A partnership in weather and climate research

Joint Weather and Climate Research Programme

Weather and Climate Modelling: highlights from PRACE-UPSCALE

Pier Luigi Vidale

Willis Professor of Climate System Science and Climate Hazards

Marie-Estelle Demory, Reinhard Schiemann,
Jane Strachan, Bryan Lawrence, Hilary Weller

University of Reading

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Malcolm Roberts

Matthew Mizielinski, Lizzie Kendon, + Met Office

(with thanks to the many MO groups

Involved in model development and elsewhere)

Met Office

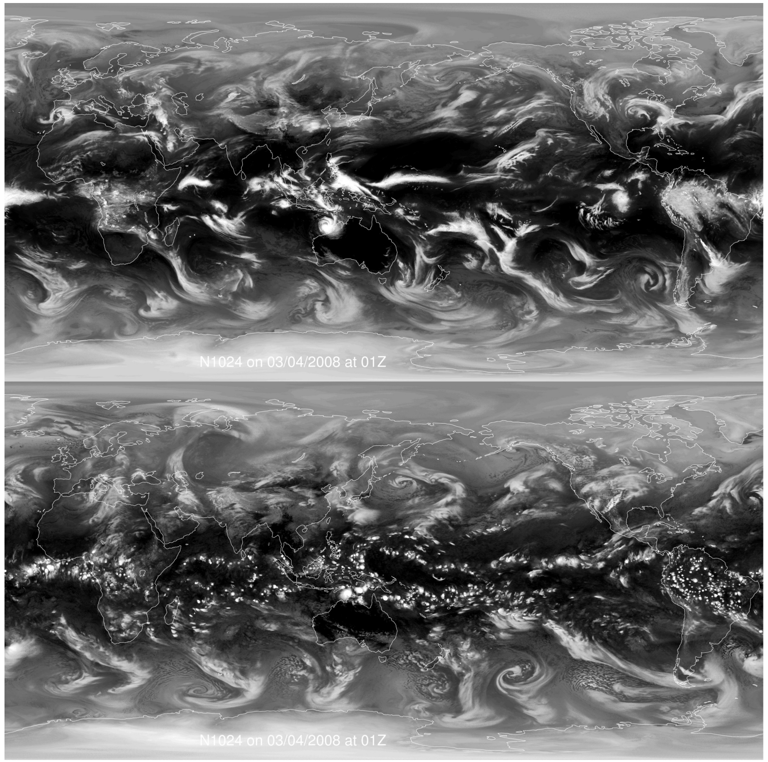

Outline

- UK High-resolution Global Climate Modelling programme:
  - Motivation
  - GCM development
- The PRACE-UPSCALE grant
  - Scalability achievements
  - Experimental design and goals
- Science highlights:
  1. European heat waves
  2. Hurricanes and typhoons in the climate system
- Enabling future GCM development
- Future directions

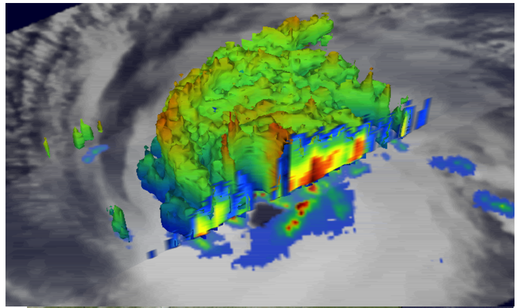
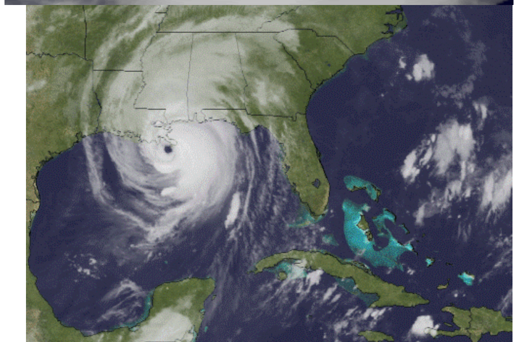

Knowledge Transfer Partnerships

Willis Research Network

Natural Catastrophes of 2011

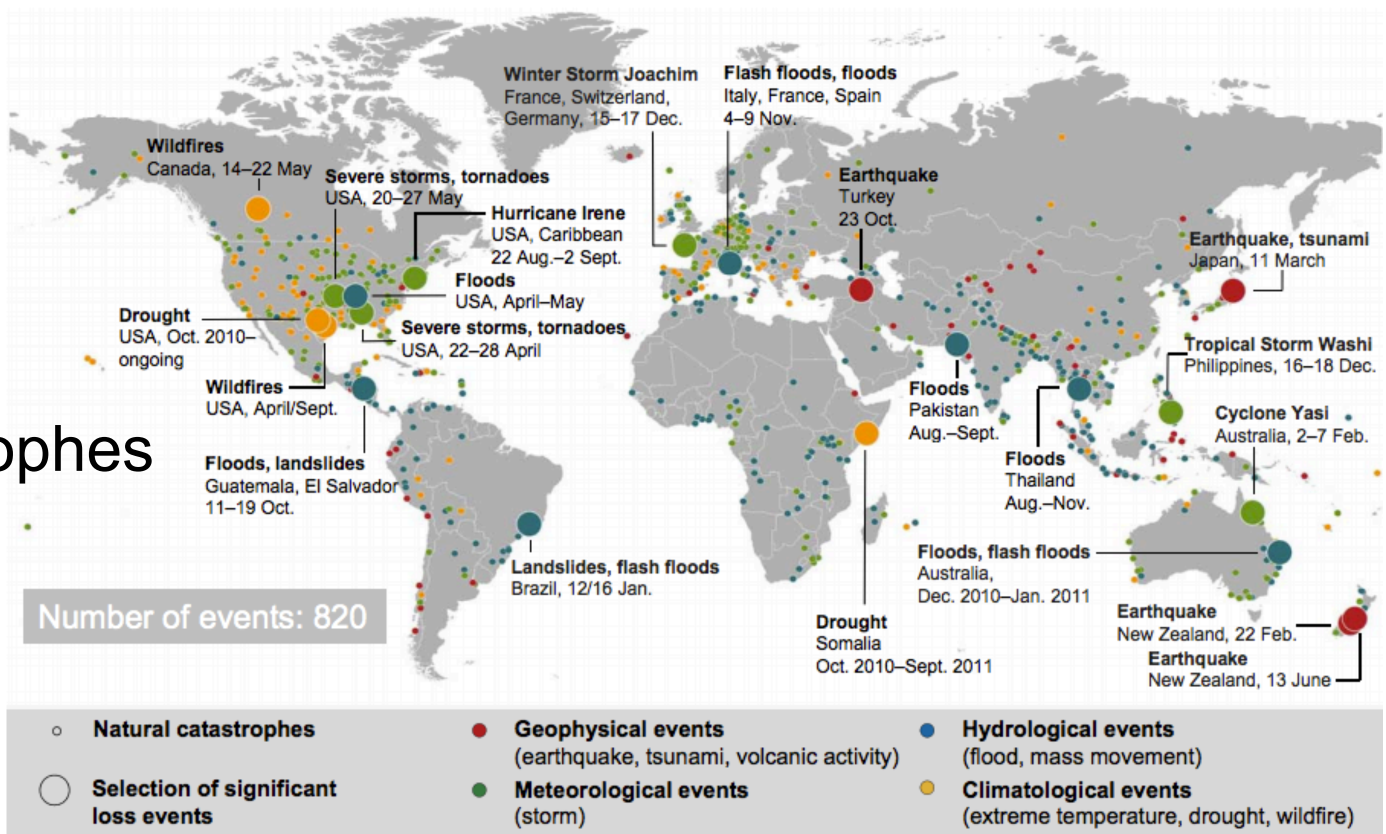

Comparing 2011 with previous years

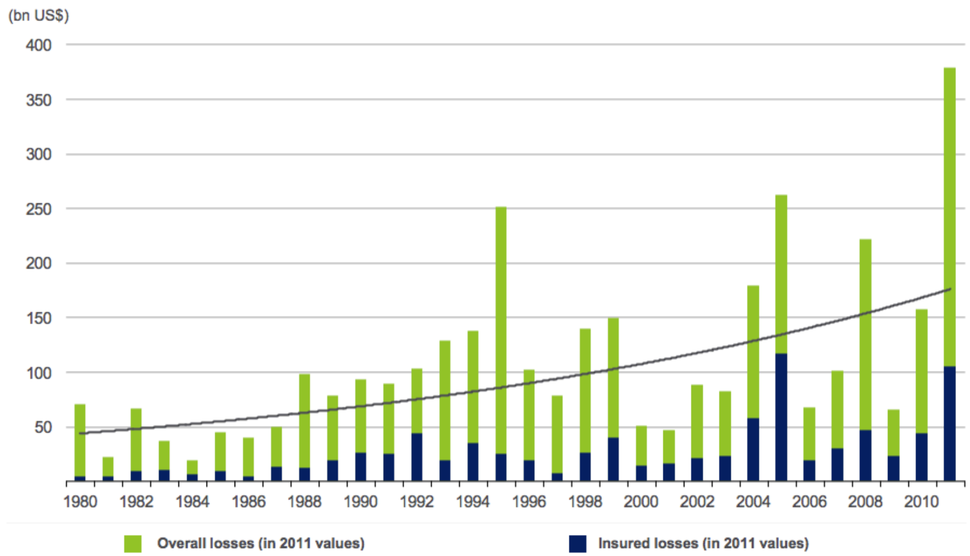

Munich Re 2012

Motivations for a **Weather and Climate Hazards Laboratory**:
Economic Impact of Weather-related Natural Catastrophes of 2011

- 2011= costliest year ever in terms of natural catastrophe
- **US$380bn** global economic losses (120bn higher than 2005)
  - Of which **only US$105bn** were insured losses
- Although earthquake dominated loss in 2011 (Japan Tsunami prominent), still **90% of the number of natural catastrophes were weather-related**

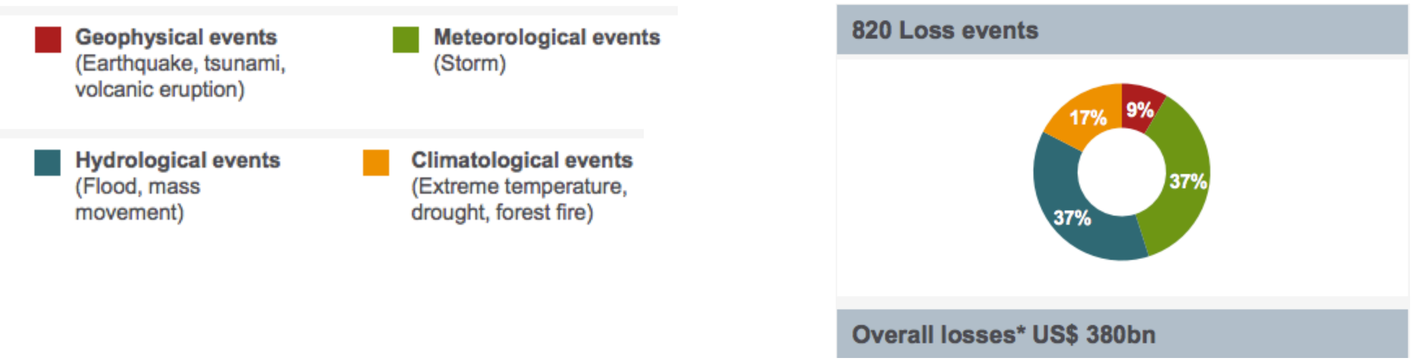

Munich Re 2012

- Observations of extreme events are rare, short, inhomogeneous. **Climate models can provide much needed complementarity; synthetic data sets for W&C extremes.**
- UK insurance and UK science share a need to understand the world around us and to understand how it is changing.
- Need for continuous engagement between the insurance industry and the scientific community to ensure industry has the best possible information to increase its market resilience to W&C risk

Knowledge Transfer Partnerships

Willis Research Network

Evolution of N. Atlantic hurricane frequency in past 100+ years: connections with C-Atlantic SSTs.

2005 was a true record year:
15 hurricanes (incl. Katrina), 27 named storms … and some of the most intense storms in US history.
Katrina damage = 1600 dead; 75-200 bn U$.
In the same region, in the Aug-Sep 2007: Dean (cat 5), Felix (cat 5)

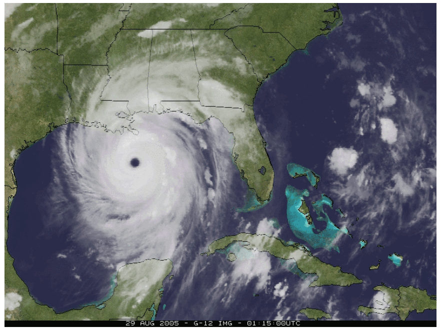

But what about before 1885, or 1950 (for the Pacific) ?

Short/inhomogeneous records of **extremely rare** events:

we need models to complement observations

And yet, coarse GCMs, especially those used for long (e.g. IPCC) integrations, cannot fully represent tropical cyclones

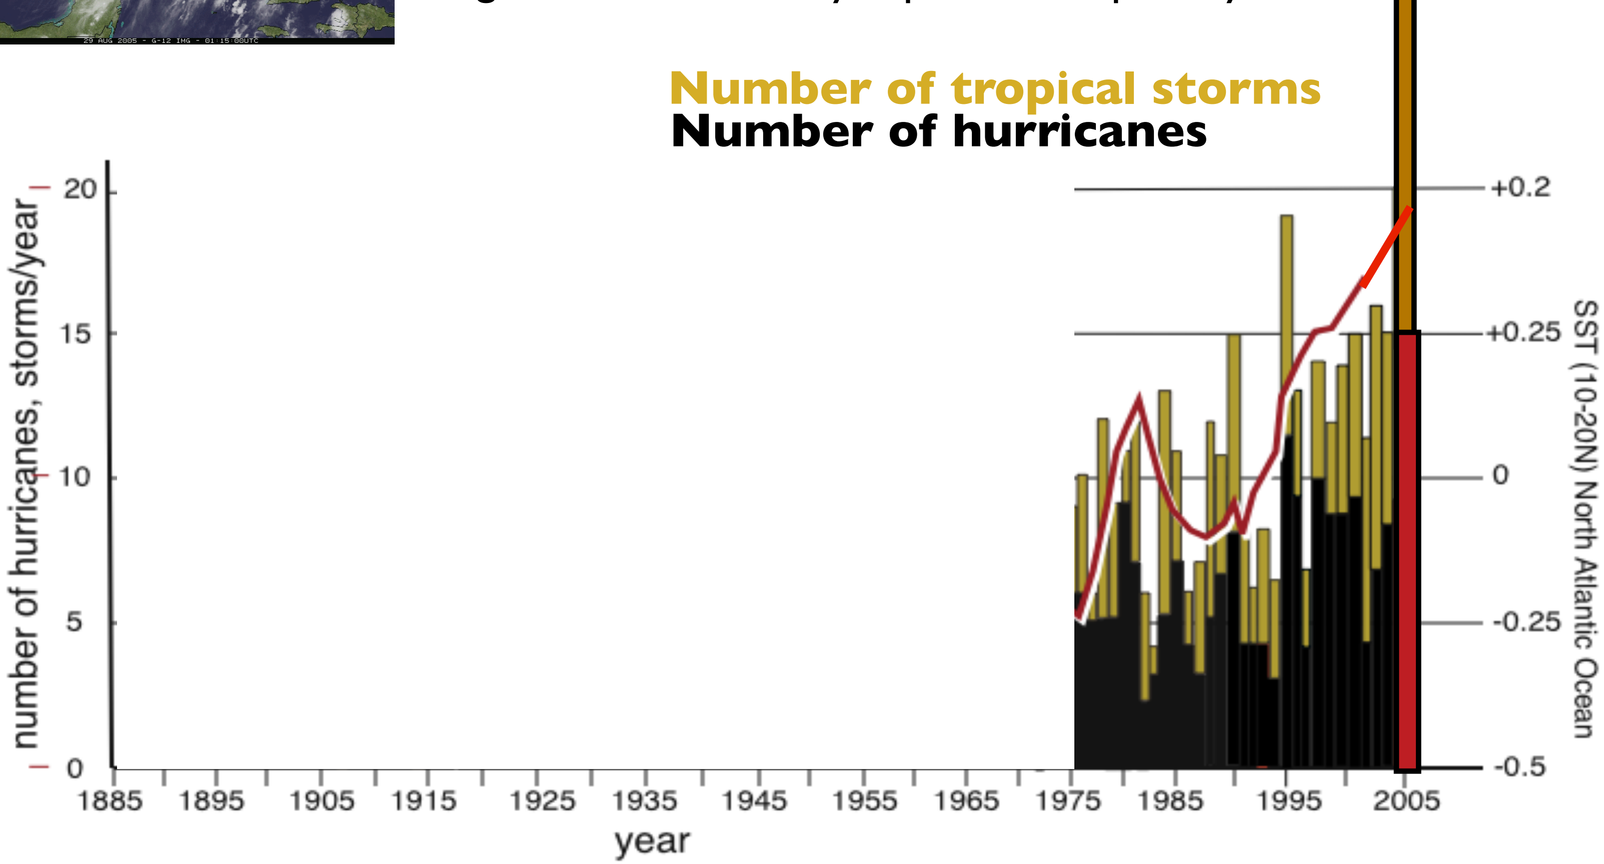

Courtesy of K. Trenberth

A **surrogate climate**:
using climate models to complement observational evidence and aid our understanding: decadal variability in hurricane frequency

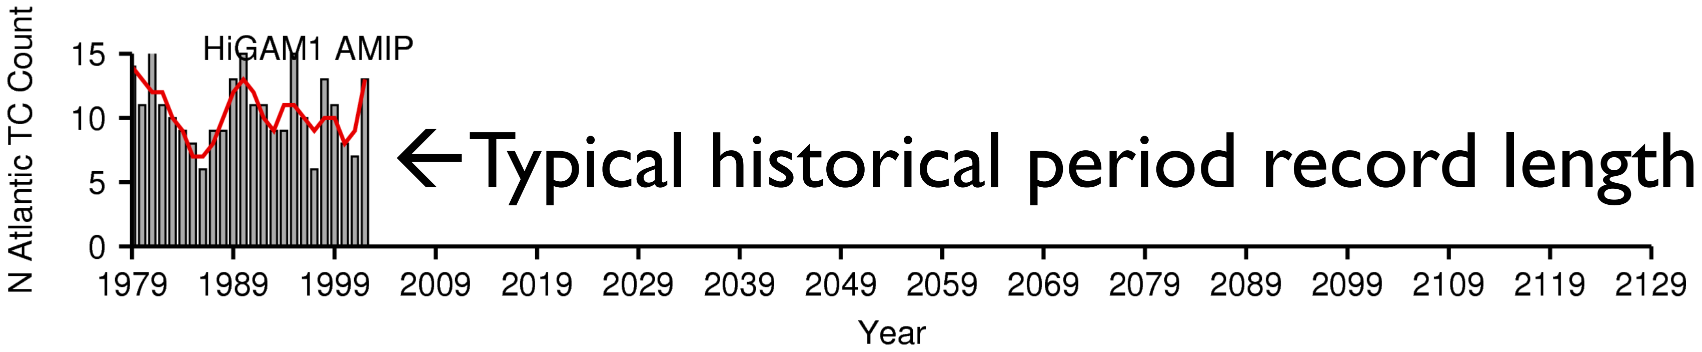

There is still large uncertainty regarding the regional details of climate change, which is what society really needs.

