

Published on *METEO 469: From Meteorology to Mitigation: Understanding Global Warming D7* (<https://www.e-education.psu.edu/meteo469>)

[Home](#) > [METEO 469 Course Outline](#) > Lesson 2 - Climate Observations, part 1

Lesson 2 - Climate Observations, part 1

The links below provide an outline of the material for this lesson. Be sure to carefully read through the entire lesson before returning to Canvas to submit your assignments.

Introduction

About Lesson 2

How do we know that climate change is taking place? Or that the factors we believe to be driving climate change, such as greenhouse gas concentrations, are themselves changing?

To address these questions, we turn first to instrumental measurements documenting changes in the properties of our atmosphere over time. These measurements are not without their uncertainties, particularly in earlier times. But they can help us to assess whether there appear to be trends in measures of climate and the factors governing climate, and whether the trends are consistent with our expectations of what the response of the climate system to human impacts ought to look like.

What will we learn in Lesson 2?

By the end of Lesson 2, you should be able to:

- discuss the various modern observational data characterizing changes in surface and atmospheric temperature over the historical period;
- discuss the nature of the uncertainties in the observational record of past climate; and
- perform simple statistical analyses to characterize trends in, and relationships between, data series.

What will be due for Lesson 2?

Please refer to the Syllabus for the specific time frames and due dates.

The following is an overview of the *required activities* for Lesson 2. Detailed directions and submission instructions are located within this lesson.

- Read:
 - [IPCC Fifth Assessment Report. Working Group 1](#) ^[1]
 - [Summary for Policy Makers](#) ^[2]
 - B Observed Changes in the Climate System: p. 4
 - B.1 Atmosphere: p. 5-8
 - Dire Predictions, v.2: p. 34-35, 38-39, 80-81
- Problem Set #1: Perform basic statistical analyses of climate data.

Questions?

If you have any questions, please post them to our *Questions?* discussion forum (not e-mail), located under the Home tab in Canvas. The instructor will check that discussion forum daily to respond. While you are there, feel free to post your own responses if you can help with any of the posted questions.

Observed Changes in Greenhouse Gases

Before we assess the climate data documenting changes in the climate system, we ought to address the question — is there evidence that greenhouse gases, purportedly responsible for observed warming, are actually changing in the first place?

Thanks to two legendary atmospheric scientists, we know that there is such evidence. The first of these scientists was [Roger Revelle](#) ^[3].



Roger Revelle.

Credit: [Cambridge Forum Speakers](#) [4]

Revelle, as we will later see, made fundamental contributions to understanding climate change throughout his career. Less known, but equally important, was the tutelage and mentorship that Revelle provided to other climate researchers. While at the Scripps Institution for Oceanography at the University of California in San Diego, Revelle encouraged his colleague [Charles David Keeling](#) [5] to make direct measurements of atmospheric CO₂ levels from an appropriately selected site.

Think About It!

Why do you suppose it is adequate to make measurements of atmospheric CO₂ from a single location as an indication of changing global concentrations?

Click for answer.

Revelle and Keeling settled on the top of the mountain peak Mauna Loa on the big island of Hawaii, establishing during the *International Geophysical Year* of 1958 an observatory that would be maintained by Keeling and his crew for the ensuing decades.



Charles David Keeling.

Credit: [NOAA](#) [6]



Mauna Loa Observatory.

Credit: [NOAA](#) [7]

From this location, Keeling and crew would make continuous measurements of atmospheric CO₂ from 1958 henceforth. Since then, long-term records have been established in other locations over the globe as well. The result of Keeling's labors is arguably the most famous curve in all of atmospheric science, the so-called *Keeling Curve*. That curve shows a steady increase in atmospheric CO₂ concentrations from about 315 ppm when the measurements began in 1958 to about 400 today (and climbing by about 2 ppm per year presently).

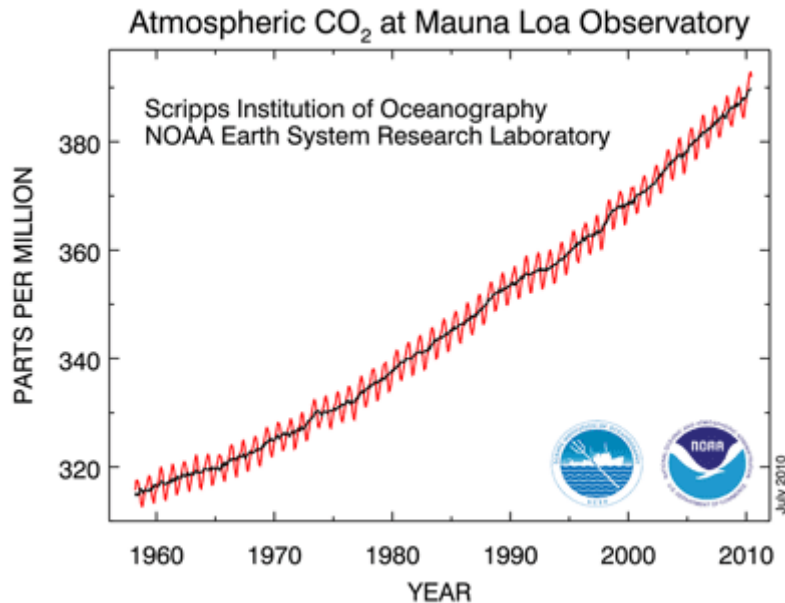


Figure 2.1: Atmospheric CO₂ at Mauna Loa Observatory

Credit: [NOAA](#) [8]

You might be wondering at this point, how do we know that the increase in CO₂ is not natural? For one thing, as you already have encountered in your readings, the recent levels are unprecedented over many millennia. Indeed, when we cover the topic of *paleoclimate*, we will see that the modern levels are unprecedented over many *hundreds of thousands* of years, and probably several million years. We will also see that there is a long-term relationship between CO₂ and temperature, though the story is not as simple as you might think.

But there is other more direct evidence that the source of the increasing CO₂ is indeed human, i.e., anthropogenic. It turns out that carbon that gets buried in the earth from dying organic matter and eventually turns into fossil fuels, tends to be *isotopically light*. That is, nature has a preference for burying carbon that is depleted of the heavier, ¹³C, carbon isotope. Fossil fuels are thus relatively rich in the lighter isotope, ¹²C. However, natural atmospheric CO₂ produced by respiration (be it animals like us, or plants which both respire and photosynthesize) tends to have a greater abundance of the heavier ¹³C isotope of carbon. If the CO₂ increase were from natural sources, we would therefore expect the ratio of ¹³C to ¹²C to be getting higher. But instead, the ratio of ¹³C to ¹²C is getting lower as CO₂ is building up in our atmosphere -- i.e., the ratio bears the *fingerprint* of anthropogenic fossil fuel burning.

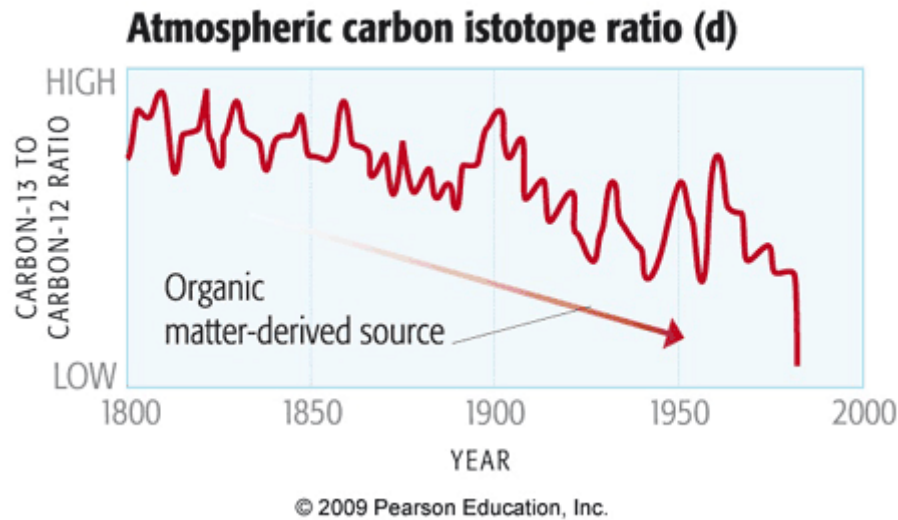


Figure 2.2: Graph of carbon-13 to carbon-12 Ratio from 1800 - 2000.

Credit: Mann and Kump, *Dire Predictions: Understanding Global Warming* (DK, 2008, 2015)

Of course, CO_2 is not the only greenhouse gas whose concentrations are rising due to human activity. A combination of agriculture (e.g., rice cultivation), livestock raising, and dam construction led to substantial increases in methane (CH_4) concentrations. Agricultural practices have also increased the concentration of nitrous oxide (N_2O).

Using air bubbles in ice cores, we can examine small bits of atmosphere trapped in ice, as it accumulated back in time, to reconstruct the composition of the ancient atmosphere, including the past concentrations of greenhouse gases. The ice core evidence shows that the rise over the past two centuries in the concentrations of the greenhouse gases mentioned above is unprecedented for at least the past 10,000 years. Longer-term evidence suggests that concentrations are higher now than they have been for hundreds of thousands of years, and perhaps several million years.

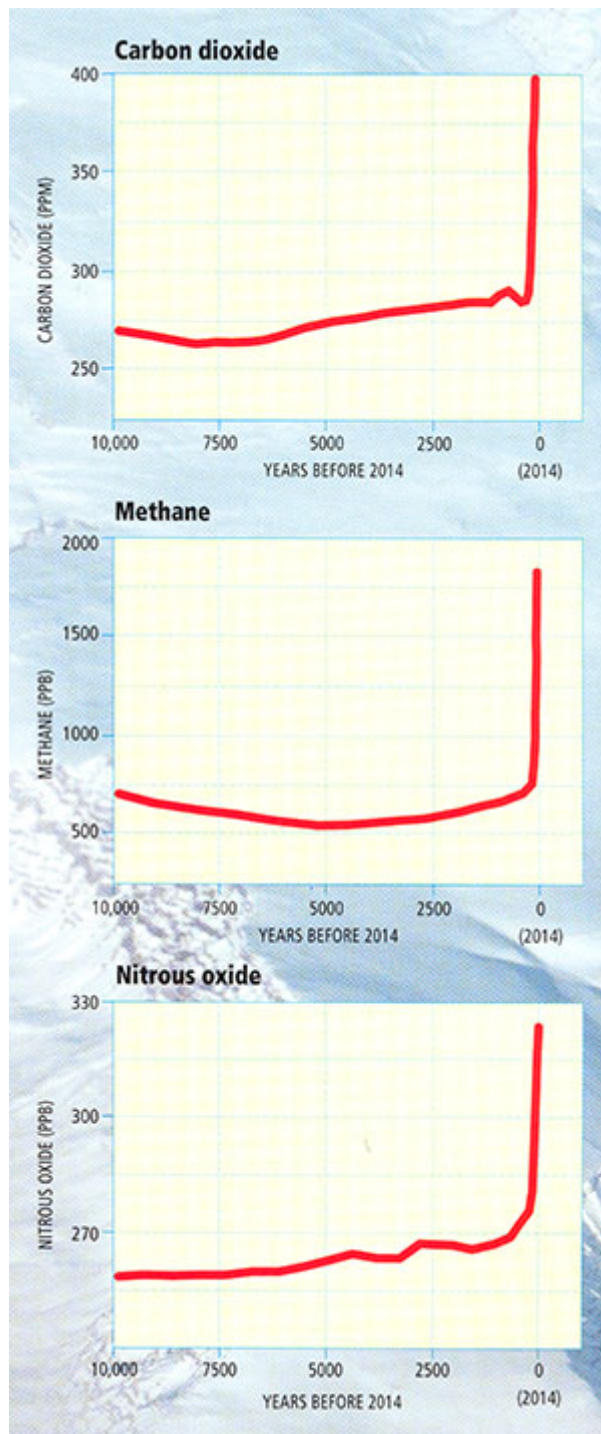


Figure 2.3: Changes in Greenhouse Gases Record in Ice Cores.

Credit: Mann & Kump, *Dire Predictions: Understanding Climate Change, 2nd Edition*
© 2015 Pearson Education, Inc.

