises the resources requested and detailed in section 6.

	Simulations (months)	Total Core-hours (millions)	Total archive (GB)
LandInit	1,680	2.5	16,200
LandClim	1,680	2.5	16,200
Total	3,360	5	32,400

Table 2: Resources requested. The estimates have been made on the basis of a cost of 0.9 wallclock hours (by using 1,616 cores) and 9.6 GB of output per month of simulation with the T511L91-ORCA025L75 configuration. These estimates have been obtained by running EC-Earth3.1 on MareNostrum III with a reduced output with respect to the CMIP5 standard.

The Autosubmit software (Asif et al., 2014), briefly described below, will be used to manage the workflow and ensure a uniform and optimal use of the resources. Our seasonal predictions offer a large computing flexibility as they are a set of 420 simulations (21 years, 2 start months, 10 members) which can be run independently and therefore in parallel for an optimal use of the PRACE computing resources. They will be managed, and packed in groups in a single big job if required, by Autosubmit to optimize the use of the machine and avoid collapsing the I/O system. The data storage and data transfer can be organized with a disk space of 10 TB in the “scratch” file system. This required scratch space is motivated by the requirement to perform many independent simulations at the same time each producing up to 80 GB of raw data. These data will be held in the workspace before being post-processed to reduce the data volume to less than a fifth of their original size, which will then be transferred to the “archive” file system and finally locally. Around 500 GB of “home” space will be required to host the code and its modified versions.

3. Describe the numerical methods and algorithms that you are planning to use, improve, or develop, the codes, packages or libraries that you need to undertake the project, and how these will enable the research to be achieved. (1 page)

EC-Earth3 comprises three major components: IFS, NEMO and OASIS3. It is essential to configure and build a separate executable for each one of them. The resolution proposed (T511: ~425,000 points, ORCA025: ~1,475,000 points) will help efficiently share calculations between 1000-1500 sub-domains, increasing the range of efficient compute-core usage per model executable. IFS and NEMO fully support a parallel environment, while OASIS3 supports a pseudo-parallel environment. OASIS3 requires Cray pointers. For IFS there is a possibility to activate an OpenMP switch but, in this case, the implemented MPI should be thread-safe. IFS generates the output in GRIB format and NEMO in NetCDF, while OASIS3 does not generate any output. At the end of a simulation the three components always generate restarts separately (IFS in binary, and NEMO and OASIS3 in NetCDF format).

For configuring and building the model executables, GNU make 3.81 or 3.81+, FORTRAN 77/90/95 complaint compiler with pre-processing capabilities and NetCDF4 deployed with HDF5 and SZIP are needed. A newly designed tool for automatic build configuration called “ec-conf” can be used. This useful tool requires Python 2.4.3 or 2.4.3+ (although it does not work yet with Python 3.0+). For NEMO, the FCM bash and perl mechanism is essential, as it is the I/O GRIB_API 1.9.9 or 1.9.9+ and GRIBEX 370 mechanism that are needed for IFS. To test the model with the run scripts, GNU date (64-bit) is also required.

The simulations will require MPI libraries and runtime facilities (MPICH2, MPICH-MX, HP-MPI, OpenMPI), optimization and data handling tools, such as BLAS, LAPACK, HDF4, HDF5, NETCDF, PARMETIS, SCALAPACK, P-NETCDF, UDUNITS, GRIB_API, CDFTOOLS v2, CDO, NCO and general configurations tools, such as PERL, PYTHON, AUTOCONF and AUTOMAKE.

4. **Explain why this project needs to run on a Tier-0 system, why the machine you have requested is suitable for the project and how the use of the system will enable the science proposed. You should describe the architecture, machine/system name and the problem sizes that have been used to test for scaling and provide supporting evidence.**

In LSIHP, most of the planned simulations are independent of each other, allowing an ensemble of simulations to be bundled into a single big job for efficient use of parallel architectures. Members of the BSC Earth Science Deparment recently tested EC-Earth3 managed by the Autosubmit system (Asif et al., 2014) on the Cray XE Jaguar in collaboration with the INCITE program of the US Department of Energy, achieving an efficient use of ~60,000 cores with a single big job wrapping an ensemble of high-resolution integrations.