PROJECTED PATTERNS OF PRECIPITATION CHANGES

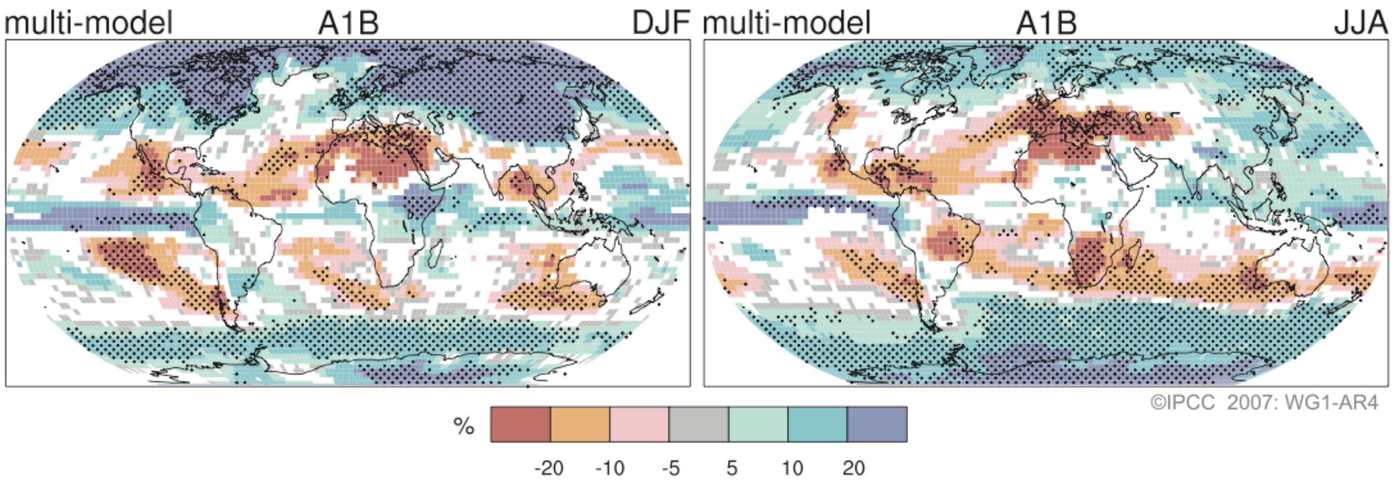

Figure SPM.7. *Relative changes in precipitation (in percent) for the period 2090–2099, relative to 1980–1999. Values are multi-model averages based on the SRES A1B scenario for December to February (left) and June to August (right). White areas are where less than 66% of the models agree in the sign of the change and stippled areas are where more than 90% of the models agree in the sign of the change. {Figure 10.9}*

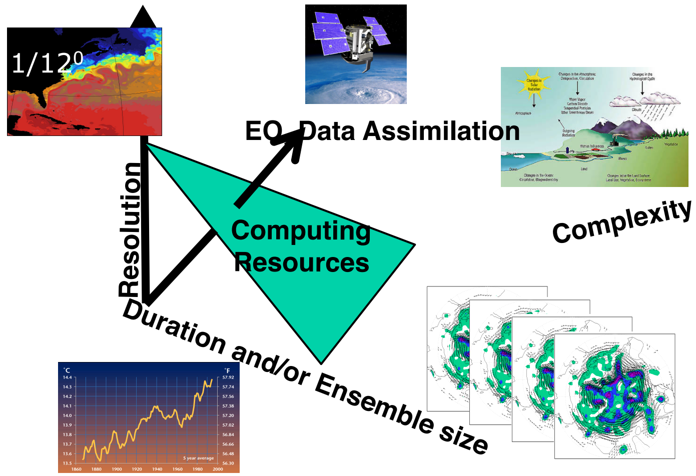

Challenge of a changing climate

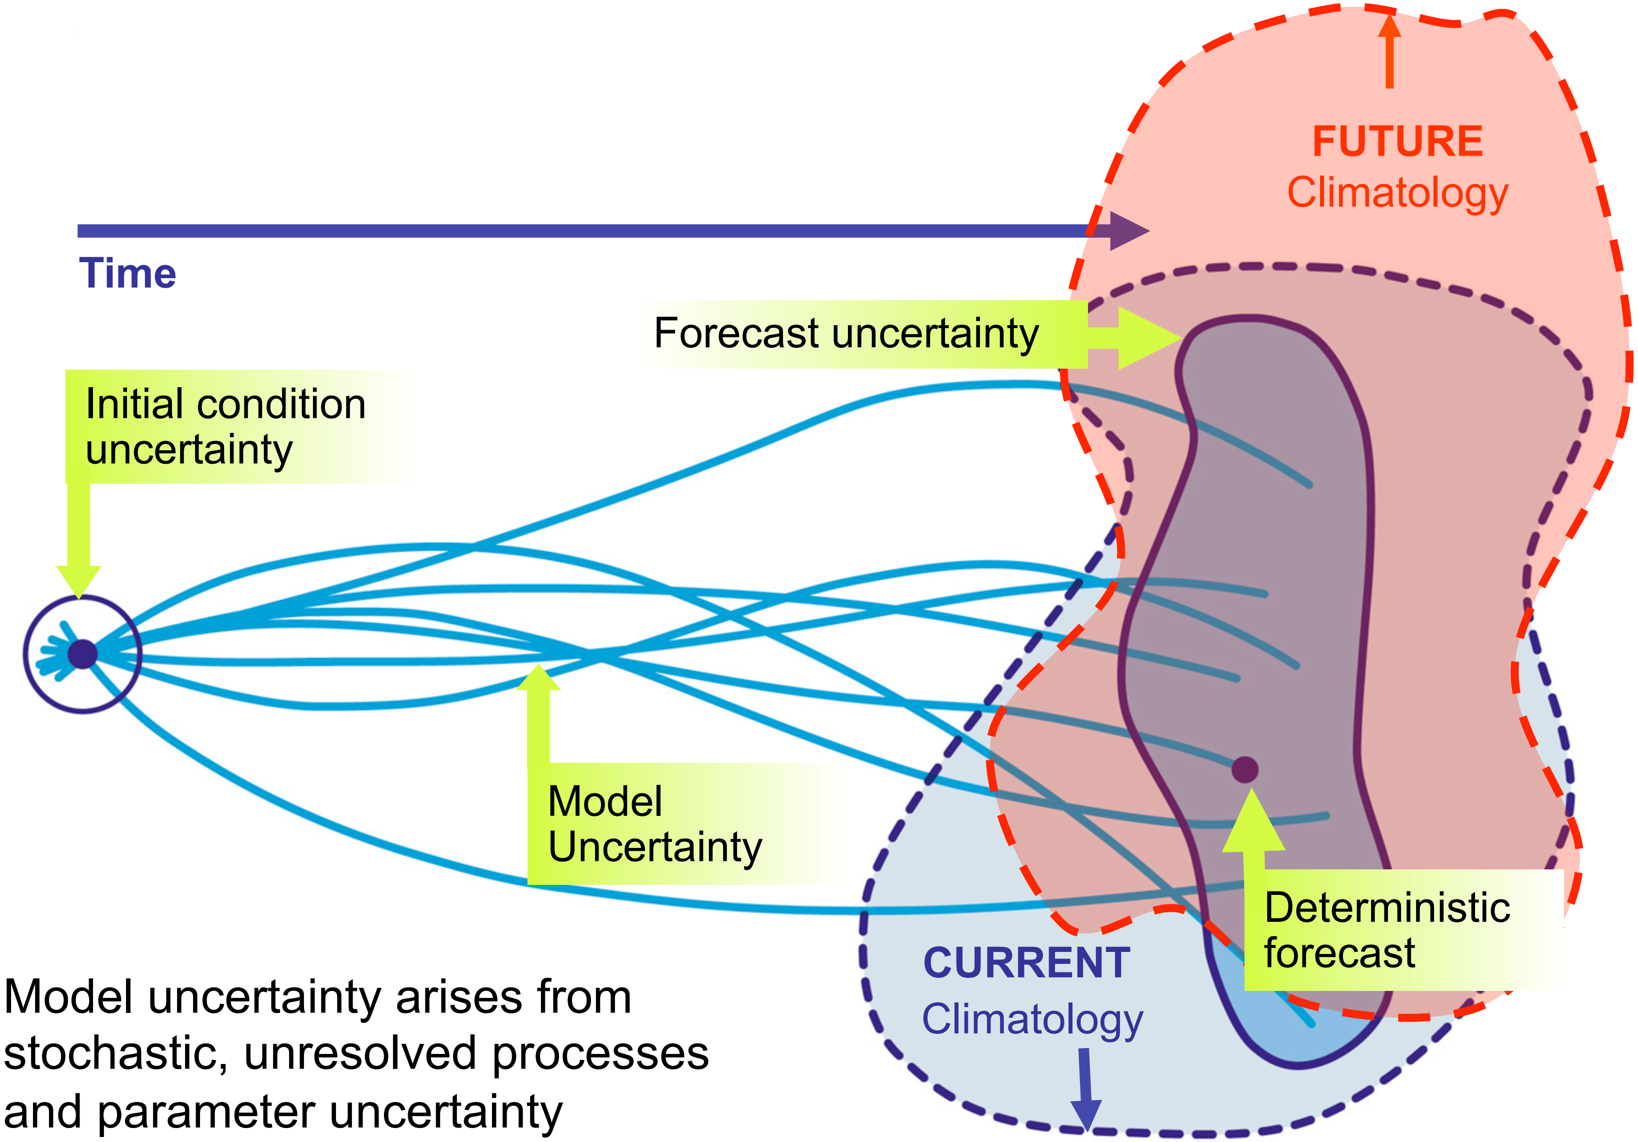

Model uncertainty arises from stochastic, unresolved processes and parameter uncertainty

How does a GCM work ?

We slice the planet in boxes… we apply the laws of physics.

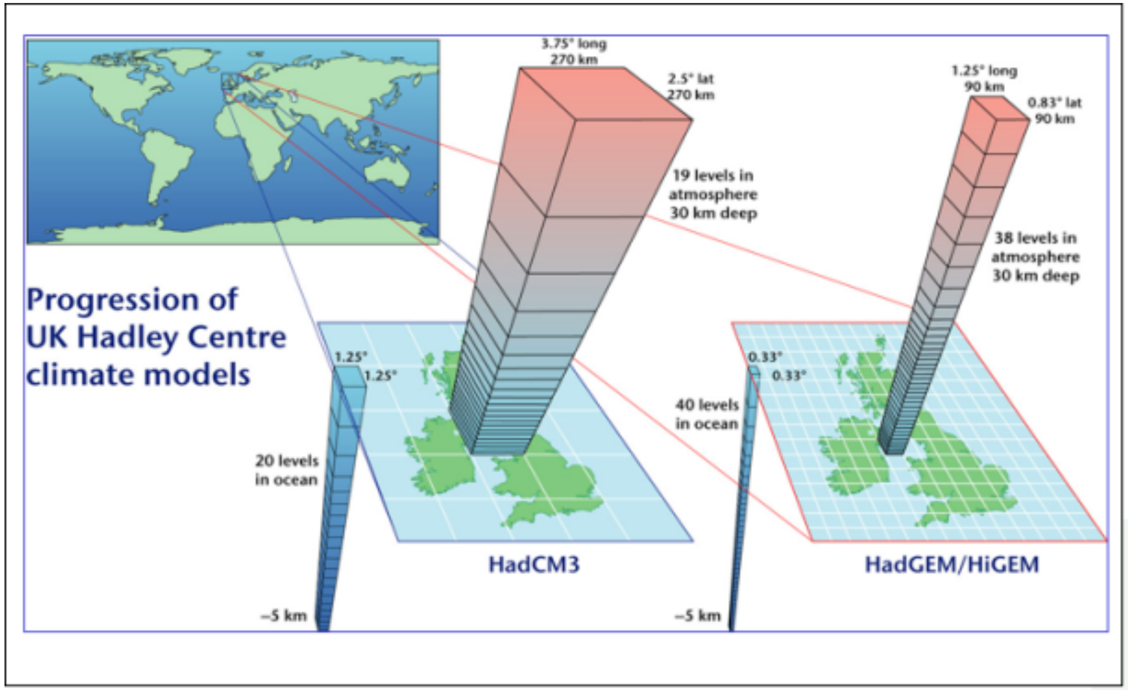

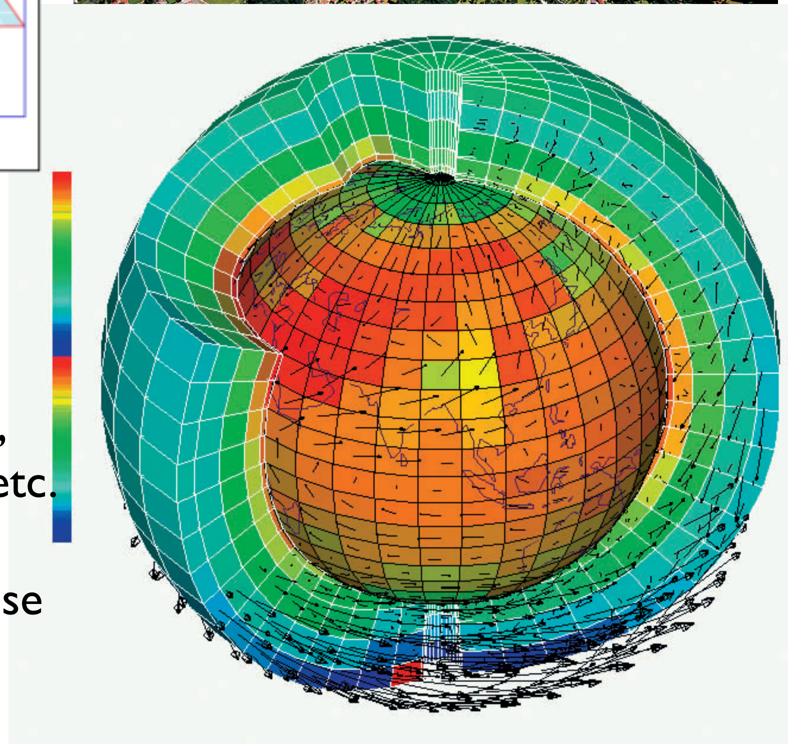

Analytical solutions are too hard: we use discretisation and numerical methods.

At every box location we compute radiation, winds, pressure, precipitation, temperature, using the laws of physics (gravitation, electromagnetism, thermodynamics, fluid dynamics, turbulence), chemistry, biology, ecology, etc.

We do all these computations every 5-30 minutes, for every single box. Every 2x resolution = 10x CPU expense

We need very large supercomputers !

Representation of orography: the importance of resolution

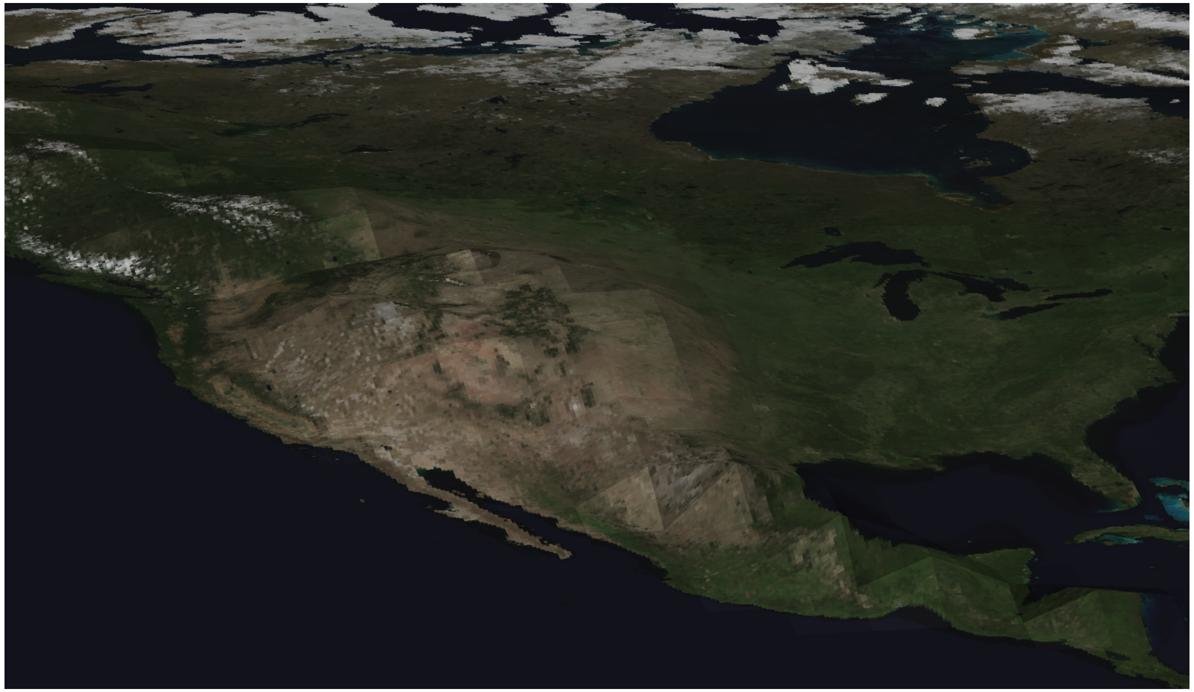

The upper figure shows the surface orography over North America at a resolution of ~300km, as in a low resolution climate model.

