What was Earth's climate like before we were measuring it?

To comprehend the full impact of climate change, now and in the future, we need a deeper understanding of past conditions on Earth. **Jason E. Smerdon** shows how palaeoclimatology – with the aid of statistical methods – can help us peer back in time

The Earth is warming, and scientists overwhelmingly agree that human activities are to blame.[1] Decades of research have established this consensus, yet there remain some who argue that we cannot know whether similar climatological changes occurred in the past, before the advent of industrialised societies. They may point to the fact that direct meteorological observations dating back more than 200 years are only available in a few locations, and global networks of observations are only available for 100–150 years. This amount of time is insufficient to fully characterise how the climate varies over centuries or how large changes to the climate system were manifest in the past. Understanding the full impact that human activities are having on our climate today, and are likely to have in the future, therefore requires estimates of past climate conditions, long before humans began measuring and altering it.

Palaeoclimatology, or the study of past climates, extends our understanding of climate back thousands to millions of years. Among other things, the science has characterised the ebb and flow of ice ages, the impacts of the long-term carbon cycle, and how large volcanic eruptions in the tropics affect climate. These insights are derived from natural archives called climate proxies – ice, mud, bone and wood – that preserve a signature of what the climate was like at the time in which they were formed. Interpreting these archives and understanding how they represent past climates is a complicated endeavour, however, and statistics is often at the centre of how such records are interpreted and how they are used to infer connections between the past and present.

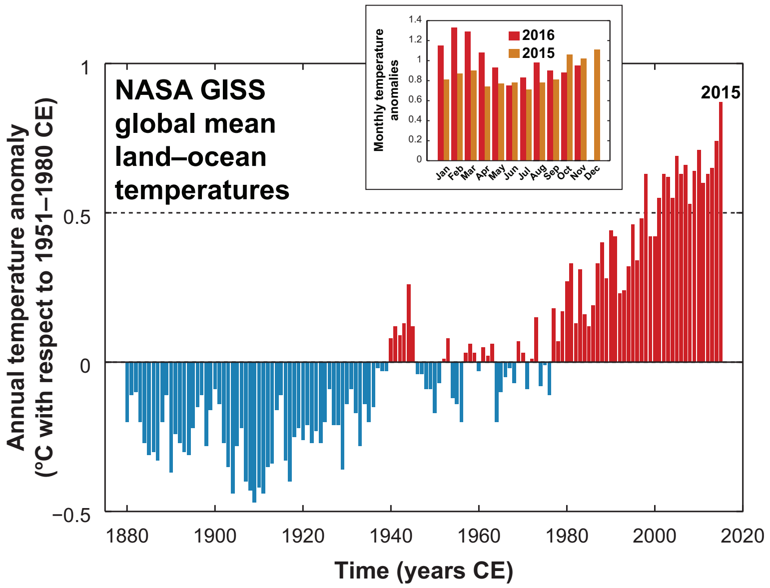

FIGURE 1 Global mean temperature anomalies, as estimated by the NASA Goddard Institute for Space Studies using land and ocean observations. The year 2015 was the warmest year on record. The inset shows the monthly anomalies for 2015 and 2016, the latter of which is on track to be even warmer than the former and therefore set a new record. Data available at data.giss.nasa.gov/gistemp

Observing the present, inferring the past

Since the mid-twentieth century, concentrations of atmospheric carbon dioxide have increased nearly exponentially to over 400 parts per million today. Geological evidence suggests that a concentration that high has not occurred for 3–3.5 million years. Meteorological records from the mid-nineteenth century to the present reveal that the mean surface temperature of the planet also has rapidly increased (Figure 1): between 1880 and 2012, global mean annual surface temperatures have warmed by an estimated 0.83–0.87°C[1], 2015 was the warmest year on record, and 2016 is expected to set a record once again. A host of additional observations of the atmosphere, ocean, biosphere and cryosphere also indicate that the Earth is warming, including melting sea ice and polar ice caps, rising sea levels, and increases in extreme weather events. As to the connection between these observations, state-of-the-art climate models demonstrate that the majority of warming since the mid-twentieth century is caused by atmospheric CO_2 increases due to human activities.[1]

