Inference using large climate simulators: HadCM3 and North American Mid-Holocene temperature anomalies

Jonathan Rougier

Department of Mathematics
University of Bristol, UK

with Tamsin Edwards, Mat Collins,
and other members of the PalaeoQUMP team

Liverpool Marine Symposium, Liverpool, 17 Jan 2011

Palaeoclimate reconstruction

1. 'Pseudo-observations' based on proxy measurements have a high spatial resolution, but sparse coverage

2. Climate simulator runs have full coverage but low spatial resolution, and there is the problem of simulator limitations

Palaeoclimate reconstruction

1. 'Pseudo-observations' based on proxy measurements have a high spatial resolution, but sparse coverage

2. Climate simulator runs have full coverage but low spatial resolution, and there is the problem of simulator limitations

... Can we construct a synthesis of these two sources of information which combines their strengths?

This is a very generic problem. A *statistical* solution emphasises the assessment and role of uncertainty, represented probabilistically.

Pseudo-observations for pointwise reconstructions

Mid-Holocene MTWA anomalies.

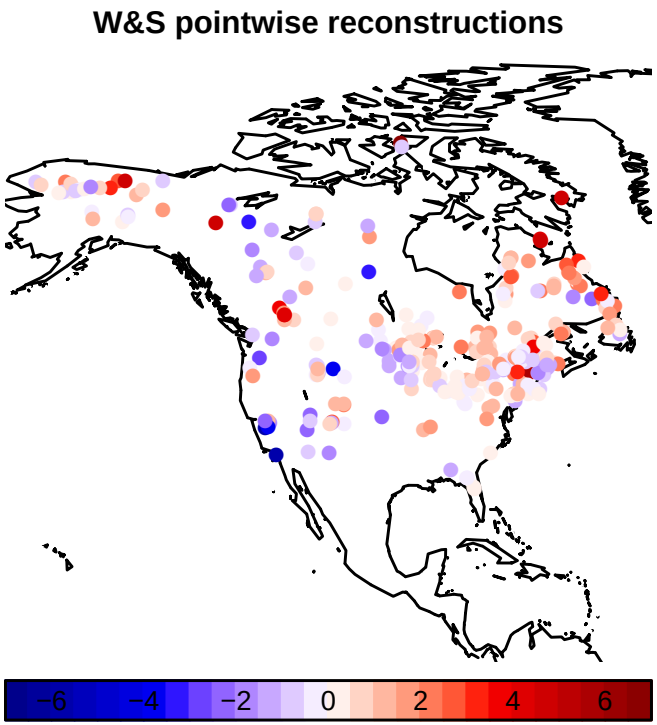

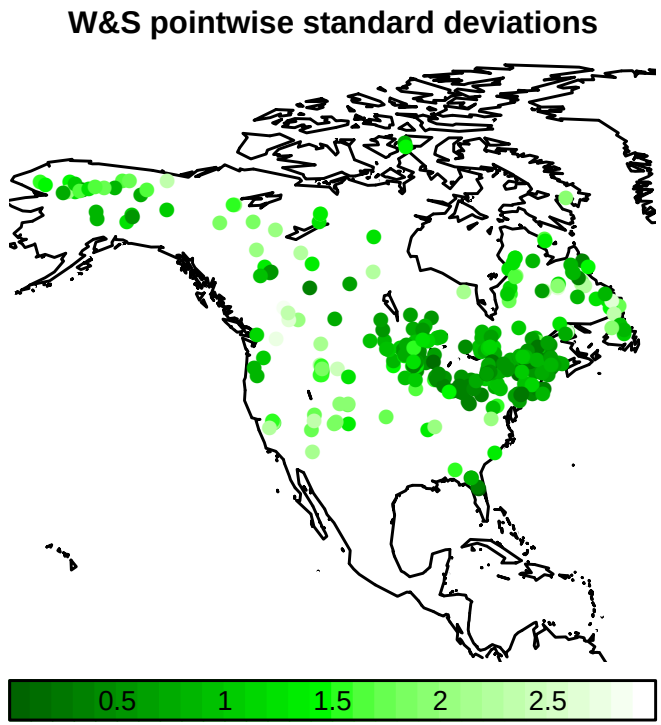

HadCM3 runs

Standard parameterisation and some of our ensemble members (n.b. different colour scale to the previous picture).