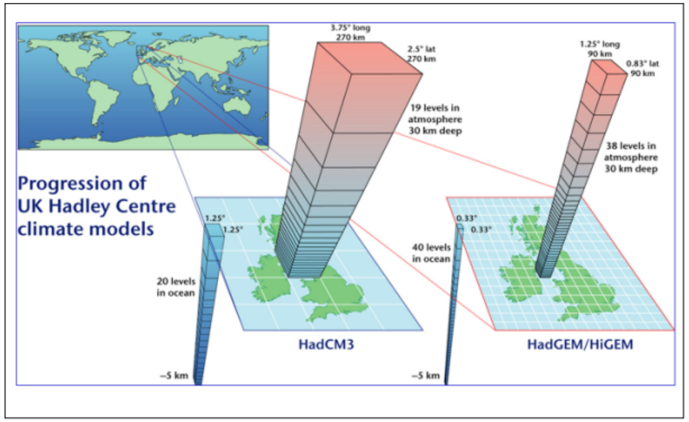

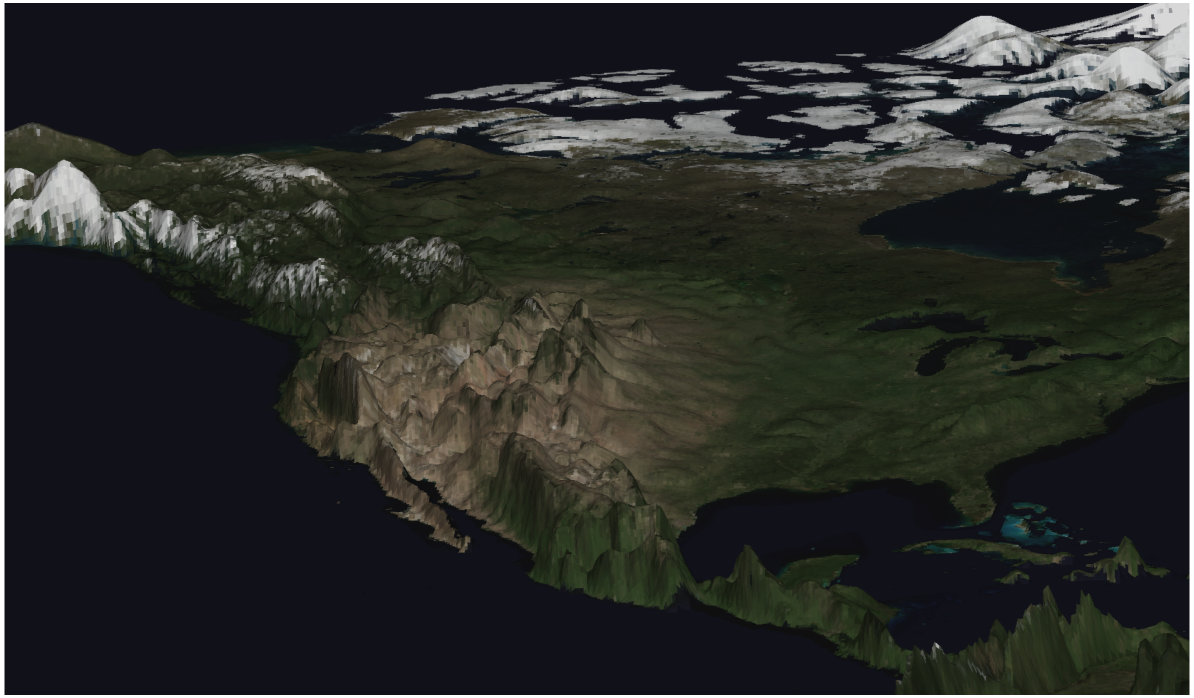

The lower figure shows the same field at a resolution of 25km, as in the climate model we are discussing today (HadGEM3-N512).

Remember that orographic processes are highly non-linear.

Resolution in itself is not a valid metric of the quality/fidelity, nor of the usefulness of a modelling **system**

16 MPixel

16 MPixel

The PRACE-UPSCALE Project

•Joint Weather & Climate Research Programme
A partnership in climate research

UK on PRACE - weather resolving Simulations of Climate for globAL Environmental risk
Current "numerical mission" of the JWCRP High-resolution climate modelling team
PI: P.L. Vidale, NCAS-Climate, Reading

In 2011 we demonstrated our capability in effectively exploiting 4'800, and up to 12'000 CRAY XE6 cores. As an ensemble of GCMs, we could **concurrently use up to 60'000 cores.**

- Cf. with Earth Simulator: we never managed to effectively use more than 88 cores (out of 5'400 cores in total)

Produced 2-4TB data/day, transferred in real time to the UK, ended up with ~400TB of data

AWARD: 144 million core hours, for 1 year.
Equivalent to: - 18x HadGEM2 submission to IPCC (= 8M core hours)
- half of the UK HECToR facility

<u>Completed</u>:

1. HadGEM3-A multi-decadal simulations at N96 (130 km) to N512 (25 km)
2. Development of a 12km (N1024) Global Climate Model

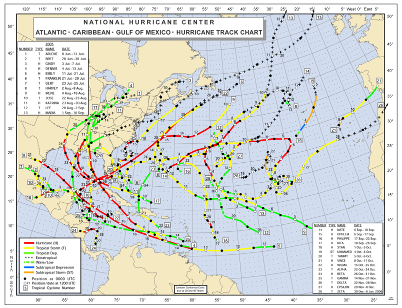

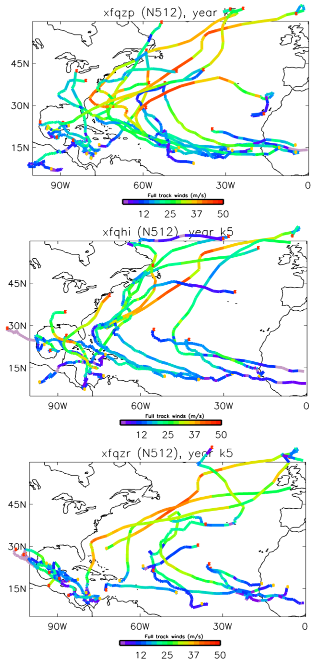

Present climate simulations

- forced with OSTIA SSTs
- 1985-2011 (27 years)
- 5 ensemble members, 27 years each

Future climate simulations

- 3 ensemble member, 27 years each
- following RCP8.5
- SST: daily OSTIA + HadGEM2-AO RCP8.5 2100 ΔSST

UPSCALE output available on JASMIN@CEDA

Model Grid Size (km) & Computing Capability

	Earth Simulator 2002-2009		PRACE-HERMIT 2012-		
Peak Rate:	**10 TFLOPS**	**100 TFLOPS**	**1 PFLOPS**	**10 PFLOPS**	**100 PFLOPS**
Cores	**1,400 (2005)**	**12,000 (2007)**	**80-100,000 (2009)**	**300-800,000 (2011)**	**6,000,000? (20xx?)**
Global NWP^0: 5-1(	**18 - 29**	**8.5 - 14**	**4.0 - 6.3**	**1.8 - 2.9**	**0.85 - 1.4**
Se 50-100					**1.3**
D 5-1(					**.2**
Change^2: 20-50 yrs/day	**120 - 200**	**57 - 91**	**27 - 42**	**12 - 20**	**5.7 - 9.1**

* Core counts above O(10^4) are unprecedented for weather or climate codes, so the last 3 columns require getting 3 orders of magnitude in scalable parallelization

teraFLOPS = 10^{12} (trillion) floating point operations per second
petaFLOPS = 10^{15} (quadrillion) floating point operations per second
exaFLOPS = 10^{18} (quintillion) floating point operations per second

Range: Assumed efficiency of 10-40%
0 - Atmospheric General Circulation Model (AGCM; 100 vertical levels)
1 - Coupled Ocean-Atmosphere-Land Model (CGCM; ~ 2X AGCM)
2 - Earth System Model (with biogeochemical cycles) (ESM; ~ 2X CGCM)

Thanks to Jim Abeles (IBM)

"Best practice" setup core counts on various architectures

Resolution		Number of cores				Notes
		NEC SX6	IBM P6	IBM P7	CRAY XE6	
N96 = 135km HG1-L38 HG3-L85		1*8 1*8	96	128		For HadGEM1: ES processors about 4x more powerful than P7
	Turnaround	1 sypd 10smo/day	3 sypd	3.5 sypd (EndGame)		
N216 = 60km HG1-L38 HG3-L85		11*8	192	3*32	1024	For HadGEM1: ES processors about 4x more powerful than P7
	Turnaround	1 sypd	8 smo/day	5 smo/day	13smo/day	
N512= 25km				64*32 40*32 (EG) 200*32 (L70 EG)	9408	
	Turnaround			7 smo/day 5.7 smo/day 2 sypd	6 smo/day	Ensemble of 5 runs, concurrent, up to 60K cores
N1024 = 12km				74*32		
	Turnaround			1.2 smo/day		

Science Highlight 1

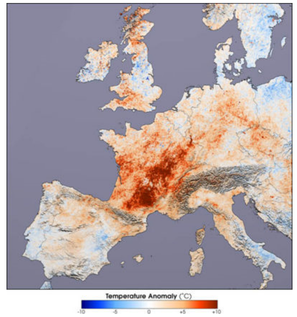

<u>European heatwaves</u>:

1. how many more summers like 2003 are likely and why ?
2. will these extreme summer episodes last longer and become more intense?

2003-type extreme summers
how frequent and why?
In the old times, only with regional models……

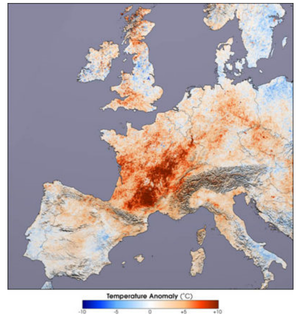

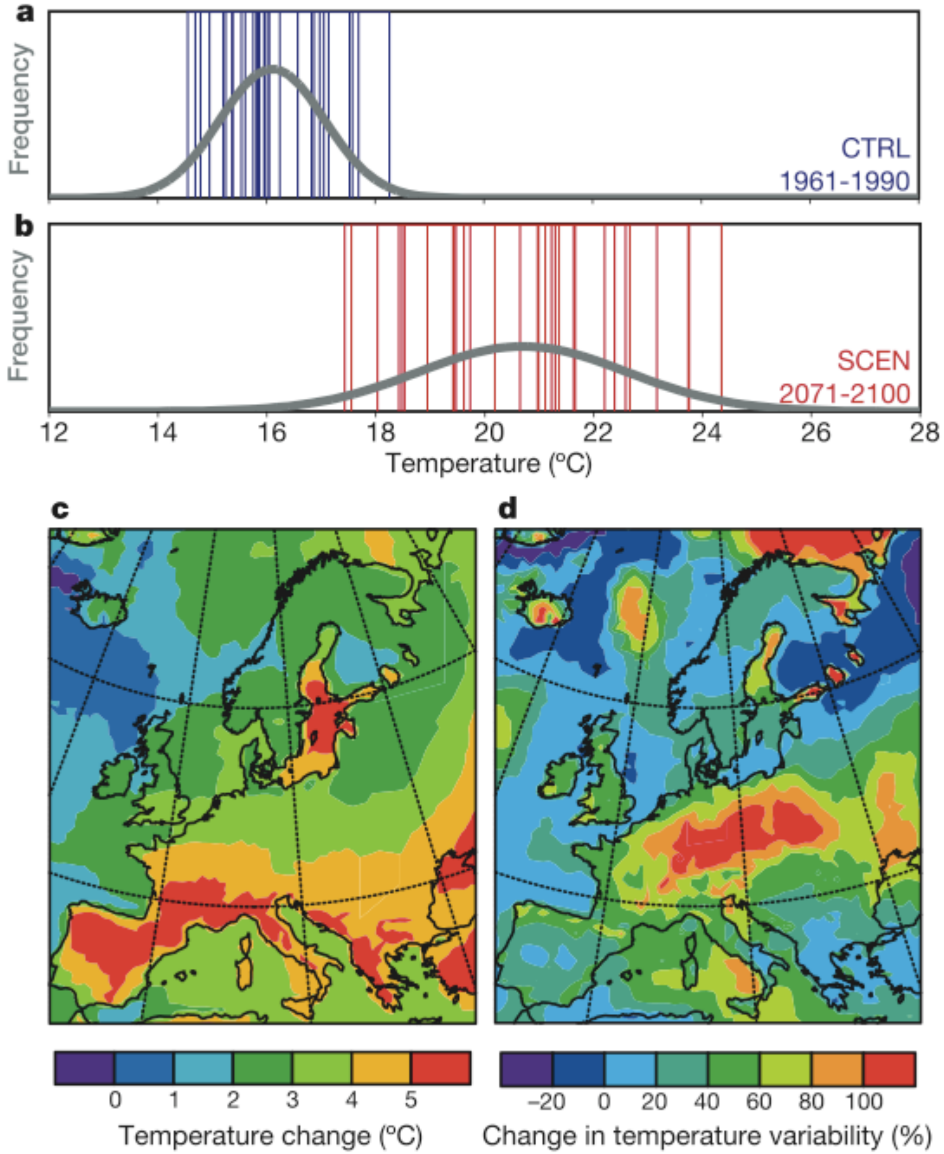

- Shift in mean <u>and</u> increase in interannual variability;
- This is crucial information for European decision makers;
- 2003 looks less special within a scenario distribution;
- Location of change in mean and location of change in variability are not the same
- How robust are these model results ?
- What are the active mechanisms ?

Schär et al., Nature 2004

2003 in the n512 simulations:
JJA Temperature anomaly in UPSCALE HadGEM3-N512 ensemble

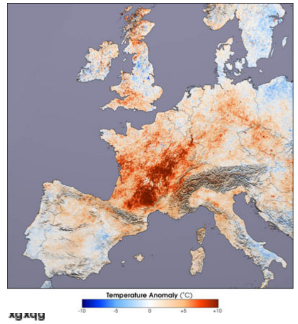

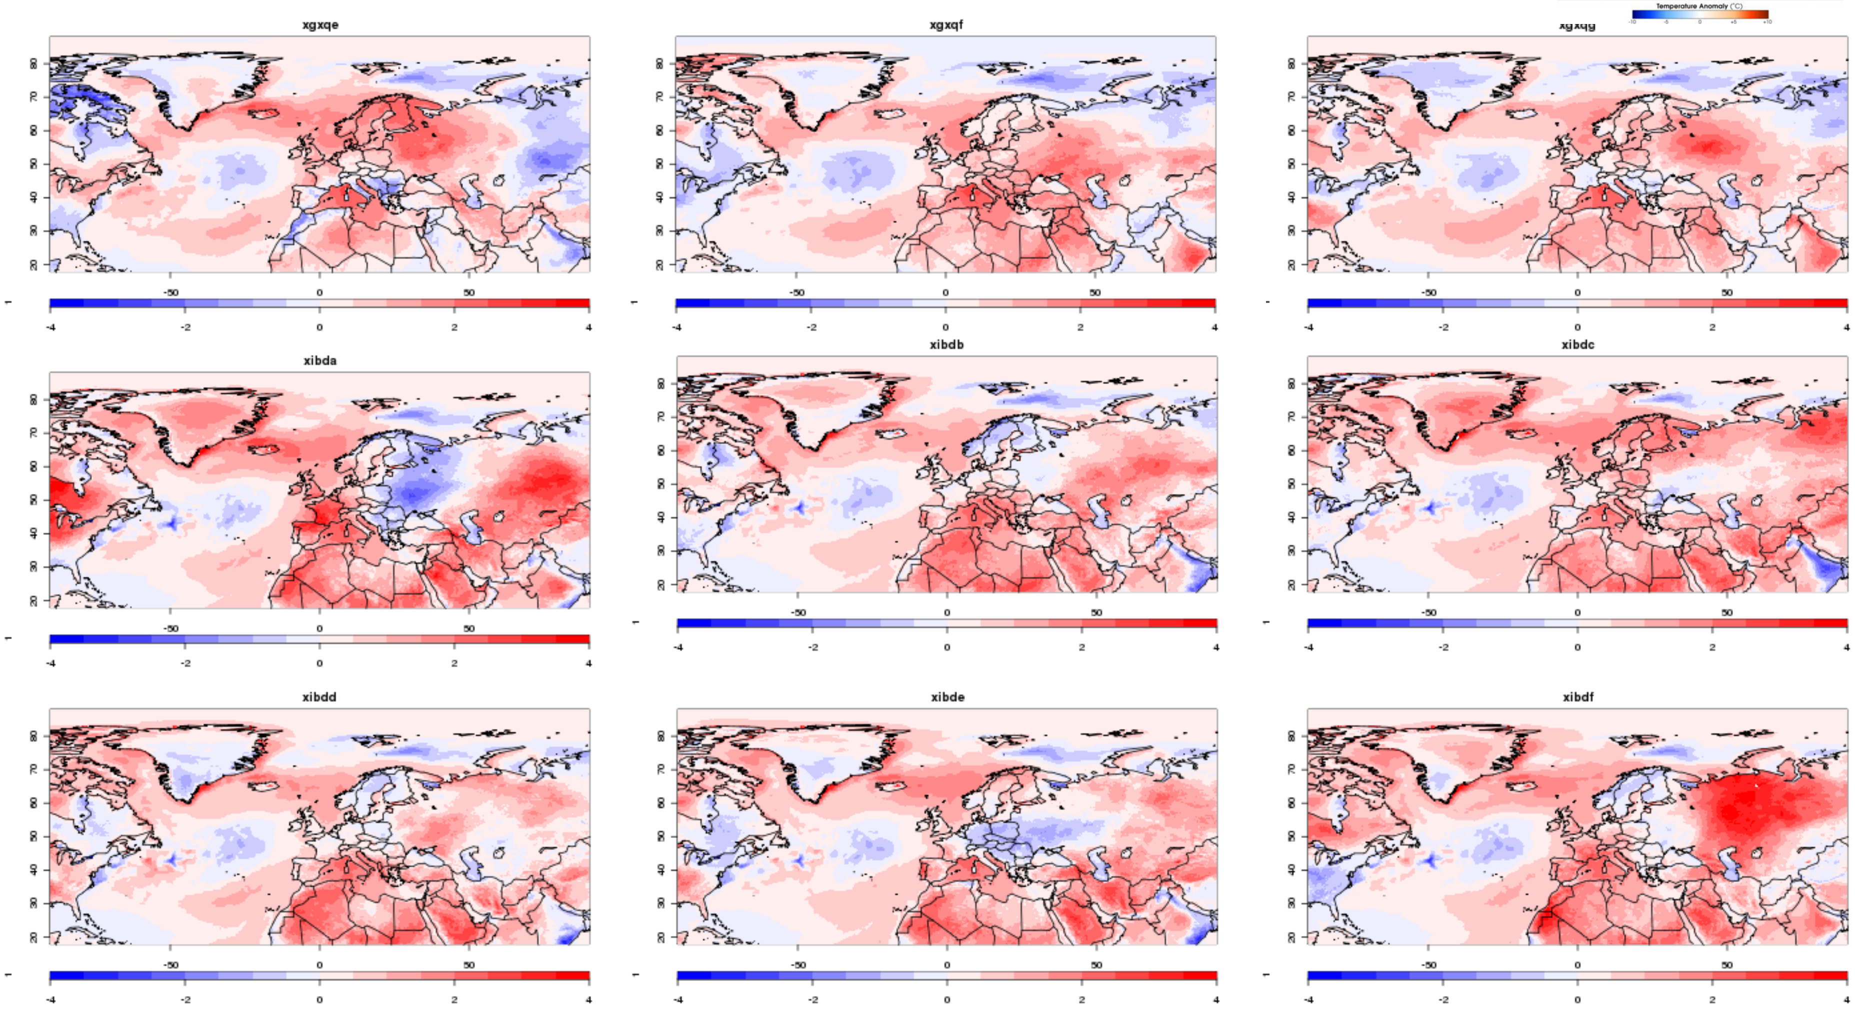

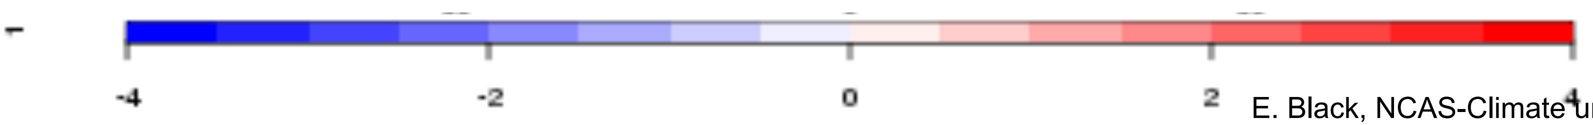

E. Black, NCAS-Climate unpublished

First place in the beauty contest…

xibda

NCEP

re-analysis,

“observations”