Total number of EC-Earth3 cores	Elapsed time (in seconds)			
	NEMO (384)	NEMO (576)	NEMO (768)	NEMO (960)
583	2,042.67			
775	1,387.67	2,042.67		
967	1,417.00	1,203.67	2,041.67	
1159	1,432.00	1,118.67	1,208.00	2,126.00
1351	1,434.33	1,104.67	1,095.67	1,234.67
1543	1,405.67	1,122.00	1,099.33	1,114.33
1735	1,434.67	1,091.67	1,119.33	1,142.00
1927	1,421.33	1,097.67	1,087.33	1,126.00
2119	1,432.33	1,086.33	1,120.67	1,130.00
2311	1,480.67	1,170.00	1,076.33	1,127.67
2503		1,131.00	1,063.33	1,184.00
2695			1,109.67	1,169.67
2887				1,296.33

Table 3: EC-Earth3 (T511-ORCA025) ran for 10 simulation days to obtain the elapsed time (average of 3 sets). The total number of cores used for EC-Earth3 is indicated in the first column while the number of cores used for NEMO (384, 576, 768 or 960) is indicated in columns 2 to 5 and OASIS always use 7 (fixed). The number of cores used for IFS is the difference between that used for EC-Earth and the sum of those used for NEMO and OASIS.

The model, its version and the highest available configuration released till date, have been chosen for LSIHP. It has already been well deployed at MareNostrum III (a tier-0 system; having 2x Intel SandyBridge-EP E5-2670/1600 8-core processor at 2.6 GHz, 8x4GB DDR3-1600 DIMMS (2GB/core) and Infiniband FDR10 network) and used within the previous PRACE projects “HiResClim” and “HiResClim2”. Members of the BSC Earth Sciences Department therefore have experience on running EC-Earth3 on MareNostrum III, and have invested a lot of effort and resources in finding the best configuration for this model in this environment.

During these previous projects, along with scientific experimentation, an extensive scaling exercise was made (Asif et al., 2014) with EC-Earth3, where the model was run, not only in coupled mode, but its components were also analyzed in uncoupled/stand-alone mode. To illustrate the problem size, Table 3 and Figure 1 depict the elapsed time obtained by running the model in coupled mode. Figure 1 also shows the respective speedup and efficiency plots.

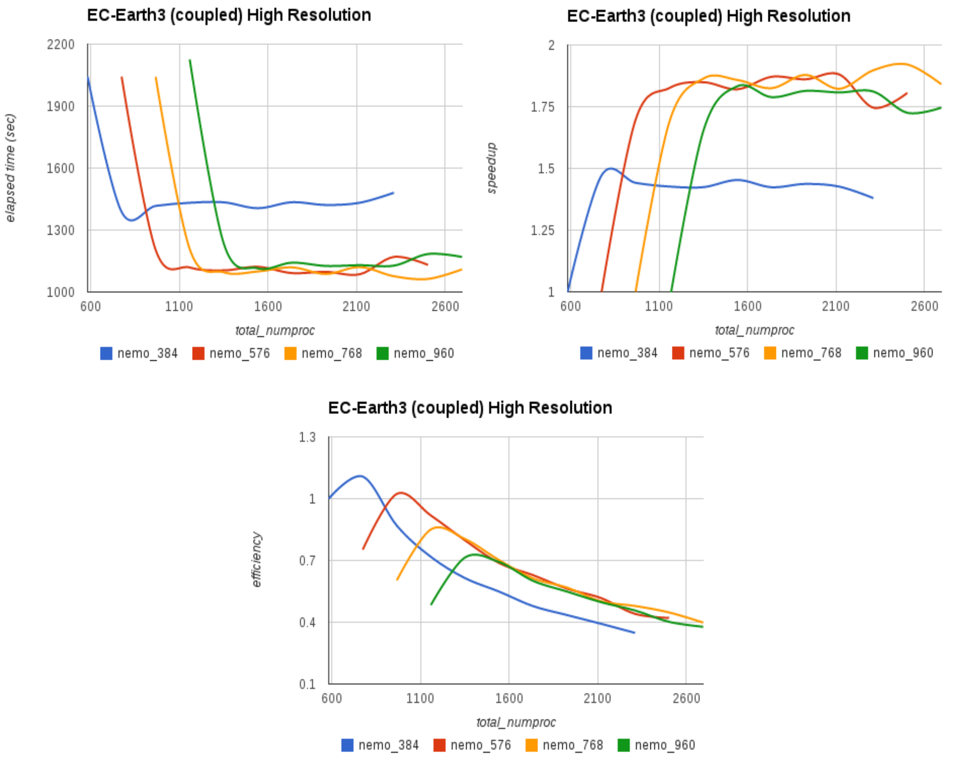

Figure 1: EC-Earth3 (T511-ORCA025-LIM3) ran for 10 simulation days to estimate the elapsed time (average of 3 sets). The blue, red, orange and green curves correspond to results for tests with NEMO using 384, 576, 768 and 960 processors while the speedup and efficiency values are relative to the reference test with 583 processors for EC-Earth.