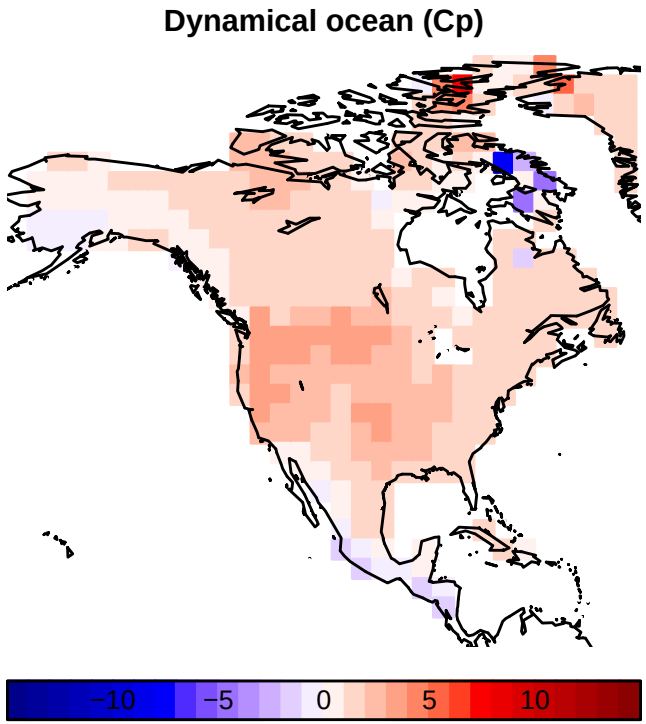

HadCM3 runs

Standard parameterisation and some of our ensemble members (n.b. different colour scale to the previous picture).

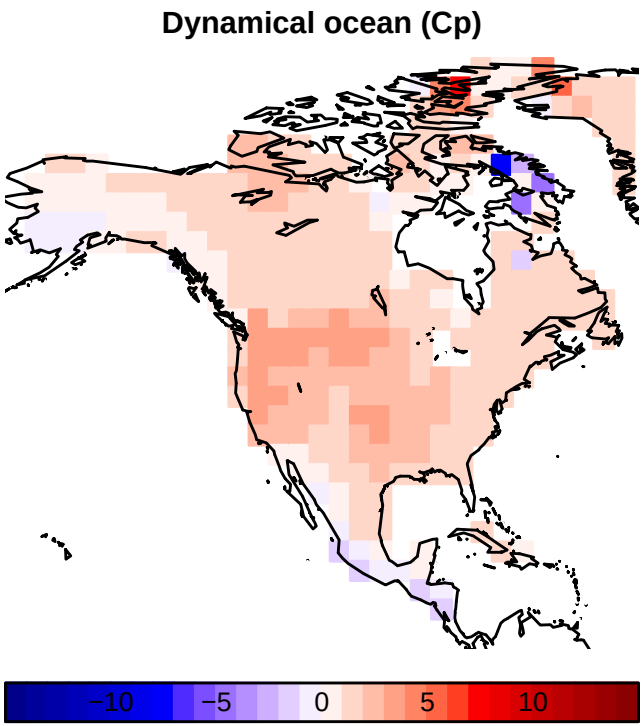

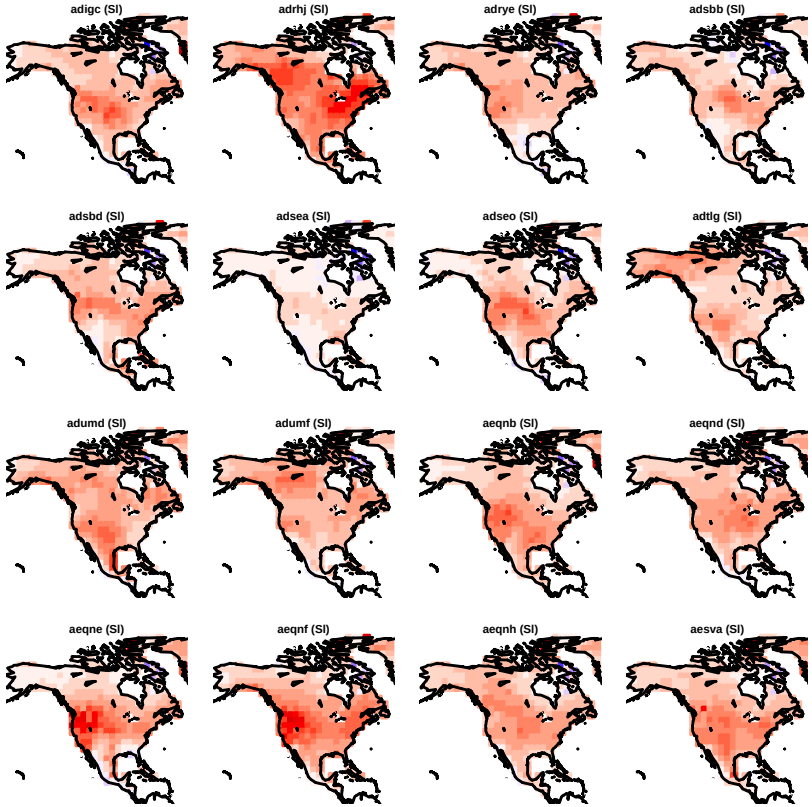

A natural method which is not quite going to work

Imagine that HadCM3 was very fast to run. We could use the following approach:

1. Sample millions of candidates for the collection of simulator parameters, and for each one:
   a. Run the simulator, and
   b. Score the result by comparison with the pseudo-observations.
2. Create a weighted average of the sample.