Jason E. Smerdon is a Lamont Associate Research Professor at the Lamont-Doherty Earth Observatory of Columbia University

But how big is the climate perturbation that we are causing? How much has the climate actually warmed and cooled in the past? How closely have atmospheric CO_2 concentrations and global surface temperatures varied together through time, and how have their changes impacted other aspects of the climate system? These questions are the provenance of palaeoclimatology and answered with data from climate proxies.

Classic examples of climate proxies include the thickness and density of annual tree rings, ice cores taken from polar ice sheets, and the shells of micro-organisms embedded in ocean sediments. But climate proxies can take more exotic forms. Fossil remains from a giant boid snake that lived 58–60 million years ago have suggested warmer temperatures in the tropics during that time because the total body size of these reptiles was controlled in part by mean annual temperatures[2] (given that the prehistoric snake was almost twice as large as a modern green anaconda, we can be thankful for the estimated tropical cooling since that time). Although this is an unusual example, it illustrates that as long as there is a discernible connection between ambient climate and a given system, and as long as components of that system are preserved as archives for scientists to measure, climate proxies can be derived from many things.

Main image: NASA

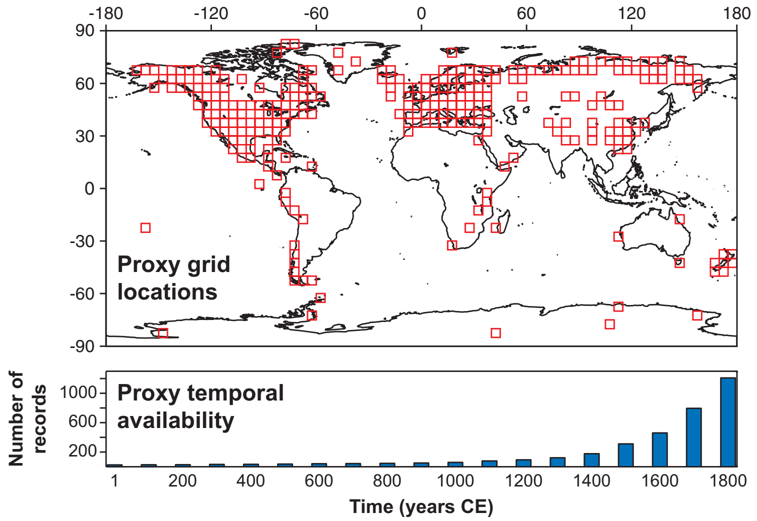

FIGURE 2 (top) Five-degree latitude–longitude grid cells that contain at least one proxy record used in a contemporary gridded temperature reconstruction[6] spanning approximately the last 1500 years. The network contains multiple proxies that include tree rings, corals, ice cores, lake sediments and cave deposits. (bottom) The temporal availability of proxies in the network at the beginning of each century illustrates the rapid reduction in the number of available high-resolution proxies back in time

