



Data you can trust

Technology that works for you

DATA61's Future Science Vision v1.4

Robert C. Williamson, 2 November 2016



Preamble

Research is at the heart of Data61. Our research is undertaken with a purpose in mind – to create a positive data-driven future. This document outlines our vision¹ regarding what we aim to achieve by focusing our research on what the world needs in areas where we have world-leading capability.

Data61 plays two complementary roles in the Australian innovation system. We are “L-shaped” (see the schematic on the right):

- 1) We conduct market driven research (end-use driven projects) in a range of industry sectors; these contribute to the horizontal part of Data61’s mission – solving problems in other CSIRO business units (and leveraging their capability and connections) and the community more broadly.
- 2) We are the home to fundamental research advancing the science and technology of data (the vertical part of the picture).

These two parts mutually support each other². Both are essential. The market component, *by definition*, is not for us to plan, but to adapt to in an agile manner. The scientific, technological and engineering research we propose to do *is* ours to plan and shape; that is what this document does.

The purpose of this document³ is to focus our work on the vertical part of the L-shaped schematic. The document captures the bold and ambitious areas of science and technology we wish to advance⁴. It should be seen as a way of focusing what we do, and allowing us to say “yes” or “no” in a more informed fashion⁵. The goal is not simply to “put more wood behind fewer arrows” but rather to get most of the arrows pointing in one direction, and to describe the target they are aiming to hit – namely the four goals listed in the callout box. This will help shape our future capability investments.

Explicitly articulating the larger technical challenges is especially important for Data61 because it is often (mistakenly) believed that data and information technology research merely supports other sciences – a sort of glorified IT helpdesk. In fact, the contrary is arguably the case, with physics⁶, chemistry⁷, biology⁸, social science⁹, and economics¹⁰ all having the science of data and information at their core, and information technology precepts, such as modularity, are essential for the understanding of many natural systems¹¹.

Ultimately, as recently witnessed by social science, any field immersed in a properly organised bath of data progressively becomes computationally based, or develops a *computational* subfield¹². The science of information and data is arguably the most fundamental research topic of the century, situated not only at the centre of mathematical research¹³, but underpinning the nature of randomness and complexity¹⁴, and situated at the very core of all the mature sciences.



Measuring the World

Improving the whole lifecycle of data capture analysis and use.

Delivering Trustworthy Analytics

Changing the way analytics is delivered; guaranteeing trust in the entire process.

Building Software you can Trust

Creating technologies that allow the construction of software that can be trusted.

Shaping Societal Transformations

Developing better data technologies through improved understanding of their potential social impact.

Context

Technologies for data are general purpose technologies¹⁵ that will have a transformative impact on Australian society, although what those impacts will be is neither predictable nor pre-determined¹⁶. These technologies are often described as “artificial intelligence”¹⁷ and include machine learning and big data analytics, automated reasoning, computer vision, natural language understanding, and robotics. Data61’s focus is on the advancement of technologies for data in a manner that provides national benefit (economic, social and environmental). Thus, a deep understanding of the *context* of technology use, the potential impacts they can have, and shaping what those impacts are, is a central part of our research vision.

Data61 lives inside an organisation dedicated to the discovery of scientific knowledge, knowledge distinguished by the high degree of *trust* one can place in it: trust in the conclusions; trust in the evidence that is derived from data; and, trust in the processes to revise the knowledge when it is found to be false. Science has always been data-driven and will remain so. We propose to exploit the scientific enterprise within CSIRO as a testbed for ideas that can, and will, have much broader impact.

General principles

The scientific vision is informed by the following five principles¹⁸

- **P1. Lead:** Strive for a greater proportion of world leading research. We should focus our efforts on areas where we are, or realistically could be, world leading.
- **P2. Multiply:** Aim for multiplicative (compositional) effects rather than additive, else we cannot scale. This implies clever “platformisation” of our technology.
- **P3. Unique:** Do what only we can do, else let others do it¹⁹.
- **P4. Bold:** Aim high. We really do want to change the world (through use-inspired fundamental research).
- **P5. Antidisciplinary**²⁰: Data traverses existing discipline boundaries. We ignore disciplinary boundaries and follow the problems wherever they take us.

Headline Visions

Data61's goal is to create our data-driven future – a future where technologies for data will play a positive role for society at large. New technologies provoke many reactions. Fear and uncertainty is common, with a belief that the precise forms of new technology are inevitable and not open to being shaped²¹. A counter to this is *trust*, which can be viewed as being at the core of all that we do. All of our work revolves around building trust in technologies for data: in automation; in security and privacy; that your software only does what it claims to do; that your personal identity is not stolen from you; and trust in all things that matter to people.

By saying “data you can trust” we do not mean that you trust it blindly, and especially we do not mean that you trust it raw – data needs to be processed and manipulated to be useful, and it is the processes of manipulation that need to be trusted. This involves both designing systems that do indeed facilitate trust in data, as well as building trustworthy technologies for doing things with the data. And in all of this “trust” itself is complex, multidimensional, and is always ultimately grounded in human needs and society²².

We are using the apparently simple notion of “trust” metaphorically²³. Without attempting to make a canonical definition of trust²⁴, we can say we have “trust” as the anchor, or point of departure, for much of what we propose to do, including:

- **Trustworthy software** – not software that you trust absolutely, but software in which you can have quantifiable degrees of trust for sound reasons
- **Trust in data** – not data you trust without cause, but data you can trust for your purpose because of the evidence provided regarding its management, provenance and what was done to it (analytics that has quantifiable effect)
- **Trust in systems** – trust that you know to what degree you can rely on data-centric systems, including communications, not that you trust it absolutely
- **Trust in data technology enabled socio-technical systems**– trust that these systems will benefit you and that any harms are manifest and controlled.

Understanding the complex interface between data, its management, manipulation and processing, and the impacts it can have on people is central to building trust around data and technologies for data. Trust in data (and its associated processes) can also underpin trust in institutions, interventions and policies.

The means of manipulating and processing data are data *technologies*. When we say “technologies that work for you” we mean they do what they are supposed to do, they don't do anything else, and they are usable and useful (and implicitly we recognise the importance of who the “you” is – technologies that help one group can harm others).

While these sentiments might be taken for granted, history shows they are often absent, and improving the degree to which the technologies we develop achieve these goals helps to shape what we do. Examples are: the construction of software that has an adequately high guarantee of securely doing only what it is supposed to do; or, statistical machine learning methods you trust because of mathematical theories that provide adequate guarantees regarding their behaviour and uncertainty.

Both these examples illustrate the necessity for deep scientific and mathematical knowledge as well as a quantitative notion of performance. This scientific depth differentiates what Data61 does from much of the data technology in the wider world.

The headline visions and scientific challenges serve as a rallying point for not only the scientific research we do, but also the shorter term end-use driven projects delivered by our engineering team. Ideally the majority of such projects, in addition to delivering on customer expectations, will further the goals below.