A natural method which is not quite going to work

Imagine that HadCM3 was very fast to run. We could use the following approach:

1. Sample millions of candidates for the collection of simulator parameters, and for each one:
   a. Run the simulator, and
   b. Score the result by comparison with the pseudo-observations.
2. Create a weighted average of the sample.

Unfortunately for us:

1. HadCM3 takes weeks/months to run
2. We have inherited an ensemble of runs that is not any kind of sample.

A natural method which is not quite going to work

Imagine that HadCM3 was very fast to run. We could use the following approach:

1. Sample millions of candidates for the collection of simulator parameters, and for each one:
   a. Run the simulator, and
   b. Score the result by comparison with the pseudo-observations.
2. Create a weighted average of the sample.

Unfortunately for us:

1. HadCM3 takes weeks/months to run
2. We have inherited an ensemble of runs that is not any kind of sample.

The solution is to use the ensemble to construct an emulator of the climate simulator, i.e. a *statistical model of the simulator* that allows us to predict its output at arbitrary parameterisations.

Two main challenges

1. High dimension. Our output field is 337 gridcells each $3.75^\circ \times 2.5^\circ$, but *HadCM3 is not that good!* Probably it is good at continental-and-a-bit-under-scales. Perhaps there is a $\sim$10-dimensional subspace which we 'trust'.

Two main challenges

1. High dimension. Our output field is 337 gridcells each $3.75^\circ \times 2.5^\circ$, but *HadCM3 is not that good!* Probably it is good at continental-and-a-bit-under-scales. Perhaps there is a $\sim$10-dimensional subspace which we 'trust'.

2. Internal variability. This is a problem because we cannot afford the very long integration times necessary to extract the mostly-smooth manifold of how HadCM3 output varies with parameterisation.

Two main challenges

1. High dimension. Our output field is 337 gridcells each $3.75° \times 2.5°$, but *HadCM3 is not that good!* Probably it is good at continental-and-a-bit-under-scales. Perhaps there is a $\sim$10-dimensional subspace which we 'trust'.

2. Internal variability. This is a problem because we cannot afford the very long integration times necessary to extract the mostly-smooth manifold of how HadCM3 output varies with parameterisation.

Our solution is to project HadCM3 output onto the column space of a matrix of 'trustworthy' linear combinations, and then to use statistical regression methods to estimate the smooth manifold of this projection using an independent estimate of the internal variability variance matrix.

