DEPARTMENT OF SCIENCE, TECHNOLOGY,
ENGINEERING AND PUBLIC POLICY

Thinking systems: how the systems we depend on can be helped to think and to serve us better

1

STEaPP Working Paper Series

Thinking systems: how the systems we depend on can be helped to think and to serve us better

The author:
Sir Geoff Mulgan CBE is Professor of Collective Intelligence, Public Policy and Social Innovation in STEaPP at University College London (UCL). He was Chief Executive of Nesta, the UK's innovation foundation from 2011–2019. From 1997–2004 Geoff had roles in the UK government including director of the Government's Strategy Unit and head of policy in the Prime Minister's office. From 2004-2011 he was first CEO of the Young Foundation. He has been a reporter on BBC TV and radio and was the founder/cofounder of many organisations, including Demos, Uprising, the Social Innovation Exchange and Action for Happiness. He has a PhD in telecommunications and has been visiting professor at LSE and Melbourne University, and senior visiting scholar at Harvard University. Geoff has advised many governments around the world and is a World Economic Forum Schwab Fellow from 2019-22. Past books include '*The Art of Public Strategy*' (Oxford University Press), '*Good and Bad Power*' (Penguin), '*The Locust and the Bee*' (Princeton University Press), '*Big Mind: how collective intelligence can change our world*' (Princeton University Press) and '*Social innovation: how societies find the power to change*' (Policy Press). Geoff has given TED talks on the future economy, happiness and education.

Author contact details:
Twitter: @geoffmulgan | Web: geoffmulgan.com

Table of contents

Acknowledgements: Huge thanks to colleagues who made comments on earlier draft and in particular Eirini Malliaraki from the Alan Turing Institute and Millie Begovic from UNDP. Of course, any errors of fact or judgement are mine alone.

About the STEaPP Working Paper Series

The STEaPP Working Paper series aims to publish emerging analysis and thought with the aim of improving links between research and policy; We will make available a diversity of outputs in different formats including pre-published versions of journal articles, drafts of longer academic or policy outputs, multi-media content and policy reports; We accept submissions from STEaPP and UCL staff, students, honorary fellows and relevant from other academic and policy colleagues.

Introduction

This draft paper describes methods for understanding how **vital everyday systems work**, and how they could work better, through improved shared cognition – observation, memory, creativity and judgement – organised as commons.

Much of our life we depend on systems: interconnected webs of activity that link many organisations, technologies and people. These bring us food and clothing; energy for warmth and light; mobility including rail, cars and global air travel; care, welfare and handling of waste. Arguably the biggest difference between the modern world and the world of a few centuries ago is the thickness and complexity of these systems. These have brought huge gains.

But one of their downsides is that they have made the world around us harder to understand or shape. A good example is the Internet: essential to much of daily life but largely obscure and opaque to its users. Its physical infrastructures, management, protocols and flows are almost unknown except to specialists, as are its governance structures and processes (if you are in any doubt, just ask a random sample of otherwise well-informed people). Other vital systems like those for food, energy or care are also hardly visible to those within them as well as those dependent on them. This makes it much harder to hold them to account, or to ensure they take account of more voices and needs. We often feel that the world is much more accessible thanks to powerful search engines and ubiquitous data. But try to get a picture of the systems around you and you quickly discover just how much is opaque and obscure.

If you think seriously about these systems it's also hard not to be struck by another feature. Our systems generally use much more data and knowledge than their equivalents in the past. But this progress also highlights what's missing in the data they use (often including the most important wants and needs). Moreover, huge amounts of potentially relevant data is lost immediately or never captured and how much that is captured is then neither organised nor shared. The result is a strangely lop-sided world: vast quantities of data are gathered and organised at great expense for some purposes (notably defense or click-through advertising) but very little for others.

So how could we recapture our systems and help them make the most of intelligence of all kinds? This paper shares methods and approaches that could make our everyday systems richer in intelligence and also easier to guide. It advocates:

- **A cognitive approach to systems** – focusing on how they think, and specifically how they observe, analyse, create and remember. It argues that this approach can help to bridge the often abstract language of systems thinking and practical action
- It advocates that much of this systems intelligence needs to be **organised as a commons** – which is very rarely the case now
- And it advocates new structures and roles within government and other organisations, and the growth of a practice of **systems architects** with skills straddling engineering, management, data and social science – who are adept at understanding, designing and improving intelligent systems that are transparent and self-aware.

The background to the paper is the great paradox of systems right now: there is a vast literature, a small industry of consultancies and labs, and no shortage of rhetorical commitment in many fields. Yet these have had at best uneven impact on how decisions are made or large organisations are run.

In the paper I show the relevance of some of the methods I suggest for governments seeking to better address challenges such as decarbonisation and care for the elderly; in relation to business I suggest that the West has tended to fall behind China in terms of designing and operating complex, interconnected systems straddling many fields, an ability which will be vital for the future of finance, energy and transport; and for citizens, I emphasise how greater influence can be achieved over systems which are now surprisingly opaque.

In summary I argue that although we live surrounded by systems we struggle to understand them let alone to guide and control them.[1] I believe we need a novel approach that focuses on how to enrich and open up the shared intelligence of the systems around us. This is the next step to take – building on the extraordinary achievements of the Internet itself, open data and other movements, and taking us to a world where there are accessible representations of the multiple systems on which we depend. My aim is to suggest some of the tools, insights and resources we could use to do this. The paper is shared as a first draft to elicit comments and improvements.

“So how could we recapture our systems and help them make the most of intelligence of all kinds? This paper shares methods and approaches that could make our everyday systems richer in intelligence and also easier to guide.”

1. The Problem

Much of our life we depend on systems: interconnected webs of activity that link many organisations, technologies and people. These bring us food and clothing; energy for warmth and light; mobility including rail, cars and global air travel; care, welfare and handling of waste.

Arguably the biggest difference between the modern world and the world of a few centuries ago is the thickness and complexity of these systems. These have brought huge gains. But one of their downsides is that they have made the world around us harder to understand or shape.

A good example is the Internet: essential to much of daily life but largely obscure and opaque to its users. Its physical infrastructures, management, protocols and flows are almost unknown except to specialists, as are its governance structures and processes.

Other vital systems like those for food, energy or care are also hardly visible to those within them as well as those dependent on them. This makes it much harder to hold them to account, or to ensure they take account of more voices and needs.

Even in their own terms however many of these systems are inefficient. Although they generally **use much more data and knowledge than their equivalents in the past, they are also much less intelligent than they could be**.[2] Most relevant data is lost immediately or never captured; much that is captured is then neither organised nor shared. The same is true of knowledge.

2. Systems and knowledge about systems

Before turning to how this problem might be solved, let me first say a word about definitions. There are many competing definitions of systems. Usually the term is used to describe some kind of whole made up of interrelated or interconnected elements – though some use the word to emphasise linked causal connections, others to emphasise boundaries or shared purpose. Here I use the word in a broad sense, primarily referring to the thick constellations of multiple elements that give us many of the things we rely on: food to eat, warmth and light, money, mobility, care, entertainment. In earlier stages of human civilization these were provided directly or with much thinner systems. Today's systems depend on complex patterns of cooperation and competition, alignment and standardization. The main thrust of this paper is to suggest ways in which this kind of thick cooperation and alignment could be enhanced.

Objective and subjective

Traditionally systems of this kind have been understood in terms of:

- their ontology – the *nature of the things* the system does, such as producing heat, mobility or income support;
- their epistemology – the *nature of the knowledge* there is about the system both outside and inside the system, and including its informational dynamics;
- their material *dynamics* – the various flows, stocks, interactions (inputs/outputs, causal links, feedback);
- their social and relational *dynamics* – the sociological patterns, psychology, culture and economics.

My interest here is just in the second part of this – the ways in which systems think about themselves and their environment, which depends on the ability of the system to represent itself, both in formal terms and in the understandings of the people working in the system. All systems combine an objective and a subjective side. They exist objectively when the behaviours of their parts are connected to a significant extent, so that changes in one cause changes in another.[3] But they are also very much subjective constructs: what we focus on, and see as the system, depends on what we believe to be important. As a result, all systems are in part shaped by their representations, which can be either endogenous or exogenous.

How these representations are organised then links to the question of how decisions are made. Discussion of systems often gets confused because of lack of clarity on this. A distinction which I have found useful distinguishes four basic types of system: **deterministic systems** where neither the parts nor the whole can display choice, such as a clock, and where the representation sits outside the system; **ecological systems** where the parts can display choice but the whole cannot, which is true of most of what we call nature, some parts of which generate representations of their bodies and states; **animate systems** where the parts cannot display choice, but the whole can, with our own bodies being a good example, or the combination of a human and a car; and **social systems**: where both the parts and the whole can display choice, which is true of organisations and nations, and where representations can usually be found at multiple levels.[4]

My main interest here is in the last category, or to be more precise the everyday social systems which also include elements that are deterministic, ecological and animate. In evolutionary terms these are systems that are both adaptive as systems and that include agents following their own adaptive strategies. The main premise of the paper is that we should want more of our systems to move from being ecological in this sense to becoming more social in this sense – ie conscious about their design and choices – and that this in turn requires more effective ways of representing their dynamics.

Intellectual resources – research, disciplines and practice

There are many useful intellectual sources to draw on to guide a more mature approach to systems. There is an emerging practice of **'systems of systems' engineering**,[5] much of it focused on the military, space, critical infrastructures and enterprise information systems within companies. Within some of these fields, particularly within many large companies, systems are well-described and are designed, managed and maintained by skilled teams. The **digital twin movement** represents one important programme to make systems more visible, particularly in relation to physical flows, and has a big influence in manufacturing, Internet of Things, risk management[6] and more recently urban planning, spurred amongst others by the Centre for Digital Built Britain. The usual definitions of digital twins describe them as a 'digital replica of a living or non-living physical entity' or 'a realistic digital representation of assets, processes or systems in the built or natural environment'.[7] However, I'm not aware of any of these that have yet integrated the social or subjective dimension of systems – psychology, sociology, power – which becomes vital for fields such as healthcare.