Reading Assignment

Before we go any further, please read the following document:

- From the IPCC Fifth Assessment Report, Working Group 1 ^[1], read section B, Observed Changes in the Climate System (p. 4), and section B.1, Atmosphere (p. 5-8) of the Summary for Policy Makers ^[2].

As you read, please pay particular attention to:

- *the variety of sources of data,*
- *internal consistency among the various data streams with regard to our changing climate.*

Modern Surface Temperature Trends

Instrumental surface temperature measurements consisting of thermometer records from land-based stations, islands, and ship-board measurements of ocean surface temperatures provide us with more than a century of reasonably global estimates of surface temperature change. Some regions, like the Arctic and Antarctic, and large parts of South America, Africa, and Eurasia, were not very well sampled in earlier decades, but records in these regions become available as we move into the mid and late 20th century.

Temperature variations are typically measured in terms of *anomalies* relative to some *base period*. The animation below is taken from the NASA Goddard Institute for Space Studies ^[9] in New York (which happens to sit just above "Tom's Diner" ^[10] of *Seinfeld* fame), one of several scientific institutions that monitor global temperature changes. It portrays how temperatures around the globe have changed in various regions since the late 19th century. The temperature data have been averaged into 5 year blocks, and reflect variations relative to a 1950-1980 base period, i.e., *warm* regions are warmer than the 1950-1980 average, while *cold* regions are colder than the 1950-1980 average, by the magnitudes shown. You may note a number that appears in the upper right corner of the plot. That number indicates the average temperature anomaly over the entire globe at any given time, again, relative to the 1950-1980 average.

Take some time to explore the animation on your own. You may want to go through it several times so you can start to get a sense of just how rich and complex the patterns of surface temperature variations are. Do you see periodic intervals of warming and cooling in the eastern equatorial Pacific? What might that be? [We will talk about the phenomenon in upcoming lessons].

Take note of any particularly interesting patterns in space and time that you see as you review the animation. You can turn your sound off the first few times so you do not hear the annotation of the animation. Then, when you are ready, turn the sound on and you can

hear Michael Mann's take.

Credit: [NASA's Goddard Institute for Space Studies](#) ^[11]

We can average over the entire globe for any given year and get a single number, the global average temperature. Here is the curve we get if we plot out that quantity. Note that in the plot below the average temperature over the base period has been added to the anomalies, so that the estimate reflects the surface temperature of the Earth itself.

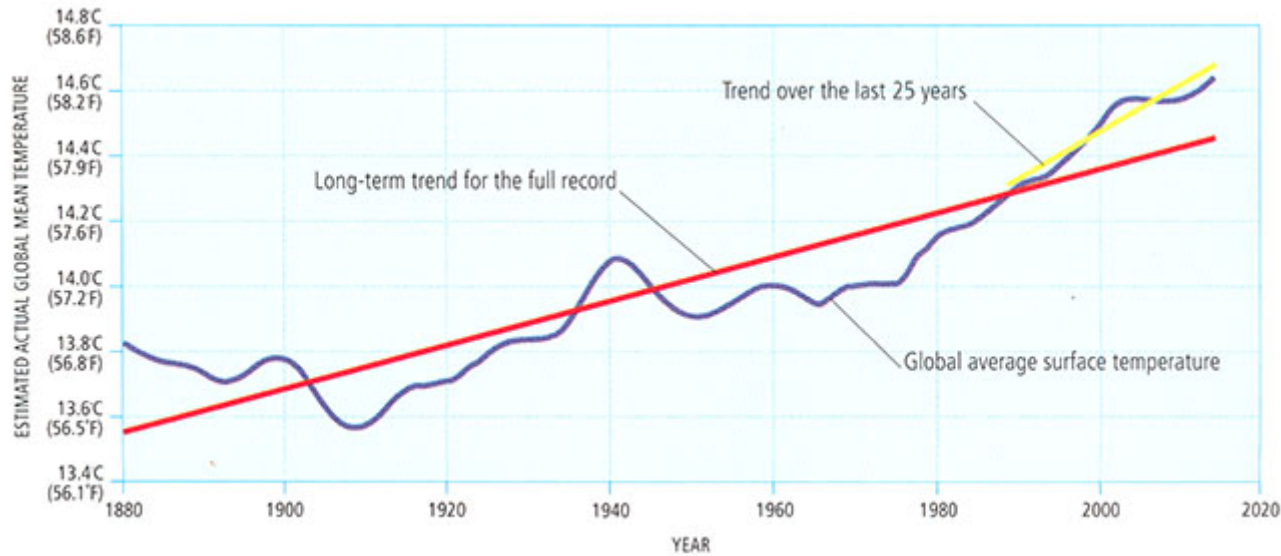


Figure 2.4: Trends in Global Average Surface Temp. 1860-2015.

Credit: Mann & Kump, *Dire Predictions: Understanding Climate Change*, 2nd Edition
© 2015 Pearson Education, Inc.

We can see that the Earth has warmed a little less than 1°C (about 1.5 F) since widespread records became available in the mid-19th century. That this warming has taken place is essentially incontrovertible from a scientific point of view. What is the *cause* of this warming? That is a more difficult question, which we will address later.

We discussed above the cooling that is evident in parts of the Northern Hemisphere (particularly over the land masses) from the 1940s-1970s. There was a time during the mid 1970s, when *some* scientists thought the globe might be entering into a long-term *cooling trend*. There was a reason to believe that might be the case. In the absence of other factors, changes in the Earth's orbital geometry did favor the descent (albeit a very slow one!) into the next ice age. Also, the buildup of atmospheric aerosols, which, as we will explore, can have a large regional cooling impact, favored cooling. Precisely how these cooling effects would balance out against the warming impact of greenhouse gases was not known at the time.

Some critics claim that if the scientific community thought we were entering into another Ice Age in the 1970s, why should we trust the scientists now about global warming? In fact, it was far from a scientific consensus [12] within the scientific community in the mid 1970s that we were headed into another Ice Age. Some scientists speculated this was possible, but the prevailing viewpoint was that increasing greenhouse gas concentrations and warming would likely win out.

We know that, indeed, the short term cooling trend for the Northern Hemisphere continents ended in the 1970s, and, since then, global warming has dominated over any cooling effects.

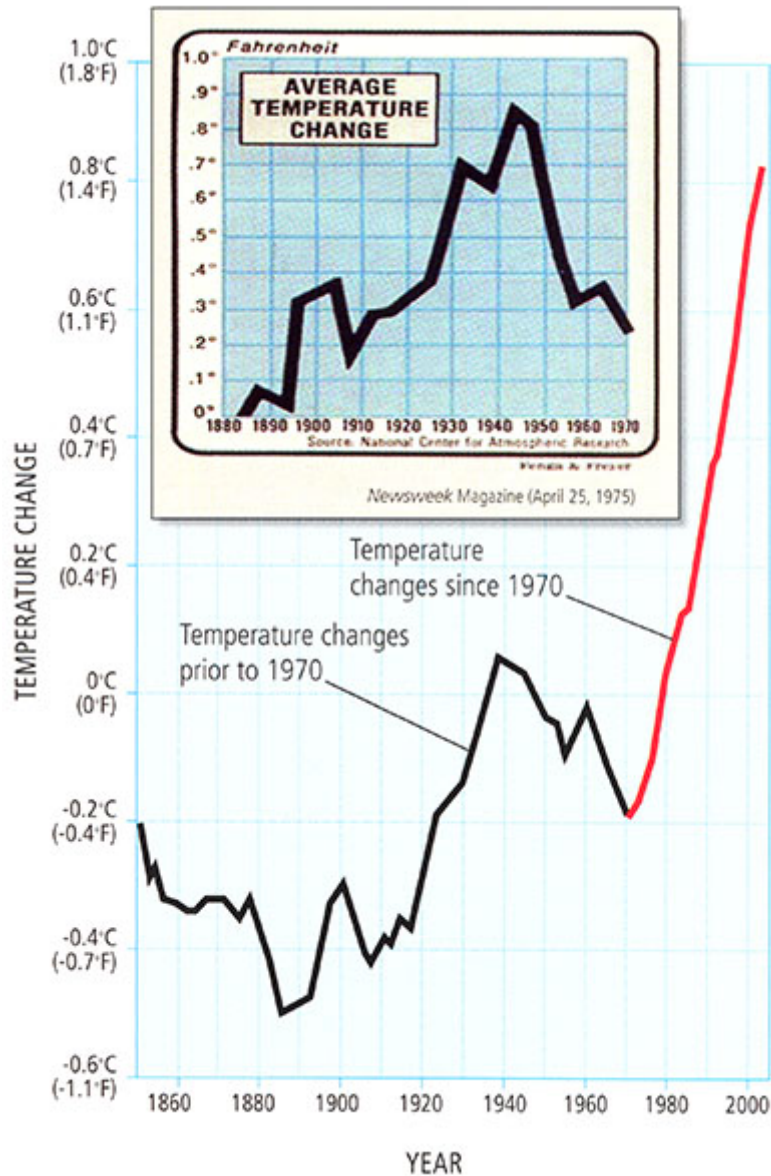


Figure 2.5: Northern Hemisphere Continental Temperature Trends.

Credit: Mann & Kump, *Dire Predictions: Understanding Climate Change*, 2nd Edition

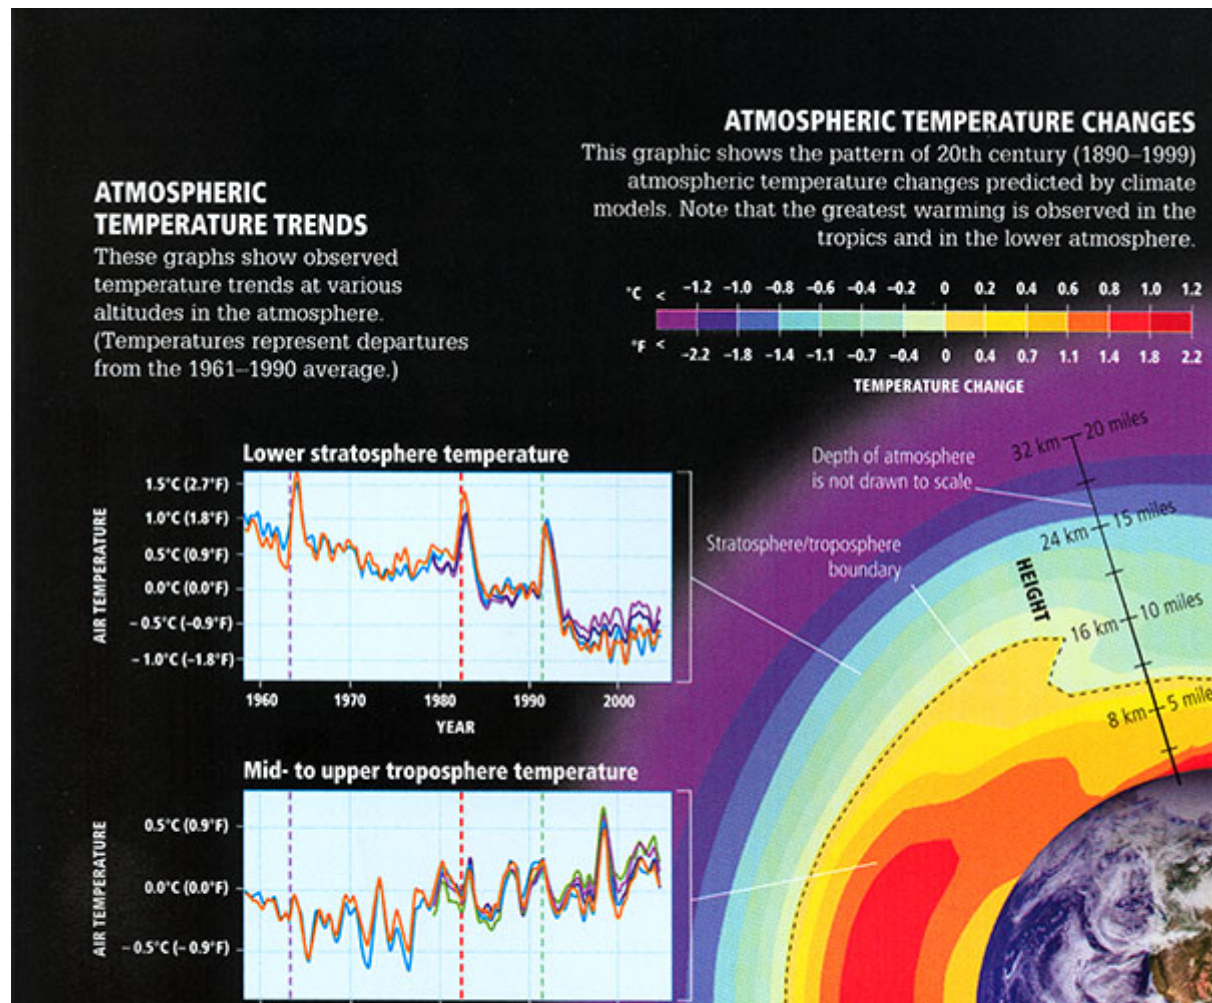
© 2015 Pearson Education, Inc.

