NETWORKS IN CLIMATE SCIENCE

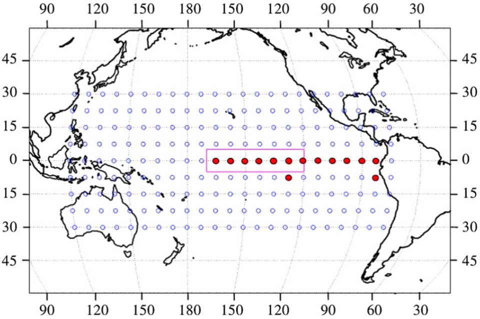

The Azimuth Project

John Baez, Jan Galkowski, Graham Jones, Nadja Kutz, Daniel Mahler, Blake Pollard, Paul Pukite, Dara Shayda, David Tanzer, David Tweed, Steve Wenner *et al*

El Niño Southern Oscillation:
a major source of temperature variability

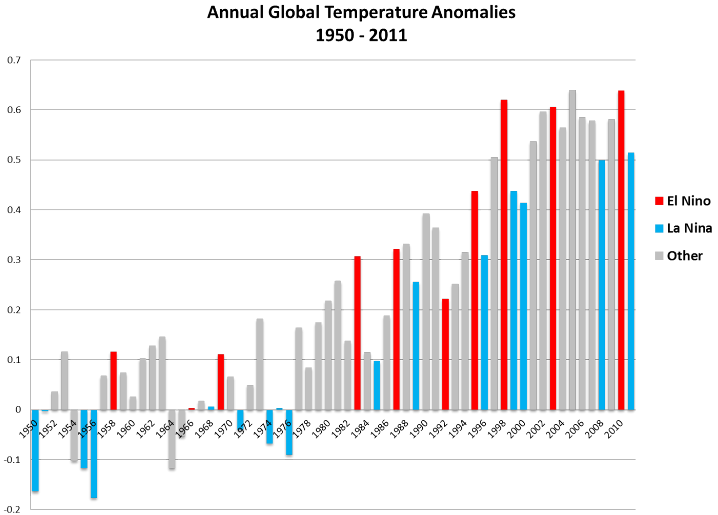

What is El Niño?

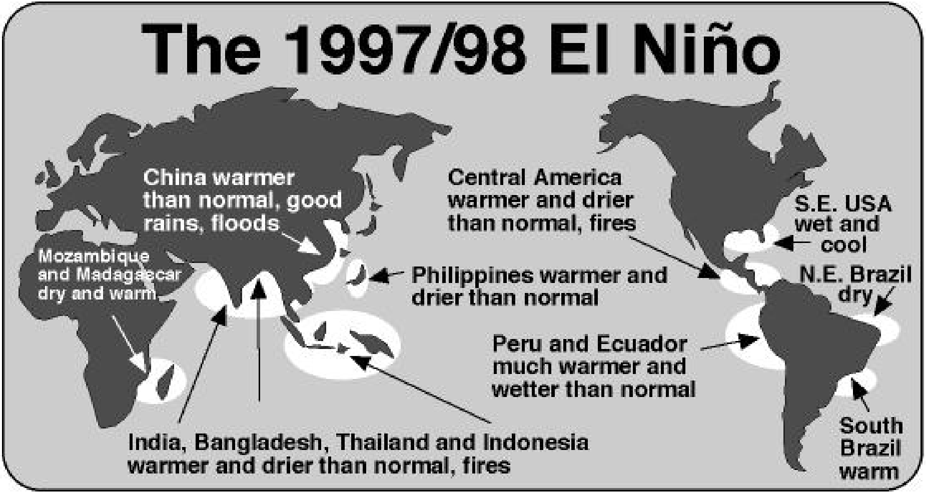

Shrimp News International

Two closely correlated signs of El Niño:

1) Increased sea surface temperatures in this region:

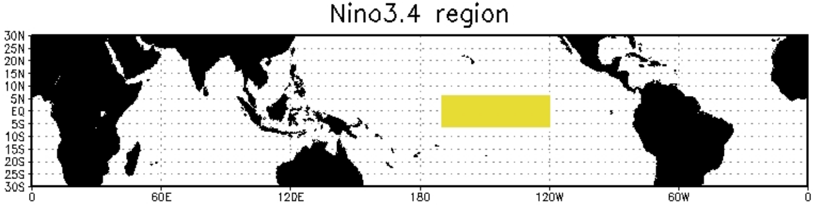

The **temperature anomaly** in this region—how much warmer it is than usual for this time of year—is called the **Niño 3.4 index.**

2) A decrease in air pressures in the *western* side of the Pacific compared to those further *east*. This is measured by the **Southern Oscillation Index** or **SOI**.