There are many **bodies of practice** which automatically think in terms of systems. Some can be found within mainstream consultancies (eg McKinsey and offshoot Quantum Black), big tech firms (particularly ones like Amazon with a strong physical as well as digital presence) and organisations rooted in structural engineering like Arup which have specialised in holistic approaches to improving the built environment, particularly focused on energy and transport.[8] Within government there are plenty of systems thinkers: the UK government, for example, has a Systems Thinking Interest Group.[9]

There is a growing body of evidence and experience around **linking data** to speed up observation and analysis. South Korea and Taiwan, for example combined data from private companies (mobile phone, credit cards etc), government, NGOs and others to respond to COVID (I discuss these in more detail later). These models are harder in jurisdictions with more restrictive rules on personal data, but point to the potential for new models that link anonymised or semi-anonymised data.

The last 20 years has seen steady advance in understanding both how to **synthesise and distil complex bodies of evidence**, and then how to make it more easily used by front-line decision-makers across a system (eg through the work of NICE and EEF in the UK; the new International Public Policy Observatory aims to further develop these models of active interation between supply and demand for knowledge). Advances in the field of 'research on research' parallel these: allowing funders to visualise funding landscapes.[10]

The rise of **knowledge graphs** on the web makes it easier than ever to visualise how sets of objects, ideas, publications, or interests connect to each other, revealing communities of interest. Google Knowledge Graphs (which appear next to every search) are the most visible version of this, and there are many more niche examples like are.na.[11]

There is long experience in how to organise the **human networks** that are often vital to ensuring that new knowledge is shared and made use of across systems, with many examples ranging from health collaboratives to study circles.

There is a large literature now on **socio-technical systems** that attempts to make sense of big, but multi-level, patterns of change, for example from a fossil fuel to a zero carbon economy.[12]

Some of these methods focus on the purely technical aspects of systems. Others recognise the links between the objective and the subjective, and the ways in which real systems generate views of:

- **What is:** the facts around it, data, observations, dynamics
- **What matters:** in terms of goals, but also which facts to attend to
- **What's possible:** the direction of travel both of the environment and the system (which in turn influences what matters in observation of what is)

These views can be implicit or they can be made explicit.[13] Richly intelligent systems make these explicit, and therefore open to challenge and development. In the language of John Warfield they encourage the combination of sciences of description, design, complexity and action, mobilising the full collective intelligence of the system, integrating 'our capacity to share meaning using words, represent causality using graphics, and model complexity using mathematics'.[14]

The crucial point is that these three dimensions are looped together: none of them exists free from the influence of the other two. Which facts are attended to depends on what is seen to matter and what is possible. What matters depends in part on the messages from the facts and from the possible future (such as potential threats). What's possible depends on the dynamics of the present and a view on what is more or less desirable.

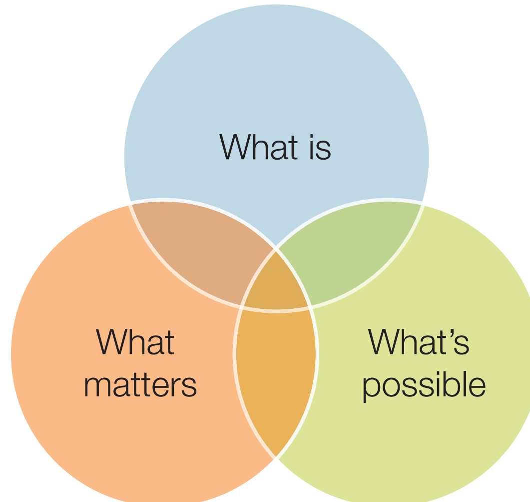

So, looking at systems in this way opens up how the system understands what is; how and who decides what matters; and how it connects actions now to possible future states. In all real human systems the answers will be fluid and contested at different points of the system. For example, views of what matters in a care system will be very different at the top of the hierarchy, amongst frontline professionals and amongst receivers of care. In energy systems there will be very different views about the relative importance of low prices, carbon reductions, national security or addressing fuel poverty. Periodic moments – such as budget setting or decisions on regulation – bring these conflicts to the surface and crystallise answers.

3. Visualising and making sense of systems

It should already be clear how important it is to address the system's representations of itself: how it addresses its own ability to decide what is, what matters or what lies ahead. Real systems are of course far more complex than our minds can grasp: but without ways to represent them we cannot hope to shape them. So we need simplified models to get hold of their dynamics.[15]

I have had many roles in governments which aimed at addressing complex problems and always tried to start any project with teams working to prepare systems maps that would describe the crucial dynamics and feedback mechanisms for such things as education, neighbourhood regeneration or reducing obesity. But we usually had to start almost from scratch.[16]

To drive a car, by contrast, there are rich tools already available: there is the view of the road ahead and behind; feedback from the speedometer; fuel gauge; engine diagnostics; GPS and Satnav and so on. The road system is visible, comprehensible and, through tools like Waze, increasingly self-aware.

The systems we were dealing with had relatively little comparable. Much the same is true now, even in those systems where there is plenty of data, statistics or organisational diagrams, or research projects that provide one-off systems maps. In public systems, there may be descriptions of responsibilities and governance roles, and often excellent statistics. There may be registers of assets. In some cases, there may be rudimentary mapping of flows, for example of young people through the education system, offenders through criminal justice or patients through healthcare. There may also be accounts of how multiple factors affect particular groups (like children at risk). But these rarely amount to systems maps.

Some of the causes of this are organisational. For example, data in the care system or energy systems is largely held within private companies, and not shared with others let alone the public. Some of the causes reflect everyday practice – such as gathering data in non-digital formats or non-machine readable forms. But another obvious reason for this gap is that it is no-one's job to fill it.

These problems became very apparent as the COVID-19 crisis hit and governments had to improvise intelligence systems to draw on data, models, predictions, experiments and scientific knowledge.[17] As I show later, one important factor differentiating the strong performers from the weak has been how well they organised systems intelligence.

So if you search for a systems map of any of the everyday systems that exist in your neighbourhood, town or city – for waste, education, water, clean air, ecosystems management – you are likely to be disappointed.

That said there are plenty of methods which *can* be used, and visualisation methods have advanced dramatically in recent years. Some approaches literally map the elements of the system: the hubs and nodes and how they link, flows of data and stores of processing power or memory. **Classic engineering** approaches emphasise the links between functions, form, structure and architecture. **Systems dynamic**s usually emphasises positive and negative feedback, and the roles of stocks and flows. **Systems engineering** approaches look at behaviour, function, functional architecture, dynamic interactions and interaction with an external environment.

Any system can be mapped in these terms, covering various flows of data, commands and resources; feedback loops and connections into shared memory or processing capacity. So, for example an energy system providing homes with electricity will have a wide range of control systems monitoring supply, peaks and outages. Increasingly there are complementary networks (not linked to the first group) managing home use, such as Hive and Nest. In an eldercare system the data and knowledge flows are thinner: instructions, regulations, market messages, inspections and rudimentary data, but little sharing of data or knowledge of any kind.

Other approaches to mapping systems see them as **interacting agents**. These are increasingly used in some fields to plan complex interactions and logistics for example. But few are used to make the system legible to the people within it or dependent on it.

More **conceptual maps** of systems are often used within leadership groups to think through dynamic interactions – but none of these are used in the way that a control panel or dashboard is.[18] The best use visualisation to bring to the surface patterns that are otherwise invisible;[19] however if these are purely conceptual they will tend to reflect the implicit assumptions of their designers rather than anything objective. There is also a familiar genre of systems maps that cover the interactions of digital technologies to guide decisions in the system.

Finally, there are parallel **bottom-up traditions**. Some focus on how systems are seen from inside and below: gossip, complaints, jokes, workarounds. All real human systems and organisations can only be fully understood by taking this informal communication seriously,[20] particularly since it has become much better organised and more visible thanks to social media providing a running commentary on almost everything. This is the relational 'dark matter' that is in practice so decisive for human systems and communities.[21] Others focus on new ways of organising democratically, formalising processes in various forms of holocracy.[22]

The diagram below shows some of the methods described above: with a classic conceptual energy system map on the top left; a conceptual map of eldercare on the top right; an agent based modelling system on the bottom left; and a digital system diagram on the bottom right (the latter three all for eldercare).

There are also many methods available to **animate systems** and enable the people within a system to understand its dynamics. Half a century ago it was assumed that these could be done with models, but most important systems also depend on a human element. So more recent methods have attempted to combine formal modelling with an experiential aspect, including simulations and games, eg for pandemics and disasters.[23] At their best these help people to move between abstraction and details, zooming in and out, and linking the organized and the self-organised aspects of systems.[24]

A good example is the labour market. I have been involved in quite a few recent projects which aimed to make the workings of labour markets more visible, and tractable: using current data to show what skills and jobs are being demanded and at what pay levels; using sophisticated forecasting to show which are most likely to grow or shrink; and then turning these into useful tools for individuals or governments to guide their decisions. This is one part of the map (opposite page), showing the likelihood of different jobs being affected by machine learning.