As mentioned earlier, we cannot deduce the *cause* of the observed warming solely from the fact that the globe is warming. However, we can look for possible clues. Just like forensics experts, climate scientists refer to these clues as *fingerprints*. It turns out that

natural sources of warming give rise to different patterns of temperature change than human sources, such as increasing greenhouse gases. This is particularly true when we look at the *vertical pattern* of warming in the atmosphere. This is our next topic.

Vertical Temperature Trends

As alluded to previously, the vertical pattern of observed atmospheric temperature trends provides some important clues in establishing the underlying cause of the warming. While upper air temperature estimates (from weather balloons and satellite measurements) are only available for the latter half of the past century, they reveal a remarkable pattern. The lower part of the atmosphere — the *troposphere*, has been warming along with the surface. However, once we get into the *stratosphere*, the temperatures have actually been decreasing! As we will learn later when we focus on the problem of climate signal *fingerprinting*, certain forcings are consistent with such a vertical pattern of temperature changes, while other forcings are not.



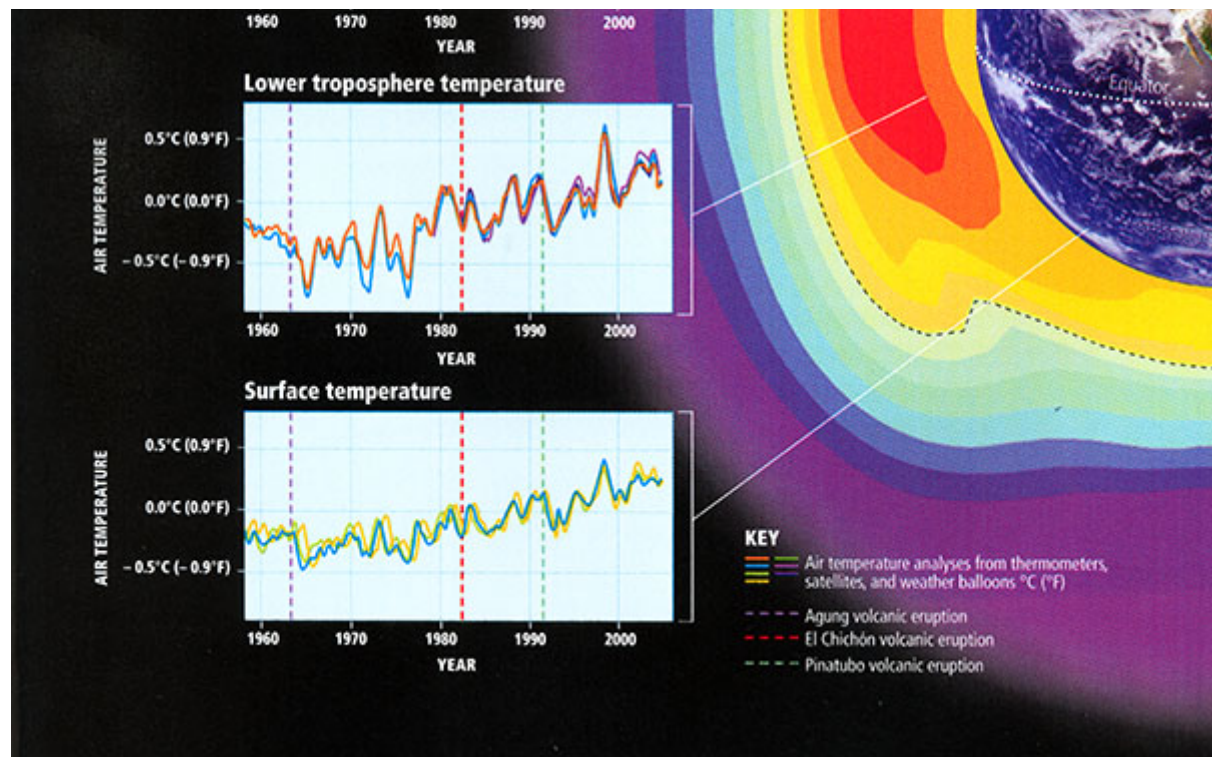


Figure 2.6: Recent Temperature Trends at Various Levels in the Atmosphere.

Credit: Mann & Kump, *Direr Predictions: Understanding Climate Change*, 2nd Edition
© 2015 Pearson Education Inc.

Think About It!

Care to venture a guess as to which forcing might be most consistent with this vertical pattern of temperature change?

Click for answer.

Historical Variations in Precipitation and Drought

Recall our discussion of the general circulation of the atmosphere ^[13] from Lesson #1.

There we learned that the circulation of the atmosphere is driven by the contrast in surface heating between the equator and the poles. That contrast results from the difference between incoming short wave solar heating and outgoing loss from the surface through various modes of energy transport including radiational heat loss as well as heat loss through convection and latent heat release through evaporation.

It, therefore, stands to reason that climate change — which in principle involves changing the balance between incoming and outgoing radiative loss via changes in the greenhouse effect — is likely to alter the circulation of the atmosphere itself, and thus, large-scale precipitation patterns. The observed changes in precipitation patterns are far *noisier* (very variable and difficult to interpret) than temperature changes, however. Regional effects related to topography (e.g., mountain ranges that force air upward leading to wet windward and dry leeward conditions), ocean-atmosphere heating contrasts that drive regional circulation patterns, such as monsoons, etc., lead to very heterogeneous patterns of changes in rainfall, in comparison with the pattern of surface temperature changes.

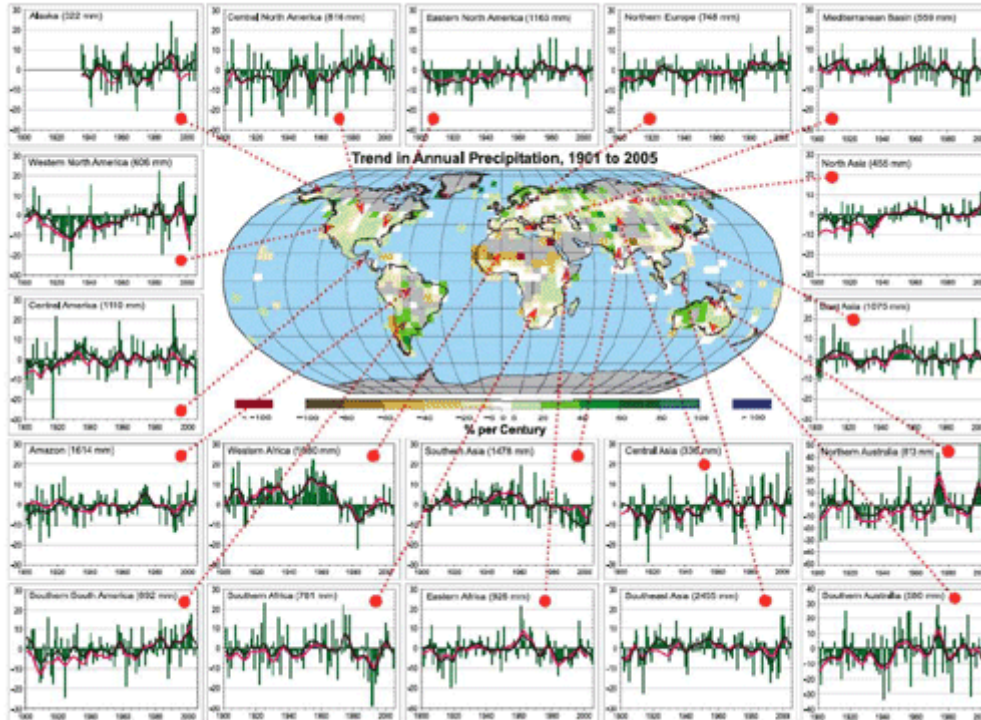


Figure 2.7: Trends in Annual Precipitation, 1901-2005 [\[Enlarge\]](#) ^[14].

Credit: IPCC Fourth Assessment Report, Chapter 3, Figure 3.14

We might expect certain reasonably simple patterns to emerge, nonetheless. As we shall see in a later lesson ^[15] looking at climate change projections, climate models predict that atmospheric circulation cells and storm tracks migrate poleward, shifting patterns of

rainfall between the equator and poles. The subtropics and middle latitudes tend to get dryer, while the sub-polar latitudes get wetter (primarily in winter). The equatorial region actually is predicted to get wetter, simply because the rising motion that occurs there squeezes out more rainfall from the warmer, moister lower atmosphere. If we average the observed precipitation changes in terms of trends in different latitudinal bands, we can see some evidence of the changes.

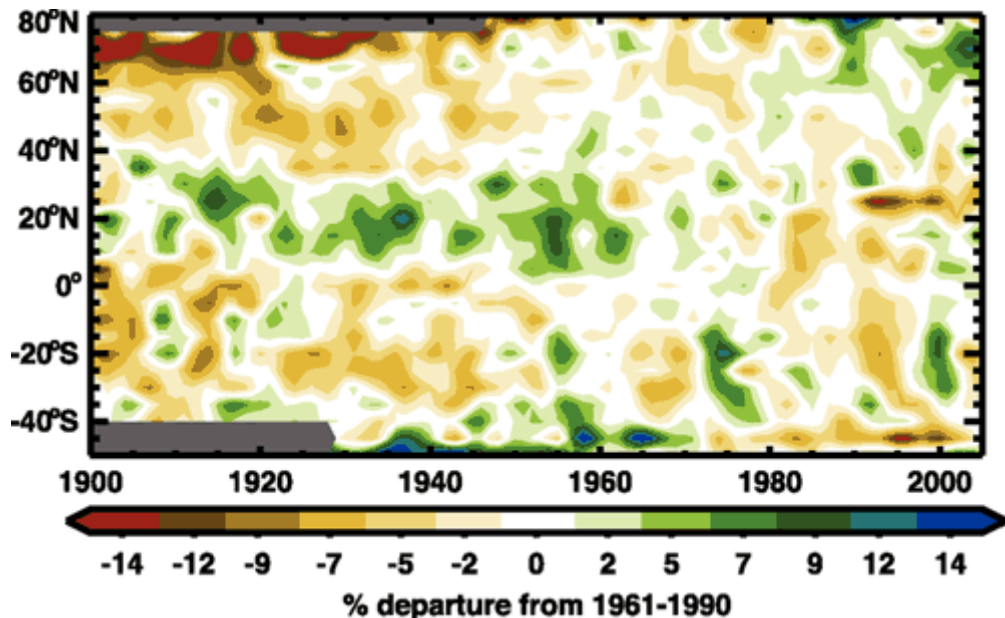


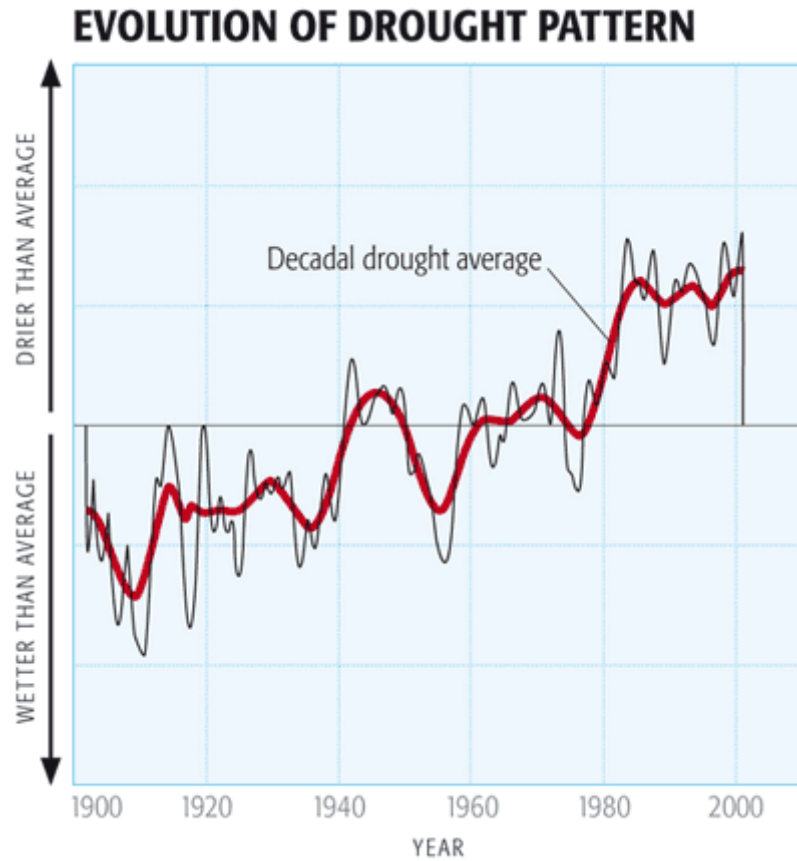
Figure 2.8: Changes over Time in Precipitation For Various Latitude Bands

Credit: IPCC Fourth Assessment Report, Chapter 3, Figure 3.15

For example, we see that over time the high northern latitudes (60-80N) are getting wetter, while the subtropical and middle latitudes of the Northern Hemisphere are getting dryer. However, there is a lot of variability from year to year, and from decade to decade, making it difficult to clearly discern whether the theoretically predicted changes are yet evident.

