HOW TO LOVE NON EXOTIC TECHNOLOGY

And use it to decarbonise, achieve cybersecurity, and improve organisational health and learning

Karl Jeffery and
Dimitris Lyras

CONTENTS

Published by Software for Domain Experts Ltd,
39-41 North Road, London, N7 9DP, UK
July 2020

I WANT YOU TO LOVE NON EXOTIC TECHNOLOGY

Introduction

Digital technology offers the most value when it gives people exactly the right information and support at the right time. But that is not the technology which gets the most attention.

The exotic technologies get the most attention - like AI/ML, blockchain, and augmented reality. But it is proving very hard for these technologies to make value, and perhaps they only provide value when handled in a very non exotic way, with careful modelling.

To make technology which can give people the right support at the right time is very difficult work - with large amounts of modelling or map making - understanding what exactly would be most helpful. And there is also a lot of data management, governance, standards and integration behind it (not covered in this book).

This is a frustration shared by many people who work with digital technology for organisations.

Many people in the digital technology industry respond to this challenge by trying to make exotic technologies sound even more exciting and relevant to the needs of organisations. Or they bury the difficult work of modelling technologies to people's needs behind fashionable concepts like agile sprints.

We propose a bold alternative approach – showing how we can love non-exotic technologies - by focussing on what we can do with them.

Talking about technology itself

Technology could do much more to help people do their work better.

But we don't talk about this much. What we talk about, when we talk about technology, is technology itself.

That is something very different. For example, technology which can analyse videos, drive cars, beat the world champion in a game, and learn how to do something new. For the purposes of this book, these can be called "exotic technologies".

Technology which can help people do their work better could be considered "non exotic technology", because it might not look like very much, in itself. It may be possible to build using old, cheap technology platforms, so not require any new software purchases. Like an old car which you use to go somewhere new.

Building non exotic technology is not easy. It requires a lot of hard dedicated work by people in the middle who understand both people and technology, and are good at reducing complex concepts to maps which other people and machines can easily understand.

But this technology could help people with expertise to make better decisions to faster decarbonise society within commercial constraints. Or help people make more effective schedules so organisations can achieve more with less resources.

Cybersecurity could be improved if people could better gain situation awareness about what is going on, rather than assuming that technology problems need technological solutions.

Online learning products could be improved if they could better work together with the expertise of a teacher who has worked in the same domain.

The purpose of this book is not to complain about the current situation, with people getting too excited by exotic technologies. Instead, the intention is to show how it can be even more exciting to work with technology which may not initially seem exotic. So exciting, we think you might fall in love with it.

We may even find that machine learning / artificial intelligence, blockchain and augmented reality can add more value to organisations when they are applied in a less exotic way. Applying them only where they can ultimately serve to improve an expert's understanding of a situation, or automate the easier decisions. Or they can help people to better fix a technical problem or learn.

Loving our tools

Let us start with the idea that we can love digital tools if they help us do our work better, the same way as we might love any other tool we use, such as a bicycle, camera, food processor, chisel, pens and laptop. We love it not for what it is, but what we can do with it.

People have put in a lot of effort since the start of humanity into perfecting tools, and we can look at digital technology in the same way. And the more love we have for the tools, the easier it is to find the motivation to do the hard work of making tools which really fit what we want to do with them.

We could be able to love our digital tools in the same way, because we feel they give us exactly what we need.

And that goes for the whole tool, not just the physical appearance of the tool, or how it feels like to hold. The digital analogy here is that we are talking about how the software is built, not just its user interface.

For example, we could have a digital tool to help us make decisions about decarbonising as part of planning an industrial process. We might expect the software to provide us all the relevant information we need to make a decision, taking into account both carbon and commercial constraints, such as whether people will travel to an offshore oil and gas platform by helicopter or by boat.

The idea of seeing technology as a tool does have a problem though. Many people are so invested in a way of looking at technology based around what it is, not what it can do for people, they may struggle to see the relevance of, or support, the approach we are advocating here.

DIGITAL TECHNOLOGY SHOULD PROVIDE SITUATION AWARENESS

When we are thinking about what digital tools should look like which help people to do something, a good starting point is to say that the prime purpose of the tool is to support somebody's situation awareness - giving people a better understanding of what is going on. Telling people what they need to know, when they need to know it.

This can apply both to people making decisions which are not related to technology (such as in the decarbonisation example above), or making decisions about technology and how it applies to the real world, such as understanding cybersecurity.