Figure below: Some visual representations of systems

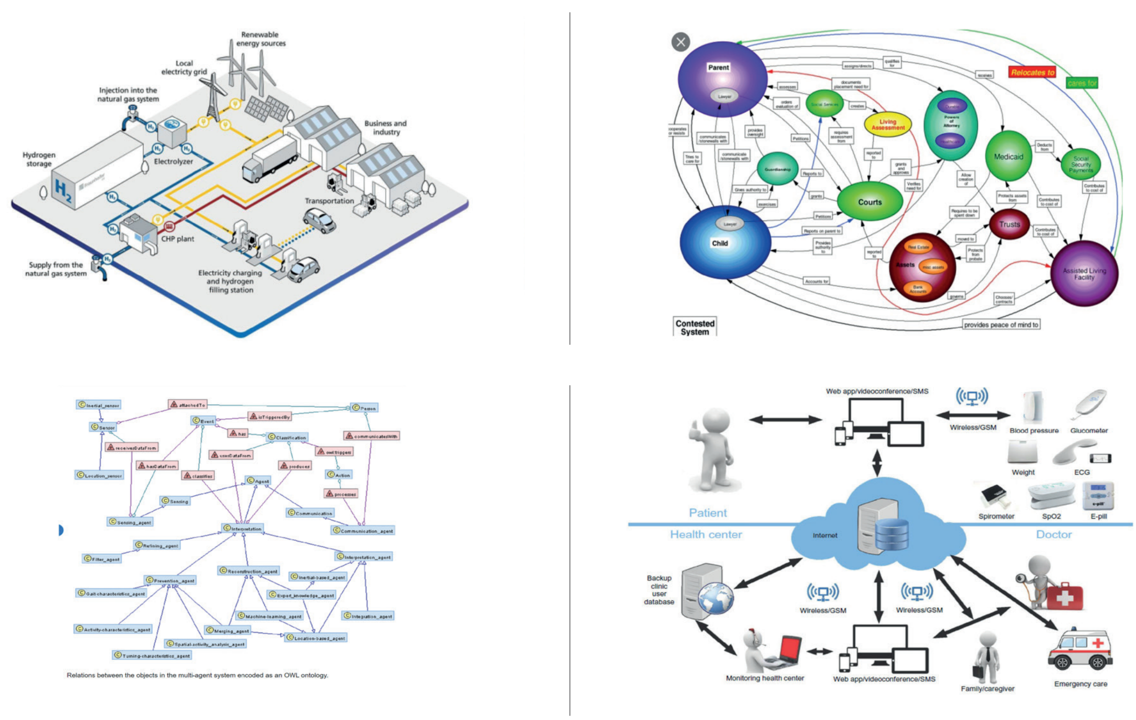

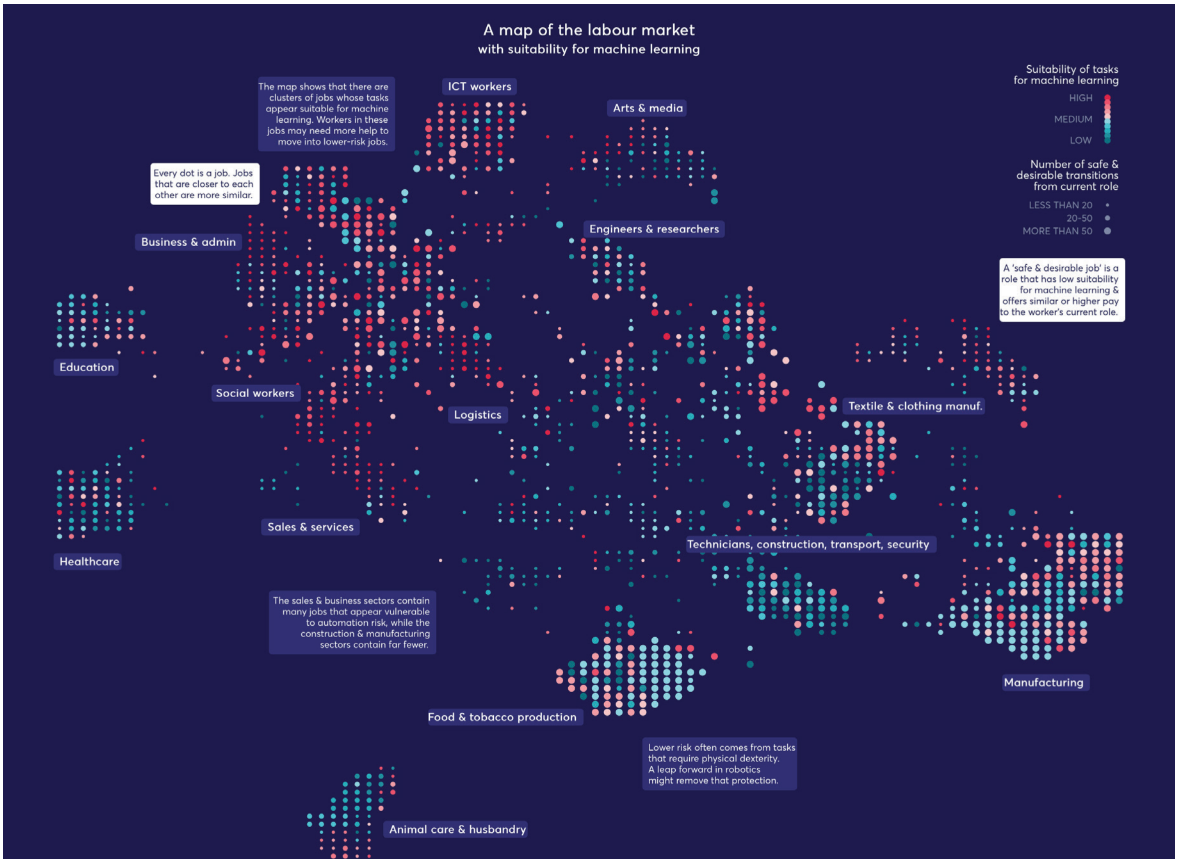

Figure above: A map showing which jobs are likely to be affected by machine learning

It can also be useful to think of these maps in less visual terms – as providing narratives of what is happening, what is important and what could lie ahead. A good recent paper analyses how the Bank of England does this using its networks of agents to provide a living picture of the state of the economy – combining data, surveys and conversation to generate a shared picture precisely of the kind I have suggested, ie covering what is, what matters and what lies ahead.[25]

The US intelligence agencies 'Intellipedia' project was another attempt at orchestrating such a shared view (though with many more constraints given the reluctance of agencies to share knowledge).

All of these different kinds of map have to decide how to handle feedback, and which to prioritise, whether:

- Observational feedback (eg what is happening to recruitment in a particular sector?)
- Performance feedback (how well is an agency achieving its goals)
- Evidence feedback (what works, where and why)
- Feedback from lived experience (from consumers, employees)
- Values feedback (are actions aligned with values?)
- Environmental feedback (eg sensors measuring air quality)
- Peripheral feedback (eg neighbouring industries, technology fields)
- Futures feedback (eg from scenarios, foresight etc)

One definition of a richly intelligent system is that it makes use of a wide range of types of feedback, across a variety of spatial, temporal and organisational scales.[26] Meanwhile its maps and visualisations will meet the four basic tests that any representation needs to satisfy:

- **Truth** – the map needs to be accurate, to represent relevant truths about how the system works, even if it has to simplify and abstract.
- **Variable granularity** – the map needs to make it easy to jump between the abstract and the concrete.[27]
- **Sense** – the map needs to make sense to its users, and there will often be a trade-off between truth and sense.
- **Use** – the map needs to be useful both for current operation and for planning ahead. This means focusing on points of action, intervention and leverage; and it usually means making the map more detailed than it was only required for sense-making.

4. Seeing shared intelligence in cognitive terms

All of the tools described so far are ways of supporting and organising *shared intelligence*. Before going further into methods it's worth acknowledging that even to talk of shared intelligence in this way can be theoretically controversial. Some fields believe that thought can only happen in the individual brain (this is true of much psychology and much economics). Contrary traditions believe that thought does meaningfully happen at larger scales – in sociology (from Durkheim onwards), anthropology (Mary Douglas et al), philosophy (John Searle), science studies (Bruno Latour and ANT), through to evolutionary biology (Joseph Henrich), and it is increasingly common to recognise that aspects of intelligence are organised at large scales both in digital technologies (databases etc) and in group minds, and indeed that most thought combines human and machine (as we use laptops, phones, pens to support us).

Such approaches shouldn't skate over the complex dynamics of shared cognition: the competition between different worldviews and interests and the pressures towards conflict; the tendencies towards deliberate disinformation, disruption and misinformation. These matter – and are why guardianship and curation of intelligence are becoming so much more important.

The cognitive approach also has one other important feature: it doesn't privilege particular scales or levels. In other words it doesn't assume that macro features of a system cause micro ones, or vice-versa. Within any real system the causal links between micro, meso and macro need to be investigated not assumed (a similar conclusion is increasingly common in neuroscience of the individual brain, and even in biology).

"Such approaches shouldn't skate over the complex dynamics of shared cognition: the competition between different worldviews and interests and the pressures towards conflict; the tendencies towards deliberate disinformation, disruption and misinformation."

5. Shared intelligence as a commons

In these and other examples the collective shared intelligence of the system has to be, to some extent, organised as a *commons* – with shared use, contribution and access to a common body of living data and knowledge, as well as the guardianship and care that is essential to all commons. This kind of openness, which also enables tapping into more sources of intelligence, may often improve overall performance. It brings the system closer to the best of the human brain which is essentially organised as a commons – albeit full of competing as well as cooperating modules.

Within some firms – such as Amazon or Alibaba – highly proprietary intelligence assemblies can be very efficient in ensuring customer service and profit, and are effectively treated as a commons within the organisation. Few if any public services have anything comparable. There are many reasons for this including legal restrictions on data-sharing (not a new issue: I oversaw a government review of how best to balance privacy concerns and the wider public interest in data sharing back in 2001), as well as authority and capability. For critical infrastructures there are complex questions of security involved – which require data to be closed, or managed through time windows, or with ontologies setting variable rules.[28] In fields where privatisation has made data and knowledge much more balkanised – firms running electricity companies, prisons, care homes or employment services – have few incentives to share. In other fields there may be incompatible systems or strong traditions of silo working.

As a general rule intelligence is helped by openness and sharing – just as the individual brain's capacity to think is enhanced by the ability to connect.

Existing systems are a long way from this vision of an interconnected commons. In the energy example observation is organised at the level of the providers, the grid and the household. The former groups are set up to spot certain kinds of patterns, though not others. Because the data is not open it's not easy for others to analyse potential patterns in consumer behaviour or opportunities for reducing use.

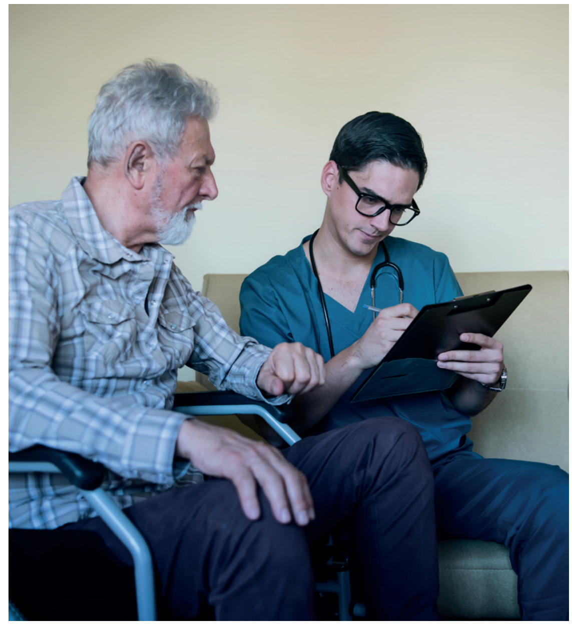

In the case of eldercare all of these elements are thin: a lot of data is collected within care homes (partly to protect against legal challenges), but almost none of this is shared, and the organisation of shared knowledge, memory and insight is rudimentary (which is why DHSC has been exploring a more systematic 'what works' centre for adult social care).