2) A decrease in air pressures in the *western* side of the Pacific compared to those further *east*. This is measured by the **Southern Oscillation Index** or **SOI**.

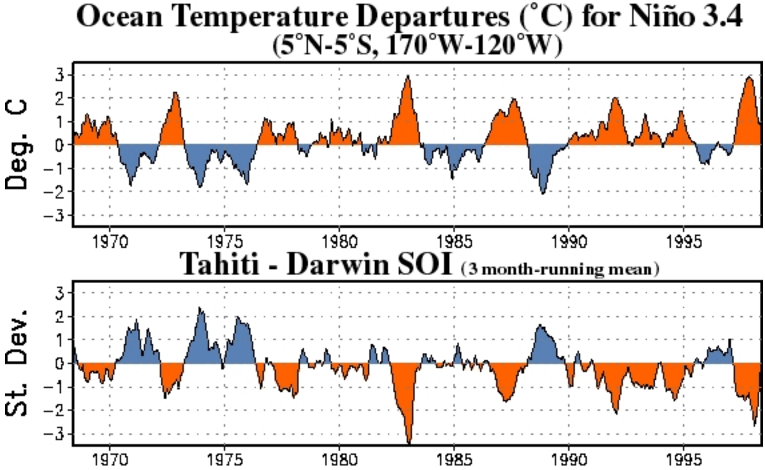

The El Niño Southern Oscillation or **ENSO** is one of many **teleconnections**: strong correlations between weather at distant locations. These can be found using principal component analysis. ENSO is the biggest on time scales between a year and a decade.

The El Niño Southern Oscillation or **ENSO** is one of many **teleconnections**: strong correlations between weather at distant locations. These can be found using principal component analysis. ENSO is the biggest on time scales between a year and a decade.

A new line of research is studying teleconnections using **complex network theory**: the study of large graphs found in nature and society.

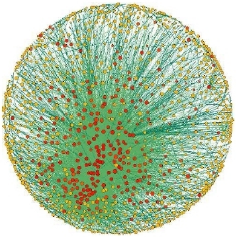

Vitali, Glattfelder and Battiston:
The network of global corporate control

In complex network theory, people often start with a **weighted graph**:

- a set N of **nodes**
- for any pair of nodes $i, j \in N$, a **weight** $A_{ij} \in [0, \infty)$.

In complex network theory, people often start with a **weighted graph**:

- a set N of **nodes**
- for any pair of nodes $i, j \in N$, a **weight** $A_{ij} \in [0, \infty)$.

We can turn a weighted graph into a graph by drawing an edge from i to j whenever $A_{ij} > c$ for some threshold c.

In complex network theory, people often start with a **weighted graph**:

- a set N of **nodes**
- for any pair of nodes $i, j \in N$, a **weight** $A_{ij} \in [0, \infty)$.

We can turn a weighted graph into a graph by drawing an edge from i to j whenever $A_{ij} > c$ for some threshold c.

In complex network theory, people try to compute quantities that reveal interesting information about large weighted graphs.

For example, the **in-degree** of a node $j \in N$ in a weighted graph is

$$\sum_{i \in N} A_{ij}$$

If lots of people have web pages with lots of links to yours, your webpage will have a high in-degree.

If lots of people like you on Facebook, *you* will have a high in-degree.

In Complex networks in climate dynamics, Donges *et al* take daily surface air temperatures on a grid — produced by a climate model — and compute the correlation between grid points.

In Complex networks in climate dynamics, Donges *et al* take daily surface air temperatures on a grid — produced by a climate model — and compute the correlation between grid points.

