

What a cholera map can teach us about trusting AI



Singapore's guidelines for AI in healthcare rightly emphasise safety, clinical effectiveness and appropriate human responsibility, but meeting these goals will require more than access to increasingly powerful models, says the writer, citing the ability to understand what those models have learnt: which relationships they preserve, which differences they overlook and how uncertainty enters their predictions. PHOTO: ST FILE

A physician spotted a consequential pattern on a street map. Working with similar principles today could make medical AI more useful and easier to question.

Zhigang Yao

In late August 1854, cholera swept through Soho in London. Within days, hundreds of people had died.

At the time, many believed that cholera was transmitted through foul-smelling air. The physician John Snow suspected contaminated water instead. While investigating the outbreak, he placed the deaths on a street map and marked the locations of public water pumps.

The deaths were not scattered randomly. They clustered around the pump on Broad Street.

The pattern became more persuasive when Snow investigated its apparent exceptions. A nearby workhouse suffered relatively few deaths because it had its own water supply. Workers at a local brewery largely escaped the outbreak because they drank beer rather than water from the pump. Some victims who lived farther away were found to have deliberately obtained water from Broad Street.

The arrangement of the deaths pointed towards something that the individual records did not reveal: a shared source of exposure.

Local officials eventually removed the pump handle. It would be neat – but historically inaccurate – to say that the map alone ended the outbreak. The epidemic was already subsiding, and a spatial pattern by itself could not prove that contaminated water caused cholera. Snow strengthened his case through interviews and comparisons of households supplied by different water companies, as described in his published investigation.

That combination was crucial. The shape of the data suggested where to look; further evidence helped establish what the shape meant.

More than 170 years later, the same principle matters for artificial intelligence. Data does not consist only of individual measurements. Their arrangement – how observations cluster, connect, branch or depart from one another – may contain information about the process

that produced them.

Shape itself can be part of the signal.

FROM A VISIBLE MAP TO A MISSING MAP

Snow had one important advantage – he possessed a street map on which every death could be located. The shape of his data could be seen.

Modern biological data may also contain meaningful shapes, but they cannot be observed directly. A single cell can be represented by the activity of thousands of genes. A biobank participant may be described by hundreds of biomarkers and clinical measurements. Each cell or person is therefore a point in a space far beyond what the human eye can visualise.

For these data, no natural street map exists. A technique that we call “manifold fitting” attempts to reconstruct this missing map – the lower-dimensional structure around which the observations are organised – so that clusters, trajectories and unusual departures can be identified and tested. The “map” is a metaphor for a structure within the original high-dimensional data, not merely a two-dimensional visualisation.

WHAT WE CAN IDENTIFY WHEN SHAPE CHANGES

Consider single-cell RNA data. Each cell is represented by thousands of gene measurements, but biological and technical noise can make cells of the same type appear far apart and different cell types appear misleadingly similar.

In work with my collaborators, we developed single-cell analysis via “manifold fitting” to bring cells closer to the hidden structure around which their measurements are organised.

The difference can be measured. In one benchmark involving three known groups of mouse embryonic stem cells, a standard clustering method correctly grouped only 57 per cent of the cells in the raw data. After manifold fitting, the same method achieved 100 per cent accuracy. Across 25 public data sets involving different species, organs and sequencing technologies, the fitting step consistently improved the separation of known cell populations.

The point was not to produce a more attractive diagram. The recovered shape made biologically different populations easier to distinguish. In this case, geometry changed what the analysis could reliably identify.

If the single-cell example shows

how shape can recover distinctions obscured by noise, a UK Biobank study shows how it can reveal differences that population averages conceal.

We analysed 251 metabolic biomarkers from more than 212,000 participants. Examining each measurement separately – cholesterol, fatty acids, proteins or measures related to blood sugar – provides only a fragmented picture. The biological meaning may lie in how these measurements vary together.

When all 251 biomarkers were viewed together, the data displayed no obvious structure. After related biomarkers were organised into seven categories and manifold fitting was applied, previously hidden structures emerged. Three of them separated participants into distinct metabolic subgroups.

These were not simply groups of people with a single “high” or “low” biomarker.

They represented different combinations of metabolic measurements moving together. One structure revealed a subgroup associated particularly strongly with severe metabolic dysfunction and its neurological and vascular complications. Another identified a broader pattern associated with cardiovascular and autoimmune conditions.

This did not diagnose individual patients or prove that a metabolic pattern caused disease. It revealed heterogeneity that individual measurements and population averages could not show. Two people may have a similar reading for one familiar biomarker yet occupy very different positions on the wider metabolic map – and belong to

When medical AI systems are evaluated in Singapore, the question should not only be whether they are accurate on average. We should also ask whether the structures they learn remain stable across age groups, disease subtypes and different segments of the population – and whether clinicians can examine those structures when a prediction is challenged.

different risk profiles.

In all three cases, shape was a starting point rather than a conclusion. Snow investigated the exceptions and gathered independent evidence about water supplies. The RNA groups were checked against known cell identities across multiple data sets, while the biobank subgroups were tested against clinical measurements and subsequent disease records. Shape became scientifically useful because it produced claims that could be tested against the real world.

This is also why shape is not merely decoration added after an analysis. It organises isolated observations into relationships that can generate a hypothesis, reveal a distinction or guide a decision.

LOOKING BEHIND AN AI-GENERATED SCORE

Modern AI systems are extraordinarily good at compressing complicated information. Millions of pixels can become a label; hundreds of medical measurements can become a prediction or risk score.

But compression always raises a question: What has been preserved, and what has been discarded?

A risk score without information about its underlying structure can be like a list of cholera deaths without a map: informative, but incomplete. It may conceal the fact that different patients arrived at the same score through very different combinations of factors. It may also rely on a correlation in the training data rather than a medically meaningful relationship.

Recovering the structure behind the score could provide an intermediate object between raw data and an AI-generated answer – something that researchers and clinicians can examine. If a system identifies a patient as high-risk, a doctor could ask where that patient lies relative to a reference population, which measurements account for that position and whether the same structure appears in other populations.

Geometry would not make the final decision or automatically make AI output explainable. But it could help turn an opaque score into questions that humans can investigate.

WHY THIS MATTERS FOR SINGAPORE

Singapore's national AI strategy places “AI for the Public Good” at the centre of the country's ambitions. Healthcare is one area in which AI could have a direct impact, from medical imaging and clinical documentation to disease prevention and personalised care.

To use these systems responsibly, predictive accuracy is not enough. A model may perform well overall while being less reliable for a particular age group, disease subtype or segment of Singapore's diverse population. It may produce a plausible answer without revealing whether it has learnt a genuine biological relationship or an accidental feature of the data.

When medical AI systems are evaluated in Singapore, the question should not only be whether they are accurate on average. We should also ask whether the structures they learn remain stable across age groups, disease subtypes and different segments of the population – and whether clinicians can examine those structures when a prediction is challenged.

Singapore's guidelines for AI in healthcare rightly emphasise safety, clinical effectiveness and appropriate human responsibility. Meeting these objectives will require more than access to increasingly powerful models.

A more durable advantage would be the ability to understand what those models have learnt: which relationships they preserve, which differences they overlook and how uncertainty enters their predictions.

Snow did not stop when he saw a cluster on a map. He investigated the exceptions and tested whether the pattern corresponded to something real.

We should demand the same of AI.

Before we trust the score, we should ask to see the shape behind it.

• The writer is with the Department of Statistics and Data Science at the National University of Singapore and holds a joint appointment in the Department of Mathematics. He co-edited the Harvard Data Science Review special theme “Shaping Data Science: A Geometric Perspective”.