The same approach applies to technology for learning, but in this case the software needs to be programmed to have an understanding of how far the student has come on their learning pathway, and what would be good to teach them next – so the situation awareness is within the software, not the student.

These examples all sound very different, but the route to achieving them comes together in one concept – that of the map, as a means of connecting machines with people intelligence.

Our concept of the map

The map is an abstract concept, and we mean it in a very abstract way, hopefully you will build an understanding of what we mean as you read through this book.

The map is where people intelligence and digital intelligence come together – because the right sort of map can bring the necessary understanding to both people and machines.

The map can be perhaps best understood as something we carry in our minds – when we understand how something works, or how to get somewhere. But if we want to share this map with someone else, or a machine, there needs to be a way to write it down.

It could be written down with a narrative like this, or by drawing lines on paper, or in an electronic document, like a PowerPoint slide. If a computer is going to read it, it needs to be in a machine readable format, which could be anything up to actual computer code, or a graph database structure or knowledge graph.

Computers can only follow very rigid instructions. A detailed map with no space for interpretation could be another name for computer code, software or algorithms.

The format for the physical manifestation of the map is less important than the idea that a map is the communications tool you need to build really useful digital tools.

Maps in our minds

Where do maps come from? We create them in our minds all the time. (They might also be known as mental models).

Think about the mental maps we form of how to navigate a deep long term relationship, such as how to behave from day to day in a way which will keep things on course, what to do when the other person behaves in various ways, signals that something may be going off course, what helped put things right last time it happened, and what we have done in the past which we thought would correct a problem but had the opposite effect, so we won’t do it again.

We hold maps in our mind at a wide range of different complexity or granularity levels – compare the map you have of a relationship with a life partner, and someone you barely know.

We can also explain our maps to others, but perhaps leaving out much of the detail, to make it easier for someone else to absorb, such as with a teacher explaining something to a novice.

Not all our mental maps are so complicated – we can have a simple map of how to reach a destination at a certain time, taking into account everything you need to do, preparing for the trip and taking a journey, leaving appropriate allowances for unplanned events.

When companies fret about the knowledge they are losing of people who are retiring, what they might really mean is the maps that these experienced people have, knowing what relates to what. This is not something you would write down in a list, but something you might draw with lines and boxes on a large piece of paper, if you could find a way to extract it from someone's mind.

People intelligence and digital intelligence

You know what people intelligence means, but when we talk about people intelligence in the context of digital intelligence, we come up with a slightly different picture, emphasising what people intelligence can do that digital intelligence cannot do.

The big difference is modelling. People can make enormously complex models in their minds using enormous amounts of diverse information and uncertainty. Computers are less good at handling diverse data than people, but can handle bigger data volumes.

If you have a situation with low volumes of data, but many diverse components, a person is ideal. This is precisely the sort of situation we were evolved to deal with, living in our tribes of 150 people but with multiple different risks to assess.

If you have a situation with high volumes of data, then only a computer could process it.

But if you have a situation with both high volumes of data and diverse data, like many situations in reasonably large organisations today, you may be better off working out how to use computers to distil the data and then having people to make decisions with it.

Many tech experts look at this scenario and say that the only possible solution is artificial intelligence, modelling large data volumes. Although in nearly all cases they have never seen AI do such a thing, and do not have much understanding of how hard and expensive it would be to build, they have just heard many promises that it can. Many technical people, it seems, have a propensity to always look for technical solutions to a problem, rather than accepting that people may have some value to add, too.

A person can easily see if an IOT sensor is doing something it shouldn't be doing – but may need a computer to first spot that the sensor is doing something which looks unusual, while monitoring 100,000 sensors at once.

A person can easily make a judgement if one approach to industrial process development will be both less expensive and better for carbon emissions than another, but may need a computer to do the legwork of serving up the information needed to make this decision.

A person can make sensitive decisions about balance – what control measures can be made which will achieve some socially desired result, such as cybersecurity or decarbonisation, without impeding people to do what they want to do, understanding that if there is any impedance people will not warm to the necessary control measures.

A person can make difficult decisions about scheduling, massing together a multitude of possible problems and obstacles and assessing which ones are and aren't worthy of consideration, a task a computer could probably never do unless specifically programmed.

Exotic and non-exotic technologies can co-exist

We defined exotic technologies as technologies people get excited about for what they can do in themselves, and non exotic technology as technologies (we hope) people can get excited about for what they can do.

Note these are not exclusive categories – in the same way you can love a car for being a car, and love a car for where it takes you.

Even the most exotic technologies like blockchain and machine learning can help people to do things.