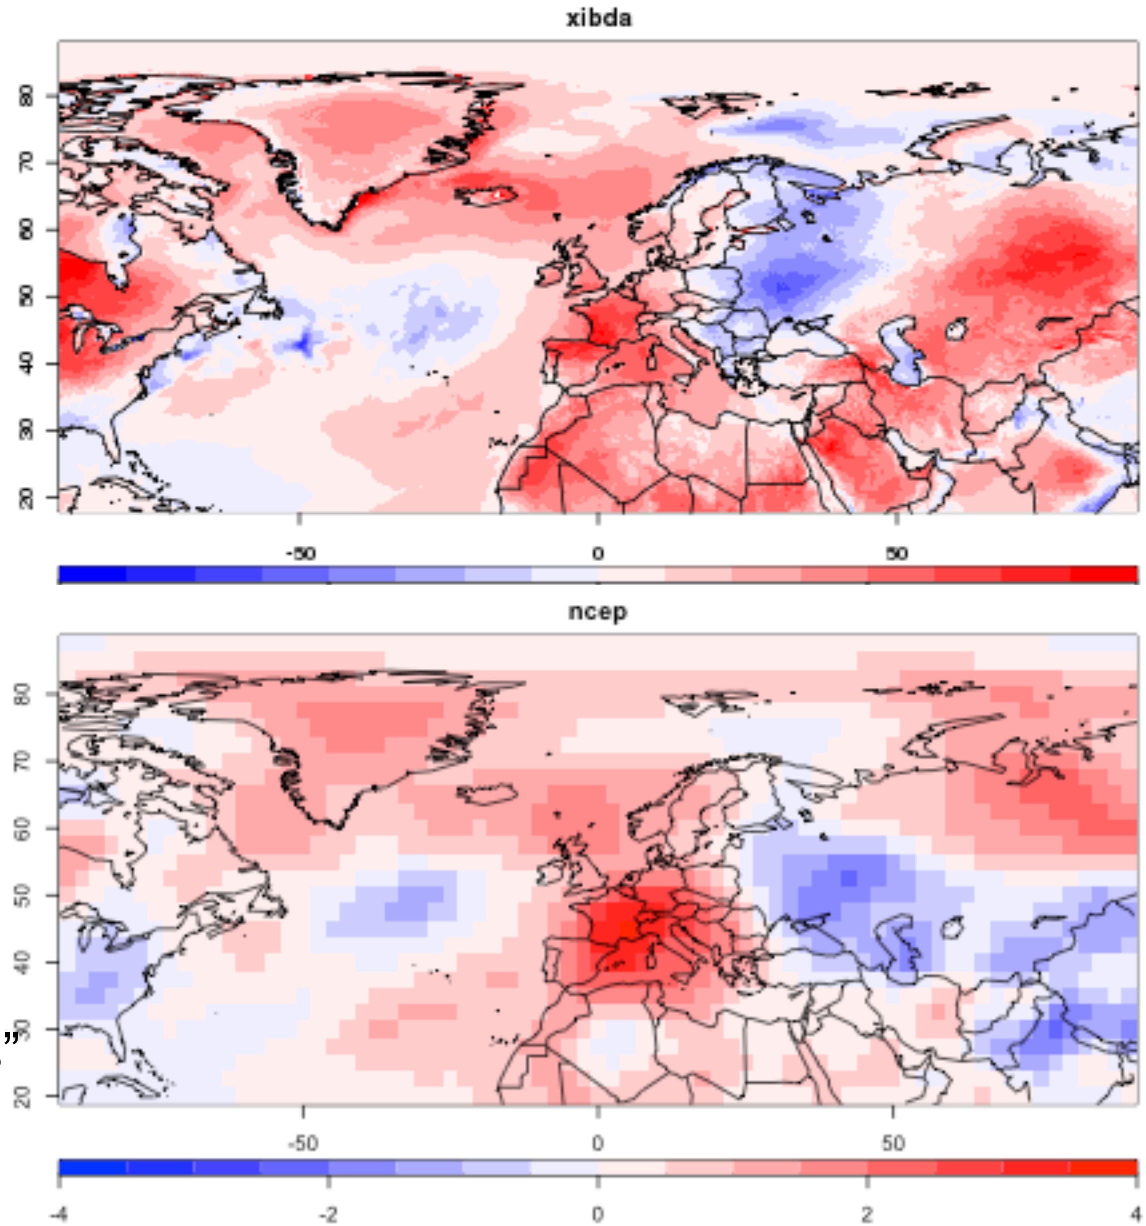

E. Black, NCAS-Climate unpublished

Hydrological pre-conditioning of European Heatwaves ?

Sequence of events in two ensemble members with identical, observed Sea Surface Temperatures (OSTIA)

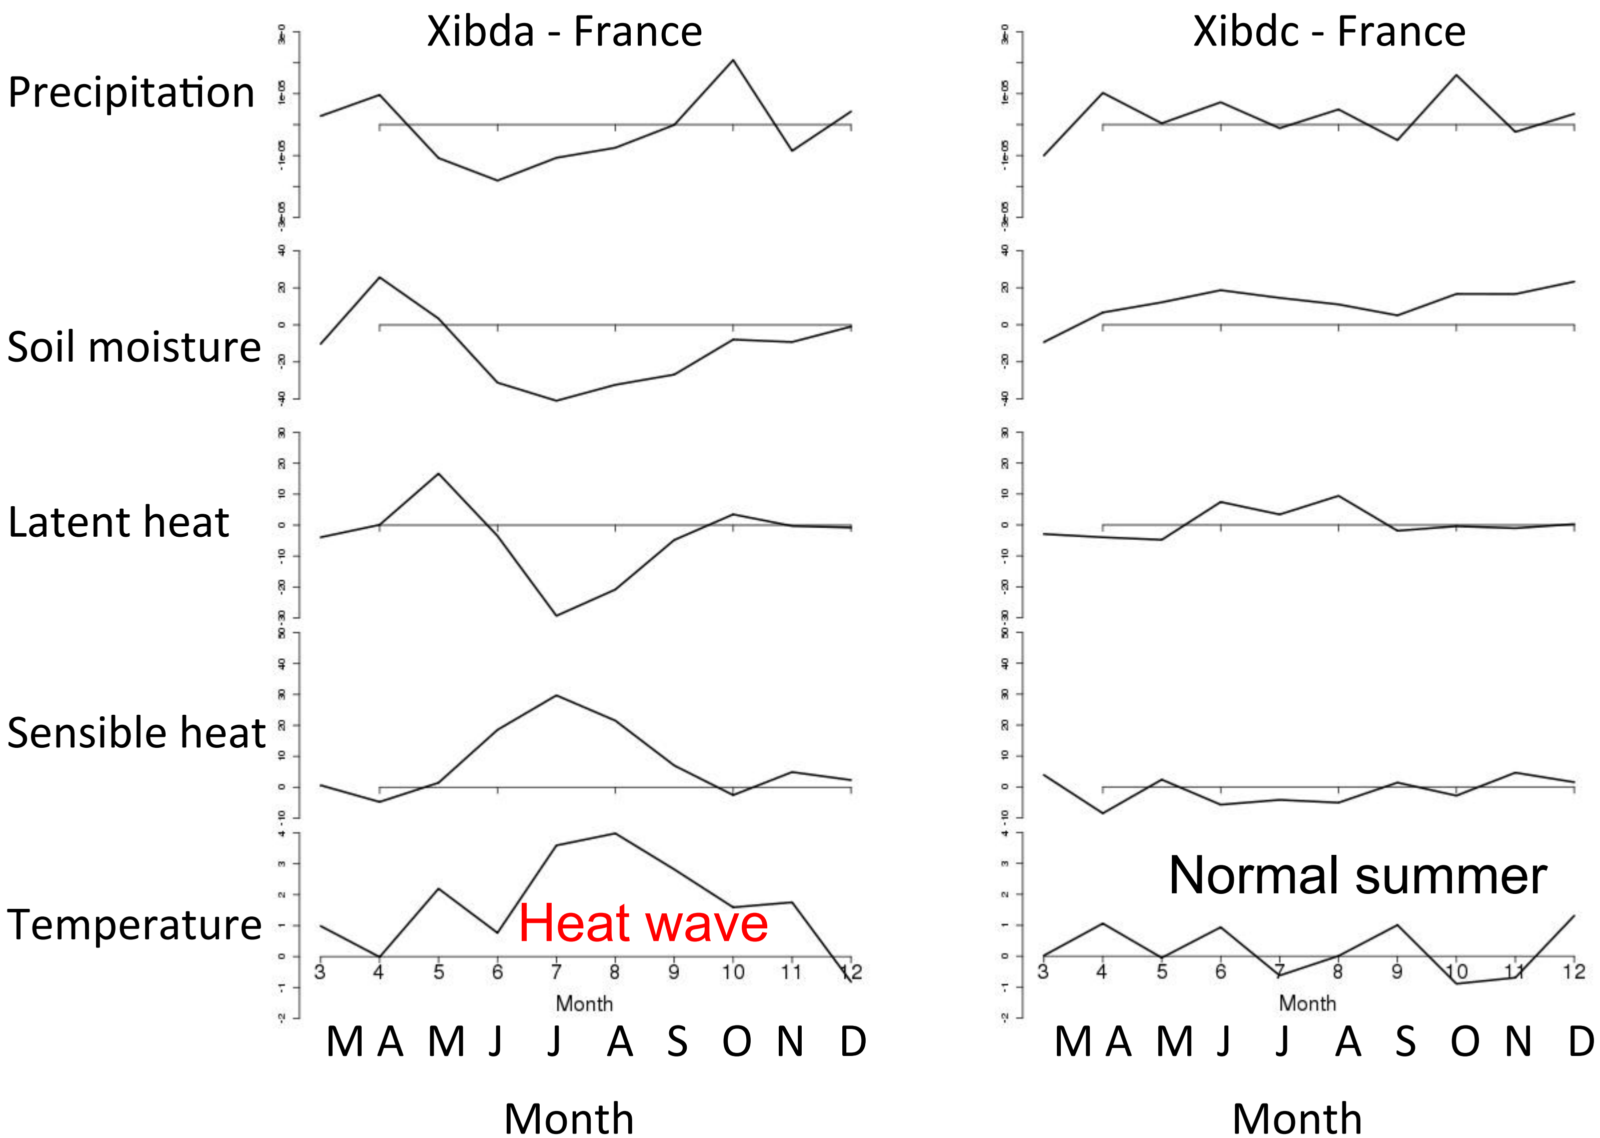

E. Black, NCAS-Climate unpublished

Science Highlight 2

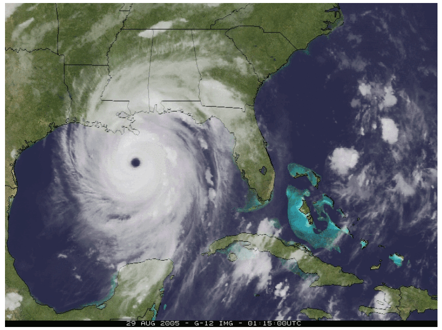

Hurricanes and Typhoons:

1. how many more years like 2005 (the year of Katrina) or 2011 can we expect?
2. will there be more/less storms in the future? more intense / less intense?
3. how many more storms like Sandy, the storm that started in the tropics and then hit New York City in late 2012?

Knowledge Transfer Partnerships

Willis Research Network

UPSCALE: emerging processes and high-fidelity GCMs

UPSCALE aims to increase the **fidelity** of global climate simulations and our **understanding** of weather and climate risk, by representing fundamental weather and climate processes more completely.

This will enhance our **confidence** in projections of climate change, including extremes such as cyclones, heat waves, floods:

- Extreme events impact society and yet most are **absent** in IPCC-class climate models (see an example in the figure)
- These are rare events and require a large sample to be studied robustly
- UPSCALE uses a fine grid, similar to that used in global weather forecasting, with a set of simulations for both current and future climates

Strachan et al., J. Clim., 2012 and Vidale et al. 2013, in preparation

Tropical cyclone track density (transits per month)

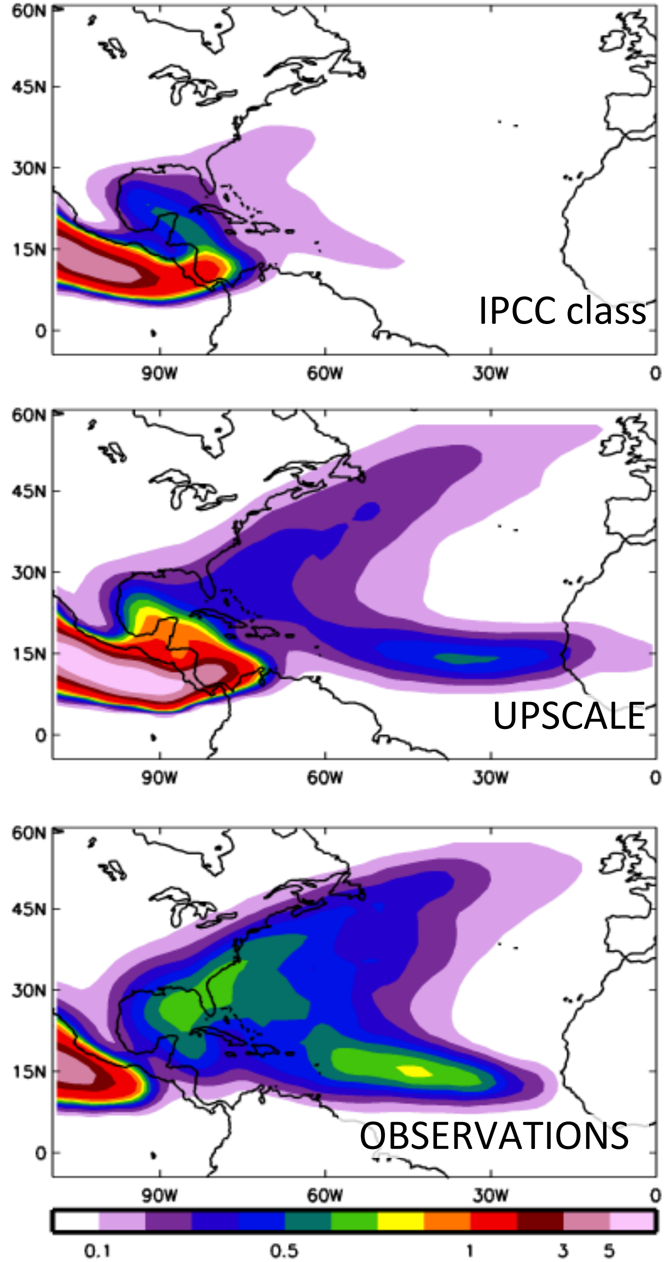

Hurricanes Ivan and Frances simulated by fvGCM (NASA)

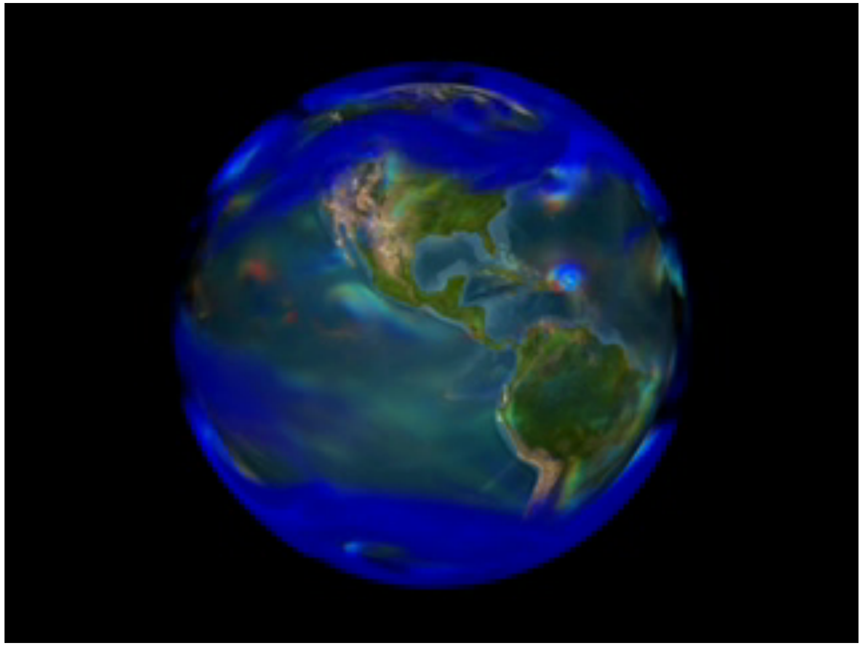

There is more than one tropical cyclone active at the same time: interactions with the atmosphere, the ocean, with each other

But, rather than simulating one hurricane/typhoon at a time, we want to simulate thousands, as well as their interactions with the climate system

First we need to find them…

Tropical storm tracking

- Tropical storms are located and tracked in GloSea5 using **TRACK** (Hodges, 1996).
- **Tracking**: maxima in 850 hPa relative vorticity on common T42 grid
- Minimum 2 day lifetime
- Includes check for a **warm core**.
- **Exactly** the same algorithm used in all of the following (no tuning)
- Obs – HURDAT + JTWC – mainly compared model to observed hurricanes

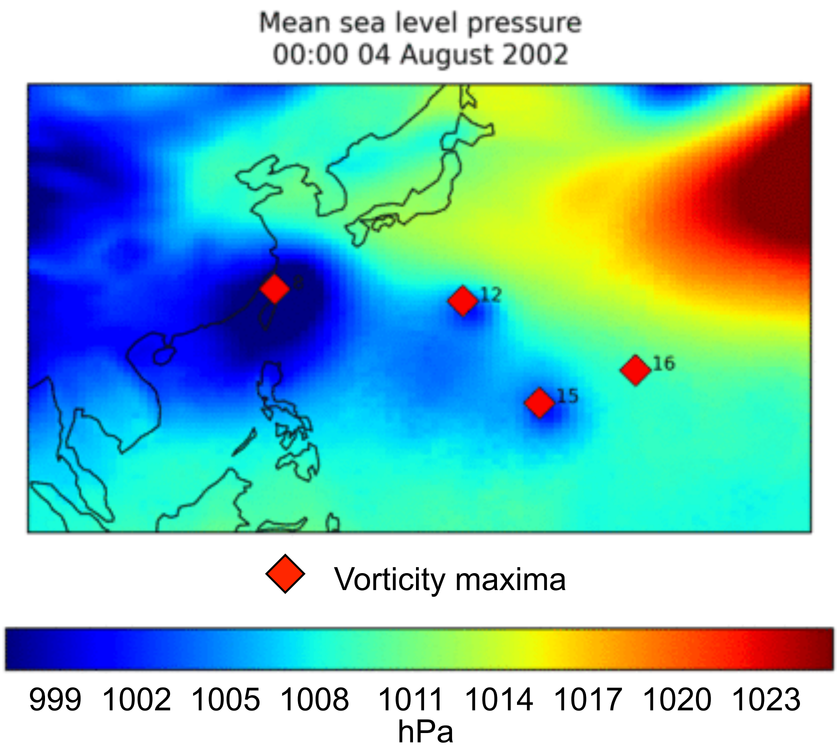

Tracking regions

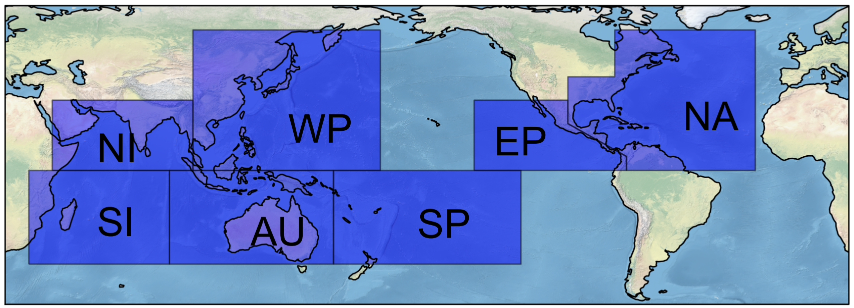

Extracting simulated tropical cyclones from model output

TRACK: a "feature tracking" methodology (Hodges,1995)
Independent of model resolution and basin