Our choice of filtering matrix

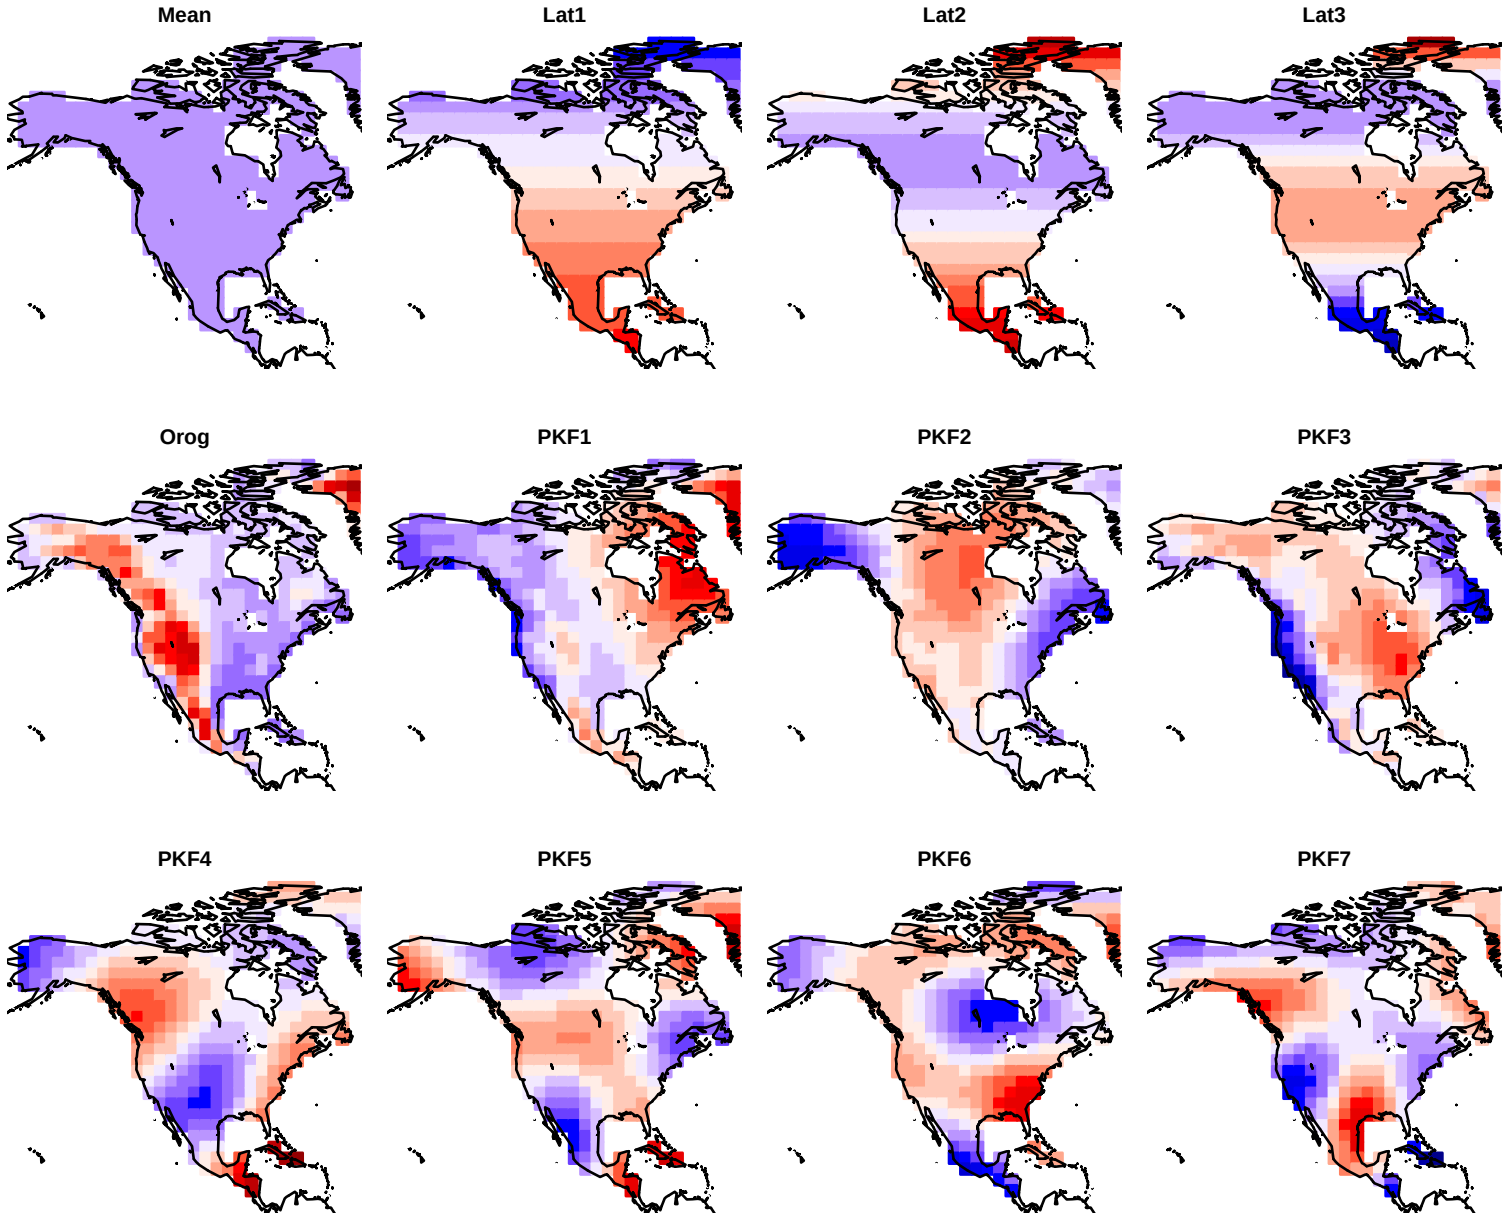

Our choice of filtering matrix

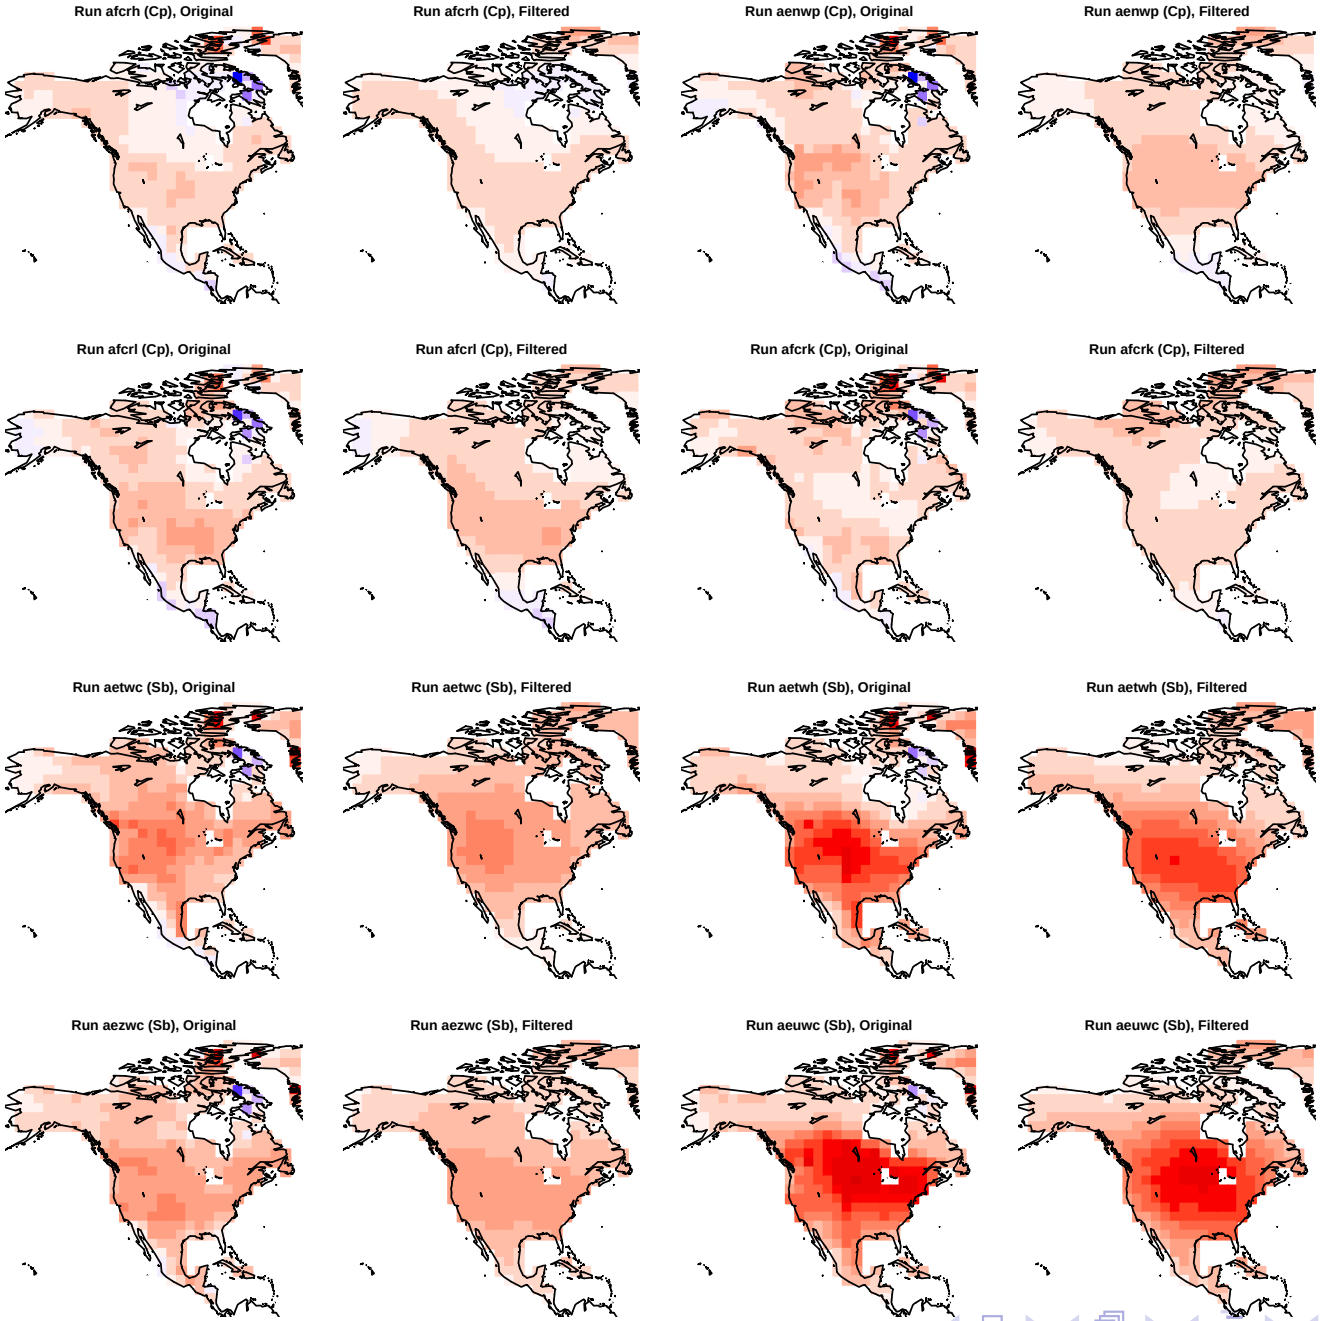

Checking the emulator

Diagnostic information based on leave-one-out; displayed as zonal means to indicate the emulator's prediction envelope.