Pragmatic improvements can be made to almost any system through diagnosing how well it undertakes the functions listed earlier – observation and memory through to creativity and judgement – and making improvements. But in far too many cases no-one has the remit, resources or responsibility to do this.

6. COVID-19 as a live test of shared intelligence

This has again become apparent during the COVID-19 crisis when the linking up of memory, observation, models and experiment has been key to effectiveness, but very uneven. Many governments around the world have had to improvise precisely these kinds of intelligence arrangements: linking multiple data sets; using models to predict patterns of spread; running quarantine systems; distributing money; and experimenting fast. And they have required an 'integrative intelligence' to make complex judgements about options and trade-offs.

These images capture some of the elements: COBRA-type central command capacities; Singapore's TraceTogether app which enabled widespread tracking in close to real time; the use of models to predict patterns; and then sophisticated digital tools such as Taiwan's digital fence.

Figure below: Responding to COVID-19

A few countries like Taiwan have been particularly adept at mobilising collective intelligence in many forms with transparent use of technology and science, very much in the spirit described above, contributing to successful results. Like South Korea they have been able to use a lot of commercial data – from mobile phone companies, banks and credit card companies – in ways that would be difficult in the more privacy-sensitive EU. But they have also opened up governance – with far more public engagement in decision-making through initiatives such as vTaiwan that's run by the government and parliament.

The next diagram summarises some of Taiwan's approach, again a hybrid of classic materials stock and flow dynamics and more attention to social factors:

Figure right: A systems map of Taiwan's response to COVID-19

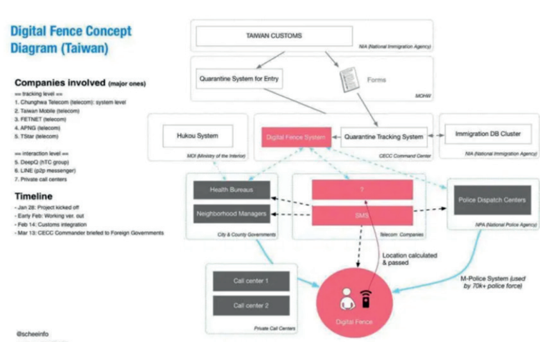

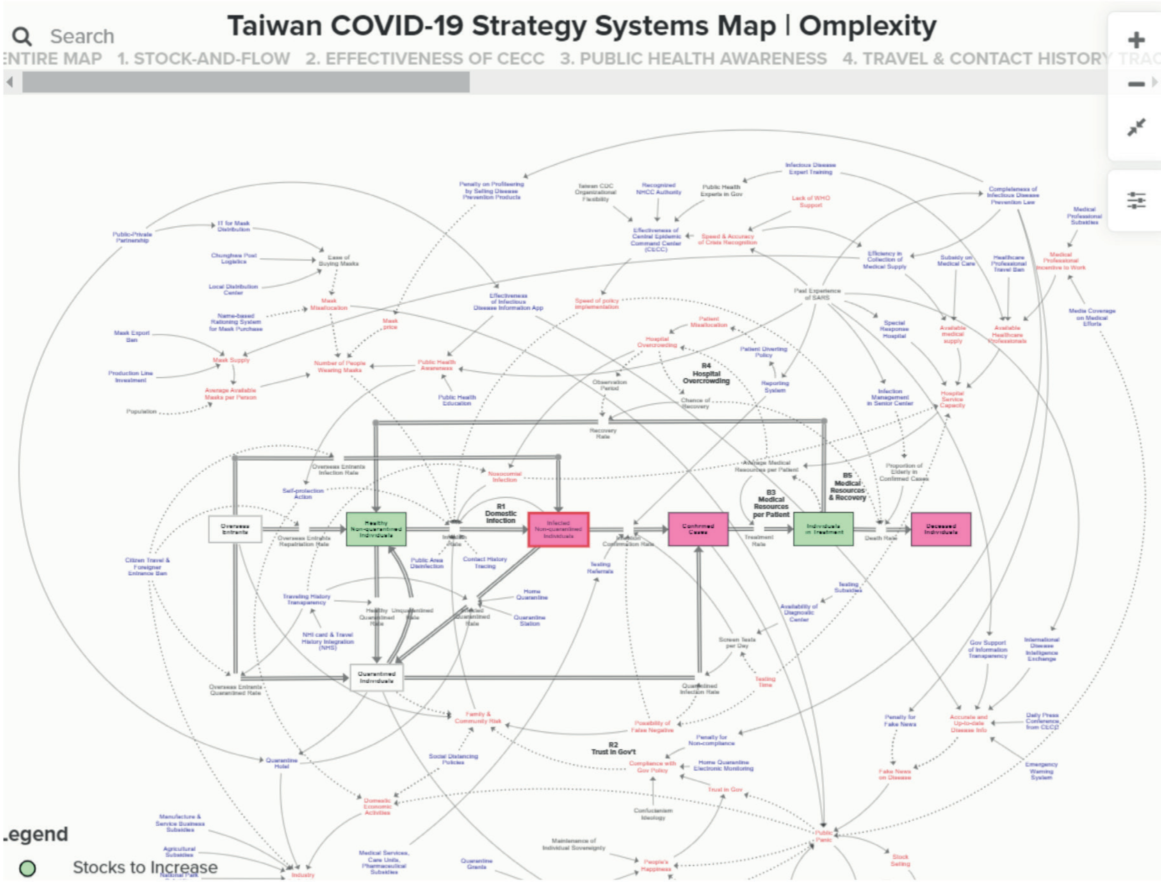

Bangladesh is another interesting example of the conscious development of intelligence assemblies to assist their response to COVID-19, led by the government's A2i team. So far the country has been relatively successful in handling both the health and economic sides of the crisis:

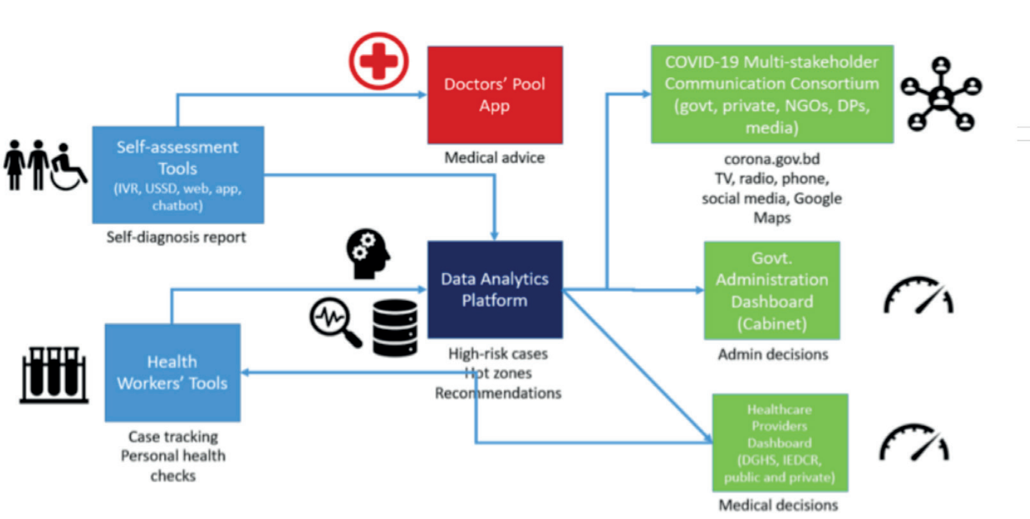

Many governments have struggled. In the UK, for example, although there are strong capabilities these are not well joined up. There is neither a responsibility nor an obvious capacity for this kind of systems approach. Responsibility sits uneasily between many different agencies – functional ones like Public Health England and the NHS; the Chief Scientific Adviser, GO-Science and SAGE; Cabinet Office; ONS; NHSX and others. No 10, Cabinet Office and Treasury are now building up a team that combines data and evaluation skills. But so far during the crisis no part of government has had the skills and methods, and the mandate, for making sense of the complex interactions between health, economic, social and other dynamics, and ensuring that key intelligence gaps are filled. The success story of vaccines actually proves the point, since it is a story of smart action but within silos: the early sourcing of vaccines on the one hand, and the very efficient mobilisation of NHS structures on the other.

Some of the problems faced during the COVID crisis were exacerbated by a more basic weakness of UK government data – the proliferation of incompatible databases, from HMRC and NI numbers to Government Gateway, NHS, drivers licenses, Verify, pupil numbers, electoral rolls and others like DBS checks. These make the UK situation very different from countries like Estonia which has a single identifiers, or India with its biometric Aadhaar. Standards that allow for data-sharing and interoperability can make it far easier for systems to become intelligent in the senses described in this paper. There has long been opposition to having any kind of single, biometric identifier in the UK because of understandable concerns about privacy and civil liberties. But there are now technology options available that allow combinations of interoperability and strong protections for personal privacy. Estonia's X-road software – which is the integrating backbone of their system – uses strong encryption and also records any sharing of data between different databases to avoid abuses.

7. Analysing and diagnosing shared intelligence

So what could be done to help our everyday systems make the most of intelligence of all kinds, whether for pandemics or other challenges like cutting carbon? My first recommendation is to break down the different aspects of cognition as a tool for diagnosis and prescription. It then becomes easier to see how these are organised; how well they are managed; and what could be done better.

As in individual brains, intelligence is an assembly of multiple elements:

- **Observation** – gathering of data, from many sources, such as administrative, open, commercial and sensor data, citizen input, lived experience, political feedback, complaints, media coverage and so on.[29]
- **Live models** – functioning models of how the system works and patterns of causation, whether for how energy might flow through a network or for how particular treatments work in a care home.[30]
- **Analysis** – spotting patterns and making sense of them, for example through research, and using many potential tools such as semantic analysis of social media or case notes, or use of neural networks and objective-oriented techniques to improve the classification of objects and distinguish noise from edge cases.
- **Prediction** – from machine learning to scenarios and simulations, or use of agent-based models.
- **Memory** – shared knowledge of what has or hasn't worked in the past and why, organised in repositories, databases, books, journals or the minds of experts.
- **Creativity and experiment** – generation of novelty, sometimes with shared real time learning, and sometimes assisted by technology (such as machine learning for discovering new proteins or recommendation engines for collaboration).