Let $T_i(t)$ be the temperature anomaly at the ith grid point at month t. They compute the correlation of these variables for each pair of grid points i, j:

$$A_{ij} = \frac{\langle T_i T_j \rangle}{\sqrt{\langle T_i^2 \rangle \langle T_j^2 \rangle}}$$

This has $-1 \leq A_{ij} = A_{ji} \leq 1$.

In Complex networks in climate dynamics, Donges *et al* take daily surface air temperatures on a grid — produced by a climate model — and compute the correlation between grid points.

Let $T_i(t)$ be the temperature anomaly at the ith grid point at month t. They compute the correlation of these variables for each pair of grid points i, j:

$$A_{ij} = \frac{\langle T_i T_j \rangle}{\sqrt{\langle T_i^2 \rangle \langle T_j^2 \rangle}}$$

This has $-1 \leq A_{ij} = A_{ji} \leq 1$.

They create a graph where there is an edge between i and j when $|A_{ij}| > c$. They adjust c so that 0.5% of the pairs i, j have an edge between them.

Then Donges *et al* plot the degree of each node, weighted by the area of its grid square:

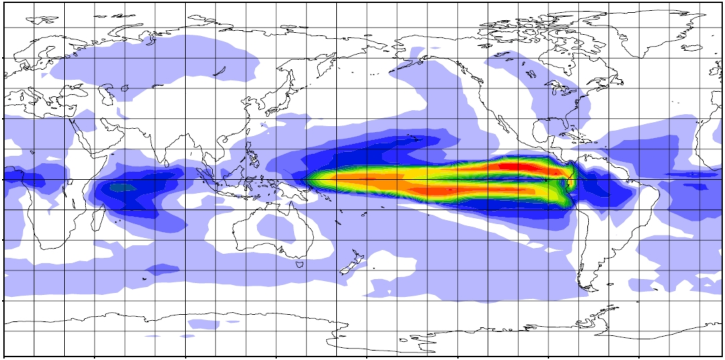

A certain patch dominates the world: the **El Niño basin.** The reddest patches are nodes connected to 5% or more of the other nodes: at least 10 times as many as average.

Is this an artifact of only looking at *linear* correlations?

Is this an artifact of only looking at *linear* correlations?

No! Repeating the procedure with mutual information, Donges *et al* get this:

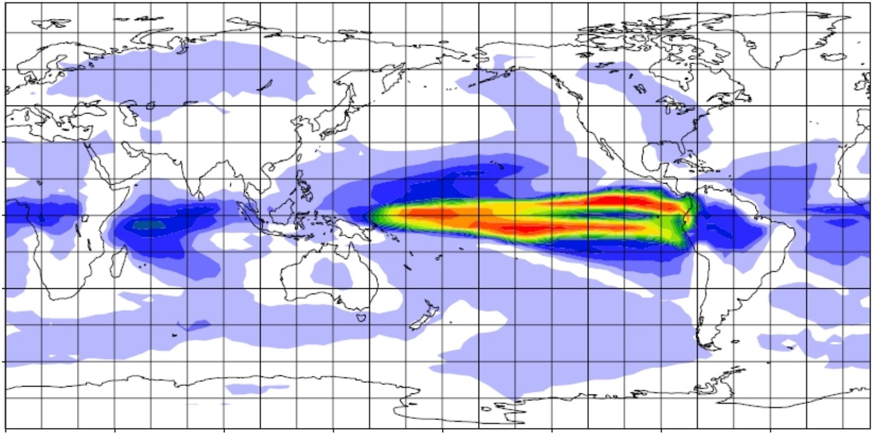

We can also look at how climate networks change with time.

We can also look at how climate networks change with time.

In Topology and predictability of El Niño and La Niña networks, Tsonis and Swanson create a climate network in a similar way, the main differences being that they:

- use actual monthly temperature data from 1950 to 2005;

We can also look at how climate networks change with time.

In Topology and predictability of El Niño and La Niña networks, Tsonis and Swanson create a climate network in a similar way, the main differences being that they:

- use actual monthly temperature data from 1950 to 2005;
- separately create climate networks for El Niño and La Niña time periods;

We can also look at how climate networks change with time.