1) Output 6-hourly global fields (e.g. pressure, winds, humidity, precipitation)

2) Locate and track all centres of high relative vorticity ⇒ **35000 / year**

3) Apply 2-day filter ⇒ **8000 storms / year**

4) Analyse vertical structure of storm for evidence of warm-core (tropical storm structure) ⇒ **120 storms / year**

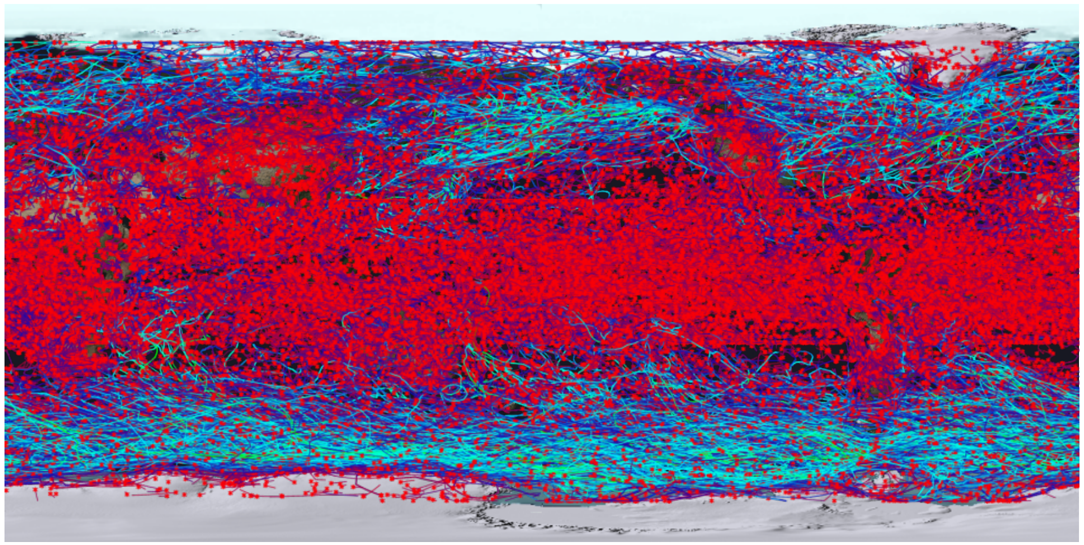

1 year of GCM simulated tropical storms

Consider the **entire life cycle** from genesis, through extratropical transition, to lysis.