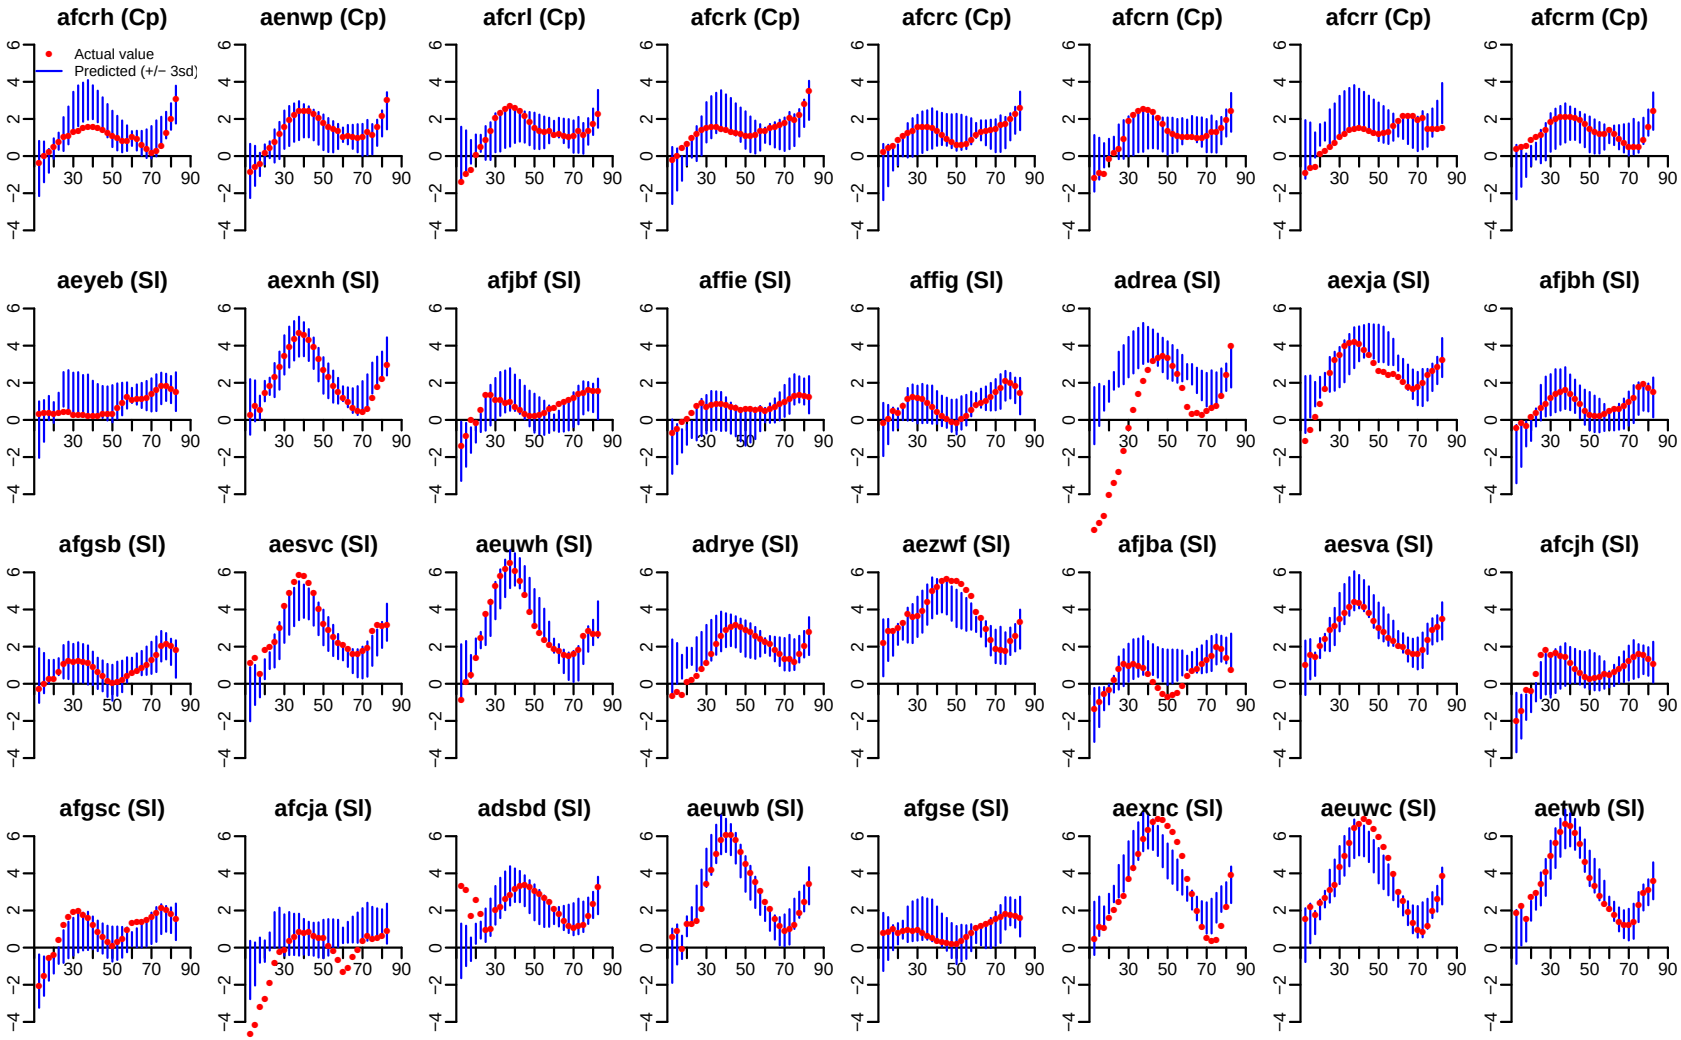

Combined reconstruction

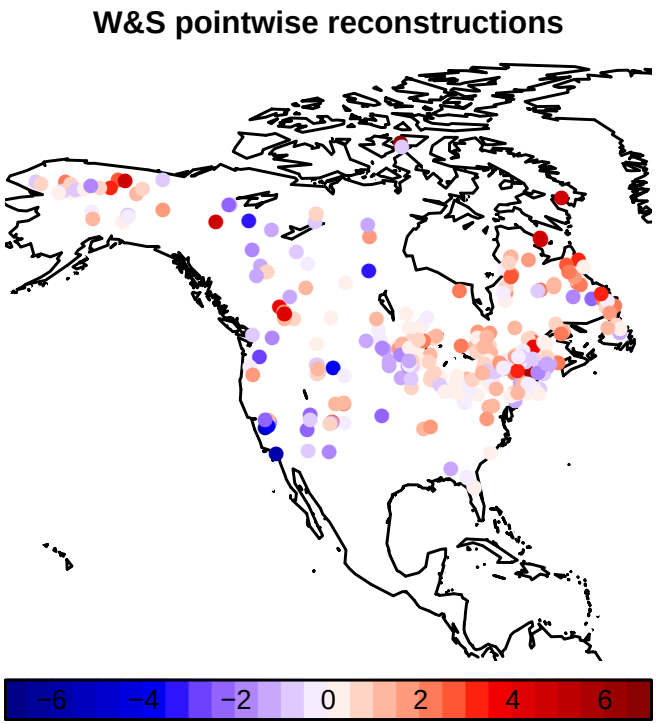

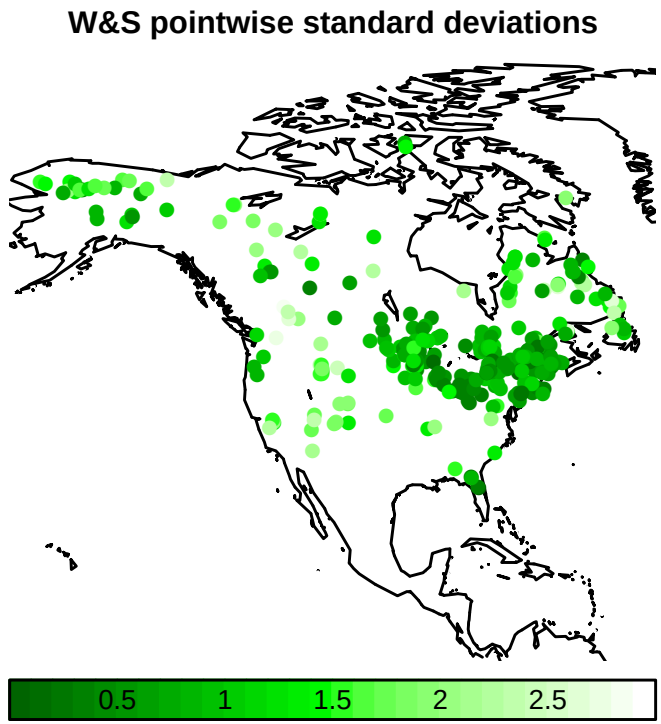

Combined reconstruction

Initial mean field

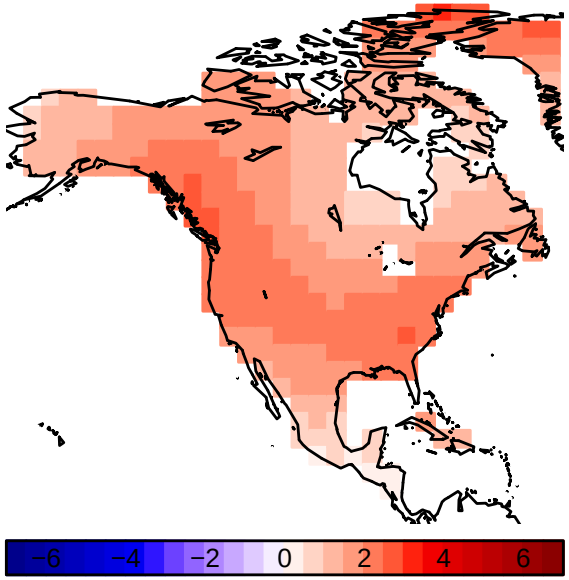

Initial SD field

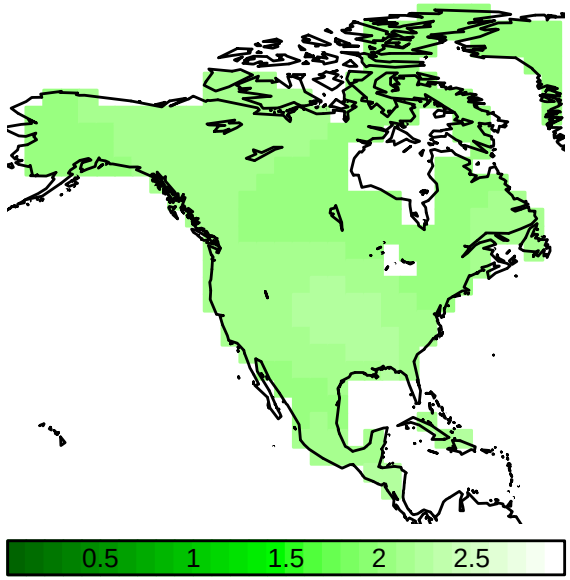

Adjusted mean field

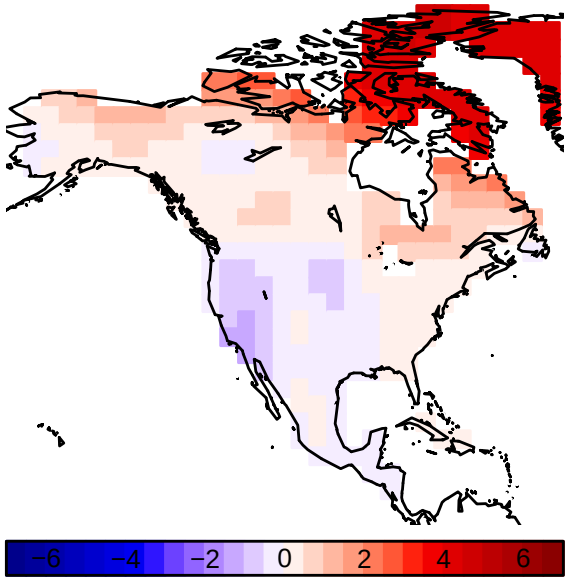

Adjusted SD field

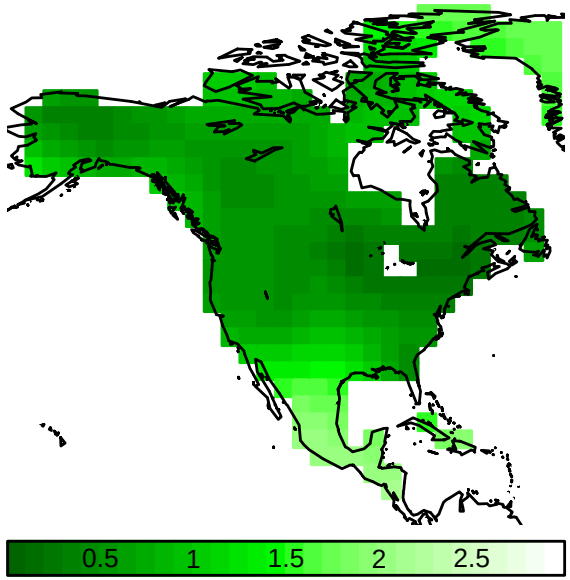

Main questions so far

`REM`: Statistics does not provide answers—it provides a framework within which You express Your judgements. One important role of this framework is to clarify the questions.

Main questions so far

`REM`: Statistics does not provide answers—it provides a framework within which You express Your judgements. One important role of this framework is to clarify the questions.

1. Emulating a climate simulator
   - What linear combinations of high-dimensional spatial outputs are 'trustworthy'?
   - How to get a good (robust) estimate of internal variability?
   - How to choose the statistical regression function describing the mostly-smooth response of the simulator mean to changes in the parameters?

Main questions so far

`REM`: Statistics does not provide answers—it provides a framework within which You express Your judgements. One important role of this framework is to clarify the questions.

1. Emulating a climate simulator
   - What linear combinations of high-dimensional spatial outputs are 'trustworthy'?
   - How to get a good (robust) estimate of internal variability?
   - How to choose the statistical regression function describing the mostly-smooth response of the simulator mean to changes in the parameters?

2. Linking HadCM3 to reality
   - What is a good probabilistic description for parametric uncertainty?
   - How to assess and quantify structural uncertainty?
   - How to present fully-probabilistic information about spatial (and spatial/temporal) reconstructions?