In Topology and predictability of El Niño and La Niña networks, Tsonis and Swanson create a climate network in a similar way, the main differences being that they:

- use actual monthly temperature data from 1950 to 2005;
- separately create climate networks for El Niño and La Niña time periods;
- create a link between grid points when their correlation has absolute value greater than 0.5;

We can also look at how climate networks change with time.

In Topology and predictability of El Niño and La Niña networks, Tsonis and Swanson create a climate network in a similar way, the main differences being that they:

- use actual monthly temperature data from 1950 to 2005;
- separately create climate networks for El Niño and La Niña time periods;
- create a link between grid points when their correlation has absolute value greater than 0.5;
- only use temperature data from November to March in each year.

Tsonis and Swanson get this map for La Niña conditions:

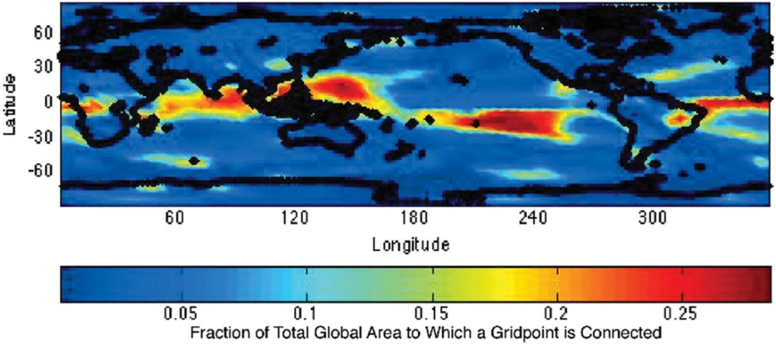

and this map for El Niño conditions:

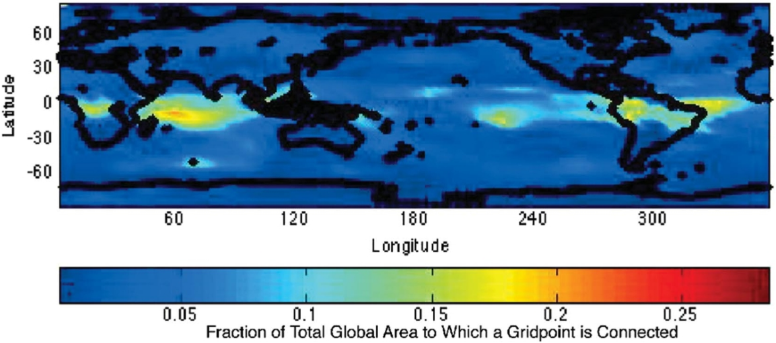

They conclude that *"El Niño breaks climate links"*.

Can we use climate networks to *predict* El Niños?

Can we use climate networks to *predict* El Niños?

People really want to predict El Niños, because they have huge effects on agriculture, especially around the Pacific ocean. But it's very hard to predict El Niños more than 6 months in advance!

Can we use climate networks to *predict* El Niños?

People really want to predict El Niños, because they have huge effects on agriculture, especially around the Pacific ocean. But it's very hard to predict El Niños more than 6 months in advance!

What would we like to predict, exactly? When the 3-month running average of the Niño 3.4 index exceeds 0.5 $^\circ$C for 5 months, we officially declare that there is an **El Niño**.

There are many teams trying to predict the Niño 3.4 index.

