

Precision maps for public health

Researchers have produced high-resolution maps of childhood growth failure and educational attainment across Africa between 2000 and 2015, to assess progress and guide policy decisions in public health. [SEE ARTICLES P.41 & P.48](#)

BRIAN J. REICH & MURALI HARAN

Maps are key to revealing public-health challenges and suggesting potential solutions. Physician John Snow's map of cholera cases and public wells in London in the 1850s was an early example of this, famously revealing that cholera is spread through contaminated water¹. Two papers^{2,3} in *Nature* now provide high-resolution maps of health and education levels for children across Africa. Their maps will draw attention to areas most in need of support, and guide future interventions to where they can have the greatest impact.

Public-health policies and decision-making are often local, and so, ideally, are based on information on a small spatial scale. For example, when analysing disparities in immunization rates between urban and rural areas, information at the level of a village is often more valuable than district-level information. Like most disciplines in the twenty-first century, public health benefits from access to an unprecedented amount of data that can provide information not only on smaller spatial scales than was previously possible, but also more frequently — an exciting prospect for scientists and policymakers.

However, individual data sets often do not provide the information needed to answer key policy-relevant questions. For example, to estimate the relationship between maternal education level and infant mortality, it might be necessary to combine data from multiple sources. In much the same way as one can make educated guesses about missing pieces in a partially completed puzzle, it can be possible to find hidden information by piecing together data from different sources. But because each data source comes with its own set of errors and complexities, sophisticated statistical methods are required to both integrate the data into usable information and reflect its intrinsic uncertainties. An example of one such set of methods is Bayesian geostatistics, which combines data from multiple locations




Figure 1 | Mapping public-health problems across Africa. Two studies^{2,3} have produced high-resolution maps of Africa, in which each 'pixel' represents 5 × 5 kilometres. The first study documented childhood growth failure between 2000 and 2015. Here, one heat map from the study shows the percentage of children classed as moderately or severely underweight (MSU) in each pixel in 2000. Grey represents areas for which data were not collected. The second study (not shown) analysed educational attainment, and produced similar heat maps. (Map taken from Extended Data Figure 2 of ref. 2.)

and data sets, exploiting their spatial correlations to predict values for regions for which information is lacking, while also providing information about the uncertainties involved in these predictions.

The current studies use new, advanced Bayesian geostatistical tools to analyse two problems Africa-wide. In the first paper (page 41), Osgood-Zimmerman *et al.*² focused on childhood growth failure. They pooled geolocated information on growth stunting, muscle wasting and weight in children under the age of 5 from several surveys across tens of thousands of villages over 15 years. They then combined this with information on local climate and geography. They carefully validated their statistical model by first fitting the model to data at a subset of locations, then comparing the predictions from this fitted model with data at a different subset of locations.

The map is split into 'pixels' of 5 × 5 kilometres across Africa, and shows changes in growth failure over time, from 2000 to 2015 (Fig. 1). The authors use their data to point out the differences in improvements across time in different regions, and show that national-level data mask nuances uncovered by their precision maps. They also provide measures of the certainty of the prediction made for each region, thereby highlighting both the uncertainty in the maps and the areas of Africa most in need of additional sampling. They then use their model to assess the likelihood of achieving 2025 global nutrition targets, and point out regions in which progress is lagging behind. The authors find that, unless there is a change in the current rates of improvement, much of the continent will fail to meet the goal of ending malnutrition by 2030.

In the second paper (page 48), Graetz *et al.*³ used a similar approach to map local variations across Africa in the number of years of education that women between the ages of 15 and 49 have received. This is of particular interest because educational attainment is linked to the health of both mothers and their children. The authors produced their maps from data gathered from geolocated household surveys and censuses. They generated maps of average attainment across regions, including changes between 2000 and 2015, and provided uncertainties for each average. In addition, they generated similar maps for men, and showed that, although there has been progress in educational attainment for both men and women across the continent, substantial differences between the sexes remain.

As these papers show, we now have sufficiently mature statistical methodology, theory and software to analyse continental-scale problems using sound methods and open-source software. This type of analysis was simply not possible ten years ago. In the past, researchers might have resorted to aggregating the data used in the current studies by country, at the expense of incorporating the local

phenomenon that really drives the science. Alternatively, they might have restricted the analysis to one country, thus failing to exploit the full power of the data at hand. But the sophisticated geospatial tools used in the current work employ clever numerical approximations to sidestep the computational bottlenecks posed by analysing so many correlated observations. These methods are applicable to much more than just the public-health domains described here, and should provide scientific insights in many disciplines. Of course, it is important that these powerful statistical tools are not applied blindly. In both papers, the authors are careful to weight data appropriately and to validate their predictions at each step.

There is much excitement these days about the way in which enormous data sets are helping us to address many hard scientific challenges. In reality, data sets are useful only when combined with a deep understanding of the relevant science, economics or sociology, such as the impact of culture in a particular region, or details about how diseases spread. A solid understanding of how data are collected is also crucial. Rigorous scientific advances emerge when interdisciplinary teams work closely together — the current papers, which involve researchers trained in epidemiology, statistics, demography and public health, are prime examples of this.