Drought, as we will see, does not simply follow rainfall changes. Rather, it reflects a combination of both rainfall and temperature influences. Decreased rainfall can lead to warmer ground temperatures, increased evaporation from the surface, decreased soil moisture, and thus drying.

Like rainfall, regional patterns of drought are complicated and influenced by a number of different factors. However, the combination of shifting rainfall patterns and increased evaporation has led to very pronounced increases in drought in subtropical regions, and even in many tropical and sub-polar regions, where rainfall shows little trend (as indicated in the earlier graphic) but warmer temperatures have led to decreased soil moisture. These broad trends are seen in measurements of the *Palmer Drought Severity Index* -- an index that combines the effects of changing rainfall and temperature to estimate soil moisture content; the more negative the index, the stronger the drought.



© 2009 Pearson Education, Inc.

Figure 2.9: Evolution of Drought Pattern.
Credit: Pearson, 2009

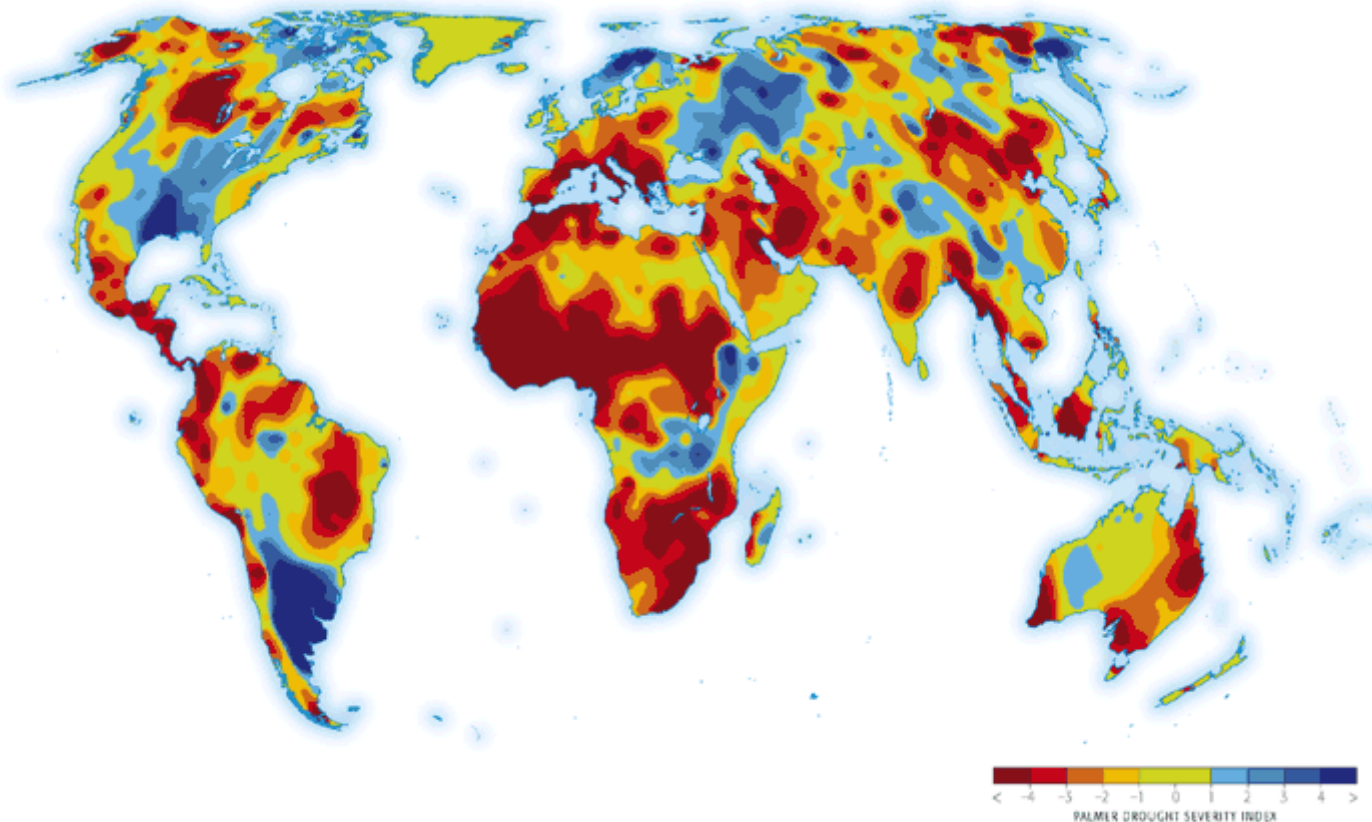
GLOBAL PATTERN OF DROUGHT, AS MEASURED BY THE PALMER DROUGHT SEVERITY INDEX

Figure 2.10: Global Pattern of Drought, as Measured by the Palmer Drought Severity Index.
Credit: Pearson, 2009

In the next lesson, we will assess evidence for changes in extreme weather events, such as heat waves, floods, tropical cyclone activity, etc. In the meantime, however, we are going to digress a bit and discuss the topic of how to analyze data for inferences into such matters as discerning whether or not trends are evident in particular data sets, and whether it is possible to establish a relationship between two or more different data sets.

Review of Basic Statistical Analysis Methods for Analyzing Data - Part 1

Now that we have looked at the basic data, we need to talk about how to *analyze* the data to make inferences about what they may tell us.

The sorts of questions we might want to answer are:

- Do the data indicate a *trend*?
- Is there an apparent *relationship* between two or more different data sets?

These sorts of questions may seem simple, but they are not. They require us, first of all, to introduce the concept of *hypothesis testing*.

To ask questions of a data set, one has to first formalize the question in a meaningful way. For example, if we want to know whether or not a data series, such as global average temperatures, display a trend, we need to think carefully about what it means to say that a data series has a trend!

This leads us to consider the concept of the *null hypothesis*. The null hypothesis states what we would expect purely from chance alone, in the absence of anything interesting (such as a trend) in the data. In many circumstances, the null hypothesis is that the data are the product of being randomly drawn from a *normal distribution*, what is often called a *bell curve*, or sometimes, a *Gaussian distribution* (after the great mathematician Carl Friedrich Gauss ^[16]):

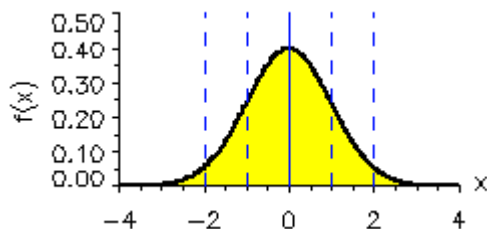


Figure 2.11: Gaussian Distribution.

In the normal distribution shown above, the average or *mean* of the data set has been set to zero (that is where the peak is centered), and the *standard deviation* (*s.d.*), a measure of the typical amplitude of the fluctuations, is set to one. If we draw random samples from such a distribution, then roughly 68% of the time the values will fall within 1 s.d. of the mean (in the above example, that is the range -1 to +1). That means that roughly 16% of the time the data will fall above 1 s.d., and roughly 16% of the time the data will fall below 1 s.d. About 95% of the time, the randomly drawn values will fall within 2 s.d. (i.e., the range -2 to +2 in the above example). That means only 2.5% of the time the data will fall above 2 s.d. and only 2.5% of the time below 2 s.d. For this reason, the 2 s.d. (or 2 *sigma*) range, is often used to characterize the region we are relatively confident the data should fall in, and the data that fall outside that range are candidates for potentially interesting anomalies.

Random Time Series

Here is an example of what a random data series of length $N = 200$ which we will call $\varepsilon(t)$, drawn from a simple normal distribution with mean zero and standard deviation one looks like (for example, you can think of this data set as a 200 year long temperature anomaly record).

$$Y_t = \varepsilon(t)$$

(1)

This sort of noise is called *white noise* because there is no particular preference for either higher-frequency or lower-frequency fluctuations. The fluctuations have equal amplitude.

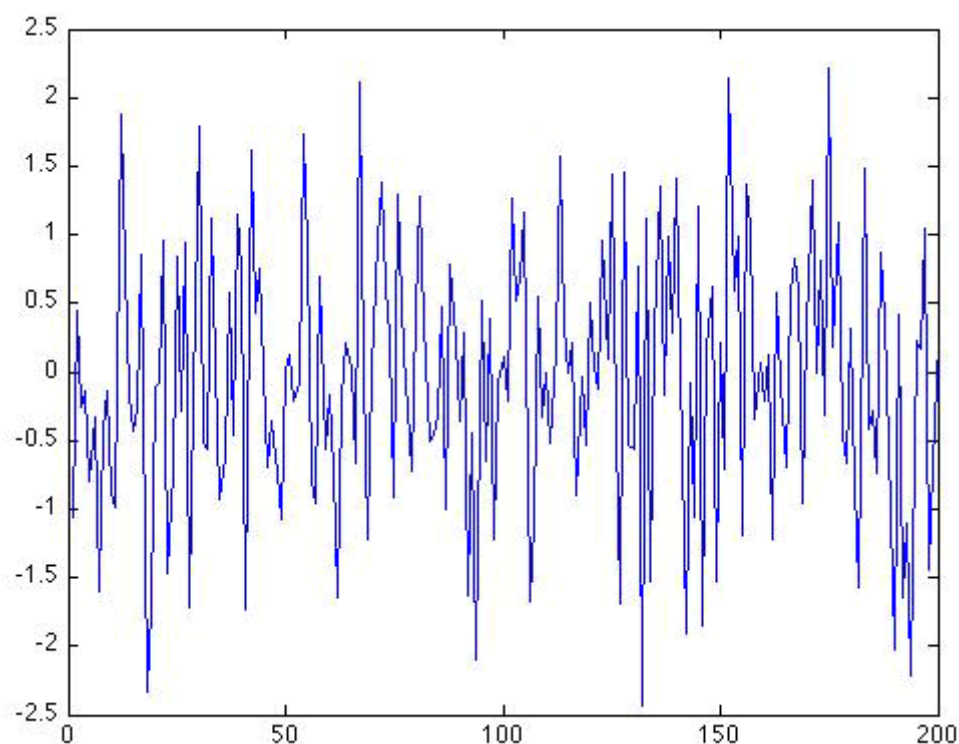


Figure 2.12(1). $N=200$ years of Gaussian White Noise.

There is another form of random noise, known as *red noise* because the long-term fluctuations have a greater relative magnitude than short-term fluctuations (just as red light is dominated by low-frequency visible wavelengths of light).