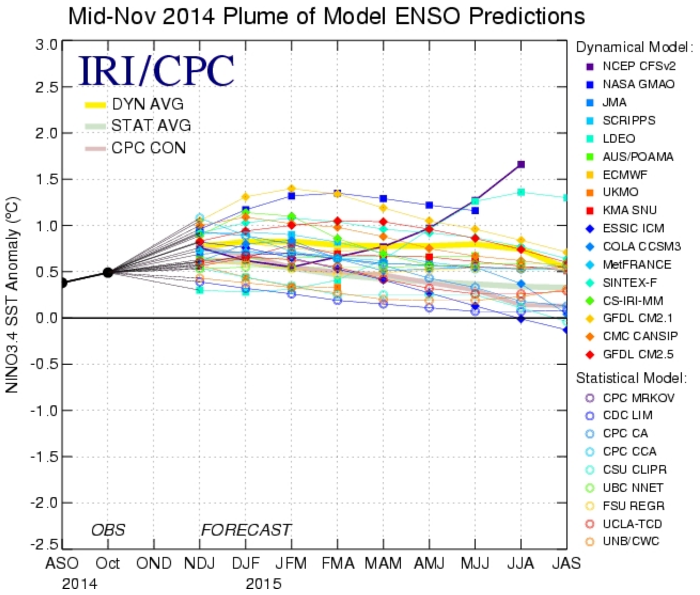
Mid-Nov 2014 Plume of Model ENSO Predictions
IRI/CPC
DYN AVG
STAT AVG
CPC CON
NINO3.4 SST Anomaly (°C)
3.0
2.5
2.0
1.5
1.0
0.5
0.0
-0.5
-1.0
-1.5
-2.0
-2.5
OBS
FORECAST
ASO
Oct
OND
NDJ
DJF
JFM
FMA
MAM
AMJ
MJJ
JJA
JAS
2014
2015
Dynamical Model:
NCEP CFSv2
NASA GMAO
JMA
SCRIPPS
LDEO
AUS/POAMA
ECMWF
UKMO
KMA SNU
ESSIC ICM
COLA CCSM3
MetFRANCE
SINTEX-F
CS-IRI-MM
GFDL CM2.1
CMC CANSIP
GFDL CM2.5
Statistical Model:
CPC MRKOV
CDC LIM
CPC CA
CPC CCA
CSU CLIPR
UBC NNET
FSU REGR
UCLA-TCD
UNB/CWC

A 2013 paper by Ludescher *et al*, called Very early warning of next El Niño, uses a climate network for El Niño prediction.

They build their climate network using correlations between daily surface air temperature data between points *inside* the El Niño basin and certain points *outside* this region:

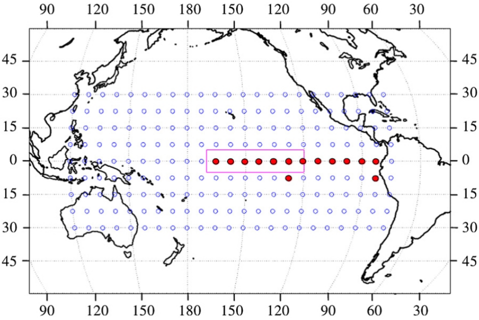

Ludescher *et al*

They use these correlations to construct a weighted graph. They predict an El Niño will occur if the average weight exceeds a certain threshold.

More precisely: they let $T_i(t)$ be the surface air temperature anomaly at the i grid point at time t. They consider the **time-delayed covariance** between temperatures at different grid points:

$$\left\langle T_i(t) T_j(t-\tau) \right\rangle - \left\langle T_i(t) \right\rangle \left\langle T_j(t-\tau) \right\rangle$$

where τ is a time delay and $\langle \ \rangle$ denotes a running average over the last year, that is:

$$\left\langle f(t) \right\rangle = \frac{1}{365} \sum_{d=0}^{364} f(t-d)$$

where t is the time in days.

They normalize this to define a correlation $C^t_{i,j}(\tau)$ that ranges from -1 to 1.

They normalize this to define a correlation $C^t_{i,j}(\tau)$ that ranges from -1 to 1.

Next, for any pair of nodes *i* and *j* and for each time *t*, they determine the maximum, the mean and the standard deviation of $|C^t_{i,j}(\tau)|$ as the delay τ ranges from -200 to 200 days.

They normalize this to define a correlation $C^t_{i,j}(\tau)$ that ranges from -1 to 1.

Next, for any pair of nodes i and j and for each time t, they determine the maximum, the mean and the standard deviation of $|C^t_{i,j}(\tau)|$ as the delay τ ranges from -200 to 200 days.

They define the **link strength** $S_{i,j}(t)$ as the difference between the maximum and the mean value of $|C^t_{i,j}(\tau)|$, divided by its standard deviation.