The ultimate goal of a spatial analysis is to design interventions for maximum impact. If we understand a spatio-temporal process, we can optimize the allocation of resources in space and time. For example, consider the spread of malaria, and the effect of interventions such as bed-net distribution. A 2016 analysis⁴ considered several malaria interventions, and determined the most cost-effective intervention for each 5-km² pixel in Africa on the basis of spatial variation in climate, mosquito populations and the current state of the disease. The results from Osgood-Zimmerman *et al.* and Graetz *et al.* should prove useful in an analogous study of optimal interventions for nutrition and education. We believe that we are entering an era in which this type of analysis can be applied broadly to improve the lives of people around the world. ■

Brian J. Reich is in the Department of Statistics, North Carolina State University, Raleigh, North Carolina 27607, USA.
Murali Haran is in the Department of Statistics, Pennsylvania State University, State College, Pennsylvania 16802, USA.
 e-mails: bjreich@ncsu.edu; muh10@psu.edu

QUANTUM COMPUTING

Qubits break the sound barrier

Quantum logic gates based on trapped ions perform more accurately than solid-state devices, but have been slower. Experiments show how trapped-ion gates can be sped up, as is needed to realize a quantum computer. SEE LETTER P.75

TOBIAS SCHÄTZ

One of the goals most eagerly pursued by physicists is the development of a universal quantum computer — a machine capable of running a superposition of many correlated tasks, and one that would offer much better performance for dedicated jobs than do conventional computers. One of the most promising approaches is to use systems of trapped ions. Such systems provide the best fidelities^{1,2} for all quantum logic operations (that is, they most reliably produce the correct outputs for a given input), including operations performed by the logic gates needed to process quantum bits (qubits).

Two-qubit gates are needed to implement quantum computers that can run any algorithm. In trapped-ion systems, two-qubit gates carry data between ions using phonons (quantum units of vibration) associated with the collective oscillation of the ions³. However, the operating speed of the gates has been limited by the oscillation frequency, which defines the speed of sound in the ensemble of ions and therefore sets the speed limit for communication. Several strategies have been proposed to overcome this limitation^{4–8}. On page 75, Schäfer *et al.*⁹ report the experimental realization of one such strategy, building on previously reported work in this area^{7,10}. The researchers find that the speed of their two-qubit gates is more than ten times that of previously reported trapped-ion gates.

The basic principle of two-qubit gates is reminiscent of classical logic gates. Let's consider a classical controlled NOT (CNOT) gate, which turns four possible input states (00, 01, 10 or 11) into four output states (00, 01, 11 and 10, respectively). In other words, the second input bit (the target bit) is flipped between the 0 and 1 states only if the first bit (the control bit) is 1.

The quantum version of the CNOT gate allows much more than these four states to be processed. For example, the control qubit can enter the gate in a superposition state, $0 + 1$. If the target qubit is 0, then running the same logic operation as for the classical CNOT flips and unflips the target at the same time. The two qubits thus end up in the final state of 00 + 11, which is a maximally entangled state: the measurement of one qubit yields a

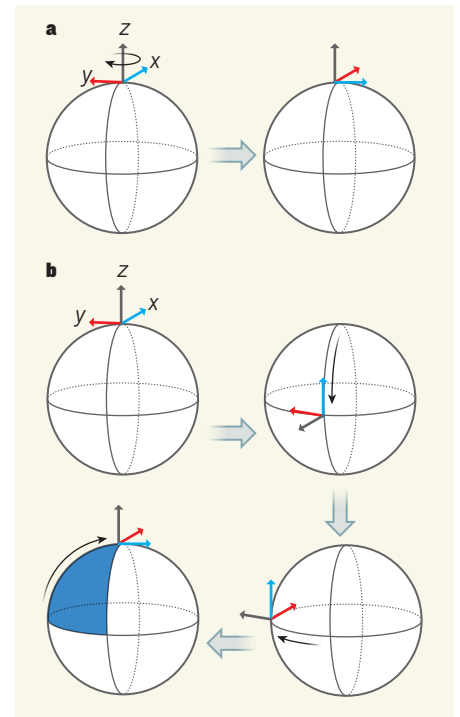


Figure 1 | Dynamical and geometric phase generation. **a**, When a physical state that can be described using x , y and z axes is placed at the north pole of a sphere, a property known as dynamical phase is generated by rotating the system around its z axis. The phase change equates to the change in the orientation of the x and y axes relative to their original ones. **b**, If the physical state moves down the surface of the sphere to the equator, then along the equator and back up to the north pole, the net result is rotation of the x axis, as in **a**. This phase change is proportional to the area enclosed by the pathway taken (blue region), and is called a geometric phase change. Schäfer *et al.*⁹ report a method for generating geometric phase changes of trapped ions, and use it to implement quantum logic gates that operate much faster than similar, previously reported trapped-ion gates.

completely random output, but instantaneously fixes the state of the second qubit to be identical to that of the first, regardless of the distance between the qubits. This correlation could form the backbone of a quantum computer — it allows many operations to be run in parallel, so that a superposition of all possible inputs yields all possible results at once.

1. Johnson, S. *The Ghost Map: The Story of London's Most Terrifying Epidemic — and How It Changed Science, Cities, and the Modern World* (Riverhead, 2008).
2. Osgood-Zimmerman, A. *et al. Nature* **555**, 41–47 (2018).
3. Graetz, N. *et al. Nature* **555**, 48–53 (2018).
4. Walker, P. G. T., Griffin, J. T., Ferguson, N. M. & Ghani, A. C. *Lancet Glob. Health* **4**, 474–484 (2016).