A simple model for Gaussian red noise takes the form

$$Y_t = \rho \cdot Y_{t-1} + \varepsilon(t)$$

(2)

where $\varepsilon(t)$ is Gaussian white noise. As you can see, a red noise process tends to integrate the white noise over time. It is this process of integration that leads to more long-term variation than would be expected for a pure white noise series. Visually, we can see that the variations from one year to the next are not nearly as erratic. This means that the data have fewer *degrees of freedom* (N') than there are actual data points (N). In fact, there is a simple formula relating N' and N :

$$N' = N \frac{1 - \rho}{1 + \rho}$$

(3)

The factor $(1 - \rho) / (1 + \rho)$ measures the "redness" of the noise. Let us consider again a random sequence of length $N = 200$ but this time it is "red" with the value $\rho = 0.6$. The same random white noise sequence used previously is used in equation 2 for $\varepsilon(t)$:

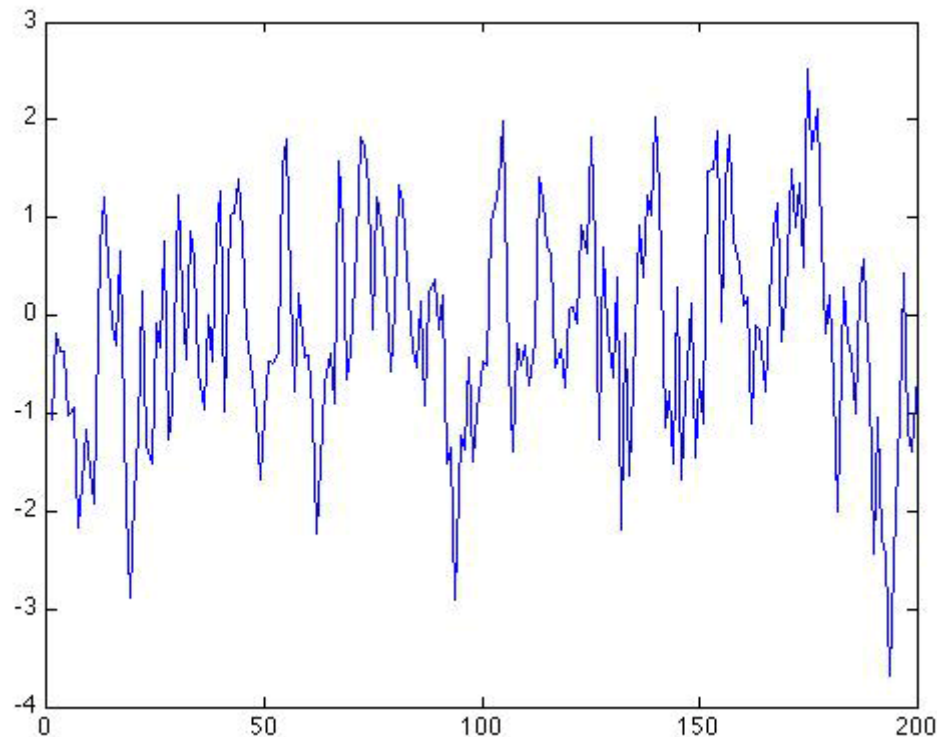


Figure 2.12(2): $N=200$ years of Gaussian 'red noise' with $\rho=0.6$

Self-Check

How many distinct peaks and troughs can you see in the series now?

Click for answer.

Self-Check

How many degrees of freedom N' are there in this series?

Click for answer.

As ρ gets larger and larger, and approaches one, the low-frequency fluctuations become larger and larger. In the limit where $\rho = 1$, we have what is known as a *random walk* or *Brownian motion*. Equation 2 in this case becomes just:

$$y_t = y_{t-1} + \varepsilon(t)$$

(4)

You might notice a problem when using equation 3 in this case. For $\rho = 1$, we have $N' = 0$! There are no longer any effective degrees of freedom in the time series. That might seem nonsensical. But there are other attributes that make this a rather odd case as well. The time series, it turns out, now has an infinite *standard deviation*!

Let's look at what our original time series looks like when we now use $\rho = 1$:

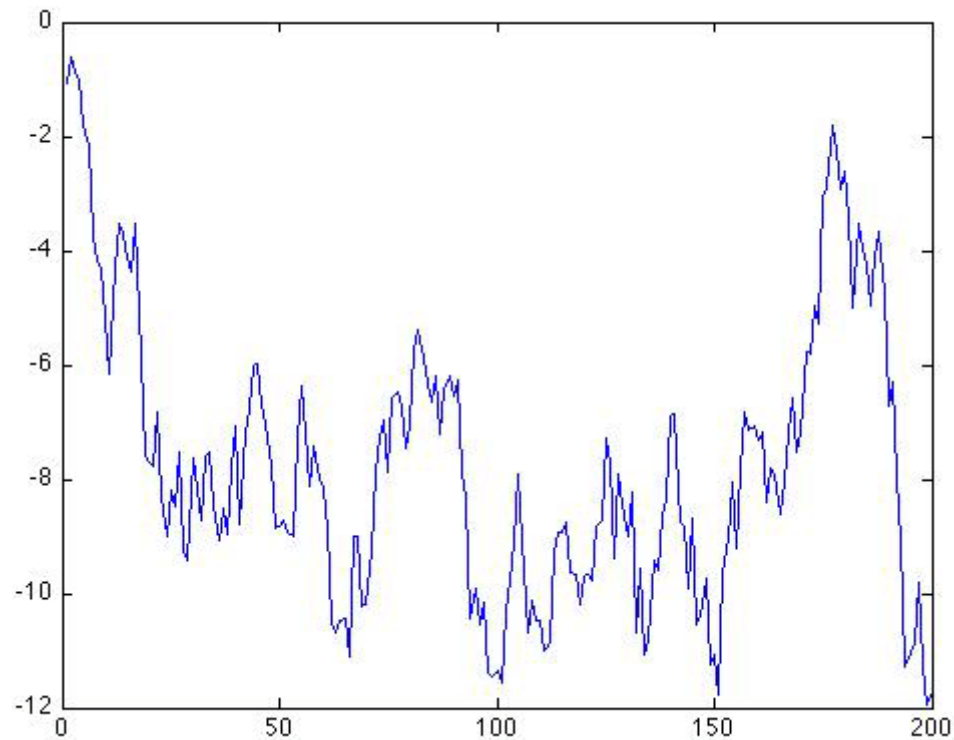


Figure 2.12(3): $N=200$ years of Gaussian 'red noise' with $\rho=1$, i.e., a 'random walk'.

As you can see, the series starts out in the same place, but immediately begins making increasingly large amplitude long-term excursions up and down. It might look as if the series wants to stay negative. But if we were to continue the series further, it would eventually oscillate erratically between increasingly large *negative and positive* swings. Let's extend the series out to $N = 1000$ values to see that:

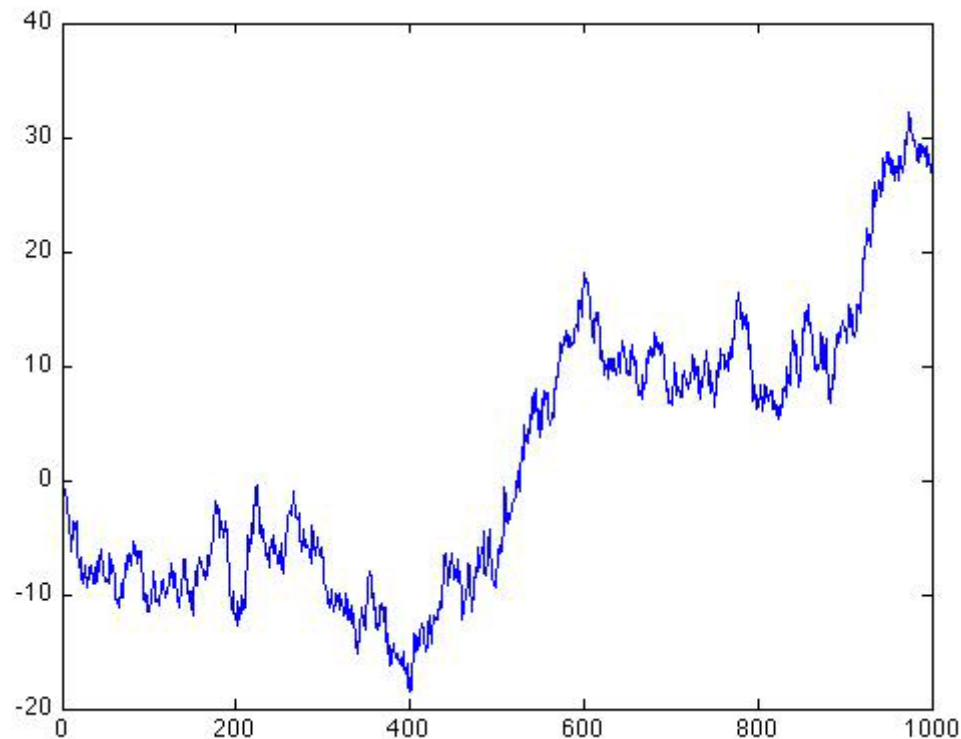


Figure 2.12(4): $N=1000$ years 'random walk'.

The swings are getting wider and wider, and they are occurring in both the positive and negative direction. Eventually, the amplitude of the swings will become arbitrarily large, i.e., infinite, even though the series will remain centered about a mean value of zero. This is an example of what we refer to in statistics as a *pathological* case.

Now let's look at what the original $N = 200$ long pure white noise series look like when there is a simple linear trend of 0.5 degree/century added on:

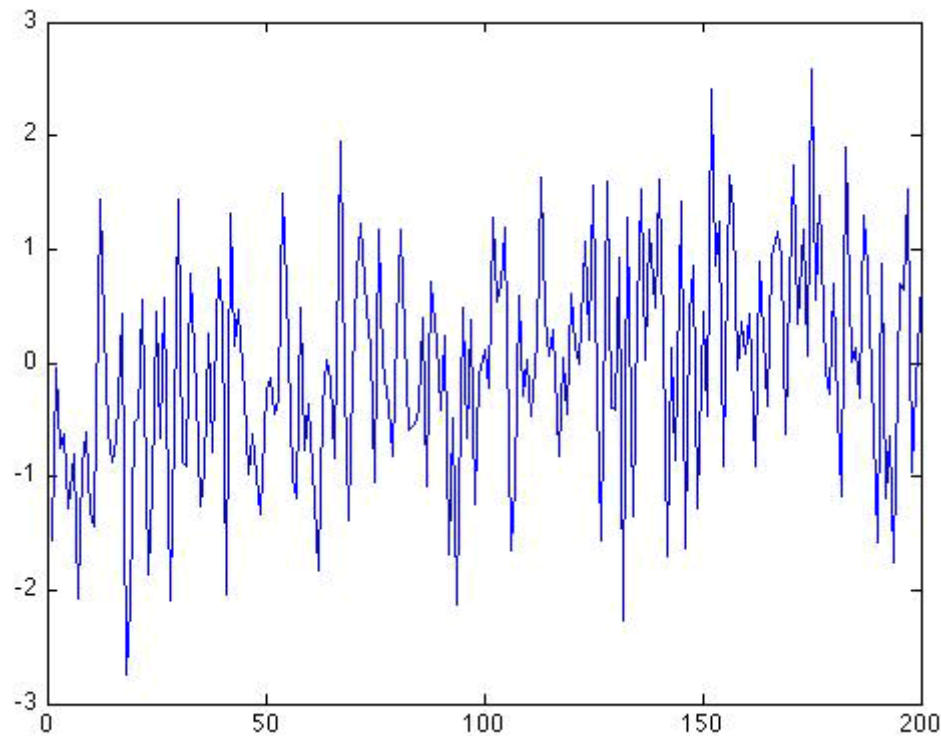


Figure 2.12(5). $N=200$ years of Gaussian White Noise with linear trend added.

Can you see a trend? In what direction? Is there a simple way to determine whether there is indeed a trend in the data that is distinguishable from random noise. That is our next topic.

Review of Basic Statistical Analysis Methods for Analyzing Data - Part 2

Establishing Trends

Various statistical hypothesis tests have been developed for exploring whether there is something more interesting in one or more data sets than would be expected from the chance fluctuations Gaussian noise. The simplest of these tests is known as *linear regression* or *ordinary least squares*. We will not go into very much detail about the underlying statistical foundations of the approach, but if you are looking for a decent tutorial ^[17], you can find it on Wikipedia.

The basic idea is that we test for an alternative hypothesis that posits a linear relationship between the *independent variable* (e.g., time, t in the past examples, but for purposes that will later become clear, we will call it x) and the *dependent variable* (i.e., the hypothetical temperature anomalies we have been looking at, but we will use the generic variable y).