H.1 Measuring the World²⁵

Thus is by geometrye mesured alle thingis

– William Caxton, *Myrrour of the Worlde* (1481)

The world becomes better understood, and thus interventions are more effective and acceptable, through the development of methods for data capture and model building that put trust at the center.

Background: Humans try to improve the world, but often fail. Their interventions don't work, or have unintended consequences. One reason for these failures is poor models of the world – it is different from what we expect. By *measuring the world* (ie capturing data about the world), one learns more about the world and thus interventions can be better designed. This is the vision of empirical science. We propose to improve how data is captured and used to advance our understanding of the world.

The world is full of data, but only a small fraction is known to us. Rather than being *given* to us (“data” comes from the Latin *dare* meaning “to give”), it is necessary to *take* the data – to actively select and gather it, and then, of course, to *do something* with it. It is thus useful to distinguish data from *capta*²⁶ (from the Latin *capere* meaning “to take, seize, obtain, get, enjoy or reap”²⁷). This terminology signals that data collection is an active process, not passive.

Data is traditionally seen as the lowest level of a hierarchy that runs from data to information to knowledge to wisdom²⁸. Implicit in this, is that in order to attain knowledge (or wisdom) one needs to start with data. While clearly true at one level, this does not capture Data61's perspective which inverts the hierarchy²⁹, and has knowledge (or the decision, action or intervention required for a particular problem) as the end point, thus focussing the needs of data collection and analytics from the reverse perspective. Data becomes useful once it is both captured (*capta*) and then made sense of through models. The models can also provide guidance regarding desirable *capta*.

Models and modelling are central to making use of *capta*. Much of the work that Data61 does is modelling based on *capta*. The distinction between models and data or *capta* is blurred³⁰; abstractly a model is *always* a function of the *capta* – whether it has a small number of “parameters” or not is irrelevant – what matters is the stability of the model (or more precisely, the stability and reliability of the conclusions drawn, and actions taken from the model) under data variations.

The important point is that it is the models that are ultimately manipulated and used for action. While much is made of a “fourth paradigm”³¹ (so called “data-driven science”) and “the unreasonable effectiveness of data”³², the fact remains that all data-driven intervention remains based upon models; they are just more complex than the models of old.

We thus embrace the “primacy of method”³³ or a “method deluge” (with methods as “first class citizens”³⁴) over a mere “data deluge”, and certainly do not envisage “making the scientific method obsolete”³⁵. For science, data alone (however it is linked or presented) is not enough³⁶. Neither data nor facts are ever entirely raw – they are constructed and theory-laden³⁷. It is indeed true that “‘Raw data’ is both an oxymoron and a bad idea”³⁸.

Some of the greatest contributions to the recent explosion of interest in data-driven everything comes from new methods³⁹ with refined notions of trust (better quantification of errors). The blurred boundary between “data” and “method” drives how methods (analysis) are being pushed towards the data (embedded analytics⁴⁰), as well as the propagation of all aspects of the data (such as its provenance) through the entire modelling process, in order to better inform interventions.

The real promise of a data-driven society is that it is an “experimenting society”⁴¹ that allows decisions, actions or interventions to be closely tied to *capta*.

We will develop new methods for achieving this universal “*capta*fication”⁴² of the physical world, the biological world and the social world:

- From modelling of materials and biological organisms at the molecular and macro level to the design of new materials and food

- From sensors measuring anything through to trusted data from those sensors and the associated trusted interventions and policy
- From all the geospatial data in the country to the rich set of services that can exploit this information
- From people's identity and reputation to systems that can guarantee the security, privacy and fairness of using this information
- From the captification of the law and public policy to make the machinery of government transparent to the user to the very development of new policy in a trustworthy evidence driven manner, and
- From transforming how science is done (tracking data and evidence and the analytical conclusions drawn) to the empiricisation of business (doing proper experiments aided by technologies for data).

Our vision is that by developing new and better *methods* we will be able to better model the world, and thus act better. Central to this is the notion of trust:

- Trust in the source of the data (collected the right *capta*) and that it was reliably captured, transmitted and not tampered with (else skeptics will challenge the result, or worse, wrong actions will be taken)
- Trust in the models underpinning the capture of the data (such models always leave something out – how does one know if the omissions do harm?)
- Trust in the methods used for analysis (that it is known what the methods actually do from a user's perspective and that the posterior uncertainty is properly calibrated)
- Trust in how the *capta* and conclusions are presented and used (if one ignores this human element, then the best methods can still lead to terrible outcomes), and
- Trust that legal and moral rights and notions of fairness are not infringed (else society will disdain the power of data analytics because of concerns regarding its abuse).

H2. Trustworthy Analytics Delivered⁴³

New methods for data analytics that offer high degrees of trust, and new methods of delivering these trustworthy methods will increase their use, reduce economic friction and speed up the process from invention to deployment. This will accelerate scientific discovery, business improvement and improve public policy outcomes.

The impact of new technologies comes from their use. We will change the way analytics is delivered to broaden its use. We will build trust into the core of how we create and deliver analytics technologies: from the mathematical foundations of trust in data-driven conclusions and the quantification of certainty; to embedded analytics at the source of data capture; and, to web services that allow the flexible composition of analysis methods in a reproducible and scalable manner, and which build in key elements of trust from the outset (provenance and traceability, management of legal and moral rights, and management and preservation of uncertainty)⁴⁴.

Background: "Data Analytics" means the computational processing of *capta* with the goal being to derive insights suitable for comprehension, decision or action. It includes mathematical or algorithmic methods as well as visualisation and presentation of the results in a manner suitable for human consumption. Analytics is not only used by a (human) statistician; many socio-technical systems have analytics embedded into their core operation, and all the points made below apply there too.

Presently analytics is implemented primarily in a manner that makes its composition (gluing together components) difficult.

The current model leads to various problems:

- Vendors of large software packages have an interest in locking in customers to their platform (so there is relatively little incentive to enable composability with other systems)
- Many of the implementations presume the data is all in one place (either local or in a cloud). Much data cannot be moved. It might be too large (the analysis has to be actually done at the source), or there is not the legal right to move it
- Provenance, traceability, legal and moral rights and uncertainty are poorly managed, resulting in outputs of analytics that lose sight of the reliability and trustworthiness of the original data (and thus the results are less trustworthy)
- It is difficult to redo analyses when mistakes are discovered (a consequence of the point above). Often not all of the “state information” is stored to enable the re-running of analyses
- Closed ecosystems make it hard to import new techniques as they are invented.