During this EC-Earth scaling exercise, the oceanic component of EC-Earth3 was identified as a clear bottleneck which limits the scalability of the coupled model. To enhance the speed of this oceanic component, the north-fold optimization technique was tested and adopted for EC-Earth3 (according to suggestions from NEMO experts). A 30% gain in performance in NEMO stand-alone mode and a 20% gain in EC-Earth coupled mode were obtained. Thanks to these improvements, the model could complete one month of simulation in 0.9 hours with IFS+NEMO+OASIS: 640+960+7=1607~1616 cores (according to MareNostrum III support, for a better optimization, the total number of cores should be divisible by the number of cores per node so instead of 1607, a total number of cores of 1616 was chosen)

Investigation on EC-Earth3 using the best parallel model performance tools are still ongoing in two directions to reach an optimum scalability, in collaboration with the BSC Computer Sciences department. These actions include not only an adjustment of the model configuration and a balance of the number of cores dedicated to each of its component, but also modifications of the code itself and work on the parallel programming models.

5. **Describe your experience of using HPC resources in the past and how you will manage using a Tier-0 system. What other experience do you and your team bring to this project? (1 page).**

EC-Earth2 made extensive contributions to the prediction protocols in CMIP5 and the BSC Earth Sciences Department was one of the most active partners of the EC-Earth consortium. Members of the BSC Earth Sciences Department have also contributed extensively to a large number of past and on-going EU projects (CLIM-RUN, QWeCi, DENFREE, PREFACE, EUCLEIA, IS-ENES2, SPECS, which is coordinated by BSC, and EUPORIAS). The BSC Earth Sciences Department has been involved in the IS-ENES/PRACE-1IP working group focusing on the EC-Earth3 adaptation to Tier-0 machines. It has tested a range of EC-Earth3 configurations, in the atmospheric resolutions T255/511/799 on several HPC systems: SGI Altix 3500, NEC-SX6, Linux cluster with Intel Xeon, Dell PowerEdge 2900, IBM pSeries 575 Power6 and IBM Power PC. The BSC Earth Sciences Department leads the development of Autosubmit, a python-based wrapper software that can manage any type of climate simulation workflow. It can also bundle numerous jobs into a single big job for performing job-control on parallel sets of simulations throughout the execution. Finally, members of the BSC Earth Sciences Department coordinated the recent HiResClim and HiResClim2 projects supported by PRACE, which demonstrates its ability to lead this kind of projects.

The BSC Earth Sciences Department has also expertise in analyzing parallel programming model codes using cutting-edge tools. This allows continuously carrying out performance analyses to reach the optimum configurations for EC-Earth and NEMO. Collaboration has been established with the NEMO developers for sharing the performance reports and code optimization tools and techniques. A similar collaboration is being established with the OpenIFS developers at ECMWF. Figure 2 illustrates an example of such performance analyses. These are two different views as shown by the Paraver tool of a NEMO 3.4 run iteration using the ORCA025 domain with 256 cores. Paraver uses different types of plots to represent the values stored in a trace file. A trace file is a file holding information of an execution to perform a subsequent analysis.

The left panel of Figure 2 offers a qualitative view of the simulation performance. It depicts a routine distribution along time, with each model function displayed with a different color along the horizontal axis, which represents the time. Different horizontal lines are used for each process, running from top to bottom. This view is obtained using a function instrumentation profiling when generating the trace file, but the information it provides is key to understand the models' bottlenecks. It is useful to compare the running time of the routines and its variability among the different processes. The ideal situation consists in straight vertical lines, meaning that each routine behaves in the same way for all the cores. However, this is not the case for NEMO, where the lines show a certain degree of tilting suggesting that there is some imbalance among the processes. The right panel of Figure 2 is a quantitative view that represents the different parts of the iteration according to their useful duration, i.e. computational bursts, or regions of the code, between two consecutive MPI calls. As in the previous panel, the horizontal and vertical axes also represent the time and the different processes. This view shows if a region of the code has parts that take longer, either because they have more instructions or because they have a low efficiency. This is complemented with other diagnostics that show the instructions per cycle or the instructions count for the different code regions.