The underlying statistical model for the data is:

$$y_i = a + b \cdot x_i + \varepsilon_i$$

(5)

where i ranges from 1 to N , a is the intercept of the linear relationship between y and x , b is the slope of that relationship, and ε is a random noise sequence. The simplest assumption is that ε is Gaussian white noise, but we will be forced to relax that assumption at times.

Linear regression determines the best fit values of a and b to the given data by minimizing the *sum of the squared differences* between the observations y and the values predicted by the linear model $\hat{y} = a + bx$. The *residuals* are our estimate of the variation in the data that is not accounted for by the linear relationship, and are defined by

$$\varepsilon_i = y_i - \hat{y}_i$$

(6)

For simple linear regression, i.e., *ordinary least squares*, the estimates of a and b are readily obtained:

$$b = \frac{[N \cdot \Sigma y_i x_i - \Sigma y_i \cdot \Sigma x_i]}{[N \cdot \Sigma x_i^2 - (\Sigma x_i)^2]}$$

(7)

and

$$a = \left(1/N\right) \cdot \Sigma y_i - \frac{b}{N \cdot \Sigma x_i}$$

(8)

The parameter we are most interested in is b , since this is what determines whether or not there is a significant linear relationship between y and x .

The sampling uncertainty in b can also be readily obtained:

$$\sigma_b = \frac{\text{std}(\varepsilon)}{\left[\Sigma(x_i - \mu(x))^2 \right]^{\frac{1}{2}}}$$

(9)

where $\text{std}(\varepsilon)$ is standard deviation of ε and μ is the mean of x . A statistically significant trend amounts to the finding that b is significantly different from zero. The 95% confidence range for b is given by $b \pm 2 \sigma_b$. If this interval does not cross zero, then one can conclude that b is significantly different from zero. We can alternatively measure the significance in terms of the *linear correlation coefficient*, r , between the independent and dependent variables which is related to b through

$$r = b \cdot \frac{\text{std}(x)}{\text{std}(y)}$$

(10)

r is readily calculated directly from the data:

$$r = \frac{(1/N-1) \cdot \Sigma(x - \bar{x})(y - \bar{y})}{\text{std}(x) \cdot \text{std}(y)}$$

(11)

where over-bar indicated the mean. Unlike b , which has dimensions (e.g., °C per year in the case where y is temperature and x is time), r is conveniently a dimensionless number whose absolute value is between 0 and 1. The larger the value of r (either positive or negative), the more significant is the trend. In fact, the square of r (r^2) is a measure of the fraction of variation in the data that is accounted for by the trend.

We measure the significance of any detected trends in terms of a p -value. The p -value is an estimate of the probability that we would wrongly reject the *null hypothesis* that there is no trend in the data in favor of the *alternative hypothesis* that there is a linear trend in the data — the *signal* that we are searching for in this case. Therefore, the smaller the p value, the less likely that you would

observe as large a trend as is found in the data from random fluctuations alone. By convention, one often requires that $p < 0.05$ to conclude that there is a significant trend (i.e., that only 5% of the time should such a trend have occurred from chance alone), but that is not a magic number.

The choice of p in statistical hypothesis testing represents a balance between the acceptable level of *false positives* vs. *false negatives*. In terms of our example, a false positive would be detecting a statistically significant trend, when, in fact, there is no trend; a false negative would be concluding that there is no statistically significant trend, when, in fact, there is a trend. A lower threshold (that is, higher p -value, e.g., $p = 0.10$) makes it more likely to detect a real but weak signal, but also more likely to falsely conclude that there is a real trend when there is not. Conversely, a higher threshold (that is, lower p -value, e.g., $p = 0.01$) makes false positives less likely, but also makes it less likely to detect a weak but real signal.

There are a few other important considerations. There are often two different alternative hypotheses that might be invoked. In this case, if there is a trend in the data, who is to say whether it should be positive ($b > 0$) or negative ($b < 0$)? In some cases, we might want only to know whether or not there is a trend, and we do not care what sign it has. We would then be invoking a *two-sided hypothesis*: is the slope b large enough in magnitude to conclude that it is significantly different from zero (whether positive or negative)? We would obtain a p -value based on the assumption of a two-sided hypothesis test. On the other hand, suppose we were testing the hypothesis that temperatures were warming due to increased greenhouse gas concentrations. In that case, we would reject a negative trend as being unphysical — inconsistent with our *a priori* understanding that increased greenhouse gas concentrations should lead to significant *warming*. In this case, we would be invoking a one-sided hypothesis. The results of a one-sided test will double the significance compared with the corresponding two-sided test, because we are throwing out as unphysical half of the random events (chance negative trends). So, if we obtain, for a given value of b (or r) a p -value of $p = 0.1$ for the two-sided test, then the p -value would be $p = 0.05$ for the corresponding one-sided test.

There is a nice [online calculator](#) [18], courtesy of Vassar college, for obtaining a p -value (both one-sided and two-sided) given the linear correlation coefficient, r , and the length of the data series, N . There is still one catch, however. If the residual series ϵ of equation 6 contains *autocorrelation*, then we have to correct the degrees of freedom, N' , which is less than the nominal number of data points, N . The correction can be made, at least approximately in many instances, using the *lag-one autocorrelation coefficient*. This is simply the linear correlation coefficient, r_1 , between ϵ , and a carbon copy of ϵ lagged by one time step. In fact, r_1 provides an approximation to the parameter ρ introduced in equation 2. If r_1 is found to be positive and statistically significant (this can be checked using the online link provided above), then we can conclude that there is a statistically significant level of *autocorrelation* in our residuals, which must be corrected for. For a series of length $N = 100$, using a one-sided significant criterion of $p = 0.05$, we would need $r_1 > 0.17$ to conclude that there is significant autocorrelation in our residuals.

Fortunately, the fix is very simple. If we find a positive and statistically significant value of r_1 , then we can use the same significance criterion for our trend analysis described earlier, except we have to evaluate the significance of the value of r for our linear regression analysis (not to be confused with the autocorrelation of residuals r_1) using a reduced, effective degrees of freedom N' , rather than the nominal sample size N . Moreover, N' is none other than the N' given earlier in equation 3 where we equate

$$\rho = r_1$$

That's about it for **ordinary least squares (OLS)**, the main statistical tool we will use in this course. Later, we will encounter the more complicated case where there may be *multiple independent variables*. For the time being, however, let us consider the problem of trend analysis, returning to the synthetic data series discussed earlier. We will continue to imagine that the dependent variable (y) is temperature T in $^{\circ}\text{C}$ and the independent variable (x) is time t in years.

First, let us calculate the trend in the original Gaussian white noise series of length $N = 200$ shown in Figure 2.12(1). The linear trend is shown below:

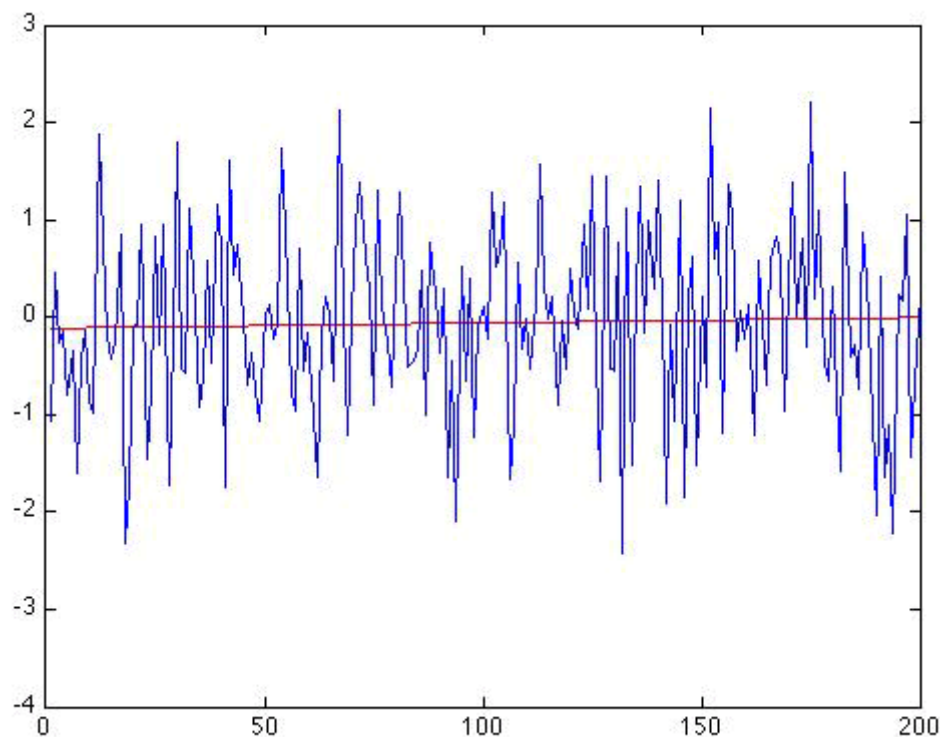


Figure 2.12(6): $N=200$ years of Gaussian White Noise with linear trend shown.

The trend line is given by: $\hat{T} = 0.0006 \cdot t - 0.1140$, and the regression gives $r = 0.0332$. So there is an apparent positive warming trend of 0.0006°C per year, or alternatively, 0.06°C per century. Is that statistically significant? It does not sound very impressive, does it? And that r looks pretty small! But let us be rigorous about this. We have $N = 200$, and if we use the online calculator link provided above, we get a p -value of 0.64 for the (default) two-sided hypothesis. That is huge, implying that we would be foolish in this case to reject the null hypothesis of no trend. But, you might say, we were looking for *warming*, so we should use a one-sided hypothesis. That halves the p -value to 0.32. But that is still a far cry from even the least stringent (e.g., $p = 0.10$) thresholds for significance. It is clear that there is no reason to reject the null hypothesis that this is a random time series with no real trend.

Next, let us consider the red noise series of length $N = 200$ shown earlier in Figure 2.12(2).

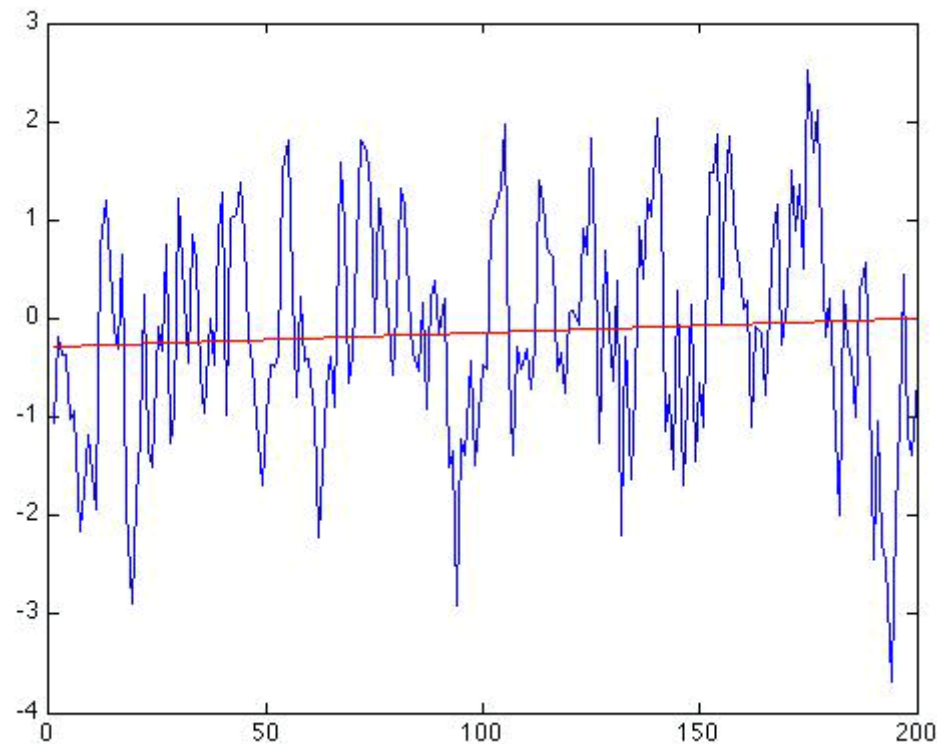


Figure 2.12(7). $N=200$ years of Gaussian 'red noise' with $\rho=0.6$ with linear trend shown.