There are potential solutions to all of these problems, all of which we envisage developing:

- By embedding analytics at the source of the data, the burden of moving large amounts of data is removed. Being able to reach all the way back to the original data source (typically embedded in a cyber-physical system) through composable data ingestion schemes allows better tracking of provenance
- Systems that deliver analytics as a RESTful web service, then it becomes more readily composable. This can remove the downside (lock-in) of proprietary systems
- By taking the computation to the data (in data centers for example), we can avoid the problem of not being able to move the data (for reasons of scale or jurisdictional constraints). This necessitates advances in not only the secure encapsulation of analytics code, but also necessitates trusted means to control information flow (so private information is not exfiltrated from the data bases)
- The ultimate delivery involves presentation to users. By improving the user experience of data analytics it will be more widely and reliably used. This requires development of visualisation as a service that represents uncertainty and provenance as first class objects
- Composable provenance of data (including legal rights such as licenses) and analytics across walled gardens allows increased trust, reliability and repeatability of analytics
- Systems that are designed to *federate* data from different sources can bypass jurisdictional and practical problems of extracting insights from distributed data
- Late binding schemas or ontologies minimise the deleterious effects of past decisions regarding data categorisation and organisation
- Systems that capture and re-execute entire workflows to facilitate late-binding, rapid prototyping and the automation of translation from exploratory to production systems

The creation of technologies as above will not only accelerate the use of data analytics for its own sake, but will play a central role in our vision for cyber-security – securing data-driven business operations through ensuring trustworthiness in the data. This is especially important for critical infrastructure protection⁴⁵.

H3. Building Software⁴⁶ you can Trust

We will develop new ways of creating software that will be the global benchmark in terms of quality, security and trust. Widespread adoption will make software companies more productive, improve cybersecurity (by addressing the root cause of one of the main problems) and enable higher degrees of trust in data-centric systems.

Technologies for data are underpinned by software, which is the means by which data is processed and transformed. Building better technologies requires building better software. We will develop the science and technology stacks to build software that provably does what it is supposed to do and nothing else – we will be able to say precisely and with strong evidence when software will be bug-free, provably secure, and will deliver guaranteed results. This will address one of the major causes of problems in cyber-security (vulnerabilities that are introduced when software does more than, or other than what it is supposed to do). We will also develop better methods to quantify risks associated with software and understand the human factors that contribute to trustworthy software.

In addition to increasing the reliability of software against attacks that cause it to do things other than which it should, the same technologies can be used to provide improved guarantees for the trustworthiness of data, whether it is that the data has not been manipulated, or that sensitive information has not been exfiltrated. Thus improving the trustworthiness of software is not only essential for making technologies that work for you, but also for ensuring that you can trust data and entrust your data to such technological systems.

H4. Shaping Societal Transformations

*Technology ... is not destiny*⁴⁷

– Jason Furman - July 2016

Technologies shape society, and technologies for data will shape the future of Australian society, but there is the opportunity to choose *what* these effects are. By developing better understandings of the complex relationships between data technology and people, we will be able to influence the development and use of technologies for data to lead to better societal outcomes. The research necessary to attain this understanding can (and needs to) be done in concert with the more narrowly technical aspects of our work.

New technologies for data will transform society, but there is much freedom regarding how. Our interest in technology does not stop with the technology itself, but extends to its use. Technologies such as UAVs and autonomous vehicles will obviously shape society, and their use will be shaped by what society finds acceptable. Collectively, as technologists and scientists, we cannot ignore the societal implications of our work. The same basic technological principles can be used in many different ways; some of which are more usable, helpful and beneficial to people than others. We will develop new ways of envisaging and influencing these societal transformations.

This will involve new approaches to the ethnography of technology (better understanding people's relationship with data-driven technology, especially in terms of trust) and deriving technological foresights. This goal aligns with strategy 2 of the recently released US *National Artificial Intelligence Research and Development Strategic Plan*⁴⁸: "Develop effective methods for human-AI collaboration. Rather than replace humans, most AI systems will collaborate with humans to achieve optimal performance. Research is needed to create effective interactions between humans and AI systems."

We will reimagine what it means to be human in a data-driven world. We will develop new technologies for ensuring rich notions of privacy and transparency in a data-driven and algorithmic world. We will develop new understandings of the complex technical tradeoffs between usability, security, privacy, efficiency and fairness. We will study how to build data-driven societal institutions that citizens can trust. We will design new computational mechanisms to enhance social welfare, enabled by pervasive technologies for data.

We will develop new methodologies that exploit data-technologies to better understand how data-technologies themselves end up being used (including the derivation of qualitative insights from quantitative data). This will extend the reach of user-experience design to new areas, and advance its state of the art. And we will develop new economic and business models enabled by data-technologies in a manner that seeks to maximise benefit for Australia as a whole.

Scientific Challenges and Foci

Theories are nets: only he who casts will catch.

–Novalis

In this section are listed some scientific⁴⁹ challenges arising from the above visions. These are not *all* the scientific challenges we will try to solve, but they capture much of what we aim to do. In all cases the timeline is roughly 5-10 years.

While each of these challenges is motivated and inspired by broader societal challenges, the particular impacts one can expect of scientific advances are notoriously difficult to predict on such a time scale (impact can be predicted more reliably for shorter term projects). Thus, apart from some rather general statements, there is no specific prediction of impact arising from the scientific challenges.

I have tried to state a high level challenge (in red) followed by some explication. It would be impossible to outline all the possibilities, and those listed are not meant to be too prescriptive.

In all cases they are stated as “How to...”. This is both a scientific challenge (development of new knowledge and understanding) as well as a technological one (development of techniques and methods and systems that achieve the goal).

Areas of Scientific Challenge

- Materials and Data
- Physical / Biological Systems and Data
- Institutions and Data
- Trustworthy Software Construction
- Architecture for composability, compartmentalisation and resilience
- Distributed Trust Mechanisms
- Analysing, Representing and Modelling Data
- Quantification of and reasoning with risk and uncertainty

S1. Materials and Data

How to turn materials into data so they can be manipulated and designed?

To understand materials (so they can be synthesised, manipulated and changed) one needs to understand them and trust that understanding (modelling and synthesis). Materials are not systems (for the purpose of this document). The question applies to both non-organic and organic materials (including for example food).

How to design materials in a data-driven manner – from quantum monte-carlo (for engineering materials) through to food designed in response to genetic information?

S2. Physical / biological systems and data

How to embed data into physical systems; understand physical systems through data-driven models; and design, build and control physical systems by using data?

This includes challenges in robotics and sensor networks and in the processing of visual data – how to embed trusted analytics into physical, biological and environmental systems. How to use data to increase trust in data-centric systems (such as the internet of things), for example by better management of privacy. How to better model physical systems using data (or more precisely, how to improve that modelling, which is the core business of all scientists, using modern technologies for data).

How to control physical systems with data in a manner that you can trust? How to turn physical or biological objects (eg scientific specimens, or aspects of living systems) into data cheaply and at scale in a manner that can be trusted? How to map the world more reliably (using spatial data as a testbed for analytics pipelines)? How to build autonomous systems for data gathering in the field. How to manage the ingestion of semi-structured sensor data? How to manage the provenance of data gathered in the world?

S3. Institutions and Data

How to represent, augment, understand, manage and control institutions better using data?