- **Judgement** – including the design of rules, policies, budget allocations, or borderline decisions on individual cases (for example on entitlement, or assessing a mistake).

- **Wisdom** – how all of these kinds of knowledge are integrated, including ethical and other dimensions (I cover what this kind of integrative intelligence means in much more detail in my recent paper on 'Loop Theory of Wisdom'[31]).

These together constitute the intelligence of the system.[32] In the case of everyday systems, each of these will be a combination of human activity and machine activity – often with greater machine intelligence around observation and prediction than the other elements. In many everyday systems their organisation is uneven; separated; not aligned; or not integrated. For example, data is collected by different professions and disciplines with very varied views both of what is and of what matters.

Seeing systems intelligence in these terms prompts useful and practical questions which avoid the traps of over-abstraction which can bedevil systems thinking. Some of the most fundamental involve data: is there reliable, comparable, open data based on good standards; is it available in machine readable form; is it stuck within organisations or shared.[33] Most everyday systems lack even these basics (often because the data is proprietary, owned for example by an electricity company or a retailer).

Then you can ask what observations are used. Are they the right ones? Are there key gaps, degraded information or misinformation? How are patterns interpreted? Whose job is it? What are the models that guide the system? How explicit are these? How reliable? What is predicted – and how successfully? What is the relationship between tacit and formal knowledge?

What memory is used – whether codified or tacit – and are there crucial gaps? How is creativity and novelty generation organised, and backed with resources and tools for assessment and scaling? How does practical insight feed back into the system? How is peripheral vision organised – spotting patterns, potential risks and non-obvious lessons? How are narratives used to support systems awareness of themselves and their options? How is intelligence integrated and by who?[34]

An odd feature of these questions is that we lack even a name for the people who are specialised in asking and answering them – as a result they tend to crop up in sporadic consultancy engagements rather than being part of the normal life of organisations.

"Seeing systems intelligence in these terms prompts useful and practical questions which avoid the traps of over-abstraction which can bedevil systems thinking. Some of the most fundamental involve data: is there reliable, comparable, open data based on good standards; is it available in machine readable form; is it stuck within organisations or shared."

8. The economics of intelligence

One of the values of a more systematic approach is to link these kinds of analysis to economics, since each aspect of systems intelligence involves costs and opportunity costs (just as the human brain is 'costly' in terms of energy use), and there may be incentives to hoard, misinform and deceive. Economics has many tools for understanding transactions costs though few for understanding the kinds of cognition described here.

However, with any system it is possible roughly to analyse what resources are devoted to different elements of intelligence – for creativity and innovation for example, memory or analysis. It can sometimes be feasible to map and measure the cost of actions taken to preserve integrity of data (with cybersecurity) or for interpretation and judgement (for example with open use of evidence).

These may point to alternative options that would raise marginal returns from investment in intelligence. Within business markets dynamics have tended to push towards high spending on data and knowledge infrastructures with the benefits seen in profit. For public services and systems, however, there are few comparable incentives and so a tendency to under-invest in these knowledge and coordination functions, or to direct money to traditionally powerful interests (this is why, for example, in health vastly more is spent doing R&D on new drugs than on digital or behavioural interventions).

There is an obvious relevance to AI. Most contemporary systems today use algorithms at multiple levels from service interfaces to infrastructure management, logistics to planning. But we lack good theories and methods[35] for understanding how to connect human and machine intelligence within an overall system (though there are some promising developments with interactive machine learning, transfer learning, crowd-assisted machine learning and other methods).[36] This creates challenges for design but also for understanding. Much of AI will never be comprehensible or legible; but situating AI within the context of how systems are made visible, may be vital in helping to make AI more democratic.

In many important systems we quickly discover that it's no one's job to attend to these questions, or we find that the roles are fragmented without connections being made. In other words, there is a missing 'integrative intelligence.' Yet connecting up the different functions of intelligence in the way that a brain does – through the kind of 'intelligence assembly' pictured below – often generates significantly greater capacity to think and learn.

Figure below: Intelligence assemblies

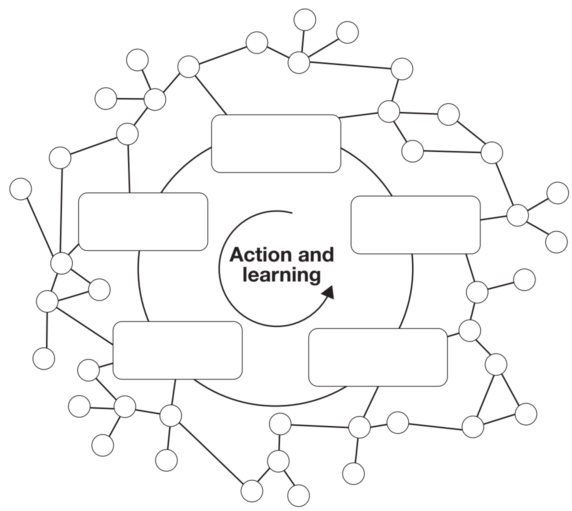

9. Developing representations and shared understanding with system stakeholders

I've already discussed the value of visualisations and representations. For any everyday system we should want there to be such shared representations in forms that can be interrogated and used: a representational twin or mirror of the world we live in.

These can never be perfect or perfectly aligned, as the representation has to be simpler than the reality. But competent systems can continuously work to ensure their representations are more richly accurate, as is their self-representation – describing back to the parts of the system how its cognitive processes are organised, and so prompting attention to how they may be improved.

This is one of many spaces where design skills can be so useful: helping with both static and dynamic visualisations that link the huge complexity of real systems to the questions that people care about most or have most potential to influence.

Next generation knowledge management tools may also move closer to this ideal: eg Project Cortex, which is trying to advance Microsoft Sharepoint and Teams to more dynamic use of knowledge, combining AI and human collective intelligence, so that the system's knowledge is visible and accessible to itself.

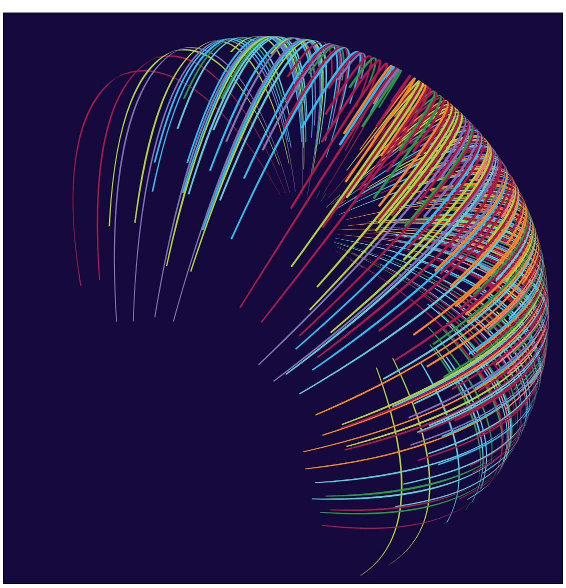

This is the ideal of some models of manufacturing going back to W Deming; of recent democratic theory; and is the animating idea behind today's digital twin projects. It requires much of the representation of the system to be opened up and treated as a commons or public good.

One ideal is a system where many of the actors within it share common understandings and knowledge, and learn in tandem. The opposite is a system where knowledge is tightly hoarded and therefore adaptation depends on central action.

There are now many methods available – many used by design and innovation labs – that bring together the stakeholders of a system to describe it together, through a combination of workshops, visualisations and live models, including social network analysis of the existing social dynamics of the system. These generally require the support and authority of a sponsor – a peak body, government or regulator. They also take time. One example I worked on was in relation to the UK fishing industry – a combination of mathematical models of fishing stock dynamics; social models of the dynamics of fishing communities; economic analysis of fishing fleets – all turned into visual form to clarify the crucial strategic choices faced by the system and feeding into face to face meetings involving the communities with most at stake.

These representations can also show the triggered hierarchies that determine how the system responds to threats. For example, the human body has many unconscious processes but represents more serious ones when they go wrong (ie pain) in order to trigger correction, or in extreme cases high fevers. Within an energy system some kinds of problem trigger very local solutions; more serious ones trigger involvement of higher tier authorities, while the most serious ones bring in authority from other fields. Ensuring these are robust is a vital part of making systems resilient against risk.

10. Making systems more inclusive

Another virtue of seeing intelligence in systems as a commons is that it opens up the scope for more inclusivity, and more democratic influence. Most everyday systems offer only limited opportunities for their beneficiaries to shape them. Most of the key design parameters are shaped top down. There may be rights of exit (in the case of utilities); or harvesting of data on consumer preferences to shape services. But there is little responsiveness to many interests and voices, and little mobilisation of insight, lived experience and other kinds of knowledge.

A key test of any system is whether there are valued outcomes or possibilities not represented in its internal representations. Recognising these is the key to improving systems – for example making energy systems more responsive to new issues such as carbon emissions; or making care systems more attuned to psychological needs, or issues such as isolation.

These can be imposed exogenously. Public policy can set obligatgions, constraints, tariffs and regulations; pressure from markets, investors or consumers can also impose new priorities on private firms. But in each case this kind of accountability depends on good representations which are often lacking. A typical UK citizen, for example, has few ways to hold the major system controllers – the companies providing electricity or broadband – to account for their actions. Customers can switch to another provider or make a customer complaint. But it's hard to find out much about their impacts on the environment, how they treat different categories of customer or what they do for the local economy.

There is also a wider issue. For most systems there will be stakeholders beyond the system's borders.
So there is a moral as well as a practical reason to want some accountability to them – or at least visibility. Such accountability also helps learning. This is very visible in some fields – such as airlines – with rules on reporting accidents, near misses, and the real-time intelligence now collected on fleets of aircraft by Airbus and Boeing. Some equivalents exist in other fields, like CROSS for buildings, and other systems for confidential reporting of problems. But again, this kind of deliberate mobilisation of a wider network of intelligence, which acknowledges the stakeholders beyond the borders of the system, remains the exception rather than the norm.