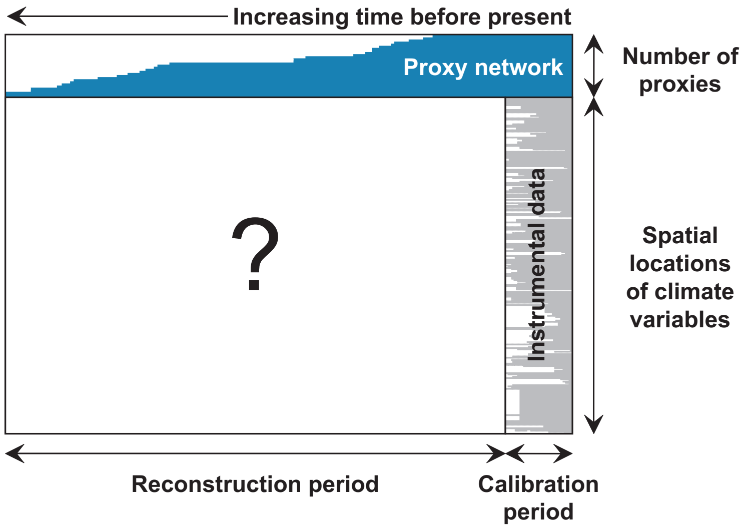

FIGURE 3 Schematic time–space data matrix representing the available proxy and instrumental data sets and their period of overlap. White areas of the matrix represent missing data in the proxy and instrumental grids, and the large white block containing the question mark is the region of the data matrix that must be estimated once a relationship is determined between the proxy and instrumental data during the calibration period

The way in which proxy archives are preserved determines how they can be used to study past climate and over which time periods. Tree cores taken from living trees are used for studying climate over the last several thousand years, but the life span of trees limits their use further back in time. In contrast, ice cores are limited by how long a given ice sheet or glacier has existed; the longest ice cores from Antarctica go back almost a million years. Ocean sediments extend back in time further still and can preserve archives of ocean conditions tens of millions of years into the past.

In general, the further back in time a proxy extends, the grainier its temporal resolution becomes. For example, tree rings can provide seasonal or annual resolution, while ocean sediments typically provide estimates of average conditions over centuries or millennia. For this reason, the Common Era (CE; the last two millennia) contains the most abundant collection of seasonally and annually resolved proxy records spread globally across land and sea.[3] When combined with data from our modern measurements, the CE is the best-documented period of climate variability in Earth's history and therefore an important target of palaeoclimatic study.

Infilling the void

It is crucial to understand that proxy archives are not direct measures of climate. They provide measurements of climate-influenced variables, but never direct measures of quantities such as average summer surface air temperatures or total annual rainfall. One consequence of this is that they are rarely sensitive to a single climatological quantity, and are often subject to site-specific causes of 'noise' that can be affected by local environmental conditions. Proxy records are also not evenly spread over time and geography, which makes them subject to sampling biases (Figure 2). The means by which measurements of proxy archives are interpreted as indicators of climatic conditions is therefore a fundamental challenge of palaeoclimatology, and one well suited to statistical methods.

A particularly apt example of statistical applications in palaeoclimatology involves the use of large networks of multiple kinds of proxy archives to derive reconstructions of hemispheric and global fields of climate during the CE. These fields can take many forms, but are easiest to imagine as gridded maps of climate variables averaged over seasons or years. Think of a series of paper maps stacked in one pile: each page represents a slice in time and the maps on the page are seasonal or annual atlases of a climate variable of interest (such as temperature or rainfall). In aggregate, the stack of maps represents a series of spatial fields incremented evenly in time over all or part of the CE.

Reconstructions of the stacked maps described above are called climate field reconstructions (CFRs), and Figure 3 schematically represents the data matrix that is characteristic of the CFR problem. The big, empty block in the lower part of the matrix is the reconstruction period. The missing data for this period can be estimated by identifying relationships between the proxy records (in the upper block of the matrix) and the instrumental data (in the bottom right block of the matrix). This objective is not unlike the famous "Netflix problem" in which viewers' ratings of a small number of movies are used to estimate their preferences for a much larger set of movies they have not seen, based on the ratings of other viewers with similar taste. In much the same way that Netflix might exploit the overlap in movie preferences between its users, the goal of CFRs is to estimate the values of a climate field in the reconstruction period using the relationship between the overlapping interval of the proxy and instrumental data during the so-called calibration period.