Strachan et al. J. Clim. 2012

Storm Track density from model ensembles and observations

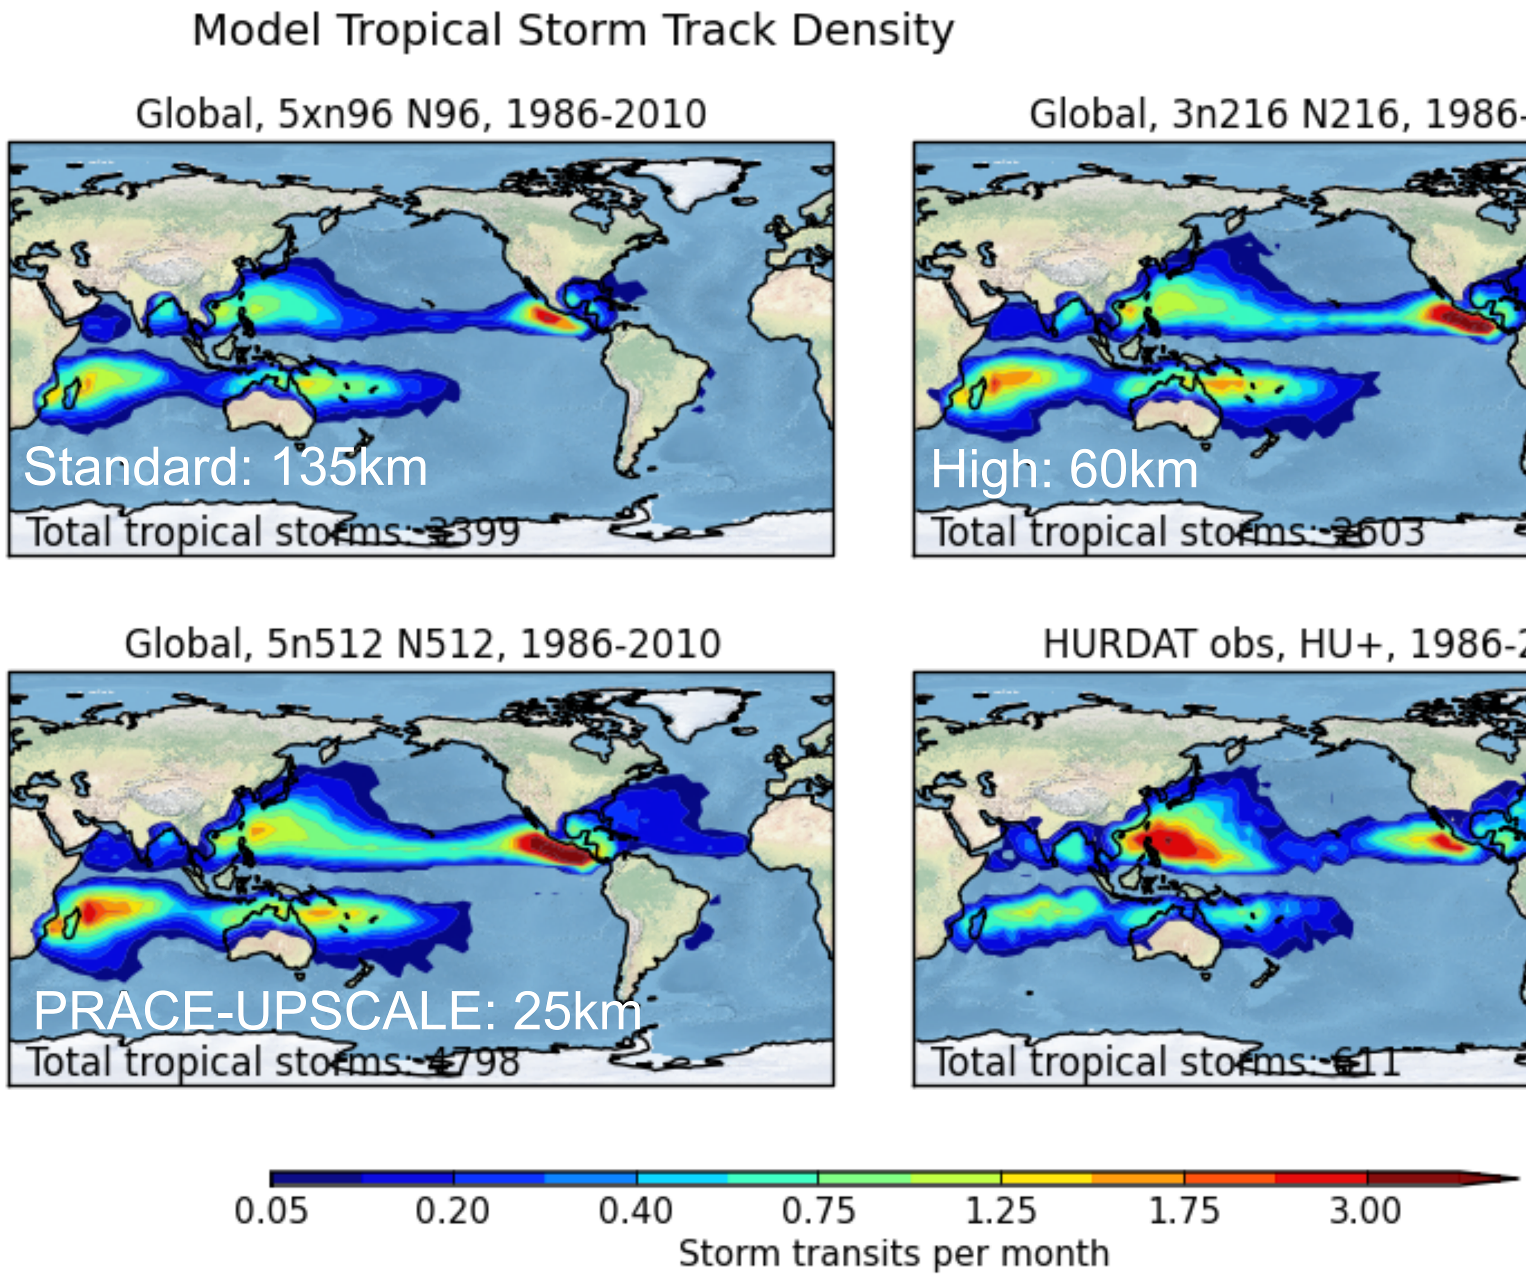

M. Roberts et al., MO, unpublished

Mean NH basin storm counts – 1986-2010

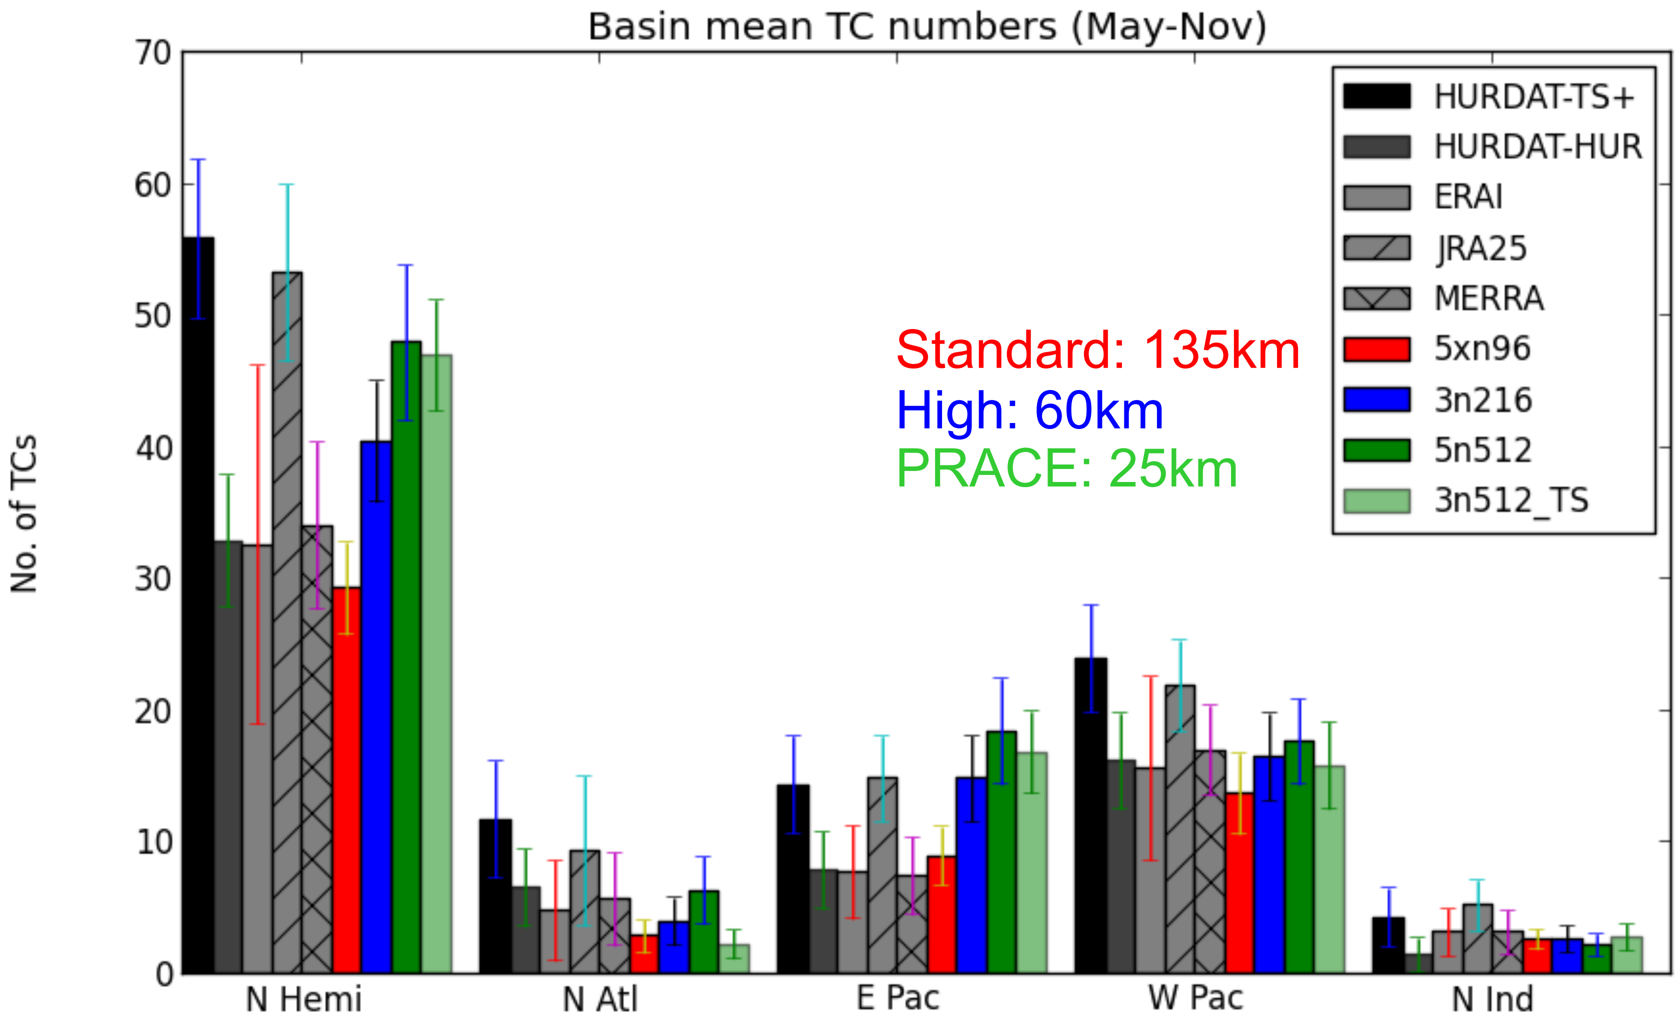

Standard deviation indicated by vertical error bar

M. Roberts et al., MO, unpublished

Solid line = ensemble mean, shading = ensemble range

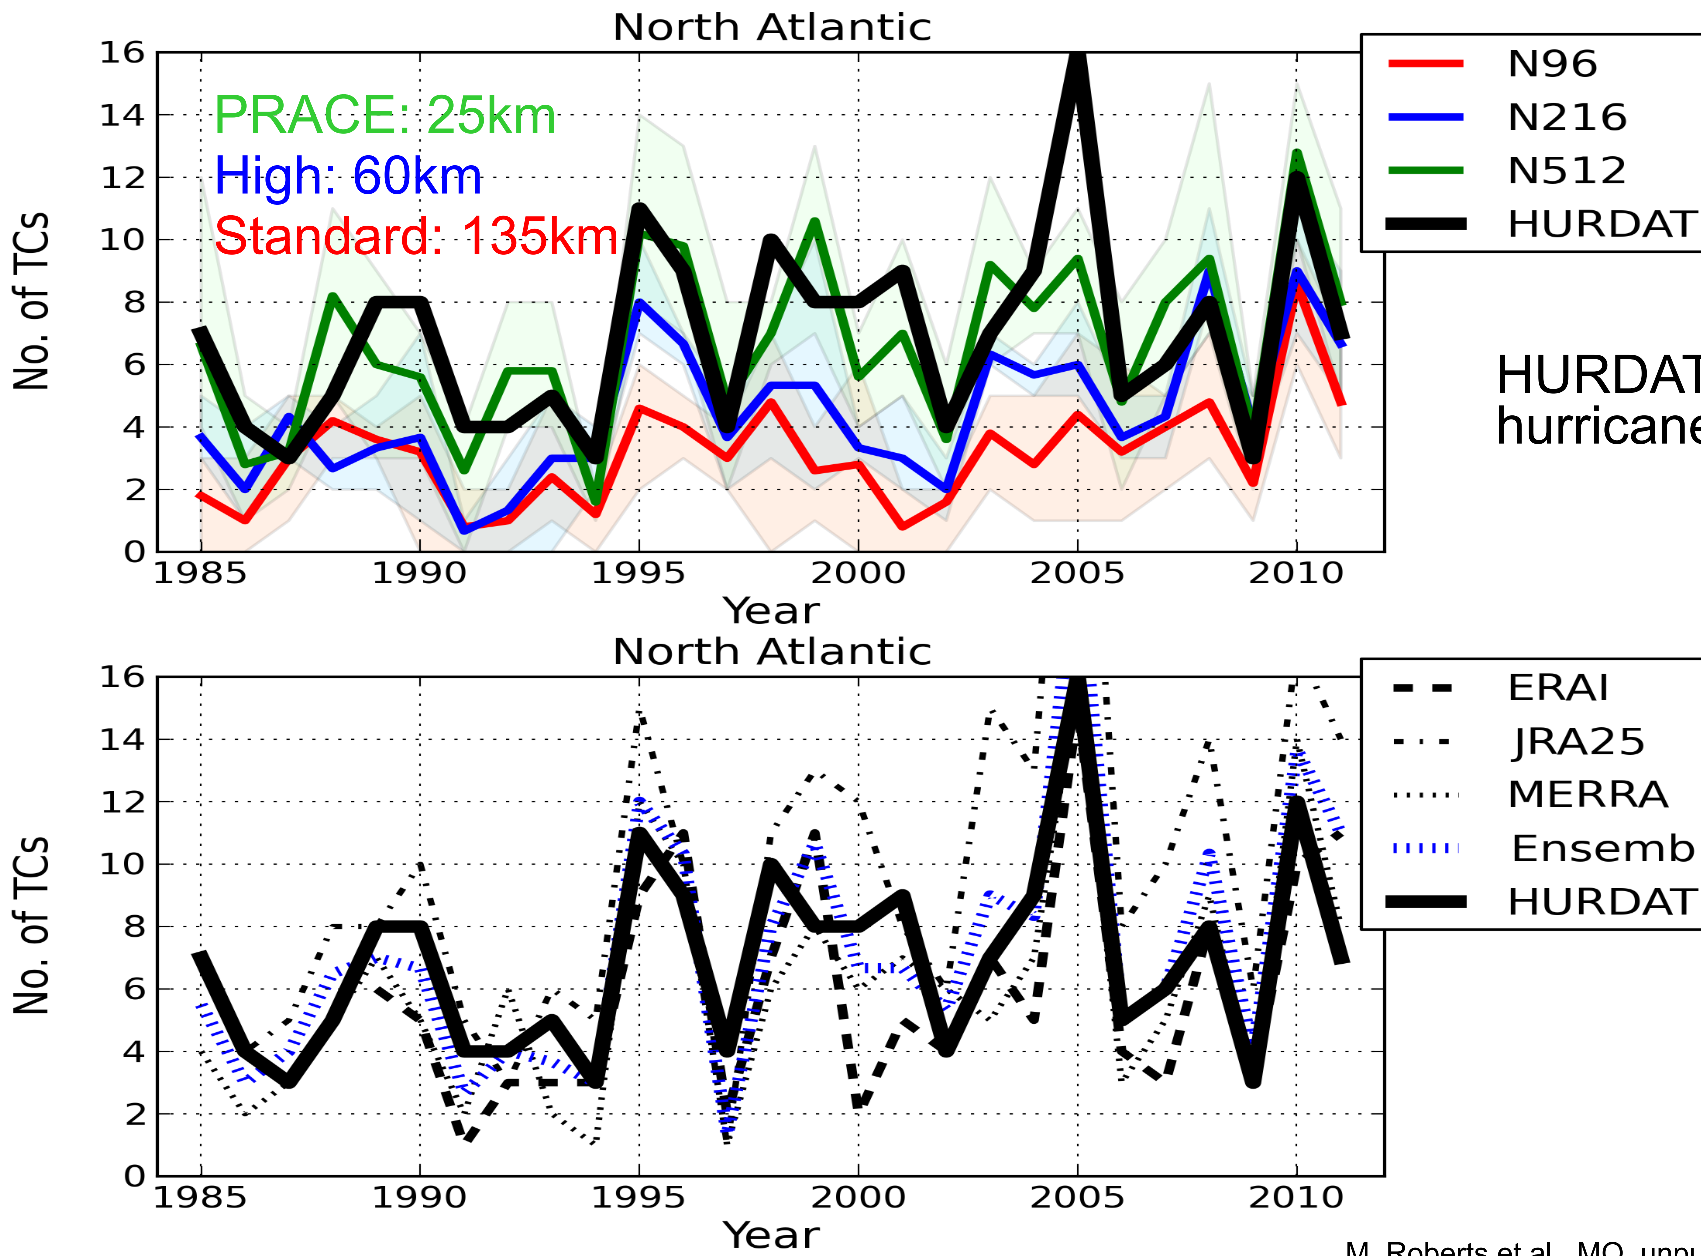

HURDAT hurricanes

M. Roberts et al., MO, unpublished

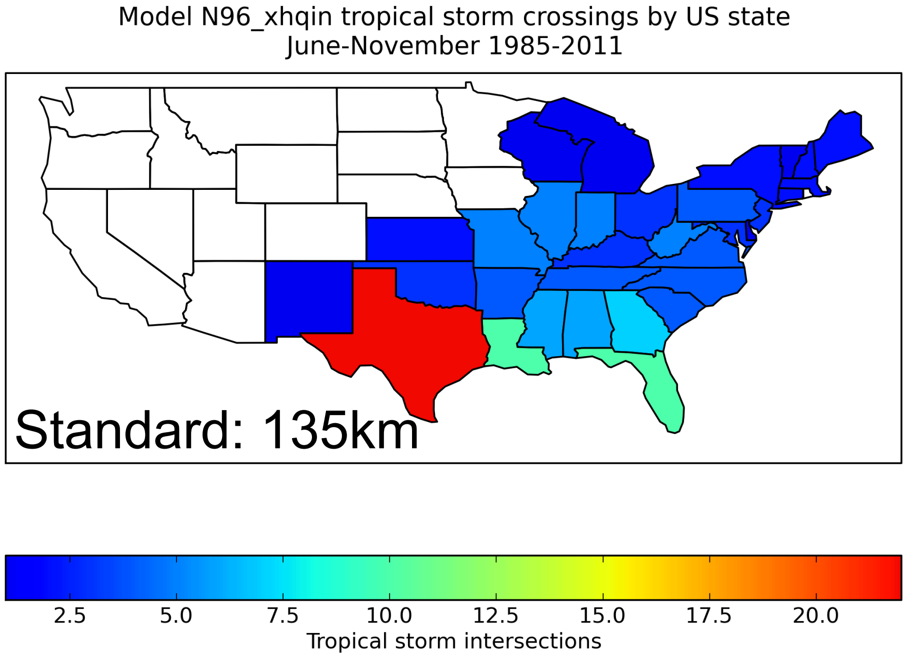

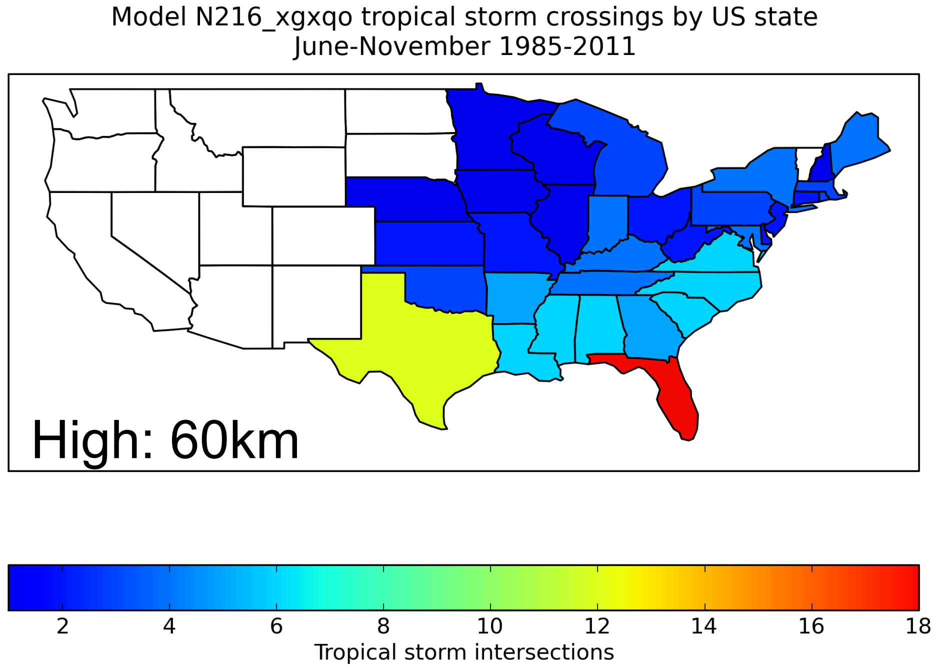

Hurricane landfall climatologies

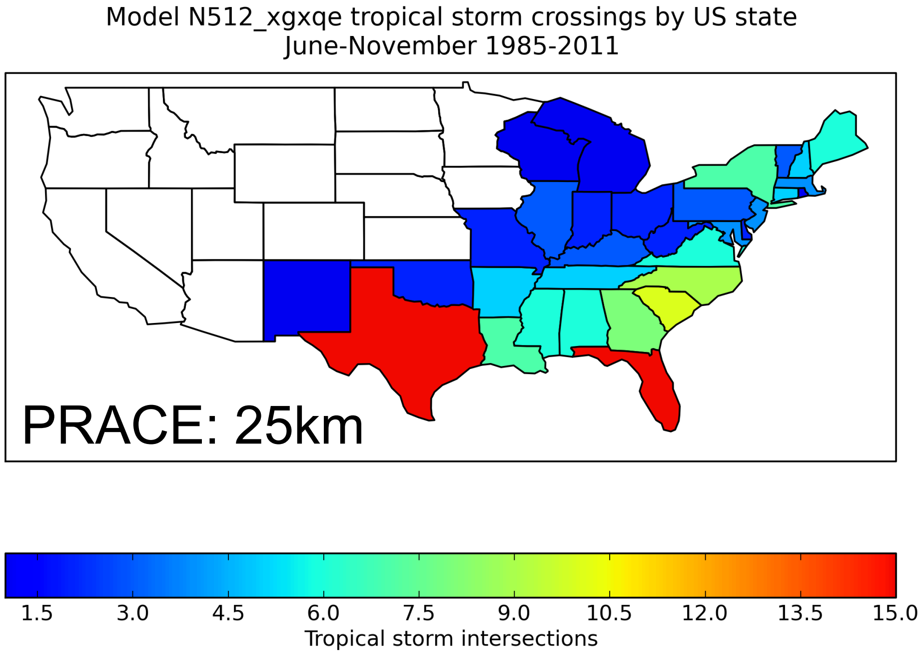

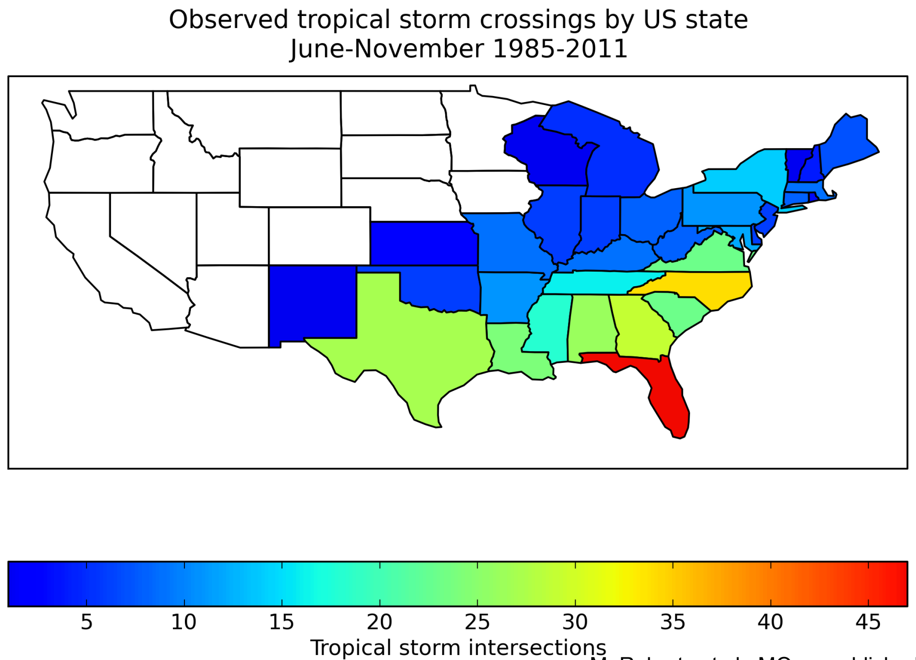

M. Roberts et al., MO, unpublished

Storm Track density difference: Niño – Niña potential for seasonal forecasting?

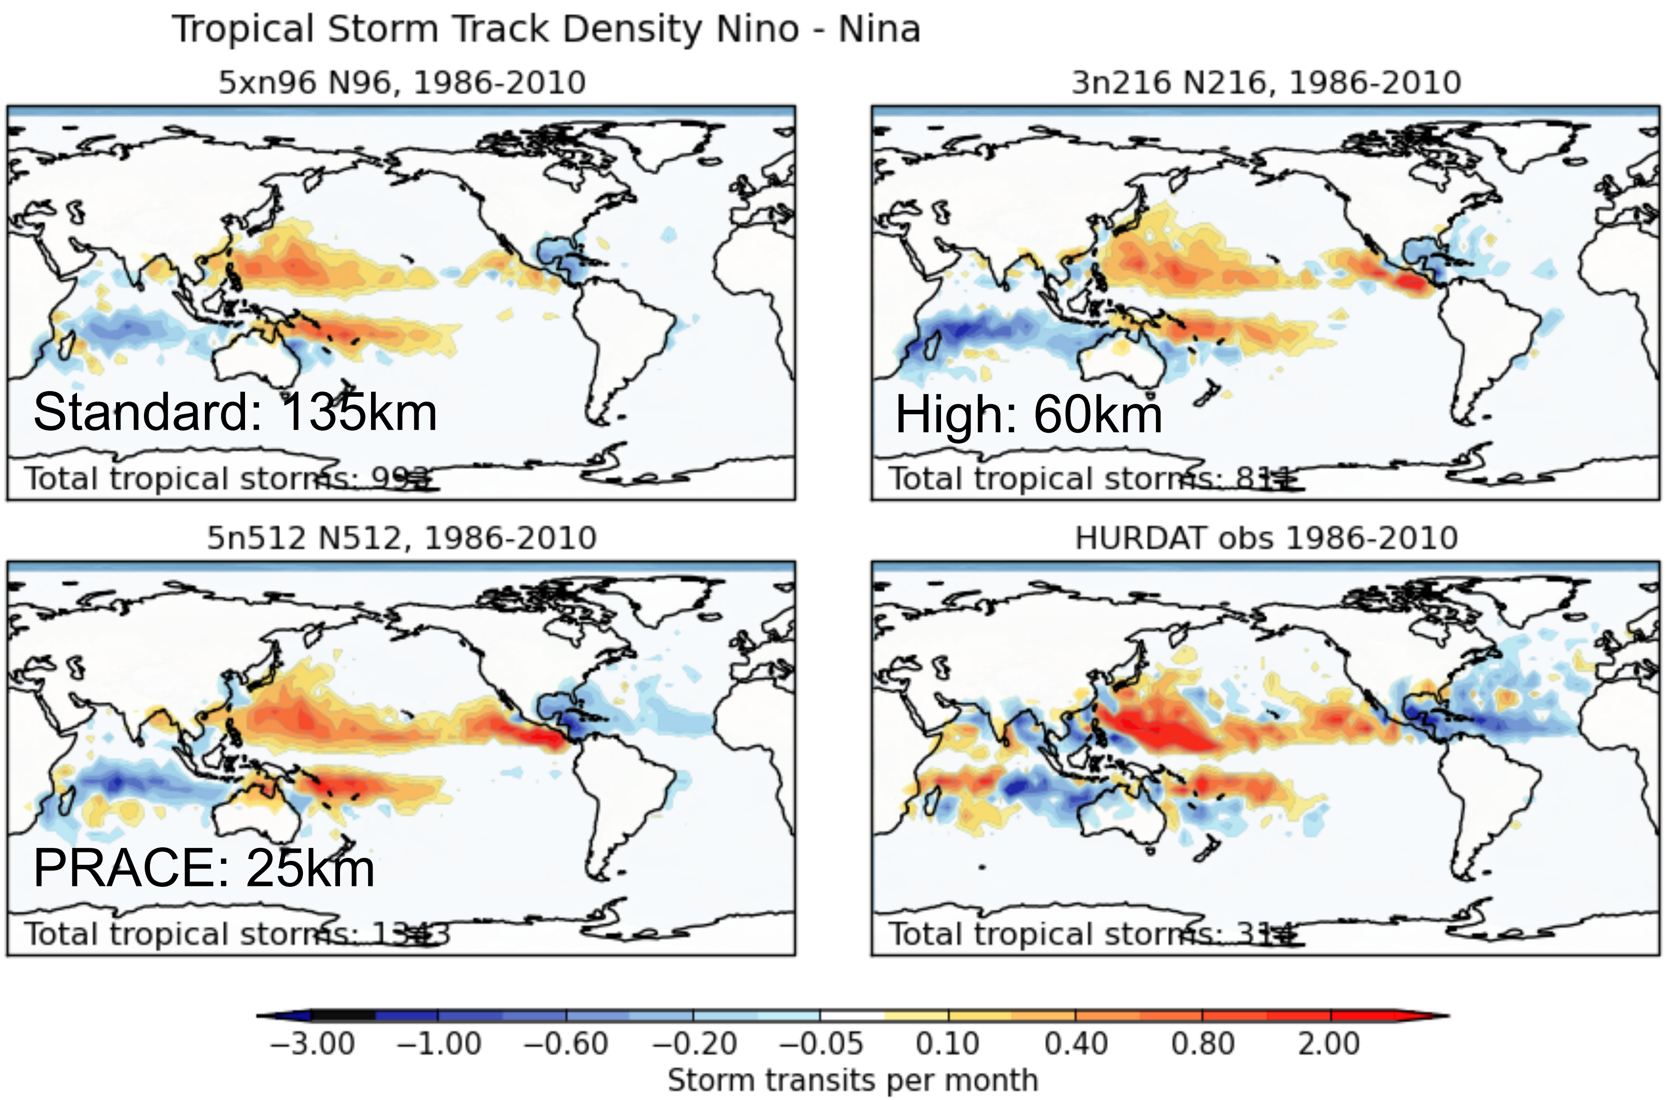

M. Roberts et al., MO, unpublished

Storm Track density: Warmer climate – present day

Joint Weather and Climate Research Programme

A partnership in climate research

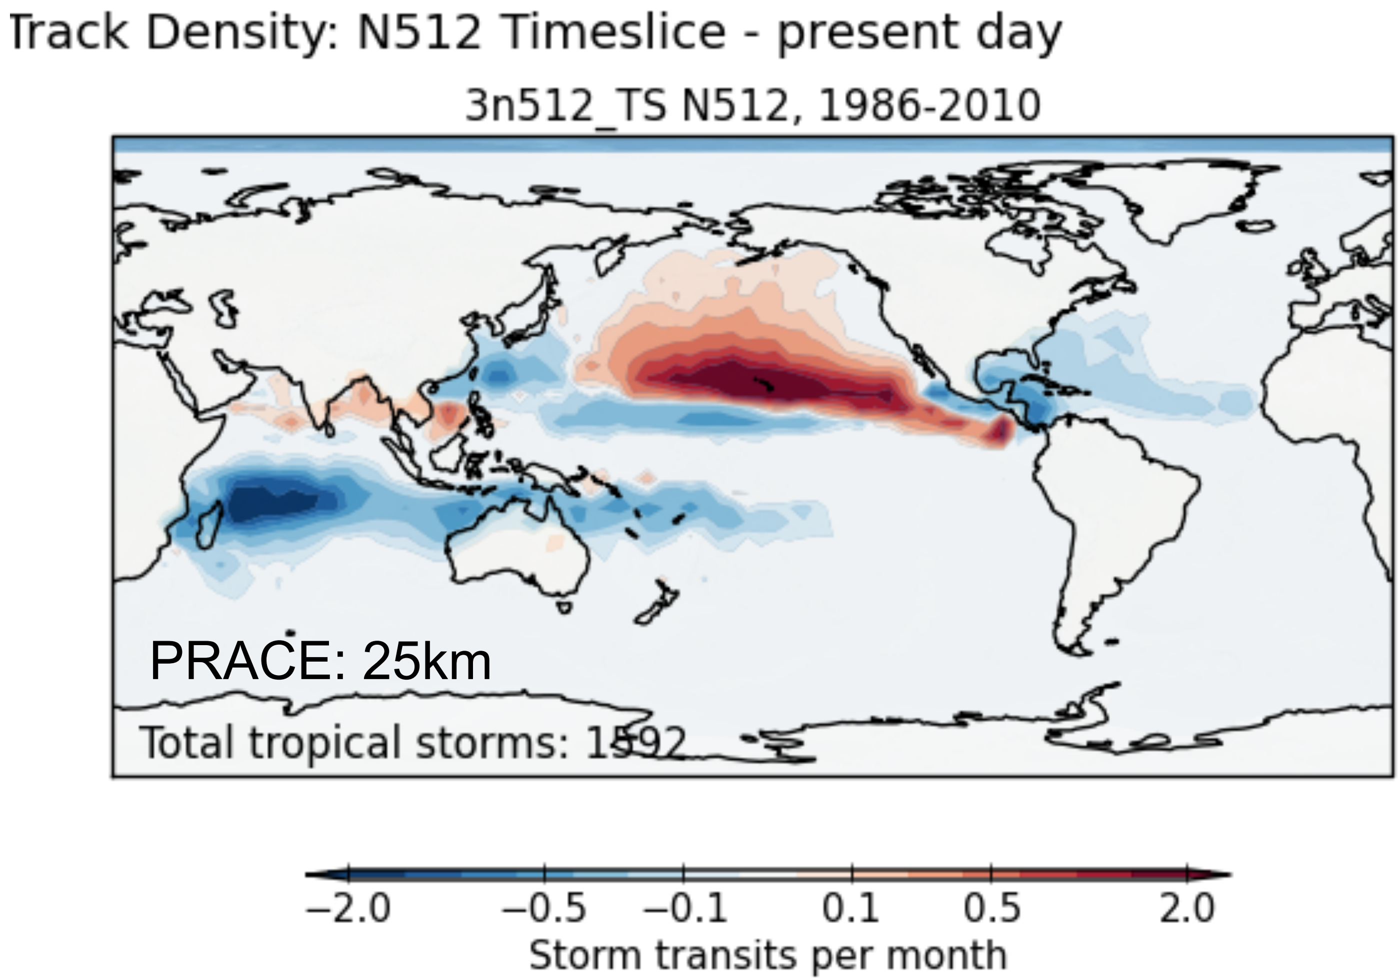

M. Roberts et al., MO, unpublished

The exceptional year of 2005 as seen by a 25km GCM.