"A typical UK citizen, for example, has few ways to hold the major system controllers – the companies providing electricity or broadband – to account for their actions. Customers can switch to another provider or make a customer complaint. But it's hard to find out much about their impacts on the environment, how they treat different categories of customer or what they do for the local economy."

11. Helping systems to learn

Here we have already touched on ways to help systems to learn better, and how this too is heavily dependent on how much intelligence is organised as a commons. Again, learning can be externalised – treated as a role for central commands, external inspection and evaluation. Or it can be internalised and embedded. As a general rule, the more that learning can be embedded the more likely it is that systems will adjust intelligently.
That learning will take at least three main forms:

- New data feeding into essentially stable models (which is the vast majority of everyday learning) but where the value may be much greater if data and lessons are shared.
- New data that challenges the models and triggers development of new models or categories (which generally has to come from outside the system as well as inside).[37]
- New challenges that force the creation of new systems of cognition (which again involves outside actors).

Some of this is classic failure mode analysis involving the design of mitigation options. But in investigating any real system we quickly come across many barriers to learning; unwillingness to admit or share problems (especially perhaps with complex subcontracting processes); lack of open data; lack of clarity on who has responsibility for orchestrating the more complex types of learning.

The most effective models are ones that institutionalise learning at multiple levels – from making sense of surprising data to encouraging use of new knowledge. For example, Study Circles in schools organise regular sessions for teachers to reflect on new observations, data and evidence. Many health services have adopted similar models, again drawing in part from the theories of Deming, Nonaka and others. However, most public services and systems lack these simple devices.

Some recent innovations in innovation itself can be seen as ways to improve learning, especially where experiments and results are made visible, including in pilots, testbeds, sandboxes and more advanced models like the Climate KIC Deep Demonstrations. Part of their role is to widen the range of options available – recognising that they are closer to ‘infinite’ than ‘finite’ games.[38] Without these tools bureaucracies tend to be bad at generating possibility.

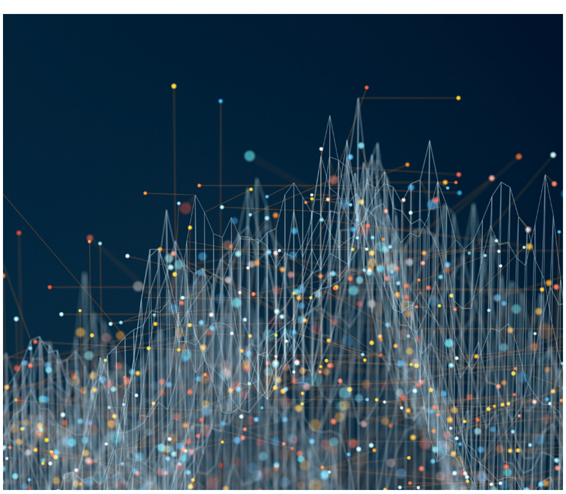

12. AI within larger systems

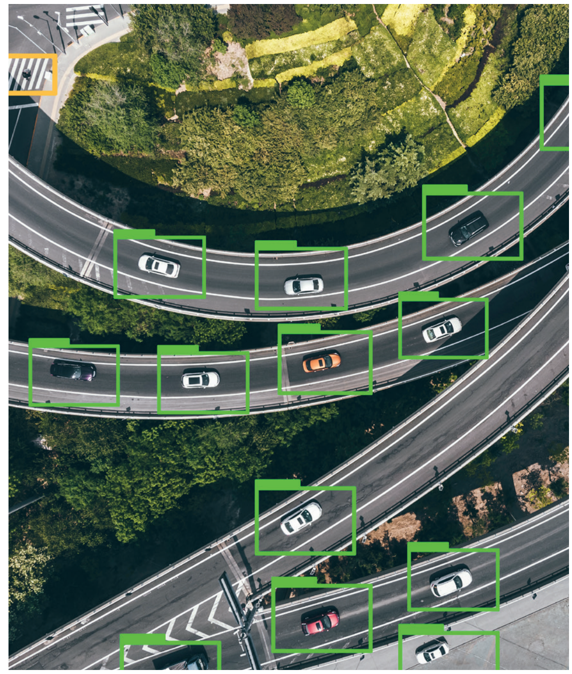

Most systems now use AI at multiple levels from chatbots interacting with the public to algorithms to allocate resources. Much of the programme already described in this paper has to include a place for algorithms of all kinds, which raises a family of related issues:

- Describing the relationship between algorithmic and human decision making in different parts of systems
- Knowing which tasks are and are not suitable to machine learning
- Knowing how to use AI to supervise human decisions (eg judges or doctors) and vice versa (eg hate speech online)
- Encouraging transparency and explainability of algorithmic decisions (as the EU intends – while recognising how hard this is in practice)
- Using new combinations of human input to guide the training process for algorithms

My view is that programmers should wherever possible be encouraged to see their work in more systemic ways – with intelligence as an outcome – rather than focusing exclusively on particular tools as inputs. For a much more detailed account of what this means for AI in relation to systems see link in this endnote.[39]

One key part of this will be advancing work to combine AI and CI, collective intelligence. Nesta has been funding a wide range of projects in this space, and commissioning more conceptual work too.[40] This was a gap in the research agendas around AI (for example of the Alan Turing Institute) but is now belatedly being filled. These considerations will be particularly important in social contexts where we have learned much more about both the potential, and limits of, machine learning, whether because of problems of bias or simply poor predictive power, as shown in the Fragile Families Challenge or the more recent work of the What Works Centre for Children's social care.[41] In systems of these kinds it is even more vital to ensure that there is sensitivity to lived experience and tacit knowledge as well as formal data and evidence.

13. Practical application: who can do this?

Systems thinking can risk being quite abstract, and impractical. Many who have become enthusiastic about the premise of working more systemically become frustrated when they have read the literature and struggle to put it into practice. Here I suggest some of the potential roles and tasks to be done in taking this work forward, ideally in relation to projects aiming to improve the operation of the many everyday systems mentioned already, at the level of cities, neighbourhoods or whole nations.

The role of governments

A primary set of users for these methods and frameworks are people with some responsibility for whole systems. These include officials and ministers in national government responsible for fields such as eldercare policy or education; officials and elected politicians in city regions with responsibility for fields like adult skills or economic policy; local authorities responsible for waste systems, clean air (and other fields such as care); regulators and utilities with responsibility for infrastructures.

Over the last few years many have become accustomed to the idea of having dashboards and datasets which help them track performance indicators. What's proposed here takes that idea further by making all of those open and shared, and linking to the broader role of government in steering, summarised in this diagram taken from a recent paper.[42]

To operationalise these ideas it is generally useful to separate out key elements:

- Options for **linking data** and organising it as more of a commons (with appropriate anonymisation and acknowledging the huge practical challenges around every aspect of management of data)
- Options for **orchestrating and sharing evidence**, as well as emergent findings (where there is growing experience and practice, see for example the new International Public Policy Observatory on COVID[43])
- Options for **peer learning** and connection
- Options for **shared foresight** and scenarios to develop better understandings of coming challenges and opportunities

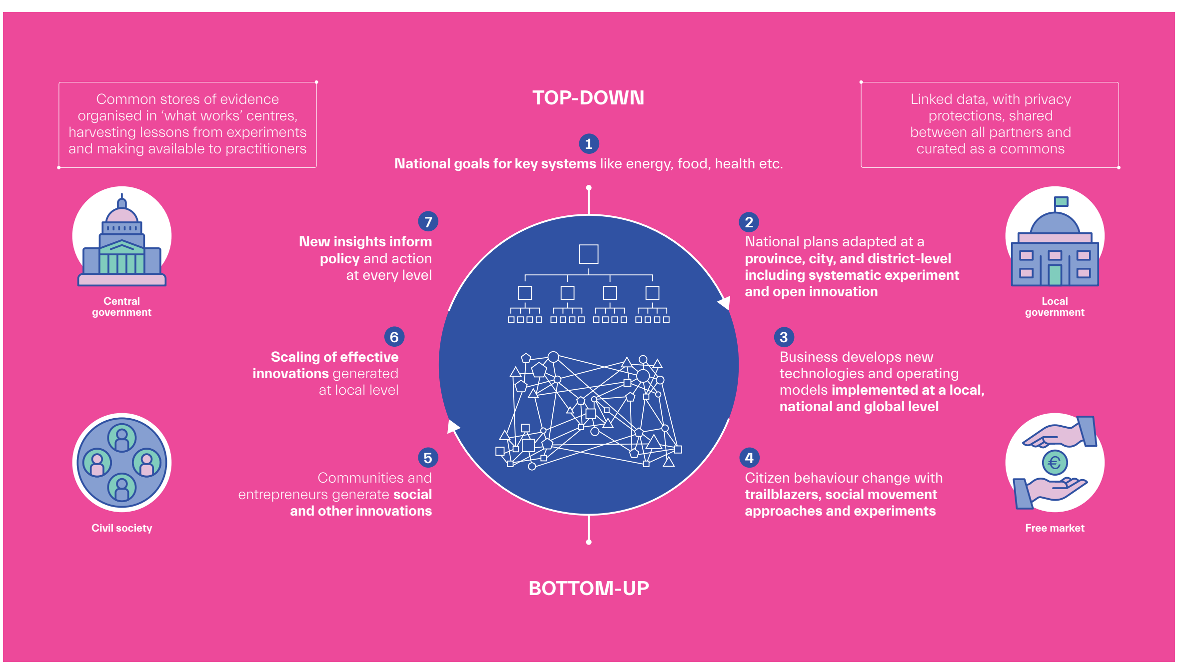

I've written elsewhere about what that might mean in practice, including for:

- **Decarbonisation and climate change**, at the level of firms, cities and nations[44]
- **Development and the SDGs**[45]
- Key fields such as **labour markets** and the urgent challenge of helping people reskill ahead of shifting patterns of job destruction and creation[46]

For government organisations the challenge is to create roles which better fit these tasks. Within the public sector Chief Digital Officers (CDOs) and other roles related to digital tend to focus on consumer-facing services – which is useful but insufficient in relation to these tasks. There are sometimes teams working on data (such as offices of data analytics), there are digital teams like GDS in the UK government, primarily focused on services; and there are often some roles focused on research (such as Chief Scientific Advisers, economists and social researchers and the Office for National Statistics).

