Check for updates

comment

Resilient government requires data science reform

Data has tremendous potential to build resilience in government. To realize this potential, we need a new, human-centred, distinctly public sector approach to data science and AI, in which these technologies do not just automate or turbocharge what humans can already do well, but rather do things that people cannot.

Ben D. MacArthur, Cosmina L. Dorobantu and Helen Z. Margetts

Resilience is the ability of an individual, organization or system to adapt to changes in circumstance or recover quickly from disturbances. In the context of government administration, resilience is also an organizational value — that is, a deliberate choice of what to prioritize — that underpins how a government designs its policymaking processes and how it makes use of technology[1]. Governments that value resilience prioritize responsiveness and adaptability: the organizational qualities that they need to withstand 'shocks' and carry on operating effectively.

The notion of resilience as a value has a long history in public administration, along with two other families of values: economy and leanness of purpose; and fairness and honesty[1]. But resilience and robustness are the organizational values most often associated with state decision-making, especially for hazard-related tasks or in times of crisis, such as war[1]. From the 1980s onwards, however, the administrative focus of many governments moved away from resilience towards economy and leanness of purpose, driven by the 'new public management' — a cohort of changes that aimed to introduce private sector management practices into the public sector[2]. This focus on economy and leanness had two important consequences for governmental uses of technology. First, agentification and outsourcing gradually led to the fragmentation of huge departments of state and gave rise to governmental structures — such as public–private partnerships and large-scale contracts — that became mismatched from the increasingly interconnected social, economic, healthcare and trade systems that they sought to serve. These new organizational boundaries were reinforced by contract relationships and privacy legislation regarding data sharing that together hindered data flow and worked against a holistic approach to governance. Second, the emphasis for government technology projects shifted from innovation to cutting costs and a focus on automation of routine tasks and staff savings via outsourcing. This shift inhibited the ability of many governments to establish in-house technological expertise and deskilled the public sector workforce more generally, which in turn meant that governments began to fall ever further behind industry in their ability to develop and use the latest data-driven technologies[2].

After struggling to serve their citizens during the COVID-19 pandemic — and faced with the next looming set of existential problems — governments have started to contemplate a move away from the administrative values of economy and leanness of purpose back towards the values of resilience and fairness[3]. Such a move should prompt them to rethink their use of technology. Instead of using it to cut costs, governments could once again use technology to strengthen decision-making processes and governmental operations. This would mark a return to the administrative values that governments prioritized in the aftermath of the Second World War. But unlike the 1940s, when computing was in its infancy, today's data-intensive technologies have the potential to radically change government for the better.

Need for public sector data science

The desire to use computers to reproduce or replace human activity is not unique to government. Indeed, it has been a central motivator for the development of data science itself, particularly in the long-standing conception of artificial intelligence (AI) as "the science of making machines do things that would require intelligence if done by [humans]"[4]. This 'intelligentist'[5] vision has been remarkably successful. It dominates research and development efforts and motivates some of the most effective machine learning methods. In recent years — particularly since the mid 2000s, when deep learning began to come of age — it has produced some extraordinary results, including super-human performance in complex strategy games and diagnosis of complex diseases[6].

However, this success has come at a cost. Notably, much progress in this area has been made (or funded) by the private sector and has been implicitly motivated by private sector concerns[7]. But the tools and methods that help companies to maximize profits are not the most appropriate for governments seeking resilience. In particular, the intelligentist vision of AI is not necessarily well placed to contribute to decision-making processes that focus on interconnected problems and require knowledge or expertise from different domains to be harmonized in a transparent manner. Such problems often do not have an unambiguous objective or 'right' answer and so are hard to approach algorithmically. Moreover, these are not the kinds of problem that any single human intelligence can solve. Yet they are precisely the kinds of problem that governments striving for resilience need to address.

For these reasons, we believe that a fresh perspective is needed. If governments are to be prepared for future shocks, then a careful re-conceptualization of data science that is tailored to the particular challenges faced by the public sector is needed.

Data science for resilience

To meet this challenge, we propose a roadmap articulated in the following three guiding principles (Box 1).

Transfer insight with integrity. Governments often have limited access to data and a fragmented approach to generating insights. The allure of 'big data' has tempted many to believe that more data will overcome all such problems. We do not believe this is the case. Rather, governments need more efficient data collection and modelling practices that are designed to derive maximal insight from sparse data

Box 1 | Three recommendations for data science in resilient government

Recommendation 1: Governments should take an inclusive and participatory approach to the design, development and deployment of data collection processes and practices, in which citizens play a central part, and use the latest advances in data science to make efficient use of sparse data resources. Doing so not only builds trust but can also substantially boost a government's ability to efficiently collect, clean and report data. As a salient recent example, Our World in Data (an academic non-profit initiative) contributed substantially to monitoring the COVID-19 pandemic and informing government decision-making. Similarly, the citizen science ZOE COVID-19 project allowed both scientists and policymakers to quickly learn about the emergence and spread of new COVID-19 symptoms[13]. Governments seeking resilience should learn from these examples and partner with reliable academic, non-profit and citizen science projects to enhance their data collection processes.