Concurrent Weather and Climate risks:

Atlantic Hurricanes <u>and</u> Amazon drought

(ensemble means)

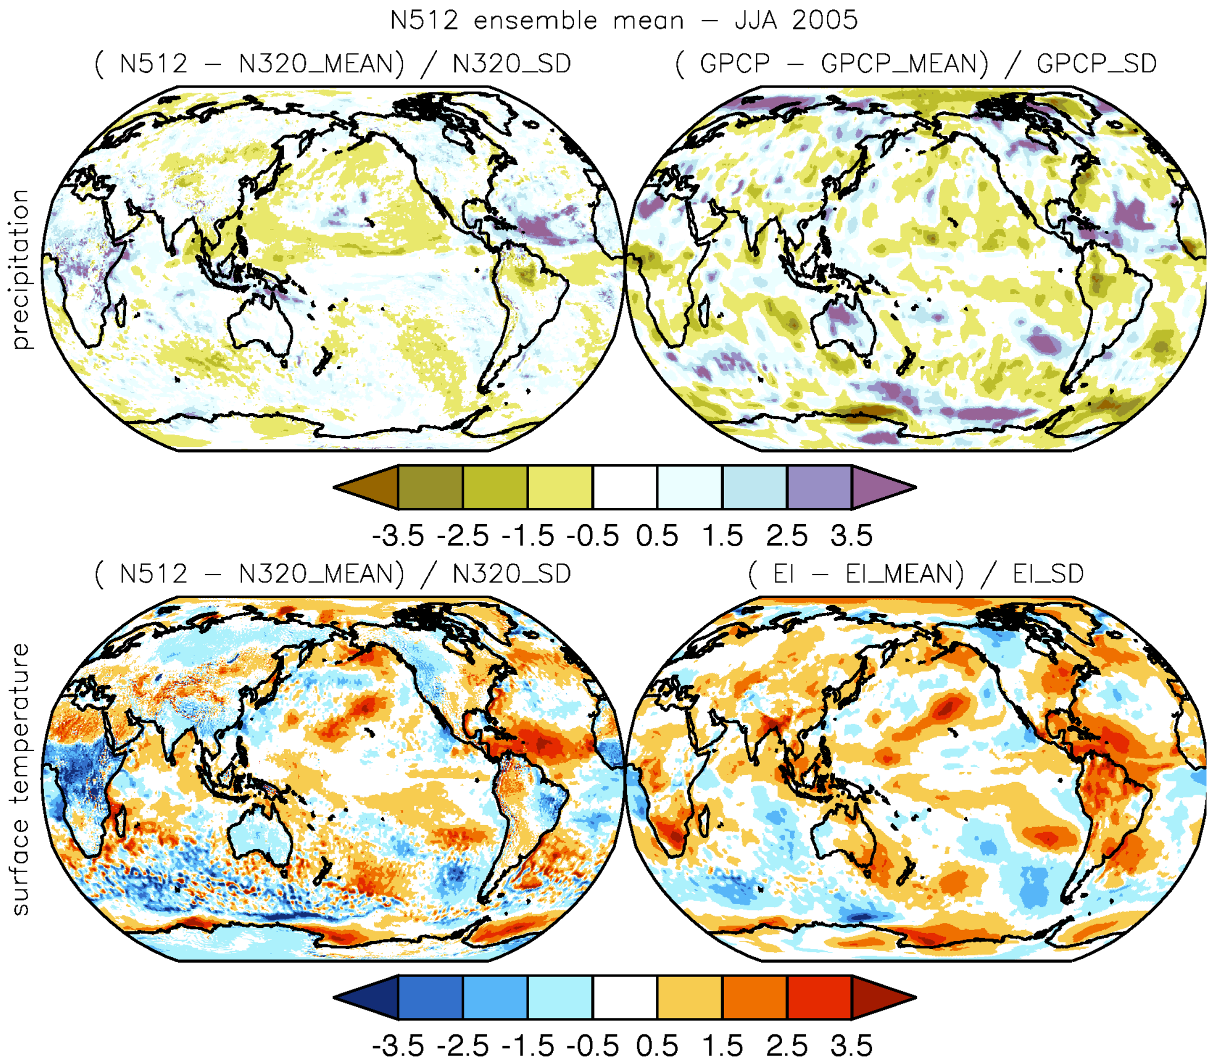

P.L. Vidale et al., unpublished

Enabling the development of next-generation forecasting systems.
N1024: a 12km GCM

- First time that a Global Climate Model leads its Numerical Weather Prediction (NWP) "parent" in resolution (current MO NWP still at 25km)
  - Obvious strategic value to the MO and this received top-priority support by all divisions
  - Originally proposed in UPSCALE and planned for the PRACE supercomputer, but it <u>required too much memory</u> and could not be ported
- We developed both a **standard HadGEM3-A version**, with parameterised convection and an **experimental version with explicit convection.**

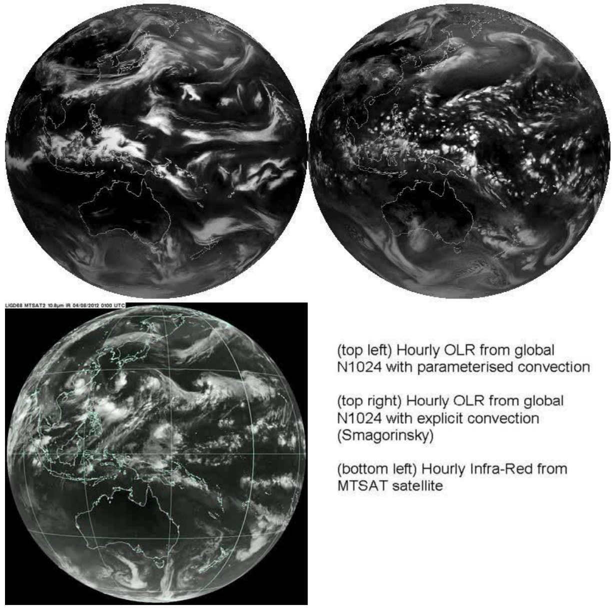

(top left) Hourly OLR from global N1024 with parameterised convection

(top right) Hourly OLR from global N1024 with explicit convection (Smagorinsky)

(bottom left) Hourly Infra-Red from MTSAT satellite

101 caveats of using explicit convection at 12km…

Consider the explicit convection version just as a process study:

- We don't represent convection at 12km (or even at 1km properly)!
- But the convective parameterisation has big issues too
- Probably the lowest resolution for which we can consider switching off the parameterisation, also for stability reasons
- And mid-latitudes may well not be as good as with parameterisation