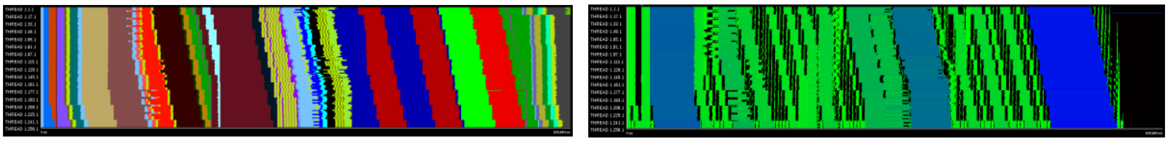

Figure 2: (Left panel) Function instrumentation view where each routine is displayed with a different color and (Right panel) useful duration view with a color gradient used to display a range of time values for the computational regions ranging between light green and dark, blue for a NEMO 3.4 model iteration. The time within the iteration is in abscissa and the different processes in ordinates

6. **Justify the number of core hours requested. This should include information such as: run type, wall clock time per step, number of jobs per run type, the number of CPU cores and the total core hours per run type. This information should take the form of a table like the one shown below with example data. Explain how the core hours requested will be used (1 page).**

Table 4 lists the two experiments described in section 2, which comprise each 420 independent simulations of 4 month length, i.e. a total of 1680 months of simulations, that can be run independently but simultaneously depending on the machine load by Autosubmit. These experiments use the HR EC-Earth3 configuration. Its cost, for a four-month simulation requires a wall-clock time of 3.6 hours. A benchmarking performed within the framework of the HiResClim project suggests that, taking into account the average load of the MareNostrum III queues, optimum performance is obtained using 1,616 procs. This configuration generates 9.6 GB of output per month of simulation.

Run type	# Runs	# Steps/Run	Walltime/Step	# CPU cores	Total core hours/Type Run
LandInit	420	1	3.6	1616	2443392
LandClim	420	1	3.6	1616	2443392
Total					4886784

Table 4: Cost of the experiments proposed.

The final estimate is for a total request of 5 million core-hours, which includes the numbers described in Table 4 plus a small buffer of 5% to account for failing jobs that will need to be repeated.

The experiments will be run using Autosubmit, the launching and monitoring solution developed by the group of the applicant that allows the remote submission of EC-Earth and NEMO experiments. Autosubmit includes in the workflow of the experiments a job that retrieves the data back to either the storage requested in the joint data pilot call or, alternatively if not granted, to the Department data storage as soon as the simulation corresponding to a member has completed. This means that the estimates for the archive are an absolute upper value in case the automatic download does not perform as expected. It is very likely that this number will approach a figure ten times smaller.

References

Alessandri, A., Borrelli, A., Masina, S., Cherchi, A., Gualdi, S., Navarra, A., Di Pietro, P., Carril, A.F. (2010) The INGV-CMCC seasonal prediction system: Improved ocean initial conditions. Mon Weather Rev, 138, 2930-2952.

Asif, M., A. Cencerrado, O. Mula-Valls, D. Manubens, F.J. Doblas-Reyes and A. Cortés (2014). Impact of I/O and data management in ensemble large scale climate forecasting using EC-Earth3. Procedia Computer Science, 29, 2370-2379, doi:10.1016/j.procs.2014.05.221.

Bellprat O, Kotlarski S, Lüthi D, Schär C (2013) Physical constraints for temperature biases in climate models. Geophys Res Lett 40:4042–4047. doi: 10.1002/grl.50737

Challinor, A., Slingo, J., Wheeler, T., and Doblas-Reyes, F. (2005). Probabilistic simulations of crop yield over western India using the DEMETER seasonal hindcast ensembles. Tellus A, 57(3).

Dee, D.P., and Coauthors (2011) The era-interim reanalysis: configuration and performance of the data assimilation system. Q J R Meteorol Soc, 137,553–59, doi:10.1175/JCLI3812.1.

Delworth, T.L., and Coauthors (2012) Simulated climate and climate
change in the GFDL CM2.5 High-resolution coupled climate model. J Climate, 25, 2755-2781, doi:10.1175/JCLI-D-11-00316.1.

Doblas-Reyes, F. J., Hagedorn, R., Palmer, T. N., & Morcrette, J. J. (2006). Impact of increasing greenhouse gas concentrations in seasonal ensemble forecasts. Geophysical Research Letters, 33(7), L07708.

Douville H (2010) Relative contribution of soil moisture and snow mass to seasonal climate predictability: A pilot study. Clim Dyn 34:797–818. doi: 10.1007/s00382-008-0508-1

Ferry, N. et al. Mercator global Eddy permitting ocean reanalysis GLORYS1V1: Description and results. Mercator Ocean Quart. Newsl. 36, 15–27 (2010).