As it happens, the trend this time appears nominally greater. The trend line is now given by: $\hat{T} = 0.0014 \cdot t - 0.2875$, and the regression gives $r = 0.0742$. So, there is an apparent positive warming trend of 0.14 degrees C per century. That might not seem entirely negligible. And for $N = 200$ and using a one-sided hypothesis test, $r = 0.0742$ is statistically significant at the $p = 0.148$ level according to the online calculator. That does not breach the typical threshold for significance, but it does suggest a pretty high likelihood (15% chance) that we would err by not rejecting the null hypothesis. At this point, you might be puzzled. After all, we did not put any trend into this series! It is simply a random realization of a red noise process.

Self Check

So why might the regression analysis be leading us astray this time?

Click for answer.

The problem is that our residuals are not uncorrelated. They are red noise. In fact, the residuals look a lot like the original series itself:

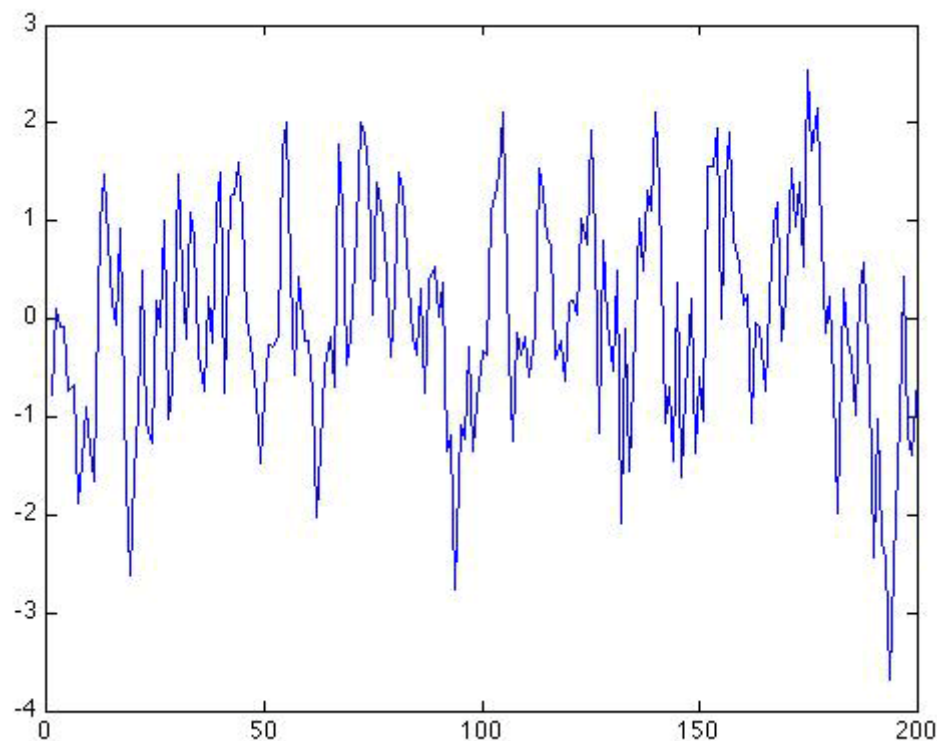


Figure 2.12(8). Residuals from linear regression with $N=200$ years of Gaussian 'red noise' with $\rho=0.6$

This is hardly coincidental; after all, the trend only accounts for $r^2 = 0.0742^2 = 0.0055$, i.e., only about half a percent, of the variation in the data. So 99.5% of the variation in the data is still left behind in the residuals. If we calculate the lag-one autocorrelation for the residual series, we get $r_1 = 0.54$. That is, again not coincidentally, very close to the value of $\rho = 0.6$ we *know* that we used in generating this series in the first place.

How do we determine if this autocorrelation coefficient is statistically significant? Well, we can treat it like it were a correlation coefficient. The only catch is that we have to use $N-1$ in place of N , because there are only $N-1$ values in the series when we offset it by one time step to form the lagged series required to estimate a lag-one autocorrelation.

Self Check

Should we use a one-sided or two-sided hypothesis test?

Click for answer.

If we use the online link and calculate the statistical significance of $r_1 = 0.54$ with $N-1 = 199$, we find that it is statistically significant at $p < 0.001$. So, clearly, we cannot ignore it. We have to take it into account.

So, in fact, we have to treat the correlation from the regression $r = 0.074$ as if it has $N' = (1 - 0.54) / (1 + 0.54) 200 = 0.30 (200) = 59.9 \approx 60$ degrees of freedom, rather than the nominal $N = 200$ degrees of freedom. Using the interactive online calculator, and replacing $N = 200$ with the value $N' = 60$, we now find that a correlation of $r = 0.074$ is only significant at the $p = 0.57$ ($p = 0.29$) for a two-sided (one-sided) test, hardly a level of significance that would cause us to seriously call into doubt the null hypothesis.

At this point, you might be getting a bit exasperated. When, if ever, *can* we conclude there is a trend? Well, why don't we now consider the case where we know we added a real trend in with the noise, i.e., the example of Figure 2.12(5) where we added a trend of $0.5^\circ\text{C}/\text{century}$ to the Gaussian white noise. If we apply our linear regression machinery to this example, we do detect a notable trend:

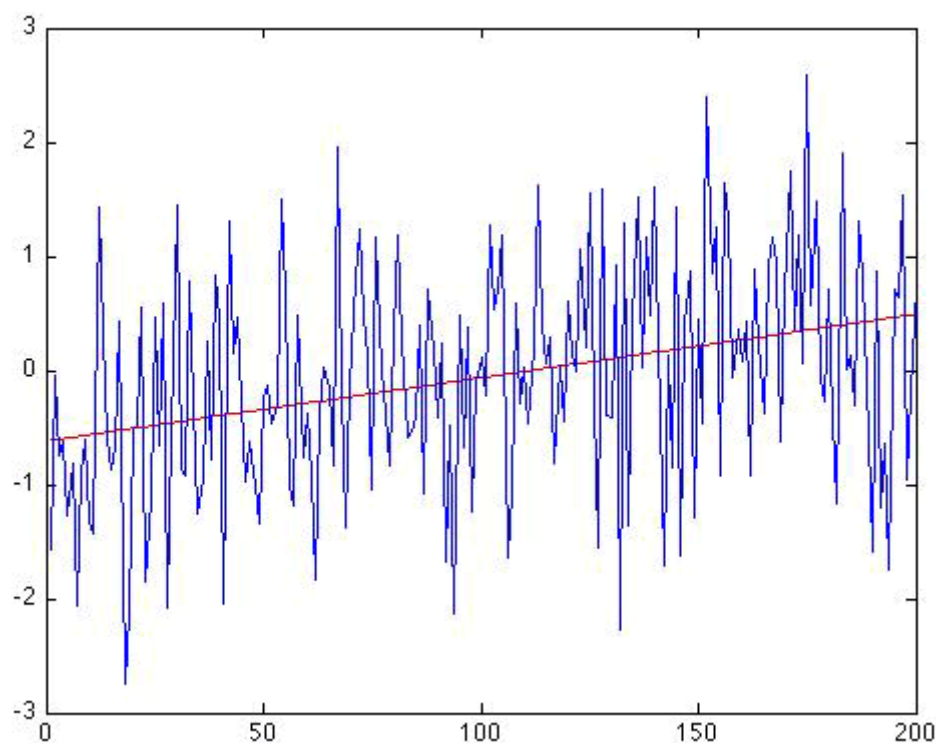


Figure 2.12(9). $N=200$ years of Gaussian white noise with added linear trend of 0.5 degrees/century; the red line shows trend recovered by the linear regression.

Now, *that's* a trend - your eye isn't fooling you. The trend line is given by: $\hat{T} = 0.0056 \cdot t - 0.619$. So there is an apparent positive warming trend of 0.56 °C per century (the 95% uncertainty range that we get for b , i.e., the range $b \pm 2 \sigma_b$, gives a slope anywhere between 0.32 and 0.79 °C per century, which of course includes the true trend (0.5 °C/century) that we know we originally put in to the series!). The regression gives $r = 0.320$. For $N = 200$ and using a one-sided hypothesis test, $r = 0.320$ is statistically significant at $p < 0.001$ level. And if we calculate the autocorrelation in the residuals, we actually get a small *negative* value ($r_1 = -0.095$), so autocorrelation of the residuals is not an issue.

Finally, let's look at what happens when the same trend (0.5 °C/century) is added to the random red noise series of Figure 2.12(2), rather than the white noise series of Figure 2.12(1). What result does the regression analysis give now?

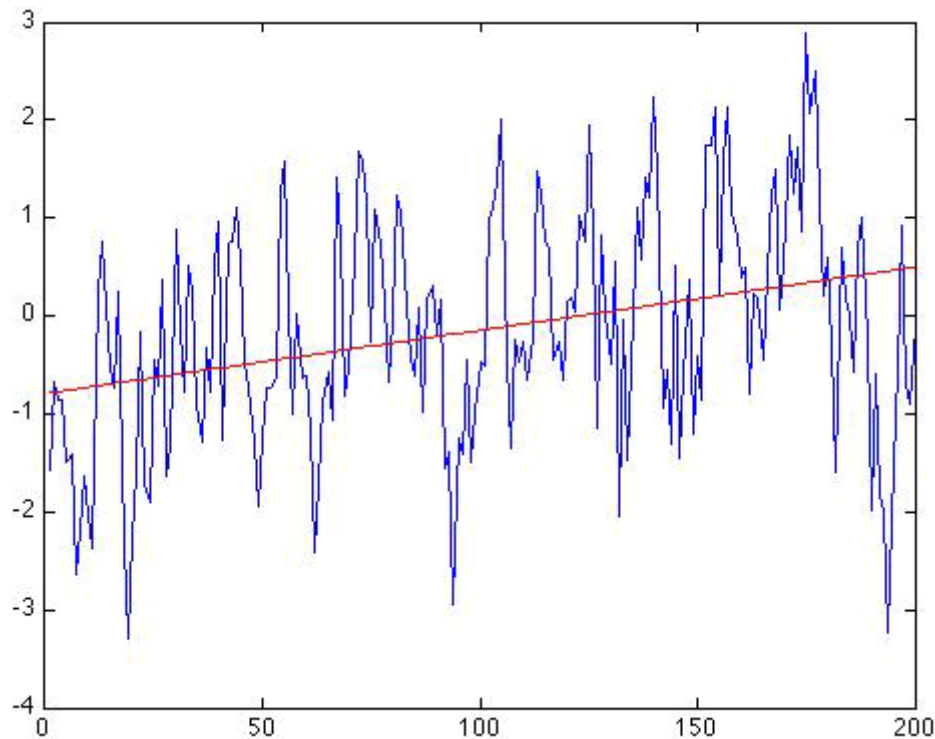


Figure 2.12(10). $N=200$ years of Gaussian red noise with $\rho=0.6$ and added linear trend of 0.5 degrees/century; the red line shows trend recovered by the linear regression.

We still recover a similar trend, although it's a bit too large. We know that the true trend is 0.5 degrees/century, but the regression gives: $\hat{T} = 0.0064 \cdot t - 0.793$. So, there is an apparent positive warming trend of 0.64 °C per century. The nominal 95%

uncertainty range that we get for b is 0.37 to 0.92 °C per century, which again includes the true trend (0.5 degrees C/century). The regression gives $r = 0.315$. For $N = 200$ and using a one-sided hypothesis test, $r = 0.315$ is statistically significant at the $p < 0.001$. So, are we done?

Not quite. This time, it is obvious that the residuals will have autocorrelation, and indeed we have that $r_1 = 0.539$, statistically significant at $p < 0.001$. So, we will have to use the reduced degrees of freedom N' . We have already calculated N' earlier for $\rho = 0.54$, and it is roughly $N' = 60$. Using the online calculator, we now find that the one-sided $p = 0.007$, i.e., roughly $p = 0.01$, which corresponds to a 99% significance level. So, the trend is still found to be statistically significant, but the significance is no longer at the astronomical level it was when the residuals were uncorrelated white noise. The effect of the "redness" of the noise has been to make the trend less statistically significant, because it is much easier for red noise to have produced a spurious apparent trend from random chance alone. The 95% confidence interval for b also needs to be adjusted to take into account the autocorrelation, though just how to do that is beyond the scope of this course.

Often, residuals have so *much* additional structure — what is sometimes referred to as *heteroscedasticity* (how's that for a mouthful?) — that the assumption of simple autocorrelation is itself not adequate. In this case, the basic assumptions of linear regression are called into question and any results regarding trend estimates, statistical significance, etc., are suspect. In this case, more sophisticated methods that are beyond the scope of this course are required.