A number of statistical frameworks have been used for the CFR problem, but the vast majority have applied forms of multivariate linear regression.[4] This methodology is further discussed in the boxout (right), but the underlying challenge of the approach is the limited amount of information available for deriving the regression coefficients and thus the proxy-based estimate of climate fields during the reconstruction period. In other words, the effort is a classic ill-posed estimation problem and various means of constraining the solution or adding information are pursued through a technique known as regularisation.

Consider a facial recognition problem in which one photo is compared to many others in search of a match. It may be both computationally impossible and yield inaccurate results to compare photos at a pixel-to-pixel level, but there are basic facial patterns that can be identified and used for the comparison. A search on these patterns alone would significantly cull the number of possible facial matches and therefore reduce the amount of information to be compared. Similarly, one form of regularisation in the CFR problem is to identify a few large-scale patterns that represent broad circulations in the climate system and explain a large portion of the spatiotemporal variability in a climate field. Targeting just these leading patterns in the regression problem, instead of the complete field, therefore constrains the information that is to be reconstructed.

The challenge with regularisation is optimising the shape and strength of the constraints to avoid overfitting (if the information is not constrained enough) or underfitting (if the information is constrained too much). The manner and amount of regularisation and the consequent impacts on derived CFRs therefore have been a subject of investigation for well over a decade. These efforts have in turn yielded reconstructions that describe important characteristics of CE climate.

Reconstructing climate fields with multivariate linear regression

The majority of CFR methods relate a matrix of climate proxies to a matrix of climate data during a common time interval, generally termed the calibration period (see Figure 3), using a linear model. If **P** is an $m \times n$ matrix of proxy values and **T** is an $r \times n$ matrix of instrumental climate records, where m is the number of proxies, r is the number of spatial locations in the instrumental field, and n is the time of overlap between the proxy and instrumental data, the linear relationship is written

$$\mathbf{T} = \mathbf{BP} + \varepsilon,$$

where **B** is an $r \times m$ matrix of regression coefficients, and ε is the residual error (the **T** and **P** matrices are assumed to be centred and normalised over the time interval of n). According to standard linear regression theory, if **B** is estimated as

$$\hat{\mathbf{B}} = (\mathbf{TP}')(\mathbf{PP}')^{-1},$$

where the prime denotes the matrix transpose, the sum of the squared residuals, ε, is minimised for all r spatial locations of the estimated climate variable. The reconstructions can then be carried out using the regression matrix $\hat{\mathbf{B}}$ during periods in which proxy data are available but observed climate variables are not.

The above formalism works best when the system is overdetermined, i.e. $n \gg m$, such that the inversion in $\hat{\mathbf{B}}$ can be reliably performed. The challenge for CFR methods is that in most practical situations this condition is not met. In such cases, the estimate requires some form of regularisation to apply additional constraints.

Two forms of regularisation are often adopted in the CFR context. The first involves covariance-based matrix factorisations (e.g. singular value decomposition or principal component analysis) of the climate and proxy matrices and subsequent rank reductions. The rank reductions are applied under the assumption that large-scale climatic circulation processes drive covariance patterns in climate fields that are expressed as leading spatiotemporal covariance patterns in both **T** and **P**. Reduced-rank representations of the two matrices therefore filter small-scale variability and noise and comprise good approximations of the full-rank versions of **T** and **P**, which can in turn be used to estimate $\hat{\mathbf{B}}$. Common CFR methods adopting these strategies are principal component regression and canonical correlation analysis.