I use “institutions” in the economist’s sense⁵⁰ which includes government, the legal system (statute law, regulation), business processes, and contracts, etc. The challenge is to represent these societal systems using data that can be processed and reasoned with by a machine. Solving this involves advancing the state of the art of natural language processing (eg, targeted at specialised uses of English, as in statute law and contracts) and the development of tools that allow the crafting of legal instruments in a manner similar to a modern programming development environment that will guarantee properties such as consistency, but will also emit human readable versions of the instruments.

Another challenge is how to use technologies for data to improve institutions, for example by data-driven experimentation for policy development⁵¹. Part of the solution is likely to be aiding the change of role of government from owner of assets, or deliverer of retail services to wholesaler and architect of modular systems.

S4. Trustworthy software construction

How to construct software that does what it is supposed to and nothing else?

How to make technologies that constructs software that guarantees its correctness, invulnerability and other properties (eg real time guarantees). One can ask similar questions regarding interaction and communication protocols. Particular challenges include: mixed-criticality, real-time, multicore, side-channels; information flow; concurrent systems verification; protocol verification (as a means to deal with composition and break the back of concurrency); automation of proof effort. How to specify and quantify dimensions of security (turning it from a binary property to a real-valued property you can reason about from a risk sensitive perspective)? How to ensure trustworthiness of mobile code (especially for analytics)?

S5. Architecture for composability, compartmentalisation and resilience

How to build data-centric systems that can be reliably composed and compartmentalised and which are resilient, robust and trustworthy?

Data-centric systems are the most complex artefacts designed by man. The challenge is to design them (including cyber-physical and cyber-societal) in a manner that facilitates composition, compartmentalisation and resilience. This is necessary in order to improve the reliability and trustworthiness of such systems.

This challenge is architectural (including questions such as how to compose trust – just because you have trusted components does not guarantee their composition can be) but includes questions such as how to monitor and manage such large systems (supervisory control and diagnostics). Examples that are worthy of attack include how to architect large distributed data analytics systems. How can trust in such systems be quantified, measured and managed?

S6. Distributed trust mechanisms

How to manage trust in distributed data-centric systems?

Trust underpins human interaction, and thus data-technologies that mediate such interactions must manage trust. The challenges include how to ensure trustworthy provenance of data and operations on data (provenance is a kind of dual to security: provenance tells you reliably where the data came from and who did what to it; data security reliably ensures where the data can go and who can do what with it). Thus we will study both provenance and security together. This needs to be done in a risk sensitive manner (see S8). How to build richer, better and more applicable distributed ledgers and allied technologies? How to understand and quantify their security and reliability? How to build social choice mechanisms that can be trusted? How to build the communications technology that underpins distributed trust?

S7. Analysing, Representing and Modelling data

How to derive insight from data that can inform action?

How do you make sense of data? How to make sense of all the methods that do so? How to build models that are usable and re-usable. How to exploit complex, structured data with all of the mess of the world in the way? How to model complex phenomena (ecologies, language, societies) using data? How to make such models trusted and reliable and composable? How to best communicate such models to people for action? How to act and decide upon models of data? How to manipulate data representations of the world? Tools for managing multiple representations of data and manipulating them (music, law, biology). How to exploit computational and algorithm advances to build better technologies for data analysis?

This all needs to be done in the context of the structure of data; data is not merely a string of bits. Many of the types of data that will have the largest impacts are highly structured (natural languages, video, social networks, etc). Advancing the stated goal with respect to these data types requires deep science and technology stacks (that can be used across diverse application domains).

S8. Quantification of and reasoning with risk and uncertainty

How to quantitatively represent the rich sources of risk and uncertainty represented by data, and how to reliably reason with this?

Whilst data can sometimes reduce uncertainty, it does not remove it; decisions still need to be made in the face of uncertainty. Furthermore, the increasing complexity of data-driven systems means that the management of partial information, uncertainty and ambiguity is essential. How can this be done in a risk-sensitive manner? How can all aspects of data technology be made resilient to uncertainty? How can different notions of uncertainty be combined (relative to the inference of decision task at hand), and how can it be reasoned with in an effective manner? How can uncertainty and risk be effectively communicated and visualised? How can legal rights, security and privacy be made risk sensitive?

S9. Fundamental limits of data

How to determine the limits of what can be done with technologies for data?

All technologies for data have limits. How can these be determined and catalogued? And how can we approach these limits? Without knowing what the fundamental limits are it is not possible to know when a technology may break down and where to put effort to prevent this from happening.

This challenge cuts across everything we do, is a fundamental differentiator, and provides credibility for our status as part of a scientific research organisation. It also sets a target for other, less “fundamental”, work by setting a gold standard to approach.

Challenges include what is possible with data analytics, optimisation, distributed trust mechanisms, and indeed all data technologies we examine. Challenges include characterising the difficulty of learning from data, inferring causality, dealing with noise, protecting privacy, transmitting and sharing data, and solving computational problems.

There are limits in terms of data, knowledge, computation, energy, time and space. As well as limits to technical components, there are also limits (which need to be determined) to composite systems (such as trust, stability, and ability to control). There are also limits to socio-technical systems built with data technologies (for example computational social choice, limits to “fairness” and other synthetic properties) and limits arising from human abilities or inabilities.

S10. Shaping data-driven society

How to understand what it means to be human in a data-driven world?

What does it mean to be human in a data-driven world? How can our humanity be enhanced by data-driven technologies; how can we prevent harm? How can we build data-technologies that are meaningful and valuable to society at large? How can we encourage and assist communities in their adoption of technologies for data to improve their lives?

Solving this challenge will require the development of new ethnographic methods for data-centric technologies. It will also require ongoing research on how people interact with data-technologies from the perspective of decision theory (social choice, bounded rationality, etc.).

Such new methods will enable the attacking of challenges such as how to design data-technologies that better protect usability, privacy, security and confidentiality. It could also provide scientific underpinnings for the practice of UX design.

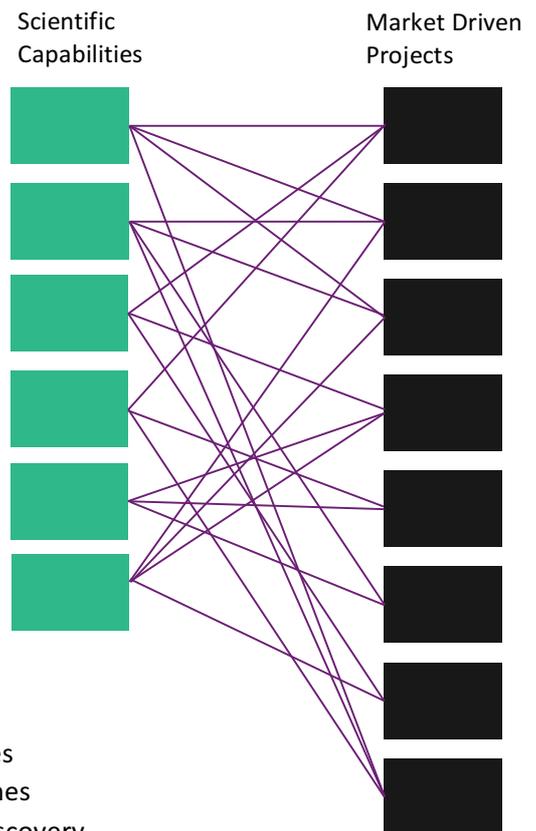
Impacts

Data61's L-shaped model (see page 1) means that our impacts are the product of our scientific capabilities with market forces and opportunities. These impacts are managed through our business development and product management processes. A given scientific capability can deliver impact in many end-use problems⁵²; a given market need can be satisfied by many different scientific capabilities⁵³ – see the schematic to the right.