Local time of peak precipitation for 12km models (diurnal cycle) – Mar-Feb 08/09

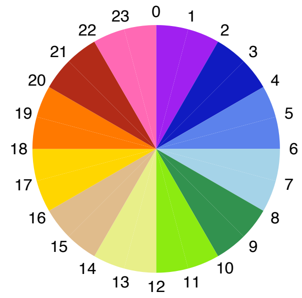

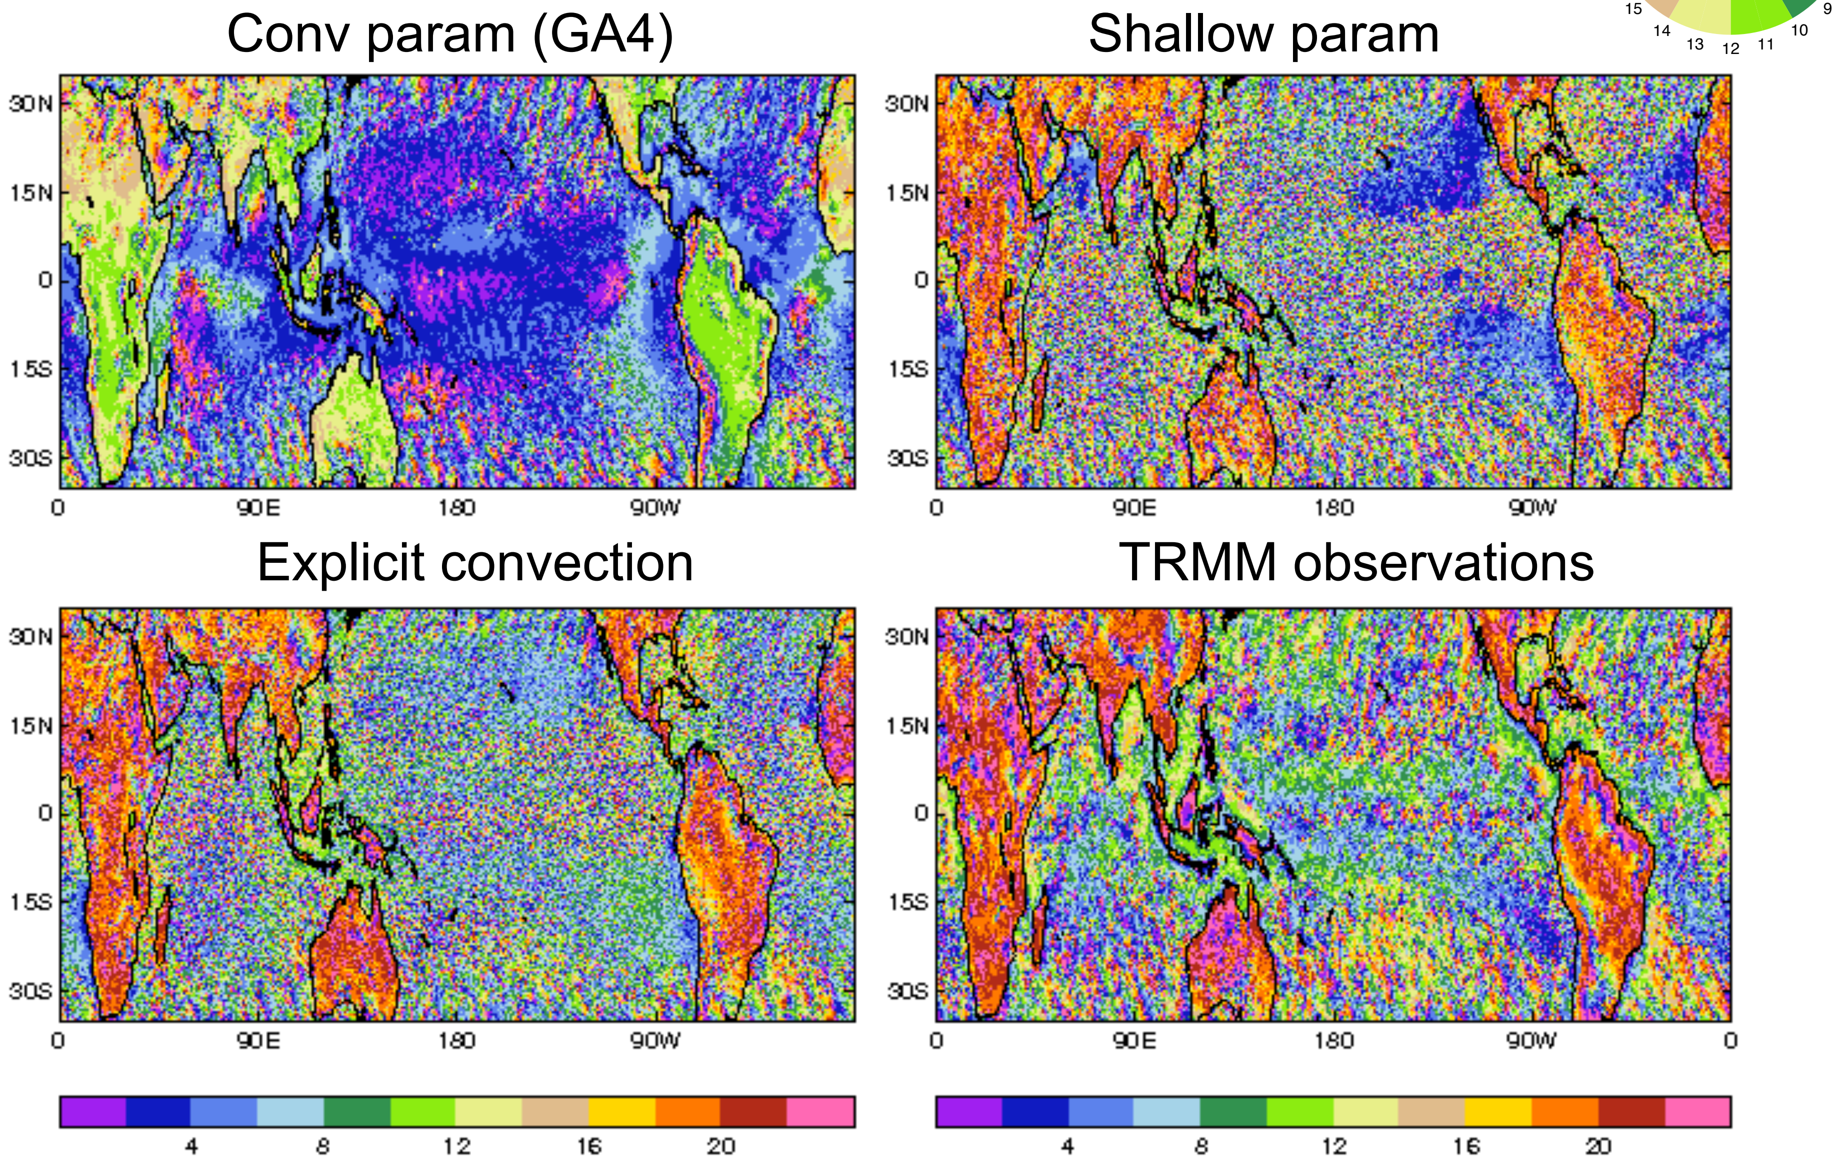

M. Roberts et al., MO, unpublished

Propagating convective systems over Africa and hurricane formation

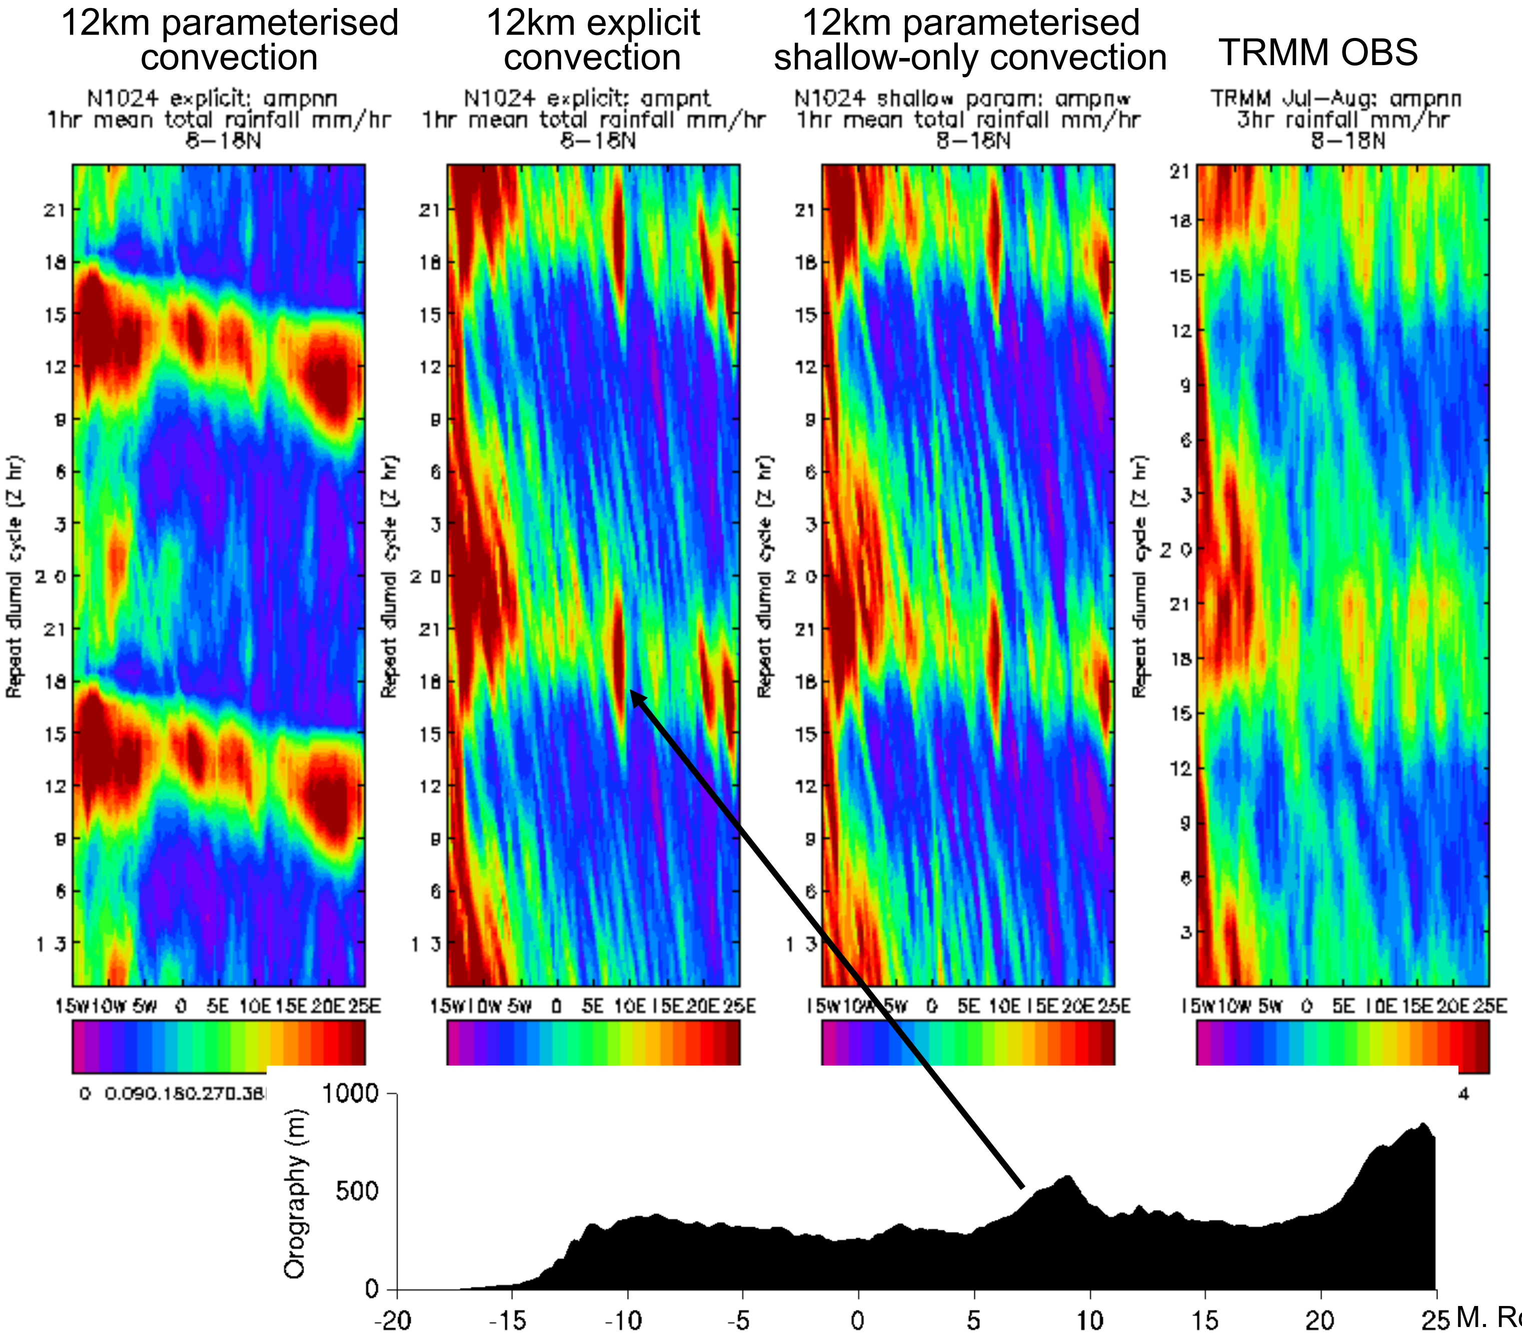

M. Roberts et al., MO, unpublished

West African Monsoon region

– mean diurnal cycle of precipitation

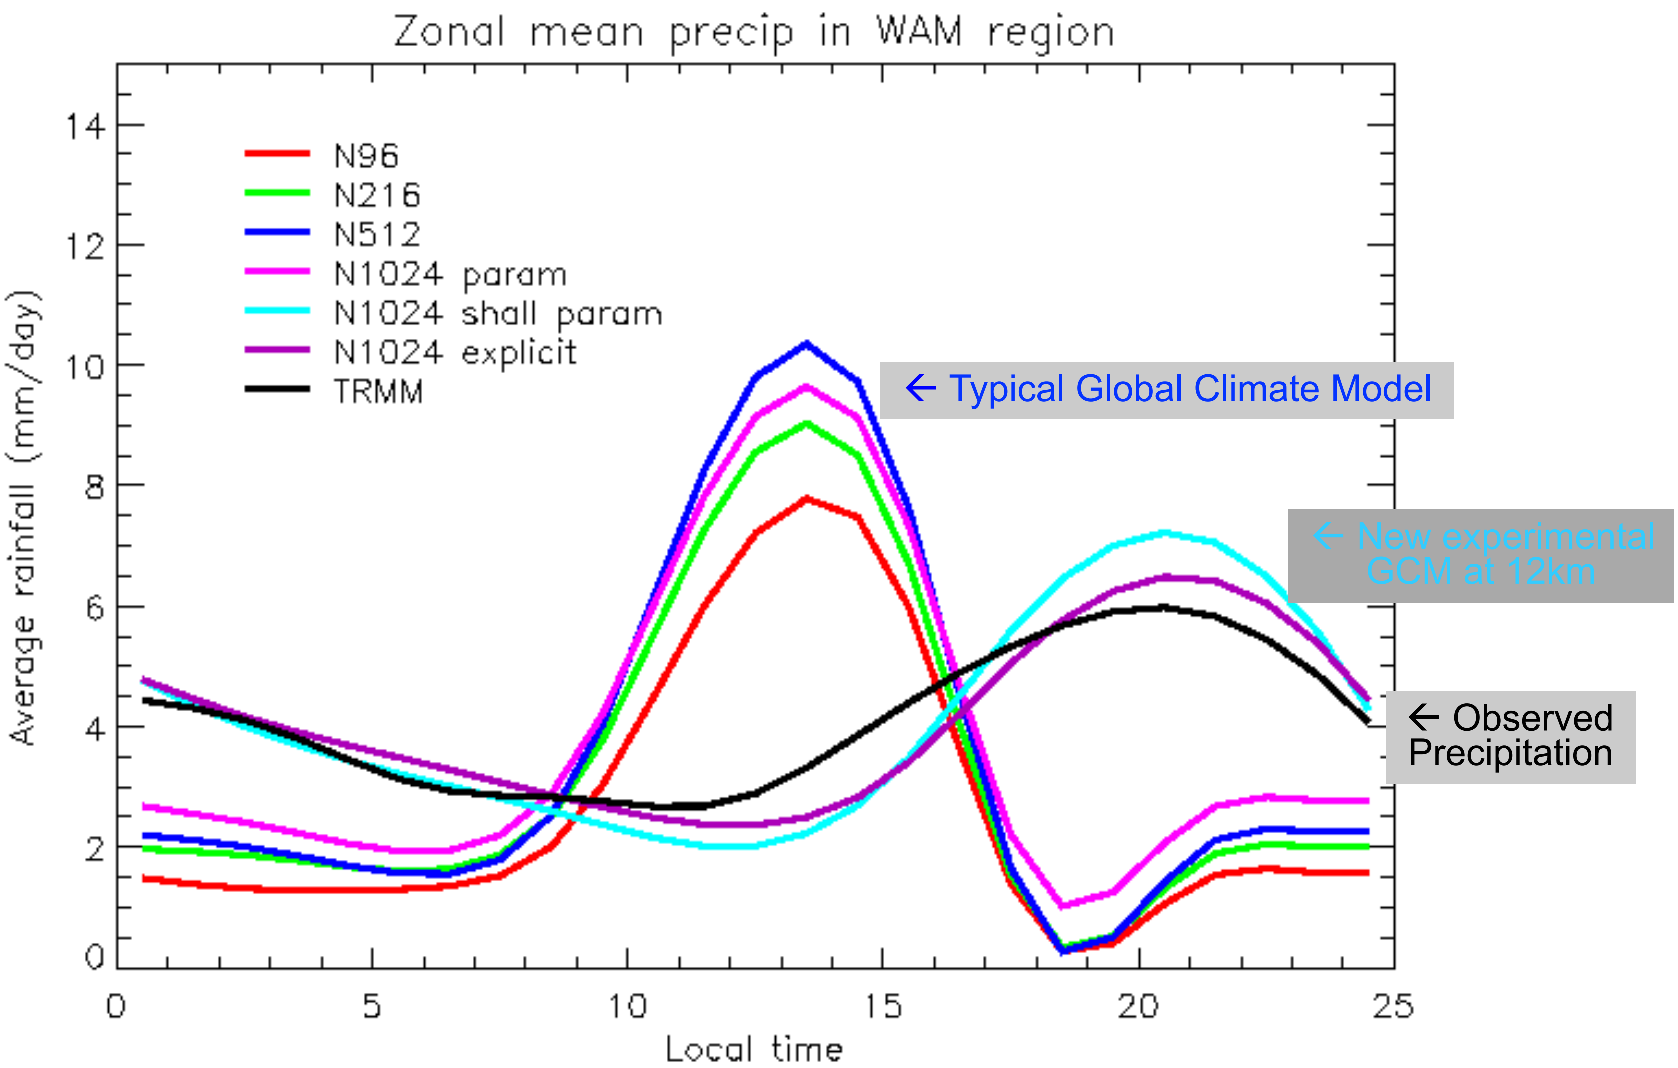

M. Roberts et al., MO, unpublished

Hurricane intensity

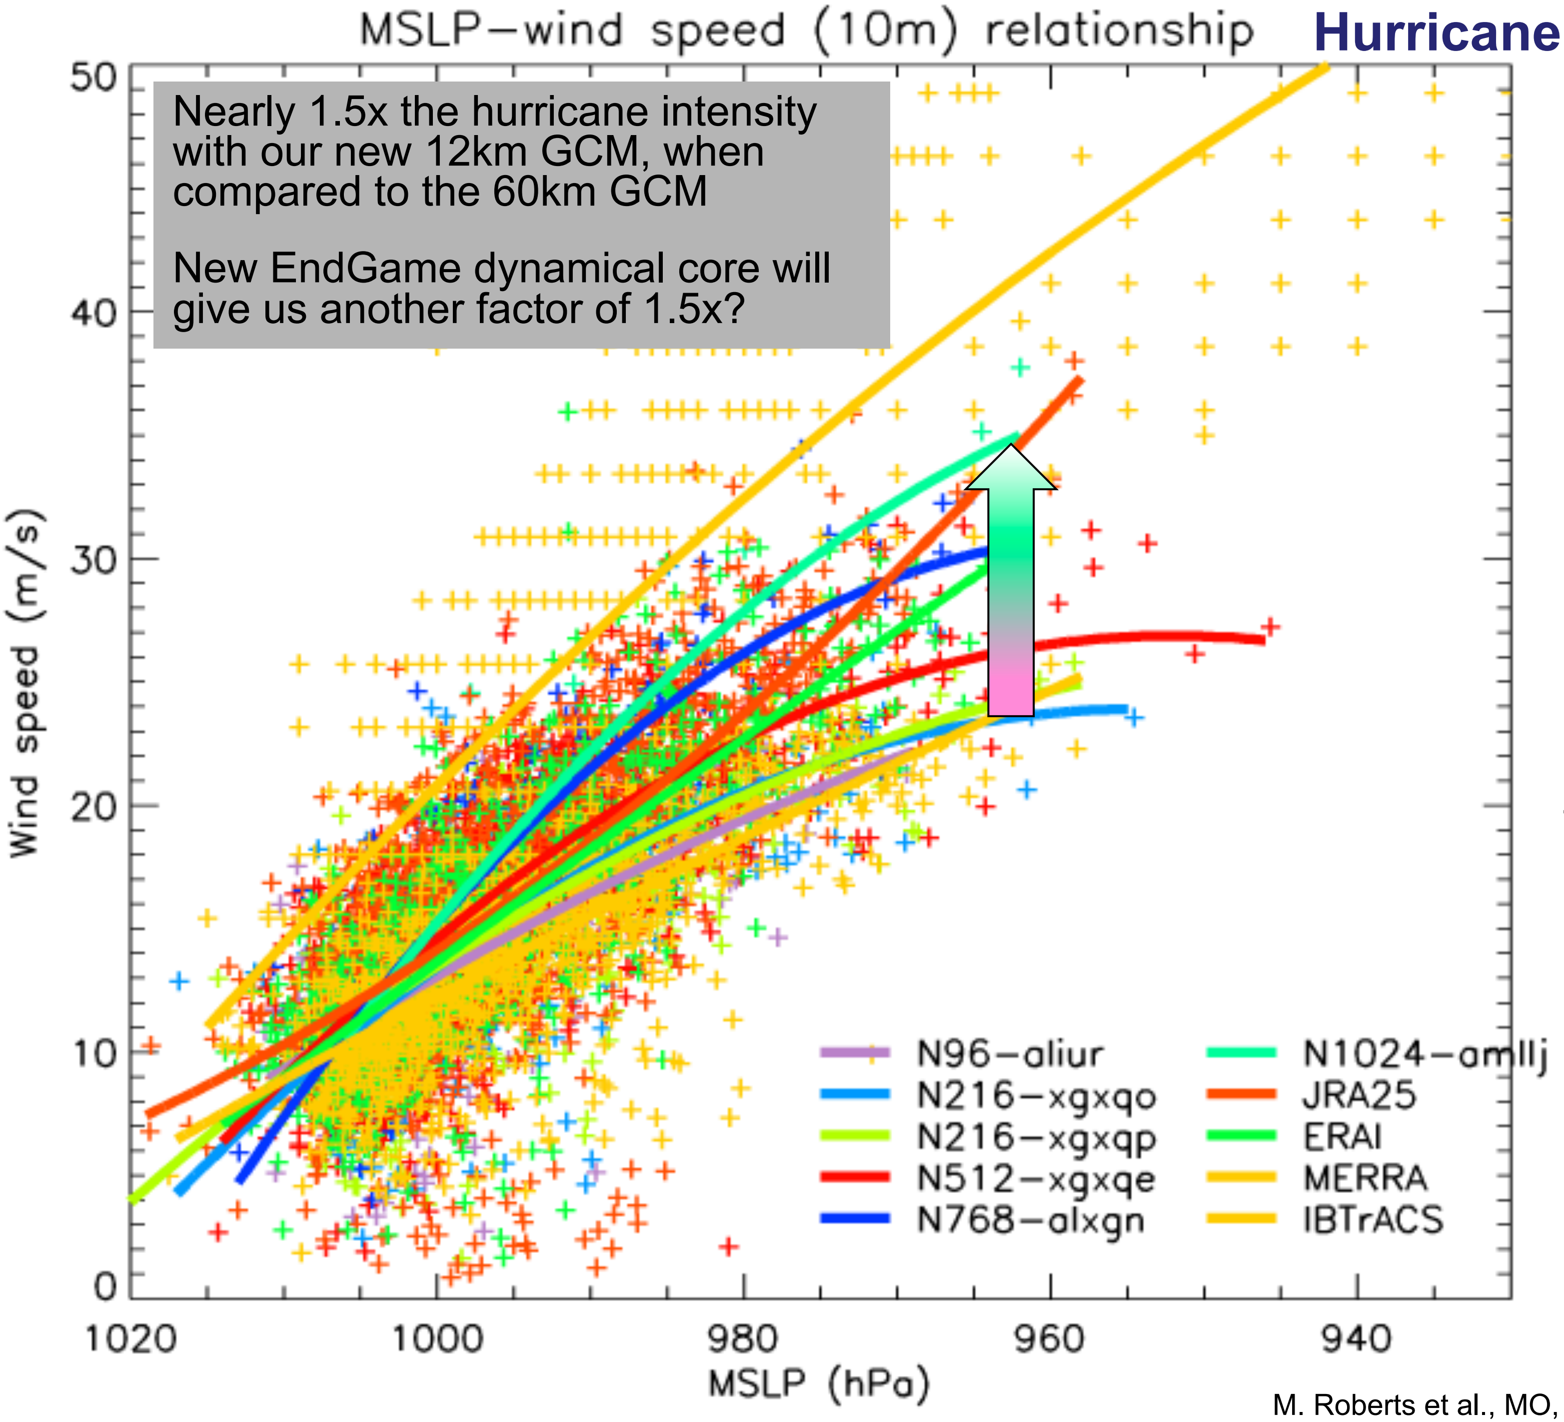

- HadGEM3-A
- GA4
- 1 year of 12km (N1024)
- 4 years of 17km (N768)

M. Roberts et al., MO, unpublished

Wisdom after PRACE-UPSCALE
1. On “Peta-scale” supercomputers

- The HLRS HERMIT service was excellent and reliable
- We successfully ran UPSCALE in 2012, exploiting the **largest High Performance Computing (HPC) allocation in science ever worldwide,** Equivalent to 18x Hadley Centre’s HadGEM2 IPCC submission (8M core hours)
  - Equivalent to 30 years of PL Vidale’s allocation on the UK’s HECToR supercomputer
- The success story stresses the importance of long-term partnerships between the Met Office and Academia: we have been collaborating since 2004 and together we have steadily built up the relevant national capability.
- The PRACE-UPSCALE grant was secured based on the quality of UK climate science and of UK Weather and Climate models:
  - The reason we have that capability is the availability of TIER-1 facilities in the UK
  - We could never have developed, nor demonstrated high-scalability, an absolute requirement for access to TIER-0 HPC, without continued access to our national TIER-1 facilities, HECToR, MO IBM Power 7, and MONSooN
- **The PRACE service can only be complementary; never a substitute for national, stable and dedicated HPC + data services**

Wisdom after PRACE-UPSCALE
2. On **data and their analysis:**
the long road from HPC to Science

- The Earth Simulator project data are still flowing into science papers now, nearly six years after the project's start. Much of this is due to:
  - the small size of the team, who are developing models and producing science at the same time
  - the daunting data analysis task.
- UPSCALE has generated 0.4PBytes of data, which **will take years for the UK community to analyse** (for reference, the size of the entire CMIP data archive is 1.8PBytes)
  - We would never have been able to store, nor analyse, that volume of data, without the existence of the national Centre for Environmental Data Archival (CEDA) service, which became available fortuitously at the same time as PRACE TIER-0 access
  - We will need to sustain the CEDA service for many years in order to maximise the UPSCALE science output.
- The PRACE funding model, providing HPC for one year, but no human resources, nor data storage/analysis facilities, is not tenable for a science enterprise, because of:
  - Too intensive use of personnel time to port, optimise, validate, run, monitor and collect data, leaving no time for science
  - Time limitation (1 year on-off) insufficient for development and understanding;
    - no time for errors or exploring emerging ideas/ hypotheses
    - can force compromises on optimal model configuration, e.g. not sufficient for coupled work, nor for centennial-scale climate integrations.