They normalize this to define a correlation $C^t_{i,j}(\tau)$ that ranges from -1 to 1.

Next, for any pair of nodes i and j and for each time t, they determine the maximum, the mean and the standard deviation of $|C^t_{i,j}(\tau)|$ as the delay τ ranges from -200 to 200 days.

They define the **link strength** $S_{i,j}(t)$ as the difference between the maximum and the mean value of $|C^t_{i,j}(\tau)|$, divided by its standard deviation.

Finally, they let $S(t)$ be the **average link strength**, calculated by averaging $S_{ij}(t)$ over all pairs i, j where i is a grid point *inside* their El Niño basin and j is a grid point *outside* this basin but still in their larger rectangle.

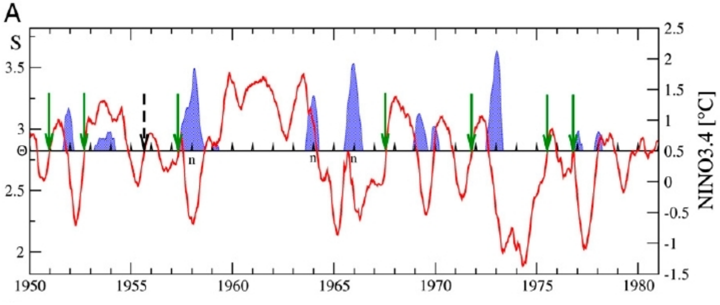

The blue peaks are episodes when the Niño 3.4 index exceeds 0.5 °C.

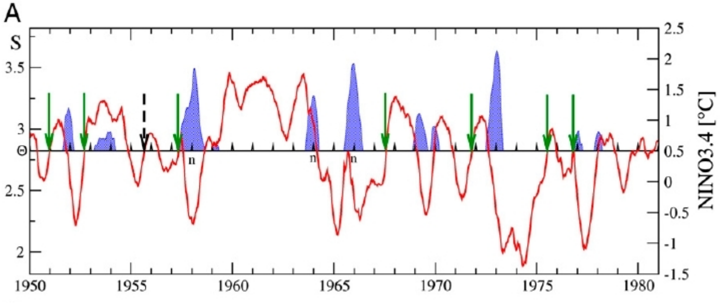

The blue peaks are episodes when the Niño 3.4 index exceeds 0.5 °C.

The red line is their 'average link strength'. Whenever this exceeds a certain threshold, and the Niño 3.4 index is not *already* over 0.5 °C, they predict it will exceed 0.5 °C in the following calendar year.

Ludescher *et al* chose their threshold for El Niño prediction by training their algorithm on climate data from 1948 to 1980, and tested it on data from 1981 to 2013.

Ludescher *et al* chose their threshold for El Niño prediction by training their algorithm on climate data from 1948 to 1980, and tested it on data from 1981 to 2013.

With this threshold, they claim their El Niño predictions were correct 76% of the time, and their predictions of no El Niño were correct 86% of the time.

Ludescher *et al* chose their threshold for El Niño prediction by training their algorithm on climate data from 1948 to 1980, and tested it on data from 1981 to 2013.

With this threshold, they claim their El Niño predictions were correct 76% of the time, and their predictions of no El Niño were correct 86% of the time.

On this basis they claimed that the Niño 3.4 index would exceed 0.5 by the end of 2014 with probability 3/4.

Ludescher *et al* chose their threshold for El Niño prediction by training their algorithm on climate data from 1948 to 1980, and tested it on data from 1981 to 2013.

With this threshold, they claim their El Niño predictions were correct 76% of the time, and their predictions of no El Niño were correct 86% of the time.

On this basis they claimed that the Niño 3.4 index would exceed 0.5 by the end of 2014 with probability 3/4.

The latest data says: yes, this happened in November!

Analysis

Graham Jones of the Azimuth Project wrote code implementing Ludescher *et al*'s algorithm and got results close to theirs, though not identical.

Analysis

Graham Jones of the Azimuth Project wrote code implementing Ludescher *et al*'s algorithm and got results close to theirs, though not identical.