The *science* driven challenges are our view of where technology needs to move. The end-use projects we do will largely be driven by the *market's* view of this. It will be primarily through these projects that the science will have its larger impact. This impact can be categorized in many overlapping ways. Three are given below:

General categories:

- Improvement in the efficiency of Australian businesses
- Improvement in the efficiency of Australian governments
- Improved reliability, safety and security of data-technologies
- Generation of new industries, especially platform centric ones
- Improvement in the speed and effectiveness of scientific discovery.



Data61 market focus categories (in partnership with other BUs where possible):

- Safety and Security
- Health & Communities
- Future Cities
- IoT/Industrial Internet
- Agri-business
- Spatial Intelligence
- Data-driven Government
- Enterprise Services + Fintech
- Defence

Whole of CSIRO categories⁵⁴

- Food security and quality
- Clean energy and resources
- Health and wellbeing
- Conservation and use of our natural environment
- Innovative industries
- A safer Australia

Data61's research in support of the scientific vision of the present document will support projects in these impact areas, and will thus find pathways to impact through them. Individual projects are responsible for analysing, shaping and articulating what those pathways and impacts will be. This needs to be done in an agile manner, adapting to opportunities, but building upon our focused scientific capability.

Endnotes

¹ It is deliberately called a “vision”, and not (metaphorically) a “roadmap” – a roadmap is a two-dimensional graphical representation of something that already exists (roads), and is rarely something inspiring and exciting; at best a “science / technology roadmap” it is a visual depiction of the expected temporal evolution of a technological product family (Ronald N. Kostoff and Robert R. Schaller, *Science and Technology Roadmaps*, *IEEE Transactions on Engineering Management*, 48(2), 132-143 (2001); Lianne Simonse, Jan Buijs & Erik Jan Hultink, Roadmap grounded as ‘visual portray’: Reflecting on an artifact and metaphor, *Helsinki EGOS 2012 Sub-theme 09: (SWG) Artifacts in Art, Design, and Organization* (2012)) which suffers by being constrained to a two dimensional visual form. Conversely, a “vision” can be of something that does not exist, and can inspire and excite and is not constrained to fit any particular format. It tells where we want to go, and outlines in broad strokes how we might get there, without actually pinning the exact path down. It is a *science* vision in the general sense of the word “science” – systematised knowledge; see endnote 4. We expect to develop more traditional technology roadmaps (i.e. temporally linear expectations and plans) for particular product and service offerings which we develop.

² At different times in computing’s evolution, either the demand (market) or the technology push side have been dominant; but it is never just one or the other; see Jan van dem Ende and Wilfred Dolfsma, Technology push, demand pull and the shaping of technological paradigms – Patterns in the development of computing technology, *Journal of Evolutionary Economics* 15, 83-99 (2005). The reality is, of course, complex, and recombination (the mixing up of different ideas) plays an essential part (Cristiano Antonelli, Jackie Krafft, Francesco Quatraro. Recombinant Knowledge and Growth: The Case of ICTs, *Structural Change and Economic Dynamics*, Elsevier, 21(1), 50-69 (2010)) and the “demand-pull” model seems to be losing favor as a satisfactory explanation (Benoit Godin and Joseph P. Lane, “Pushes and Pulls”: *The Hi(story) of the Demand Pull Model of Innovation*, Project on the Intellectual History of Innovation, working paper No 13 (2013); Benoit Godin, *Innovation Contested: The Idea of Innovation over the Centuries*, Routledge (2015)).

³ The document has multiple intended audiences:

- DATA61 talent (existing and potential future) – to align what we do, to help us say “no” to opportunities that do not align, and to achieve large impact multiplicatively.
- Rest of CSIRO and external partners – to articulate our own longer term research goals to serve as one of the filters we will apply in considering engaging in joint projects.
- Wider public – to explain what we do.

⁴ It would be unfortunate, and unhelpful, to get hung up on the distinction between science, engineering and technology. This document presents an aspiration for the new knowledge we will create – *novum scientia*. While engineering knowledge *is* different from scientific knowledge (Walter G. Vincenti, *What Engineers Know and How They Know I: Analytical Studies from Aeronautical History*, The Johns Hopkins University press (1990)) and technology is more than mere scientific knowledge (W. Brian Arthur, *The Nature of Technology: What it is and How it Evolves*, Simon and Schuster (2009)), the essence of engineering research (the improvement of technology) remains the production of new knowledge (Edwin T. Layton Jr, Technology as Knowledge, *Technology and Culture* 15(1), 31-41 (January 1974)). The research Data61 does spans all of these headings, and more, such as “design-driven innovation” – the phrase is from Roberto Verganti’s book *Design-Driven Innovation: Changing the Rules of Competition by Radically Innovating What Things Mean*, Harvard Business Press (2009) – new business models, and ethnographic approaches to data technologies.

We should aspire to seek new knowledge (motivated by real problems and the desire to improve our current technologies) wherever it takes us, in the spirit of the great researchers of the past (Lisa Jardine, *Ingenious Pursuits: Building the Scientific Revolution*, Little Brown, London, 1999; Jenny Uglow, *The Lunar men: The Friends Who Made the Future*, Faber and Faber 2002). Our inspirations and role models should be polymaths such as Robert Hooke (Lisa Jardine, *The Curious Life of Robert Hooke: The Man who Measured London*, HarperCollins (2003); Stephen Inwood, *The Man Who Knew Too Much: The Strange and Inventive Life of Robert Hooke 1635-1703*, MacMillan (2002); Robert D. Purrington, *The First Professional Scientist: Robert Hooke and the Royal Society of London*, Birkhauser (2009); Jim Bennet, Michael Cooper, Michael Hunter and Lisa Jardine, *London’s Leonardo – The Life and Work of Robert Hooke*, Oxford University press (2003)) or Charles Babbage (Laura J. Snyder, *The Philosophical Breakfast Club: Four Remarkable Friends who Transformed Science and Changed the World*, Broadway Books (2011)) both of whom freely moved between science and technology.

As noted long ago (Robert P. Multhauf, The Scientist and the “Improver” of Technology, *Technology and Culture* 1(1), 38-47 (1959)), there is no perfect word for the *improver of technology*: “engineer” is widely used, but it still primarily refers to the expert practitioner and not necessarily the *improver*. Perhaps we, as improvers of technologies for data, should not worry whether what we do is adequately described as “science”, “engineering” or anything else, and just refer to ourselves by Hilary Cinis’ elegant neologism: “datanauts”.