Recommendation 2: Governments should use ensembles of models that each take a different view of the world, or tackle different aspects of a hard problem, to harmonize different sources of advice. The practical advantages of a collective approach to forecasting are well-established and have been demonstrated in numerous areas, from climate modelling to macroeconomic predictions. As a salient current example, the US Centres for Disease Control and Prevention uses ensemble forecasting to predict COVID-19 cases, hospitalizations and deaths, drawing on an established multi-model forecasting methodology for seasonal influenza predictions[14]. Governments seeking resilience should learn from these examples and use robust multi-model ensembles to inform their decision making.

Recommendation 3: Governments should develop models that incorporate essential causal mechanisms into their analysis and thereby allow the effects of interventions to be rigorously assessed. Such models should be preferred over black boxes, no matter how accurate. Doing so will not only bolster the long-term robustness of government decision making, but also can facilitate effective decision making in rapidly changing circumstances by providing a principled basis for weighing interventions. For example, researchers supporting the UK's health security agency developed a causal analysis framework that provided fine-scale spatiotemporal 'nowcasts' of COVID-19 prevalence and robust estimates of the epidemiological parameters needed for real-time policymaking[15].

Notably, the frameworks needed to take full advantage of citizen science; ensemble forecasting and causal inference methods were not established until later in the pandemic response. Governments seeking resilience in public sector governance, day-to-day decision making and crisis management should learn from these examples and establish robust analysis frameworks in advance of their need.

resources without compromising their citizens' right to privacy. These practices should take advantage of the latest advances in data science to build resilience into decision making by facilitating the flow of information within government and the transfer of insight between policy domains.

Two considerations are key. First, to derive maximal insight from sparse data resources, governments should take advantage of the fact that policy domains are often inherently interconnected, and insight gained from one domain may be used to improve understanding in another. In such cases, tools from transfer learning — the branch of machine learning that concerns passing information from one domain to another[8] — may be particularly useful. Although not yet widely used in policy settings, we anticipate that transfer learning may provide powerful tools to policymakers, for example to transfer insight between or within countries (to extrapolate insight gained from geographic areas for which data are abundant to those for which it is not, for instance) or between related healthcare or economic domains.

Second, collection of socio-economic, healthcare and behavioural data inevitably means collection of information about individuals, each of whom should have the right to decide how their data are collected and used. Unprincipled data collection practices, without scrupulous regard for individuals' privacy and autonomy, risk becoming intrusive and undermining public trust. Because resilience requires trust, it is vital that data collection practices are conducted with citizen input and support and provide informative data without intruding into citizens' lives. To approach these issues, governments should make use of emerging new tools, such as privacy-enhancing technologies and synthetic data, to maximize the insight they gain from multiple data sources while maintaining privacy.

Integrate diverse perspectives. To address complex multi-sector problems, decision makers typically seek counsel from advisors with different areas of specialism. Although there are practical ways to improve the accuracy of specialist advisors' judgements[9], there are few ways for policymakers to harmonize disparate sources of specialist advice or to weigh the effects of policy choices that may be beneficial in one area, but costly in another. Data science for resilient government should not aim to replace such human advisors, but rather should aim to augment and connect different areas of human expertise and to harmonize different viewpoints.

One way to approach this issue is to take a collective modelling approach, in which an ensemble of models (or 'learners') — each of which may be informed by specialist domain expertise or make different basic assumptions about the world — is developed, and decisions are based on the output of the collective. This simple idea is the basis of so-called ensemble learning methods, which have proven to be among the most powerful tools in modern machine learning and predictive modelling[10].

Ensemble methods are particularly useful to policymakers for three reasons. First, they are beneficial whenever multi-modal data are available but hard to fuse (for example, combining data from different government, health or economic sectors) or when data can be partitioned into disparate pieces, each with different characteristics. In this case, a divide-and-conquer approach can be taken in which specialist models are trained on different subsets of the data before being combined for output, thereby providing a way to integrate different sources of expertise.

Second, because ensemble models gain their power from their ability to combine diverse perspectives, individual learners do not always need to be highly accurate or refined and therefore can be quickly and easily trained. Thus, ensemble methods may allow new learners to be easily added and old ones removed, and so provide a natural way to produce models that are able to adapt to new data as it arrives and design interventions on the basis of the latest knowledge.

Third, they can be used alongside other powerful mathematical modelling and machine learning tools, for instance to integrate models that make use of different styles of learning or make different causal assumptions. For governments striving for resilience in policy-making systems, using an ensemble of models therefore represents a pragmatic approach that cautions against the search for the one 'right' model and makes the most of the latest machine learning advances, available evidence and any disciplinary insight to inform decisions that appropriately account for diverse perspectives.