Steve Wenner did a statistical analysis and estimated the mean success probability is 80% using Ludescher's method — but 71% if we always predict "no El Niño".

Analysis

Graham Jones of the Azimuth Project wrote code implementing Ludescher *et al*'s algorithm and got results close to theirs, though not identical.

Steve Wenner did a statistical analysis and estimated the mean success probability is 80% using Ludescher's method — but 71% if we always predict "no El Niño".

Their method has an 61% expected posterior probability of successfully predicting an El Niño initiation when one actually occurs — but the 95% confidence interval is huge: 39% to 80%.

How much does average link strength say about whether an El Niño is coming?

How much does average link strength say about whether an El Niño is coming?

There's a problem with small sample size. Ludescher *et al* only make a single yes-or-no prediction each year for 65 years.

31 years were used for training their algorithm, leaving just 34 retrodictions and one actual prediction.

How much does average link strength say about whether an El Niño is coming?

There's a problem with small sample size. Ludescher *et al* only make a single yes-or-no prediction each year for 65 years.

31 years were used for training their algorithm, leaving just 34 retrodictions and one actual prediction.

We need more ideas! Daniel Mahler and Dara Shayda of the Azimuth Project compared simple linear regression models to more complex models and found some interesting results.

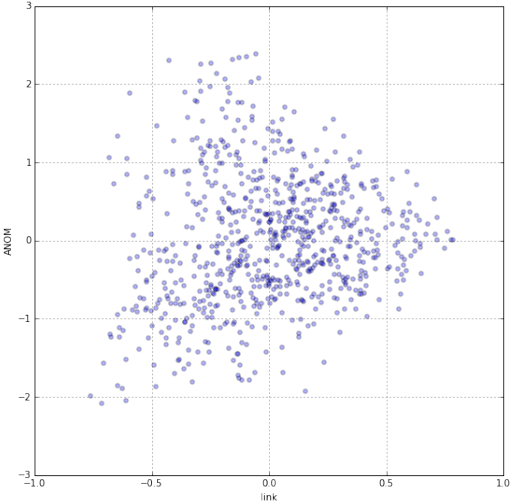

Niño 3.4 as a function of average link strength — no time lag

$R^2 = 0.0175$

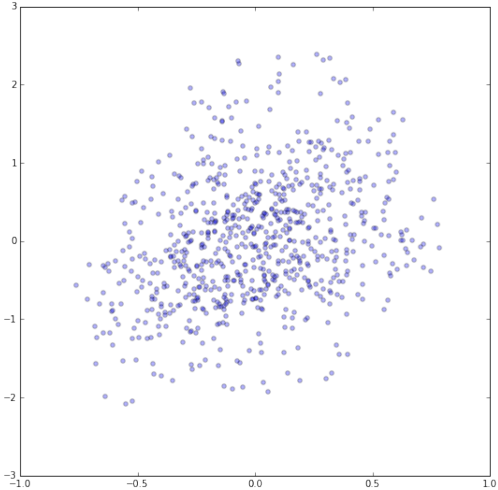

Niño 3.4 as a function of average link strength 6 months earlier

$R^2 = 0.088$

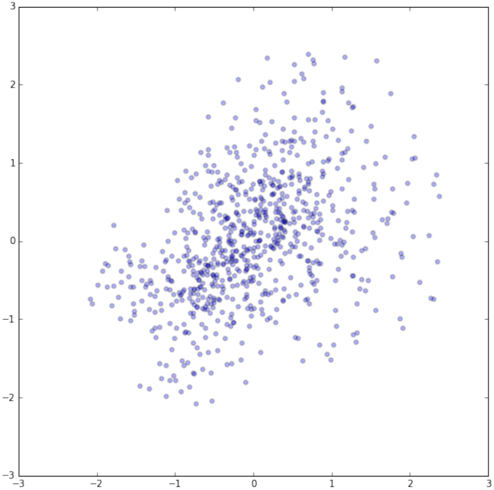

Niño 3.4 as a function of Niño 3.4 six months earlier

$R^2 = 0.162$

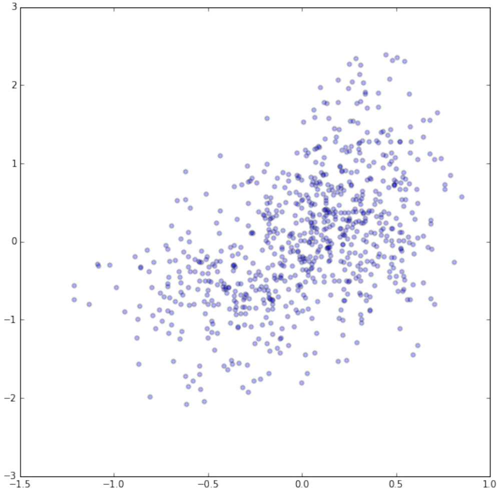

Niño 3.4 as a function of
average link strength and Niño 3.4 six months earlier

$R^2 = 0.22$

Summary

Using a linear model, the average link strength on a given month accounts for 8% of the variance of Niño 3.4 index 6 months in the future.

Summary

Using a linear model, the average link strength on a given month accounts for 8% of the variance of Niño 3.4 index 6 months in the future.