⁵ It is common that vision statements become all-encompassing, excluding nothing. That the present vision does not aim to cover everything can be tested by comparing it to the substantially broader set of goals in *Future Science – Computer Science: Meeting the Scale Challenge*, Australian Academy of Science (2013), or President’s Council of Advisors on Science and Technology, *Report to the President and Congress. Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology*, Executive Office of the President (December 2010).

⁶ See John Archibald Wheeler, Information, Physics, Quantum: The Search for Links, in *Proceedings of the 3rd International Symposium on the Foundations of Quantum Mechanics*, Tokyo, (1989); Hector Zenil (Ed.), *A computable universe: understanding and exploring nature as computation*, World Scientific (2013); Rolf Landauer, Uncertainty principle and minimal energy dissipation in the computer, *International Journal of Theoretical Physics* 21(3/4), 283-297, (1982); Rolf Landauer, The physical nature of information, *Physics Letters A*, 217, 188-193 (1996); Antonie Berut et al., Experimental verification of Landauer’s principle linking information and thermodynamics, *Nature* 483, 187-190, (8 March 2012); Juan M.R. Parrondo, Jordan M. Horowitz and Takahiro Sagawa, Thermodynamics of Information, *Nature Physics*, 11, 131-139, (February 2015); Gilles Brassard, Is information the key? *Nature Physics* 1, 2-4, (October 2005).

⁷ Jean-Marie Lehn, Perspectives in Supramolecular Chemistry—From Molecular Recognition towards Molecular Information Processing and Self-Organization, *Angewandte Chemie International Edition in English*, 29(11), 1304–1319, (November 1990); Jean-Marie Lehn, Supramolecular chemistry – scope and perspectives – molecules – supermolecules – molecular devices, *Nobel Prize Lecture*, (8 December 1987).

⁸ John Maynard Smith, The concept of information in biology, *Philosophy of Science* 67(2), 177-194 (2000); confer Ladislav Kovac, Information and knowledge in biology: time for reappraisal, *Plant Signalling and behaviour* 2(2), 65-73 (2007).

⁹ David Easley and Jon Kleinberg, *Networks, crowds and markets: reasoning about a highly connected world*, Cambridge University Press (2010).

¹⁰ Friedrich A. Hayek, The use of knowledge in society, *The American Economic Review*, 35(4), 519-530 (1945); George J. Stigler, The Economics of Information, *The Journal of Political Economy* 69(3), 213-225 (1961); Joseph E. Stiglitz, Information and the change in the paradigm in economics, *Nobel Prize Lecture* 8 (December 2001).

¹¹ Werner Callebaut and Diego Raskim-Gutman, *Modularity: Understanding the development and evolution of natural complex systems*, MIT Press, (2005); Jeff Clune, Jean-Baptiste Mouret and Hod Lipson, The evolutionary origins of modularity, *Proceedings of the Royal Society (series B)*, 280, 20122863 (2013)

¹² David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-Lazlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy and Marshall Van Alstynr, Computational Social Science, *Science* 323, 721-723 (2009).

¹³ Committee on the Mathematical Sciences in 2025, Board on Mathematical Sciences and Their Applications, Division on Engineering and Physical Sciences, National Research Council of the National Academies, *The Mathematical Sciences in 2025*, The National Academies Press, (2013).

¹⁴ Cristian S. Calude (Ed), *Randomness and Complexity: From Leibniz to Chaitin*, World Scientific, (2007).

¹⁵ Richard G. Lipsey, Kenneth I. Carlaw and Clifford T. Bekar, *Economic Transformations General Purpose Technologies and Long-Term Economic Growth*, Oxford University Press (2005).

¹⁶ Robert C. Williamson, Michelle Nic Raghnaill, Kirsty Douglas and Dana Sanchez, *Technology and Australia’s future: New technologies and their role in Australia’s security, cultural, democratic, social and economic systems*, Australian Council of Learned Academies, September 2015.

¹⁷ National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The National Artificial Intelligence Research and Development Strategic Plan*, (October 2016).

¹⁸ These complement other broader principles underpinning everything we do, such as national benefit; see the Data61 operating model document.

¹⁹ “We” here refers to the broader Data61+ network. This principle implies avoiding NIH (Not Invented Here) syndrome; we do not need to invent everything ourselves. We should focus on the things that we, and we alone, can

do; and then network with others in a rich and complex manner. It would be supremely ironic if our organisation that underpins the information society does not embrace all of its implications (Manuel Castells, *The Rise of Network Society* (2nd Edition), Wiley-Blackwell (2010)).

²⁰ The word is pinched from a suitably inspiring institution: The MIT media lab, which so describes itself <https://www.media.mit.edu/about>. The principle, of course, implies much collaboration with other disciplines, but goes beyond the traditional “multi-disciplinary” to a stronger problem-oriented perspective – “There are no subject matters; no branches of learning – or, rather, of inquiry: only problems and the urge to solve them. A science such as botany or chemistry ... is, I contend, merely an administrative unit” (Karl Popper, *Realism and the Aim of Science*, Rowman and Littlefield (1983)). Such a stance implies widespread collaboration without fear of crossing boundaries. It does not imply a lack of “canon” or core; our canon is *primarily* that of cybernetics broadly construed.

²¹ This viewpoint is given the fancy name of “technological determinism” with the concomitant fear of “autonomous technology” (Langdon Winner *Autonomous technology: Technics-out-of-control as a theme in political thought*. MIT Press, 1978). The counter is that technologies can be, and are, shaped by society. The reality is that while technologies do indeed have “momentum” (Thomas P. Hughes “The evolution of large technological systems.” Pages 51-82 in Wiebe E. Bijker et al. (eds), *The social construction of technological systems: New directions in the sociology and history of technology* (1987)) and “drive history” (Merritt Roe Smith and Leo Marx. *Does technology drive history? The dilemma of technological determinism*. MIT Press (1994)) there remains a huge freedom of choice in terms of how they are used and their precise form. Like all technologies of the past, technologies for data can also be shaped for social and national benefit.

²² Russell Hardin, *Trust and Trustworthiness*, Russell Sage Foundation, New York, (2002); Frances Fukuyama, *Trust: The Social Virtues and the Creation of Prosperity*, Simon and Schuster (1995); Eric M. Uslaner, *The Moral Foundations of Trust*, Cambridge University Press (2002). An excellent short summary of the social side of trust is chapter 21 of Jon Elster, *Explaining Social Behaviour: More Nuts and Bolts for the Social Sciences*, Cambridge University Press (2007). People’s trust in technology is a complex matter (Karen Clarke, Gillian Hardstone, Mark Rouncefield and Ian Sommerville, *Trust in Technology: A Socio-Technical Perspective*, Springer (2006); Meinolf Dierkes and Claudia von Grote (eds), *Between Understanding and Trust: The Public, Science and Technology*, Routledge (2000)); and trust in *technological experts* (as opposed to the technology itself) is surprisingly weakly correlated with perceptions of risk (Lennart Sjöberg, Limits of Knowledge and the Limited Importance of Trust, *Risk Analysis* 21(1), 189-198 (2001)).