But none has a remit to address systems intelligence. For example, during the COVID crisis if you asked who in the UK government was responsible for domestic intelligence – ie really knowing what is happening on the ground – there was no good answer. The responsibilities were split between many individuals and organisations, with only a very thin integrative capacity in the centre of government to pull these together and make sense of them in a holistic way.

Over the next ten years I hope we will see Prime Ministers, Mayors and others create teams with a broad remit to improve the intelligence infrastructures – covering data and knowledge in all its forms – to underpin more effective steering., and to build up better capacities for integrative intelligence at their core. Their jobs should be to focus on the outcome – better system intelligence – rather than privileging any particular input.

The results should not be a single plan; or complete consensus; but rather what I have called integrative intelligence,[47] and Dave Snowden describes as coherent heterogeneity[48]: sufficient alignment which still allows for diversity.

Philanthropy

Philanthropic foundations have a unique freedom to support the organisation of intelligence in more effective ways. Precisely because they lack power they can be neutral intermediaries, connecting governments, NGOs and business.

But philanthropy has been remarkably slow to collaborate[49] or put in place shared infrastructures. There are some exceptions – covered in an excellent survey[50] from SIX. A small group of foundations are becoming more engaged in data (thanks to the leadership of figures like Stefan Verhulst at NYU[51]), using new tools in their own work,[52] supporting data commons of different kinds and pooling evidence. Some of the biggest, such as the Wellcome Trust, are able to take on a systems leadership role in part because of their scale and relationships.[53] I hope that more foundations will recognise the useful role they can play in supporting the data and knowledge infrastructures of vital systems – the essential plumbing that is so often missing while far more resources go into more glamorous but less effective one-off projects.[54]

Business as system

Some parallel challenges face business. If the rhetoric around the Fourth Industrial Revolution is to be believed, the next generation of business will involve much more combination of data, processing power and physical networks. This will be relevant to housing, transport of all kinds, energy and also to healthcare.

The question then is how to do this well. First is a challenge of business models. The ability to operationalise in a holistic and systems way has become very evident in China in recent years in the rise of Alibaba, Tencent, Baidu, Meituan and others, offering families of interconnected products and services. It shows up in their very different approach to driverless cars – offering a combination of infrastructures and vehicles rather than only vehicles – and in projects like Alibaba's 'City Brain'. And it has shown up in the COVID-19 response with the creation, at great speed, of smart health surveillance infrastructures. This may reflect differences of culture (more attention to wholes than is normal in the more atomistic philosophy of the West), legacies of central planning as well as lax competition law. But its net effect may be to give China significant advantages in terms of both designing and implementing the 4IR. Some Western companies have elements of a similar mindset – from

Amazon to Ocado, Siemens and Schneider – but most do not.

A second challenge is how to reshape regulation to support 4IR systems integration in business. I have long believed that regulators would need to force sharing of data to unlock these potentials. The experiments around open data in banking have shown how this can be done,[55] and are likely to be followed up with some requirements in the European Union on big platforms to open up their data. My hope is that in a few years time it will be obvious that some of the data from things like smart meters or bus services, drones or smart health devices, will be made open in suitably privacy-enabled form, in order to allow for systems coordination, and enable the big productivity jumps that the 4IR makes feasible.

A third aspect of this is better government support for innovation in systems. I have been an advocate of systems approaches to R&D which deliberately connect the funding of R&D with policy and regulation, rather than organising missions separate from these key levers of power. This was the idea of 'Advanced Systems Agencies'. This is not a universal recipe but is arguably a better way to steer R&D linked to fields like transport or energy than general purpose research funders or DARPA variants[56].

System architects: equivalents to architecture and planning for a world of knowledge and data

Both government and business need new skills to do this work well. At present the capabilities described in this paper are divided up. Parts sit within data teams; others in knowledge management, product development, research, policy analysis or strategy teams, or in the various professions dotted around government, from economists to statisticians. In governments, for example, the main emphasis of digital teams in recent years has been very much on service design and delivery, not intelligence. This may be one reason why some aspects of government intelligence appear to have declined in recent years – notably the organisation of memory.[57]

What we need is a skill set analogous to architects. Good architects learn to think in multiple ways – combining engineering, aesthetics, attention to place and politics. Their work necessitates linking awareness of building materials, planning contexts, psychology and design. Architecture sits alongside urban planning which was also created as an integrative discipline, combining awareness of physical design with finance, strategy and law.

So we have two very well-developed integrative skills for the material world. But there is very little comparable for the intangibles of data, knowledge and intelligence. What's needed now is a profession with skills straddling engineering, data and social science – who are adept at understanding, designing and improving intelligent systems that are transparent and self-aware[58]. Some should also specialise in processes that engage stakeholders in the task of systems mapping and design, and make the most of collective intelligence.

As with architecture and urban planning supply and demand need to evolve in tandem, with governments and other funders seeking to recruit 'systems architects' or 'intelligence architects' while universities put in place new courses to develop them.

The academic challenge

Universities have a crucial role to play in training these systems architects, and in the parallel task of developing better knowledge to guide them, drawing on complementary advances, such as those being made around data, computer science, and AI; the evidence movement learning much more about how to make evidence used and useful; the many parallel fields using the words 'systems' or 'complexity'; and work on better understanding the causal links between micro, meso and macro phenomena.

There are also developments that are more squarely aligned with the approach proposed here, analysing cognition at the level of whole systems. The Collective intelligence field is growing fast, now with several professorships, centres (MIT, Carnegie Mellon; Copenhagen; Cardiff and Huddersfield), journals (in particular the new Collective Intelligence journal launched by Sage and ACM in 2021) and books. Other relevant fields and sub-disciplines include: implementation science; web science; brain science and cognitive sciences.

However, work in universities is organised in its own silos which means that even the language used in this paper means very different things, typically, to engineers, policy analysts and computer scientists. As a result there is no obvious centre in any university that is yet able to do the kind of systems analysis and design proposed here.

The world badly needs a new integrative discipline that goes beyond cross-disciplinarity and is focused on the 'how' of organising intelligence at a large scale to help solve big challenges. This has some historical echoes. As indicated there are parallels with the development of urban planning a century ago (integrating across architecture, engineering, sociology etc) and the development of business studies after WW2. Both were hybrid, integrating disciplines that had a strong link to practice and helped to make large bodies of knowledge more useful. Today's task also parallels the emergence of new fields like climate science that integrated many disciplines: '…meteorology, oceanography, geography, hydrology, geology and glaciology, plant ecology and vegetation history—to mention only some' which had 'made it impossible to work … with common and well-established definitions and methods."[59]

"The world badly needs a new integrative discipline that goes beyond cross-disciplinarity and is focused on the 'how' of organising intelligence at a large scale to help solve big challenges."

14. Dilemmas and hypotheses

A key aim of thinking about cognition and systems intelligence in these ways is to provide a framework into which other bodies of knowledge can be integrated, and then to test out hypotheses, since there is so much that we don't know. Here are a few:

- The best combinations of AI and large-scale human intelligence and how to organise it in relation to tasks such as de-carbonisation;
- How to understand the critical trade-offs, for example in prioritising different functions of intelligence, and the opportunity costs;
- The most effective methods for engaging stakeholders in mapping and design processes – ie how to handle trade-offs of breadth and depth, cost and time;
- The optimum degrees of openness and collaboration for different tasks and timescales, given that there are likely to be inverted U shaped patterns of collaboration, where too much sharing and collaboration can become as inefficient as too little;
- The incentives for citizens or front line workers to contribute data, insights and ideas, and how to understand the power laws of voluntary contribution to knowledge platforms,
- The roles of curation and relevant skills and powers needed to help systems think well
- How to handle IP, property rights, open source and creative commons elements of new knowledge
- The relative roles of institutions and relationships in systemic change (in other words how much should we focus on the formal structures and how much on informal networks);
- The characteristics that make intelligence assemblies useful and used.
- How to model and understand the effects of scale in terms of aspects of cognition and action, ie the relationship between micro, meso and macro dynamics.

15. A brief conclusion

I hope that the frameworks set out here are plausible and point to how the rhetoric of thinking and acting systemically can be turned into action. At heart much of what I suggest echoes the best in art: more attentive seeing and listening, but this time applied to the systems around us. My main claim is that this depends on:

- Making systems visible and graspable
- Making much of their cognition open, and organised as a commons
- Implementing explicit processes for learning
- Growing skills and structures that can do these well, including a new breed of systems architects

Many of these ideas will be familiar to some, even common sense. But they are quite unfamiliar to many more and are still very rare in mainstream practice. None of what I cover is offered as a panacea: but without better systems cognition all ambitions for the kinds of systems change we badly need in the next few decades are likely to fall short.

References

[1] Of course the systems that serve us can also turn us into their servants – this is the much bigger story of rationalisation and control.

[2] There are many definitions of intelligence. I tend towards ones that focus on the ability to think in ways that lead to surviving and thriving and both generating goals and meeting them. In the case of a collective intelligence or a system the definition turns to the ability of the system to serve its stakeholders.

[3] Some systems also have clear boundaries. Most of the ones discussed in this paper do not.

[4] Russell L. Ackoff, in *Re-Creating the Corporation*. See also David Sloan-Wilson's work on complex adaptive systems, which distinguishes 'between two meanings of the phrase *complex adaptive system*: a system that functions adaptively as a system (CAS1) and a system composed of agents that *separately pursue their own adaptive strategies* (CAS2).' I have attempted various frameworks for understanding in more detail the interactions between these, particularly the 'triggered hierarchy' framework set out in my book Big Mind which describes, for example, the relationship between 'invisible hand' markets, conscious steering and governance, both in normal times and during crises.