The second form of regularisation is a family of methods commonly referred to as penalised regressions or maximum likelihood methods such as ridge regression or lasso. These methods estimate modified versions of $\hat{\mathbf{B}}$ using a trade-off between the loss (error) and likelihood or fit. This latter penalty assumes that regression coefficient values that are closer to zero are most likely, and therefore prioritises estimates of $\hat{\mathbf{B}}$ that include coefficients meeting this criterion. The magnitude of this penalty is determined by the value of regularisation parameters, which determine the degree of the trade-off between loss and fit. Overall, selecting rank reductions of the climate and proxy matrices and the regularisation parameters associated with penalised regressions are steps of technical importance and have been an important part of the work to improve CFRs.[4]

Progress and challenges

Figure 4 (page 28) is a summary of state-of-the-art reconstructions of Northern Hemisphere mean temperatures, either derived from direct reconstructions of the mean temperature index or aggregated from temperature CFRs. The differing estimates are the result of specific methodological and proxy selection choices made by various research teams. While uncertainties remain, the results in aggregate have placed contemporary global warming in the context of a longer climatic history and indicated that the most recent past decades are likely to have been the warmest in the Northern Hemisphere in the last 800–1400 years. Similarly, the spatial

Palaeoclimatology tells us about climate variability in the past, but can also help characterise risks that are yet to come

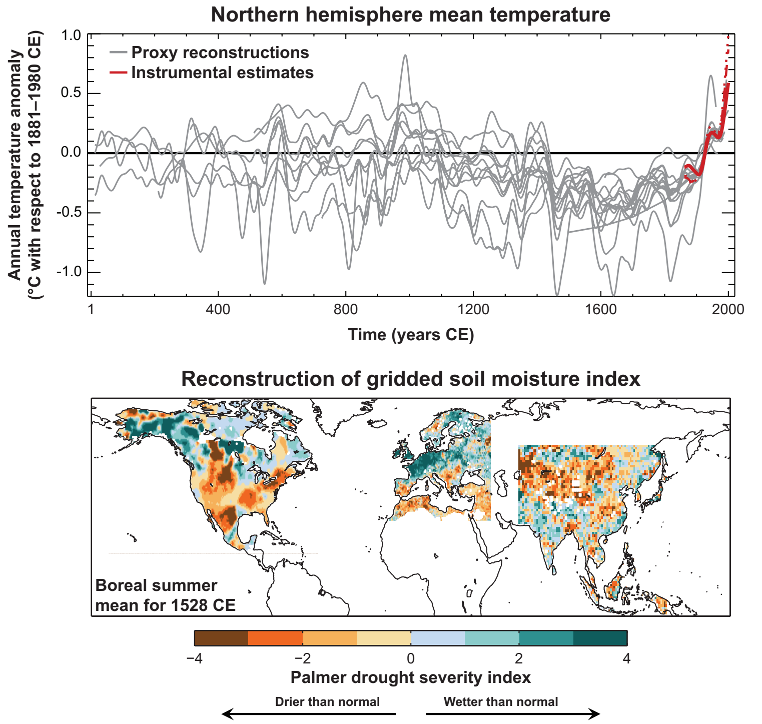

FIGURE 4 Example reconstructions of (top) Northern Hemisphere mean temperatures. The figure is modified from Figure 5.7 in the *Fifth Assessment Report* of the Intergovernmental Panel on Climate Change and collects the range of contemporary reconstructions based on different methods and proxy network selections. All series represent anomalies from the 1881–1980 mean (horizontal line) and have been smoothed with a filter that reduces variations on time scales less than about 50 years. (bottom) A spatial map of a reconstructed soil moisture index in the year 1528 CE (data courtesy of E. Cook at the Lamont-Doherty Earth Observatory of Columbia University). The Palmer Drought Severity Index is a normalised soil moisture metric in which positive (negative) deviations represent wetter (drier) than normal conditions

information in CFRs has characterised important past patterns of variability. Figure 4 also includes a single-year example from a CFR of normalised soil moisture anomalies that have been used to characterise drier and wetter periods in Earth's history. This information has been used to understand climate impacts on past societies such as the Ancestral Puebloan civilisations of the American Southwest and the civilisation of Angkor that existed between the ninth and fifteenth centuries in the area of modern-day Cambodia.