²³ In the sense of George Lakoff and Mark Johnson, *Metaphors we Live By*, The University of Chicago Press (1980) – not as a mere rhetorical flourish, but as an essential way in which to make sense of what we do.

²⁴ Trust is a very complex notion, and means different things to different people: (D. Harrison McKnight and Norman L. Chervany, *The Meanings of Trust*, University of Minnesota, (1996); Donna M. Romano, *The Nature of Trust: Conceptual and Operational Clarification*, PhD thesis, Louisiana State University (2003)).

The complexity is illustrated follows:

Trust has not only been described as an “elusive” concept, but the state of trust definitions has been called a “conceptual confusion”, a “confusing potpourri”, and even a “conceptual morass”. For example, trust has been defined as both a noun and a verb, as both a personality trait and a belief, and as both a social structure and a behavioral intention. Some researchers, silently affirming the difficulty of defining trust, have declined to define trust, relying on the reader to ascribe meaning to the term. (D. Harrison McKnight and Norman L. Chervany, *Trust and Distrust Definitions: One Bite at a Time*, in R. Falcone, M. Singh, and Y.-H. Tan (Eds.): *Trust in Cyber-societies*, LNAI 2246, pp. 27–54, Springer-Verlag (2001)).

Perhaps, like “culture” (confer Kroeber’s 164 definitions of culture: Alfred L. Kroeber and Clyde Kluckhorn, *Culture: A critical review of concepts and definitions*, Peabody Museum of American Archeology and Anthropology, (1952) or “technology” (confer Robert C. Williamson, Michelle Nic Raghnaill, Kirsty Douglas and Dana Sanchez, *Technology and Australia’s future: New technologies and their role in Australia’s security, cultural, democratic, social and economic systems*, Australian Council of Learned Academies, (September 2015)), it makes little sense to attempt to *define* trust, but rather we should focus upon the technological and scientific problems we want to solve (as done in the main text).

The notion of trust as a concept in computing has had attempts to formalise it for some time, starting at least 20 years ago (Stephen Paul Marsh, *Formalising Trust as a Computational Concept*, PhD thesis, University of Stirling, (1994)), with conferences on the topic starting over decade ago (Sokratis Katsikas, Javier Lopez and Gunther Pernul (eds), *Trust and Privacy in Digital Business: First International Conference, Trustbus 2004*, Springer (2005); Thorsten Holz and Sotiris Ioannidis, *Trust and Trustworthy Computing: 7th International Conference TRUST 2014*, Springer (2014)).

One reason for the complexity is because of the many threats to trust (in the same way there are many threats to security, which need to be explicitly declared or modelled: Adam Shostack, *Threat Modelling: Designing for Security*, Wiley (2014)). But primarily the complexity comes simply from the diverse elements to trust in data-centric systems including, but not limited to:

-
- Trust in the reliability of software (never absolute: see Donald MacKenzie, *Mechanizing Proof: Computing, Risk and Trust*, MIT Press (2001); Juan C. Bicarregui and Brian M. Matthews, *Proof and Refutation in Formal Software Development*, 3rd *Irish Workshop on Formal Methods* (1999));
 - Trust in security (e.g. Jeffrey J.P. Tsai, Philip S. You (eds), *Machine Learning in Cyber Trust: Security, Privacy, and Reliability*, Springer (2009));
 - Trust in data management (Milan Petkovic and Willen Jonker (eds), *Security, Privacy, and Trust in Modern Data Management*, Springer (2007));
 - Trust in the credibility of information, such as which scientific results one can rely upon: (Christine L. Borgman, *Scholarship in the Digital Age: Information, Infrastructure and the Internet*, MIT Press (2007)) and what sensor measurements one can trust (J.C. Wallis, C.L. Borgmann, Matthew Mayernik, Alberto Pepe, Nithya Ramanathan and Mark Hansen, Know thy Sensor: Trust, Data Quality, and Data Integrity in Scientific Digital Libraries, *11th European Conference on Research and Advanced Technology for Digital Libraries*, September 16–21, 2007, Budapest, Hungary (2007)). This is already front-of-mind in work such as “bees with backpacks” that Data61 has done. It is hardly a new concern – the (apparently simple) notion of a scientific measurement is deeply entangled with notions of trust, as is evident from the history of Victorian science (Graeme J.N. Gooday, *The Morals of Measurement: Accuracy, Irony, and Trust in Late Victorian Electrical Practice*, Cambridge University Press (2004)).
 - Trust that social mechanisms built with data-technologies cannot be manipulated (See Eric Friedman, Paul Resnick and Rahul Sami, Manipulation-Resistant Reputation Systems, Chapter 27 in Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay V. Vazirani, *Algorithmic Game Theory*, Cambridge University Press (2007));
 - Trust that sensitive information is not leaked (Guillermo Lafuente, The big data security challenge, *Network security* 2015., 12-14 (2015);
 - Trust that data-analytics are fair (Solon Barocas and Andrew D. Selbst. Big data's disparate impact. *California Law Review* 104 (2016); Danah Boyd and Kate Crawford, Six provocations for big data. In *A decade in internet time: Symposium on the dynamics of the internet and society* (pp. 1-17). Oxford Internet Institute, (September 2011));
 - Trust in the communication system underpinning data technologies (White House: "Cyberspace policy review: Assuring a trusted and resilient information and communications infrastructure." *White House, United States of America* (2009)). There is no perfectly trustable communication system, and so like all other elements of the trust chain, a risk sensitive approach will be warranted.
 - Trust that the overall systems constructed can be sufficiently relied upon (Piotr Cofta, *Trust, Complexity and Control: Confidence in a Convergent World*, John Wiley and Sons (2007)).

²⁵ The phrase alludes to an admirable novel about two famous scientists who are further (in addition to Hooke and Babbage – see endnote 4) great role models for Data61 – Alexander von Humboldt and Carl Freidrich Gauss (Daniel Kehlman, *Measuring the World*, Pantheon (2006)). Humboldt is one of the most important creators of modern science, who undertook outstandingly painstaking data gathering and analysis (Andrea Wulf, *The Invention of Nature: The Adventure of Alexander von Humboldt, Lost Hero of Science*, John Murray, (2015)). Gauss is famously credited as the originator of least squares data analysis (Stephen M. Stigler, Gauss and the invention of least squares, *The Annals of Statistics*, 9(3), 465-474 (1981)) and thus one of the fathers of modern data analytics.

In an earlier version of this document, I used the awkward polysyllabic neologism “datafication”, apparently coined in the article by Kenneth Cukier and Viktor Mayer-Schoenberger: [The Rise of Big Data](#), *Foreign Affairs* 28–40, May/June, (2013). It is already widely used, but it is an ugly word that many Data61 folks reacted negatively to, and, crucially, it misses the distinction between data and capta (see below).