Tackle questions of causality. Machine learning is perhaps best known for its capacity for prediction. But in a policy-making context, understanding causal principles is arguably more important than prediction: it is what enables decision makers to understand the key drivers that influence the outcomes of their decisions, to identify and assess the effects of their policy measures in the real world, and to prepare and adapt for the future.

Economists and social scientists have made substantial progress towards understanding causality in specific policy settings (two out of the last three Nobel memorial prizes in economic sciences have been awarded for methodological advances in this area, for instance). But much more work remains. The next challenge is to combine causal modelling with advances in machine learning. Recent years have seen tremendous advances in this area, much of it building on the work of Judea Pearl, who proposed a three-rung 'ladder of causation'[11] — climbing from purely associative models (that describe what is), to those that can explore the effects of intervention (what could be) and finally to those that can explore counterfactuals (what could have been).

Because policy is fundamentally about making interventions, it is inherently associated with rungs two and three of Pearl's ladder. Models that operate at the first level may therefore be useful for understanding patterns in administrative data but cannot design reliable interventions, and should not be used to determine policy or to inform sensitive or high-stakes decisions — even if retrospectively interpreted[12]. Rather, effort needs to be directed at building models that clarify how myriad socio-economic factors affect each other and enable policymakers to properly understand the effects of interventions before implementing them in the real world. These efforts should capitalize not only on the latest advances in machine learning but also on the vast literature within the social sciences on using empirical evidence to inform policy.

Conclusion

Faced with myriad healthcare, social, economic and environmental challenges, governments the world over are seeking resilience. When approaching such challenges, understanding patterns of interconnection between sectors is vital to robust decision making. For this reason, we have argued that to build resilience, a reform of data science for government — explicitly designed to tackle complex multidisciplinary public sector challenges — is needed. Rather than focusing on reducing costs through automation of what humans can already do well, such a reform should focus on doing what humans cannot do well: addressing interrelated problems that require the harmonization of data, knowledge and expertise from different domains. Rethinking data science in this way is a substantial challenge that will require government investment and citizen support, and be characterized by strong collaborative interactions among data scientists, ethicists, domain experts — including and especially social scientists — and decision-makers. Although it may not appear as glamorous as some AI developments, developing this vision is equally challenging, exciting and societally transformative. ❐

Ben D. MacArthur [1,2,3,5] ✉, Cosmina L. Dorobantu [1,4,5] ✉ and Helen Z. Margetts [1,4,5] ✉

[1]*The Alan Turing Institute, London, UK.* [2]*Mathematical Sciences, University of Southampton, Southampton, UK.* [3]*Faculty of Medicine, University of Southampton, Southampton, UK.* [4]*Oxford Internet Institute, University of Oxford, Oxford, UK.* [5]*These authors contributed equally: Ben D. MacArthur, Cosmina L. Dorobantu, Helen Z. Margetts.*
✉*e-mail: bmacarthur@turing.ac.uk; cdorobantu@turing.ac.uk; hmargetts@turing.ac.uk*

Published online: 8 August 2022
https://doi.org/10.1038/s41562-022-01423-6

References

1. Hood, C. *Public Adm.* **69**, 3–19 (1991).
2. Dunleavy, P., Margetts, H., Bastow, S. & Tinkler, J. *Digital Era Governance—IT Corporations, the State and e-Government* (Oxford Univ. Press, 2008).
3. *The National Resilience Strategy: A Call for Evidence* (UK Cabinet Office, 2021).
4. Minsky, M. E. *Semantic Information Processing* (MIT Press, 1968).
5. Leslie, D. *Nature* **574**, 32–33 (2019).
6. LeCun, Y., Bengio, Y. & Hinton, G. *Nature* **521**, 436–444 (2015).
7. Klinger, J., Mateos-Garcia, J. & Stathoulopoulos, K. Preprint at https://doi.org/10.48550/arXiv.2009.10385 (2020).
8. Zhuang, F. et al. *Proc. IEEE* **109**, 43–76 (2021).
9. Sutherland, W. J. & Burgman, M. *Nature* **526**, 317–318 (2015).
10. Sagi, O. & Rokach, L. *Data Min. Knowl. Discov.* **8**, e1249 (2018).
11. Pearl, J. & Mackenzie, D. *The Book of Why: The New Science of Cause and Effect* (Basic Books, 2018).
12. Rudin, C. *Nat. Mach. Intell.* **1**, 206–215 (2019).
13. Menni, C. et al. *Nat. Med.* **26**, 1037–1040 (2020).
14. Reich, N. G. et al. *PLoS Comput. Biol.* **15**, e1007486 (2019).
15. Nicholson, G. et al. *Nat. Microbiol.* **7**, 97–107 (2022).

Acknowledgements

This work was funded by the Engineering and Physical Sciences Research Council (EPSRC, grant ep/w006022/1). The funder played no role in the decision to publish or preparation of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Peer review information *Nature Human Behaviour* thanks Paul Henman and Baobao Zhang for their contribution to the peer review of this work.