Despite the many successes of CE reconstruction efforts, some methodological studies have nevertheless indicated that the fundamentally data-limited CFR problem is still subject to uncertainties. For instance, some experiments suggest that many CFR methods estimate global or hemispheric mean temperature indices rather accurately, but the spatial characteristics of the fields are subject to large uncertainties.[5] Reducing these uncertainties is therefore an important contemporary research focus.

CFR improvements rest, first and foremost, on the acquisition of more high-quality proxy information from undersampled regions (Figure 2, page 26). Methodological advances are, however, also important for moving forward. New CFR methods are including information about how proxies respond to multiple climatological influences, based on growing efforts to numerically model how proxies respond to climate. Simulations from climate models are also being used to provide independent physical estimates of how climate varies in space and time; this information can in turn be used with proxy data in CFR methods. Such developments are at the forefront of efforts to reconstruct CE climate and driving new and exciting insights into the climate of the past.

A vital perspective

Palaeoclimatology is vital for understanding how the climate naturally varies when instrumental observations are not sufficient for the task – a motivation with practical implications in areas such as infrastructure and policy planning. Consider, for instance, the water allocations drafted in the 1922 Colorado River Compact, which were based on absolute (not proportional) amounts of water. The agreement might have been very different if the stakeholders had realised that the early twentieth century was an anomalously wet period in the American Southwest and that decades-long droughts unlike anything they had seen firsthand were also a feature of the region.

Image: NASA

Palaeoclimatology can therefore inform us about the range of natural climate variability in the past, but it can also help us characterise risks that are yet to come.

In addition to natural variability, the climate of the future will be impacted by human activities. Climate models are critical tools used to characterise risks associated with these activities and they are continuously being improved and updated as computing technology evolves and scientific understanding expands. Evaluating these evolving tools over palaeoclimatic intervals is one means of refining them, and robust palaeoclimatic reconstructions are the bedrock information by which these evaluations are possible. Advances in the application of statistical methods that are used in these reconstructions are therefore a central aspect of how the field is moving forward, and it is through the joint application of palaeoclimatic and statistical methods that the past can ultimately serve as an essential and quantitative guide for the future. ■

Note

I am grateful to Alexey Kaplan and Stacy Morford at the Lamont-Doherty Earth Observatory of Columbia University and the *Significance* editorial board for their insightful comments on early drafts of this manuscript.

References

1. Stocker, T. F., Qin, D., Plattner, G. K., Tignor, M., Allen, S. K., Boschung, J., Nauels, A., Xia, Y., Bex, V. and Midgley, P.M. (eds) (2014) *Climate Change 2013: The Physical Science Basis. Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. New York: Cambridge University Press.

2. Head, J. J., Bloch, J. I., Hastings, A. K., Bourque, J. R., Cadena, E. A., Herrera, F. A., Polly, P. D. and Jaramillo, C. A. (2009) Giant boid snake from the Palaeocene neotropics reveals hotter past equatorial temperatures. *Nature*, **457**, 715–717.

3. Smerdon, J. E. and Pollack, H. N. (2016) Reconstructing Earth's surface temperature over the past 2000 years: the science behind the headlines. *WIREs Climate Change*, **7**(5), 746–771.

4. Tingley, M. P., Craigmile, P. F., Haran, M., Li, B., Mannshardt-Shamseldin, E. and Rajaratnam, B. (2012) Piecing together the past: statistical insights into paleoclimatic reconstructions. *Quaternary Science Reviews*, **35**, 1–22.

5. Smerdon, J. E. (2012) Climate models as a test bed for climate reconstruction methods: pseudoproxy experiments. *WIREs Climate Change*, **3**(1), 63–77.

6. Mann, M. E., Zhang, Z., Rutherford, S., Bradley, R. S., Hughes, M. K., Shindell, D., Ammann, C., Faluvegi, G. and Ni, F. (2009) Global signatures and dynamical origins of the little ice age and the medieval climate anomaly. *Science*, **326**, 1256–1260.