²⁶ This distinction is quite old, but rarely used. See Rob Kitchin, *The Data Revolution: Big data, open data, data infrastructures and their consequences*, Sage, Los Angeles (2014); this explains some of the history of the word; Christopher Chippindale, Capta and data: on the true nature of archaeological information, *American Antiquity* 65(4), 605-612 (2000); Bettina Berendt, *Big Capta, Bad Science? On two recent books on “Big Data” and its revolutionary potential*, Department of Computer Science, KU Leuven, <https://people.cs.kuleuven.be/~bettina.berendt/Reviews/BigData.pdf> (March 2015).

²⁷ Quoted from the entry for *captus* in *A Latin Dictionary. Founded on Andrews' edition of Freund's Latin dictionary. revised, enlarged, and in great part rewritten by. Charlton T. Lewis, Ph.D. and. Charles Short, LL.D.* Oxford. Clarendon Press (1879).

²⁸ The traditional view is widespread; e.g. Paul Cooper, Data, information and knowledge, *Anaesthesia and Intensive Care Medicine*, 11(12), 505-506 (2010).

²⁹ Ashley Braganza, Rethinking the data-information-knowledge hierarchy: towards a case based model, *International Journal of Information Management*, 24, 347-356 (2004); Ilkka Tuomi, Data is more than Knowledge: Implications of the Reversed Knowledge Hierarchy for Knowledge Management and Organizational Memory, *Journal of Management Information Systems* 16(3), 103-117 (1999).

³⁰ It is sometimes claimed to be a clearer distinction than it really is: Sreenivas Rangan Sukumar, *Machine learning for data-driven discovery: thoughts on the past, present and future*, Oak Ridge National Laboratory, (2014).

³¹ Tony Hey, Stewart Tansley and Kristin Tolle, *The Fourth Paradigm: Data-intensive scientific discovery*, Microsoft Research, (2009).

³² Alon Halevy, Peter Norvig and Fernando Pereira, The Unreasonable Effectiveness of Data, *IEEE Intelligent Systems Magazine*, 8-12 (March/April 2009).

³³ Carole Goble and David De Roure, *The impact of workflow tools on data-centric research*, <http://www.myexperiment.org/files/215/download/workflows-v8-05May2009.pdf> (May 2009).

³⁴ David De Roure and Carole Goble, Anchors in Shifting Sand: The Primacy of Method in the Web of Data, *Web Science Conference*, (April 2010).

³⁵ This (entirely wrong) phrase is due to Chris Anderson: "The end of theory: the data deluge makes the scientific method obsolete," *Wired* (23 June 2008). It does no such thing! It simply allows for more sophisticated models.

³⁶ Sean Bechhofer, Iain Buchan, David De Roure, Paolo Missier, John Ainsworth, Jiten Bhagat, Philip Couch et al., Why linked data is not enough for scientists, *Future Generation Computer Systems* 29(2), 599-611, (2013).

³⁷ Ludwick Fleck, *Genesis and Development of a Scientific Fact*, University of Chicago Press (1979); Bruno Latour and Steve Woolgar, *Laboratory Life: The Construction of Scientific Facts*, Sage Publications (1979); Karl Popper, *The Logic of Scientific Discovery*, Hutchinson, (1959).

³⁸ Geoffrey C. Bowker, *Memory Practices in the Sciences*, MIT Press, (2005).

³⁹ Mark Stalzer and Chris Mentzel, A preliminary review of influential works in data-driven discovery, *SpringerPlus* 5:1266, (August 2016).

⁴⁰ There are other reasons that serve to push for embedding analytics, especially latency and bandwidth limitations.

⁴¹ William N. Dunn (ed), *The Experimenting Society: Essays in honour of Donald T. Campbell*, Transaction Publishers, (1997); Donald T. Campbell, Methods for the Experimenting Society, *American Journal of Evaluation* 12, 223-260, (1991); Donald T. Campbell, Reforms as Experiments, *American Psychologist*, 24, 409-429, (1969).

⁴² As explained elsewhere in this document, such a phrase ("universal captafication") does not imply it is done once, without a theoretical stance, and the data "speak for themselves." What is meant here is simply the push towards more pervasive (hence approaching "universal") translation of the data in the world into capta that can be manipulated.

⁴³ "Delivered" in the title of this headline is the right word – we propose to change the delivery modality, and to actually build systems that literally deliver the results.

⁴⁴ Confer strategy 4 of the National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The National Artificial Intelligence Research and Development Strategic Plan*, October 2016: it articulates the need for explainable and transparent systems that are trusted by their users, perform in a manner that is acceptable to the users, and can be guaranteed to act as the user intended.

⁴⁵ Patrick McDaniel et al, Towards a Secure and Efficient System for End-to-End Provenance. *2nd workshop on the theory and practice of provenance* (2010).

⁴⁶ Data technologies are made up of hardware and software, the boundary of which is somewhat blurred. Our primary (but not exclusive) focus here is on the software because it is with regard to that that we have a global competitive advantage. One could use the more general phrase "systems you can trust" but that misses the specificity that I currently have. And all of the research I am alluding to here is indeed on software.

⁴⁷ Jason Furman, Is this Time Different? The Opportunities and Challenges of Artificial Intelligence, Remarks at *AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term*, New York University, (July 7, 2016).

⁴⁸ National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee, *The National Artificial Intelligence Research and Development Strategic Plan*, (October 2016).

⁴⁹ "Scientific" is meant in the broad sense described in endnote 4.

⁵⁰ E.g. Nathan Rosenberg and L.E. Birdzell Jr., *How the West Grew Rich: The Economic Transformation of the Industrial World*, Basic Books (1986).

⁵¹ Huw T.O. Davies, Sandra M. Nutley and Peter C. Smith, *What Works? Evidence-based policy and practice in public services*, The Policy Press (2000).

⁵² Pleiotropy (genetically), or non-injectivity of the inverse map (mathematically).

⁵³ Genetic heterogeneity or non-injectivity of the forward map.

⁵⁴ Elizabeth Eastland, *Future Australia – Market Vision: Unlocking a more prosperous and sustainable future for all Australians*, Powerpoint presentation (2 November 2016).

CONTACT US

t 1300 363 400
+61 3 9545 2176
e csiroenquiries@csiro.au
w www.data61.csiro.au

AT CSIRO WE SHAPE THE FUTURE

We do this by using science and technology to solve real issues. Our research makes a difference to industry, people and the planet.

FOR FURTHER INFORMATION

Bob Williamson
Chief Scientist, Data61
t +61 2 6218 3712
m +61 404 053 877
e Bob.Williamson@data61.csiro.au
w www.data61.csiro.au

Adrian Turner
CEO, Data61
t +61 9372 4202
m + 61 475 981 219
e Adrian.Turner@data61.csiro.au
w www.data61.csiro.au