García-Morales, M. B., and Dubus, L. (2007). Forecasting precipitation for hydroelectric power management: how to exploit GCM's seasonal ensemble forecasts. International Journal of Climatology, 27(12), 1691–1705.

Guemas, V., Doblas-Reyes, F.J., Mogensen, K., Keeley, S., Tang, Y. (2014) Ensemble of sea ice initial conditions for interannual climate predictions. Climate Dynamics, doi:10.1007/s00382-014-2095-7.

Hazeleger, W., and Coauthors (2012) EC-Earth V2.2: description and validation of a new seamless earth system prediction model. Climate Dynamics, 39, 2611-2629, doi:10.1007/s00382-011-1228-5.

Hirschi M, Seneviratne SI, Alexandrov V, et al. (2011) Observational evidence for soil-moisture impact on hot extremes in southeastern Europe. Nat Geosci 4:17–21. doi: 10.1038/ngeo1032

Jung, T., and Coauthors (2012) High-resolution global climate simulations with the ECMWF model in Project Athena: Experimental design, model climate and seasonal forecast skill. J Climate, 25, 3155-3172, doi:10.1175/JCLI-D-11-00265.1.

Kinter III, J.L., and Coauthors (2013) Revolutionizing climate modeling with Project Athena: A multi-institutional, international collaboration. Bull Am Meteorol Soc, 94, 231-245, doi:10.1175/BAMS-D-11-00043.1.

Koster RD, Mahanama SPP, Yamada TJ, et al (2010) Contribution of land surface initialization to subseasonal forecast skill: First results from a multi-model experiment. Geophys Res Lett 37:L02402. doi: 10.1029/2009GL041677.

Koster RD, Mahanama SPP, Yamada TJ, et al (2011) The Second Phase of the Global Land–Atmosphere Coupling Experiment: Soil Moisture Contributions to Subseasonal Forecast Skill. J Hydrometeorol 12:805–822. doi: 10.1175/2011JHM1365.1

Prodhomme, C., Doblas-Reyes, F, Bellprat, O, Dutra, E. (2015) Impact of land-surface initialization on sub-seasonal to seasonal forecasts over Europe. Climate Dynamics, doi:10.1007/s00382-015-0094-3.

Rodwell, M., Doblas-Reyes, F.J. (2006) Predictability and prediction of European monthly to seasonal climate anomalies. J Climate, 19, 6025–6046.

Saha, S., and Coauthors (2006) The NCEP Climate Forecast System. J Climate, 19, 3483-3517, doi:10.1175/JCLI3812.1

Scaife, A., and Coauthors (2014) Skillful long-range prediction of European and North American winters. Geophys Res Lett, 41, 2514-2519, doi: 10.1002/2014GL059637.

Schär C, Lüthi D, Beyerle U, Heise E (1999) The soil-precipitation feedback: A process study with a regional climate model. J Clim 12:722–741.

Seneviratne SI, Corti T, Davin EL, et al (2010) Investigating soil moisture-climate interactions in a changing climate: A review. Earth-Science Rev 99:125–161. doi: 10.1016/j.earscirev.2010.02.004

Seneviratne S, Wilhelm M, Stanelle T, et al. (2013) Impact of soil moisture‐climate feedbacks on CMIP5 projections: First results from the GLACE‐CMIP5 experiment. Geophys Res Lett 40:5212–5217. doi: 10.1002/grl.50956

Shukla, J., Kinter, J.L., Palmer, T.N., Hagedorn, R. (2006) Predictability of seasonal climate variations a pedagogical review. Predictability of Weather and Climate. Cambridge, UK:Cambridge University Press.

Stockdale, T.N., D.L.T. Anderson, M.A. Balmaseda, F.J. Doblas-Reyes, L. Ferranti, K. Mogensen, T.N. Palmer, F. Molteni and F. Vitart, (2011) ECMWF seasonal forecast System 3 and its prediction of sea surface temperature. Climate Dynamics, 37, 455-471

Taylor K E, R J Stouffer, G A. Meehl (2012) An Overview of CMIP5 and the Experiment Design. Bull. Amer. Meteor. Soc., 93, 485–498. doi: http://dx.doi.org/10.1175/BAMS-D-11-00094.1

Thomson, M. C., Doblas-Reyes, F. J., Mason, S. J., Hagedorn, R., Connor, S. J., Phindela, T., Morse, A. P., and Palmer, T. N. (2006). Malaria early warnings based on seasonal climate forecasts from multi-model ensembles. Nature, 439(7076), 576–9.