Now, let us look at some *real temperature data*! We will use our very own custom online [Linear Regression Tool](#) ^[19] written for this course. The demonstration how to use this tool has been recorded in three parts below (click each link to open a new window and then the arrow to begin the demonstration):

[Part 1](#) ^[20]

[Part 2](#) ^[21]

[Part 3](#) ^[22]

You can play around with the temperature data set used in this example using the [Linear Regression Tool](#) ^[19]

Review of Basic Statistical Analysis Methods for Analyzing Data - Part 3

Establishing Relationships Between Two Variables

Another important application of OLS is the comparison of two different data sets. In this case, we can think of one of the time series as constituting the independent variable x and the other constituting the independent variable y . The methods that we discussed in the previous section for estimating trends in a time series generalize readily, except our predictor is no longer time, but rather, some

variable. Note that the correction for autocorrelation is actually somewhat more complicated in this case, and the details are beyond the scope of this course. As a general rule, *even* if the residuals show substantial autocorrelation, the required correction to the statistical degrees of freedom (N'), will be small as long as either one of the two time series being compared has low autocorrelation. Nonetheless, any substantial structure in the residuals remains a cause for concern regarding the reliability of the regression results.

We will investigate this sort of application of OLS with an example, where our independent variable is a measure of El Niño — the so-called *Niño 3.4 index* — and our dependent variable is December average temperatures in State College, PA.

The demonstration is given in three parts below (click each link to open a new window and then the arrow to begin the demonstration):

[Demonstration - Part 1](#) [23]

[Demonstration - Part 2](#) [24]

[Demonstration - Part 3](#) [25]

You can play around with the data set used in this example using this link: [Explore Using the File testdata.txt](#) [26]

Problem Set #1

Activity: Statistical Analysis of Climate Data

NOTE: For this assignment, you will need to record your work on a word processing document. Your work must be submitted in Word (.doc or .docx), or PDF (.pdf) format.

For this activity, you will use the application below to perform basic statistical analyses of climate data. The data we will use are global temperature anomalies and Niño 3.4 index, both measured in °C. You need to:

- determine historical trends in global temperatures and determine if there has been an increase in the trend.
- analyze the influence of El Niño on global temperatures.

[Link to Linear Regression Tool](#) [27]

Directions

1. First, save the [Problem Set #1 Worksheet](#) [28] to your computer. You will use this word processing document to electronically record your work in the remaining steps.
 - Save the worksheet to your computer by right-clicking on the link above and selecting "Save link as..."

- o The worksheet is in Microsoft Word format. You can use either Word or Google Docs (free) to work on this assignment. You will submit your worksheet at the end of the activity, so it must be in Word (.doc or .docx) or PDF (.pdf) format so the instructor can open it.
 - o Please show your work! When you are explicitly asked to create plots in a question, please cut-and-paste graphics and the output from the screen (e.g., by first printing the output to a pdf file and then directly inserting into the worksheet) to submit along with your discussion and conclusions.
- 2. Use the plotting tool to create a line plot of global temperature anomalies vs. time over the full 160 year period (1850 - 2009). Determine the basic statistics of the time series, i.e., the mean and standard deviation. If the mean is not zero, can you guess the 30 year base period that was used to calculate the anomalies? To do that, note that by definition, the mean of the base period used to calculate anomalies should be zero. Visually determine the year where the temperature anomalies graph crosses zero. Zoom in on different 30-year intervals around that year and use Viewport Data under the Statistics tab to view the mean for each selected 30-year sub-set of data, until you find a period for which the mean is closest to zero.
- 3. Evaluate the linear trend in global temperature anomalies over the full 160 year period (1850 - 2009) following the steps below. (A) Add a trend line to the plot of global temperature anomalies vs. time over the full 160 year period: use the Trend Lines tab. Determine the slope of the linear regression line, b , in $^{\circ}\text{C}/\text{century}$ and the correlation coefficient, r . As a side exercise, calculate the overall warming trend in temperature over the 1850-2009 period by multiplying the slope by the number of years in the period. (B) Assess statistical significance of the linear trend in global temperature anomalies without considering autocorrelation: use the [online Statistical Calculator tool](#) ^[18] from Lesson 2, Statistical Analysis Part 2 to calculate the p-value for the number of samples, N , and the correlation coefficient, r , from part A. Interpret the p-value. Use the standard error of the slope, S_b , to calculate 95% confidence interval for the slope as $b \pm 2S_b$, and report the calculated warming range in $^{\circ}\text{C}/\text{century}$. (C) Determine lag-one autocorrelation coefficient for the residuals, ρ . To do that, run the regression model using the Regression Model tab: select Model Parameters = year, Target Observation = Temp Anom. Plot the residuals: in the Plot Settings tab, select year for X and Model Residuals for Y. Check whether the autocorrelation is statistically significant using the Statistical Calculator. Remember that lag-one autocorrelation is simply correlation between the original data and an exact copy of the data but shifted by one time step, so the number of samples, N , is decreased by one; and the correlation coefficient, r , should be substituted by autocorrelation coefficient, ρ . Interpret the p-value. (D) Reassess statistical significance of the linear trend in the global temperature anomalies, taking into account autocorrelation of the residuals. Remember that in the presence of a significant autocorrelation, the actual number of samples, N , must be replaced by the degrees of freedom, N' : use formula (2) from Lesson 2 to calculate N' . Use the Statistical Calculator to calculate the p-value for $N = N'$ and r from part (A). Interpret the p-value. (E) Look closely at the plot of the residuals that you created in part (C). Do you see evidence of heteroscedasticity (additional structure superimposed on the random fluctuations)? Do you think that the hypothesis of a simple linear warming trend in this data series is appropriate or not?
- 4. Test the hypothesis that there is a difference in trend over two sub-periods: the first 130 years and the final 30 years. Follow the steps below. (A) Use the Plot Settings tab to plot the entire data series and zoom in to select 1850-1979 sub-period. Select Viewport in the Trend Lines tab to perform linear regression for the sub-period. Determine the slope of the linear regression line, b , in $^{\circ}\text{C}/\text{century}$ and the correlation coefficient, r . Use the standard error of the slope, S_b , to calculate 95% confidence interval for the slope in $^{\circ}\text{C}/\text{century}$. (B) Repeat the analysis for the 1980-2009 sub-period. (C) Determine whether trends are statistically different: you need to check whether the 95% confidence intervals for the two sub-periods overlap. Based on your results, has global warming accelerated over the past 30 years? What important caveat about the 95% confidence intervals was not taken into account in our analysis?

5. Is there a statistically significant influence of El Niño on global temperatures? (A) Use the Plot Settings tab to plot global temperature anomalies (Temp Anom) over time (Plot #1) and then Niño 3.4 index (Niño) over time (Plot #2) on the same plot. Determine the number of years, N , in the time interval over which the two data series overlap. This is the number of samples you will use in the analyses below. (B) Plot the relationship between Niño 3.4 index and global temperature anomaly: use the Plot Settings tab to plot Niño 3.4 index (Niño) on the X-axis and global temperature anomalies (Temp Anom) on the Y-axis. Determine the slope of the linear regression line, b (in $^{\circ}\text{C}$ change in global temperature per a unit change in Niño 3.4 index) and the correlation coefficient, r . Assess statistical significance of the linear trend without accounting for autocorrelation. (C) Plot model residuals over time and discuss whether heteroscedasticity of the regression model residuals is a point of concern in your analysis.
6. Given that there are La Niña year conditions this year, how do you expect ENSO to influence this year's global temperature? ENSO – El Niño-Southern Oscillation – is a climate pattern of oscillation between El Niño events (the positive phase, i.e., Niño 3.4 index above average) and La Niña events (the negative phase, i.e., Niño 3.4 index below average). Using the most negative observed value of the Niño 3.4 index during the historical time period used for the regression equation from question (5), determine the temperature anomaly predicted by the regression model for this most negative value. By comparing this temperature anomaly to the temperature anomaly expected in neutral (i.e., Niño 3.4 index = 0) conditions, estimate the largest perturbation of global temperatures expected in the most extreme La Niña events.
7. Save your word processing document as either a Microsoft Word or PDF file in the following format: PS1_AccessAccountID_LastName.doc (or .pdf).

For example, student Elvis Aaron Presley's file would be named "PS1_eap1_presley.doc". This naming convention is important, as it will help the instructor match each submission with the right student!

Submitting your work

Upload your file to the "Problem Set #1" assignment in Canvas by the due date indicated in the Syllabus.

Grading rubric

The instructor will use the general [grading rubric for problem sets](#) ^[29] to grade this activity.

Solutions will be uploaded to *Lesson 2: Climate Observations, Part 1* in Canvas after the due date.

Lesson 2 Summary

In this lesson, we reviewed key observations that detail how our atmosphere and climate are changing. We have seen that:

- greenhouse gas concentrations, including atmospheric CO_2 and methane, are increasing dramatically and these increases are associated with human activity;

- the surface of the Earth is warming and certain regions (e.g., the Arctic) are warming faster than others, consistent, as we will see, with expectations from climate model projections;
- the vertical pattern of the warming indicates that the surface and lower atmosphere (troposphere) are warming, while the atmosphere is cooling at altitude (in the stratosphere), a pattern that is consistent with greenhouse warming, but not with the natural factors such as solar output changes;
- there is a complicated pattern of changes in rainfall patterns around the globe, with some regions becoming wetter while other regions become drier;
- despite the heterogeneous pattern of changes in rainfall, there is a trend towards more widespread drought, consistent with the additional impact of warming on evaporation from soil.

We also learned how to analyze basic relationships in observational data, including:

- how to assess whether or not there is a statistically significant trend over time in a data series;
- how to assess whether or not there is a statistically significant relationship between two distinct data series.

In our next lesson, we will look at some additional types and sources of observational climate data, and we will explore some additional tools for analyzing data.

Reminder - Complete all of the lesson tasks!

You have finished Lesson 2. Double-check the list of requirements on the first page of this lesson to make sure you have completed all of the activities listed there before beginning the next lesson.

Source URL: <https://www.e-education.psu.edu/meteo469/node/118>

Links

- [1] <https://www.ipcc.ch/report/ar5/wg1/>
- [2] https://www.ipcc.ch/pdf/assessment-report/ar5/wg1/WG1AR5_SPM_FINAL.pdf
- [3] <http://earthobservatory.nasa.gov/Features/Revelle/revelle.php>
- [4] <http://www.harvardsquarelibrary.org/biographies/roger-revelle/>
- [5] http://scrippsco2.ucsd.edu/history_legacy/charles_david_keeling_biography
- [6] <http://www.esrl.noaa.gov/>
- [7] <http://www.esrl.noaa.gov/research/themes/carbon/>
- [8] <http://www.esrl.noaa.gov>
- [9] <http://data.giss.nasa.gov/gistemp/>
- [10] <http://www.giss.nasa.gov/about/>
- [11] <http://data.giss.nasa.gov/gistemp/animations/>
- [12] <http://www.realclimate.org/index.php/archives/2008/03/the-global-cooling-mole/>
- [13] <https://www.e-education.psu.edu/meteo469/203>
- [14] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu/meteo469/files/lesson02/PrecipTrendsIPCC_large.gif
- [15] <https://www.e-education.psu.edu/meteo469/node/152>
- [16] http://en.wikipedia.org/wiki/Carl_Friedrich_Gauss
- [17] http://en.wikipedia.org/wiki/Ordinary_least_squares
- [18] <http://faculty.vassar.edu/lowry/tabs.html#r>

- [19] https://courseware.e-education.psu.edu/courses/meteo101/stat_plotting_tool.swf?filename=STCDecTemp.txt
- [20] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/part_01.html
- [21] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/part_02.html
- [22] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/part_03.html
- [23] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/demo_part_01.html
- [24] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/demo_part_02.html
- [25] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/demo_part_03.html
- [26] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/from_DBabb/stat_plotting_tool.swf
- [27] https://courseware.e-education.psu.edu/courses/meteo469/stat_plotting_tool.swf?filename=assignment1.txt
- [28] https://www.e-education.psu.edu/meteo469/sites/www.e-education.psu.edu.meteo469/files/lesson02/PS1_worksheet_Sp2017.doc
- [29] <https://www.e-education.psu.edu/meteo469/node/243>