[5] See for example this useful overview from IEEE – A systematic mapping of the research literature on system-of-systems engineering – https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7151918

[6] https://www.arup.com/perspectives/digital-twin-managing-real-flood-risks-in-a-virtual-world

[7] The National Academy of Engineering in the US has just published a collection on systems which is close in spirit to this paper https://www.nae.edu/244761/Editors-Note-Systemic-Vistas

[8] A few years ago I tried to persuade some of the Arup leadership to specifically target the kinds of 'intelligence design' and systems architecture described here. They haven't done so but are as close as any organization to the mindsets and methods that are needed.
See for example this recent publication with Alibaba: https://www.arup.com/perspectives/publications/research/section/empowering-urban-design-and-planning-with-dynamic-data

[9] Adam Jones is one of the coordinators of this https://systemsthinking.blog.gov.uk/author/adam-jones/

[10] See https://rori.figshare.com/articles/report/Supporting_priority_setting_in_science_using_research_funding_landscapes/9917825 and the work of the research on research programme.

[11] https://www.are.na/

[12] See for example Johan Schot's work on deep transitions http://dx.doi.org/10.1016/j.eist.2018.07.006 or Frank Geels work, such as Geels F. W. (2004). From sectoral systems of innovation to socio-technical systems: Insights about dynamics and change from sociology and institutional theory. *Research Policy*, 33(6–7), 897–920. https://doi.org/10.1016/j.respol.2004.01.015 ; Geels, F. W. (2010). Ontologies, socio-technical transitions (to sustainability), and the multi-level perspective. *Research Policy*, 39(4), 495–510. https://doi.org/10.1016/j.respol.2010.01.022 ; Geels, F. W. (2011). The multi-level perspective on sustainability transitions: Responses to seven criticisms. *Environmental Innovation and Societal Transitions*, 1(1), 24–40. https://doi.org/10.1016/j.eist.2011.02.002; Geels, F. W., & Schot, J. (2007). Typology of sociotechnical transition pathways. Research Policy, 36(3), 399–417. https://doi.org/10.1016/j.respol.2007.01.003

[13] In homeostatic systems the third is essentially set to neutral – to repetition in time rather than change. This is also true in systems with no ambition but survival – in these the third part of the triad is flat.

[14] See the interesting recent paper from Michael Hogan and others on collective intelligence and systems change: https://aran.library.nuigalway.ie/bitstream/handle/10379/6917/Collective_Intelligence_Design_and_a_new_Politics_of_System_Change_Oct17_%281%29.pdf?sequence=1&isAllowed=y

[15] Representations rest on ontologies which in computer science means "an explicit specification of a conceptualization", described by Tom Gruber as "an abstract, simplified view of the world that we want to represent."

[16] I cover some of this in my book 'The Art of Public Strategy' from Oxford University Press.

[17] My 2020 lecture for the Cambridge University Bennett Centre set out how the crisis was casting new attention on the cognitive processes of government: https://www.bennettinstitute.cam.ac.uk/publications/bennett-2020/#PAPER1

[18] There are many such competing approaches to conceptual maps – see eg https://www.chora.foundation/ – though it is rarely clear why they privilege particular factors and not others.

[19] For example, the Cynefin framework distinguishes fields that *are* simple, complicated, complex, chaotic, and a centre of disorder, and in the hands of its designer David Snowden can be very useful. It built on other 'soft systems' methods that try to incorporate a sense of how people learn in different contexts. I drew on this and similar approaches in my book 'The Art of Public Strategy' back in 2007.

[20] See for example Michael Thompson's work on organisation and disorganisation in his classic book 'Organising and Disorganising: A Dynamic and Non-linear Theory of Institutional Emergence'.

[21] I describe my take on this in more detail in my paper on wisdom – the unevenly distributed ability to fix problems, solve disputes, avoid unnecessary conflict, stop bullying or taking offence too easily.

[22] See for example the collaboration platform for holocracy https://www.glassfrog.com/ which was used to generate a dynamic map for roles/actors/projects for Extinction Rebellion at a time when they were facing the acute strains that often hit movements trying to drive radical change.

[23] Other interesting methods include the work of the Synthesis Center at the Arizona University on simulating heat-scapes in cities and alternative economies and the work of Roger Kneebone's Center for Performance Studies at Imperial College.

[24] Carlos Gershenkon's work on traffic is a excellent example.

[25] https://www.bankofengland.co.uk/working-paper/2020/monetary-policy-and-the-management-of-uncertainty-a-narrative-approach

[26] See 'Methods and approaches to modelling the Anthropocene" https://www.sciencedirect.com/science/article/pii/S0959378015300285

[27] Dave Snowden writes well on this topic, and the challenge of getting the kinds of abstraction right: using geographical maps as an analogy. https://www.cognitive-edge.com/flexuous-landscapes/

[28] The practical challenges of organising data in any organisation, let alone a system, are vast: readiness, quality, recording processes, encoding protocols, extraction procedures, velocity, levels of granularity, generalizability, and governance rules to determine who or what can access continuous feeds, security. Within large firms and military organisations heavy investment helps to ensure these challenges are overcome. But we lack any comparable investment at the systems level. So the pragmatic answer is to start relatively simple with key elements of data that have the most potential to add value when shared.

[29] There are many challenges in data gathering. Just in relation to sensor networks for example these include how to organize in-field processing, data transmission and network security, calibration of sensors, understanding where to put sensors, upgrades, relevance determination; data integration; moving from semi-manual to fully automated; standards.

[30] Just as many challenges face anyone shaping models: causality; dynamics; timescales; model structure – what are the important processes to capture? Parameterization – how are processes represented mathematically, what are the functional relationships between model components?

[31] https://www.geoffmulgan.com/post/a-loop-theory-of-wisdom-how-do-we-respond-to-foolish-times

[32] This approach to the self-representation and cognition of systems echoes Maturana's ideas about system autopoeisis & cognition https://monoskop.org/images/3/35/Maturana_Humberto_Varela_Francisco_Autopoiesis_and_Cognition_The_Realization_of_the_Living.pdf

[33] This is becoming ever more relevant to AI too through initiatives like Open Climate Fix https://openclimatefix.org/ and privacy enhancing open technologies like https://www.openmined.org/

[34] See for example the work of David Tuckett on decision-making in uncertainty: https://journals.sagepub.com/doi/full/10.1177/0959354317713158

[35] https://www.nesta.org.uk/blog/intelligence-outcome-not-input/

[36] There are some promising developments. In interactive machine learning (IML), there is a human-in-the-loop that passes knowledge to a learning system. This knowledge can take different forms. These include sampling (selecting a useful example of something that a learning system has not seen before so that it can learn from it); labelling (providing the learning system with sources of truth); featuring (identifying or selecting the properties that improve the learning system's representation of the concept one wants to model. Machines can help humans too through 'active learning', 'transfer learning', 'machine teaching' interactive extraction of knowledge from a human teacher/expert – who selects examples and counter-examples towards the building of a model, while not knowing details about the learning algorithm. This echoes but is distinct from the expert systems methods of a generation ago and is useful for the growing set of ML problems where unlabeled data is plentiful and domain knowledge to articulate a concept is essential. On the boundaries of AI and CI there is also 'crowd assisted machine learning' to help with low quality data (missing values, dirty data etc), insufficient training labels crucial for classification tasks, discovering discriminative features with limited data, uncertain reliability of a machine's decisions.
See https://participatoryml.github.io/

[37] I set out in some detail a framework for triggered hierarchies of learning, and for what I call first, second and third loop learning, in my book Big Mind.

[38] https://eight2late.wordpress.com/2020/01/21/complex-decision-making-as-an-infinite-game/

[39] https://www.nesta.org.uk/blog/intelligence-outcome-not-input/

[40] https://www.nesta.org.uk/feature/ai-and-collective-intelligence-case-studies/

[41] https://whatworks-csc.org.uk/research-report/machine-learning-in-childrens-services-does-it-work/

[42] My recent paper for Demos Helsinki set out in more detail what this could mean: https://www.demoshelsinki.fi/wp-content/uploads/2020/12/dh_steering-through-capability2020.pdf

[43] IPPO is the new International Public Policy Observatory on social aspects of COVID which is supported by the ESRC and run out of UCL STEaPP.

[44] https://www.geoffmulgan.com/post/net-zero-integrating-data-digital-radical-carbon-reduction-strategies

[45] https://www.geoffmulgan.com/post/the-vital-missing-support-that-s-needed-for-the-sdgs-and-how-it-could-be-provided

[46] https://www.nesta.org.uk/blog/open-jobs-making-labour-markets-smarter-and-empowering-jobseekers/

[47] https://www.geoffmulgan.com/post/a-loop-theory-of-wisdom-how-do-we-respond-to-foolish-times

[48] https://www.cognitive-edge.com/coherent-heterogeneity-1-of-2/

[49] Four years ago I shared a paper on how philanthropy could use AI: https://www.nesta.org.uk/blog/philanthropy-and-innovation-how-could-open-data-and-artificial-intelligence-help-funders-do-better/

[50] https://socialinnovationexchange.org/sites/default/files/uploads/the_role_of_philanthropy_in_using_data_to_address_complex_challenges-_a_global_scan.pdf

[51] https://www.thegovlab.org/stefaan-verhulst.html

[52] A good recent example of this is: https://medium.com/wellcome-data-labs/what-research-fields-have-emerged-in-the-last-years-a-machine-learning-approach-9eaf63d329c3 which shows how they employed text clustering to help grant managers discover portfolio in innovative ways.

[53] See for example: https://medium.com/wellcome-data-labs/data-science-at-the-wellcome-trust-a-year-in-review-e73ea0360625

[54] See my piece for Alliance Magazine: https://www.alliancemagazine.org/blog/fixing-the-plumbing-from-shared-data-to-shared-imagination/

[55] https://www.wired.co.uk/article/nesta-open-up-challenge

[56] https://www.nesta.org.uk/blog/mission-oriented-innovation-seven-questions-search-better-answers/ and my colleague Stian Westlake's note on https://www.nesta.org.uk/blog/if-not-a-darpa-then-what-the-advanced-systems-agency/

[57] https://www.nesta.org.uk/blog/fighting-memory-loss-in-government/

[58] This note builds on practical work I've been involved in over many decades trying to reshape systems both top down through governments (on topics such as decarbonisation and skills) and bottom up through working with grassroots innovators. It also draws on my past writings on systems, including my book Connexity (Harvard Business Press, 1997) which examined systems dynamics in the emerging Internet, and Big Mind (Princeton University Press, 2017) which showed how collective intelligence methods can be used to reframe systems.

[59] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3586608/

1

STEaPP Working Paper Series

Contact us

steapp.communications@ucl.ac.uk