The Niño 3.4 index explains 16% of its own variance 6 months into the future.

Summary

Using a linear model, the average link strength on a given month accounts for 8% of the variance of Niño 3.4 index 6 months in the future.

The Niño 3.4 index explains 16% of its own variance 6 months into the future.

Taken together, they explain 22%. So, these two variables contain a fair amount of *independent* information about the Niño 3.4 index 6 months in the future.

Furthermore, they explain a surprisingly large amount of its variance for just 2 variables.

For comparison, Mahler used ExtraTreesRegressor to predict Niño 3.4 six months into the future from much larger datasets. Out of the 778 months available he trained the algorithm on the first 400 and tested it on 378.

For comparison, Mahler used ExtraTreesRegressor to predict Niño 3.4 six months into the future from much larger datasets. Out of the 778 months available he trained the algorithm on the first 400 and tested it on 378.

The result: using a full world-wide grid of surface air temperature values at a given moment in time explains only 23% of the Niño 3.4 six months into the future: $R^2 = 0.23$.

For comparison, Mahler used ExtraTreesRegressor to predict Niño 3.4 six months into the future from much larger datasets. Out of the 778 months available he trained the algorithm on the first 400 and tested it on 378.

The result: using a full world-wide grid of surface air temperature values at a given moment in time explains only 23% of the Niño 3.4 six months into the future: $R^2 = 0.23$.

A full grid of surface air pressure values explains 34% of the variance. Using twelve months of the full grid of pressure values only gets around 37%.

So, explaining 22% of the variance with just two variables doesn't look so bad!

Moreover, the average link strength is maximally correlated with Niño 3.4 *10 months into the future*:

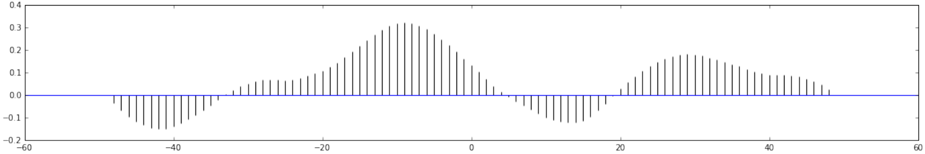

(The lines here occur at monthly intervals.)

Moreover, the average link strength is maximally correlated with Niño 3.4 *10 months into the future*:

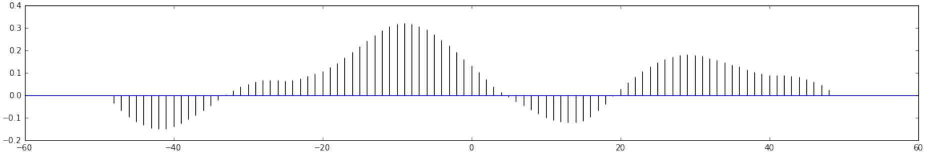

(The lines here occur at monthly intervals.)

So, *something like average link strength has predictive power for El Niños. But what is it,* ***exactly*** *— and why?*