So the core goal of this book could be more precisely defined as finding ways to love technology for where it can take us, not for what it is – and exploring more what we can do if we take this approach.

There are many analogies in parenting, where a parent tries to get a child enthused in a long term journey by giving them something they value having right now. Buying an exciting computer game which also teaches maths. Encouraging effort in a sport or musical instrument using certificates and prizes. Encouraging the child to eat food which will provide long term nutrition by covering it with cheese.

Making non exotic technology work is hard

While we can get interest in exotic technologies based on their exotic characteristics, getting interest and maintaining focus in non exotic technologies is really hard.

We need to understand what information people need to see and when, and structure the digital technology to do that. We need to make it easier for people to understand what the technologies are doing and the various mechanisms they have to stop them getting hacked or giving us wrong information. We need to understand the processes which work getting these technologies used in our organisations, and we need to understand the roles that people do, to get this done.

We are really amateur at this, as of 2020. Companies are treating digital technology as a big experiment. The dominant way of working with digital technology is to throw all the expensive experts into a room for 2 weeks to see if they come up with something useful. Meanwhile many companies still don't know or much respect what their data or digital managers actually do.

Exotic technologies alienate the non-fans

It is worth reflecting on how much our talk about exotic technologies might be alienating people who could otherwise be happily engaged with making technology work.

There are many people who have reasonable technical skills, although not at the level of an IT expert, but also better people skills than a typical IT expert. They could combine these skills to do the work of modelling digital technology to what people need.

And understanding how to really make machine learning and blockchain add value does take a large amount of technical competence, and people who do not have this can easily feel alienated from a discussion about them.

Technology vendors can make the problem worse, when they have a marketing plan based on talking about exotic technologies as much as they can. It makes it hard to determine what they are actually selling, and if it has anything of value or not.

Making enemies

So the ideas in this book hold potential to make many enemies.

We will make enemies of the people whose job it is to sell exotic technologies, and everyone else involved in these processes, including investors, other employees, people on the buying side who have a role to promote it internally within their organisations, and consultants making a living out of implementing it or telling people they need it.

We will make enemies of the people who believe that artificial intelligence can make a big change to everything, and that 'artificial general intelligence' is just around the corner, so any advantages to people intelligence will only last for a few years.

We will struggle to make much money from these ideas ourselves, since selling technology which does not sound exciting has a tough job beating something which sounds exciting. And non-exotic technology, by definition, does not sound that exciting, in itself.

WHAT YOU CAN DO WITH NON EXOTIC TECHNOLOGY

Introduction

Here are some ideas about how we could start thinking about how non exotic technologies could help us make the world better. We'll look at how they can help fix carbon emissions, fix cyber security, make learning more equal, support developing countries, help us have better healthcare, recruit more black police officers, and fix more potholes in the road, with less temporary traffic lights.

How to reduce carbon emissions

The question of reducing carbon emissions is normally thought of like this – those nasty oil companies need to be stopped, people should fly and drive less, we need more renewable energy and less use of plastic.

We are starting here with a different way of looking at the problem – to say that we can reduce CO2 emissions when we find ways to do what we were usually doing, at a price we are used to paying, but making much less CO2 in how we go about it.

Governments will usually only make new rules or taxes for carbon when they see that it does not impose much higher costs on anyone. For example they will encourage investment in wind power but only if it is not too much more expensive than coal. So we cannot rely on regulation and policy alone to decarbonise.

The challenge comes down to decision making – finding better ways to incorporate carbon emission predictions in every financial decision. This can be anything from a design for a new school, our choice of family holiday, and a maintenance program for a manufacturing plant. And when government sees that it is possible to make choices with much lower CO2 but at similar costs, they won't mind forcing people to choose the low CO2 option. The role that digital technology could play is supporting this by providing the right data at the right time.

The second part of this challenge is that carbon data is generally enormously diverse, which makes it much more suitable for 'model' based thinking which people do well, rather than 'box' based thinking which computers do well.

A computer could calculate the CO2 emissions from a single flight if it knew the fuel consumption. But allocating those emissions to individual passengers, who may have bought different sized seats is hard. Allocating it to cargo, which filled space which would otherwise have been empty, is also hard.

When we want to answer a question like "how does the CO2 emission compare if I buy flowers from Kenya and flowers from Netherlands", where the Dutch flowers grew in a heated greenhouse and were carried to the UK by road and ship, to Kenyan flowers brought by plane in cargo space which would otherwise have been empty, you can see how hard it gets.

People can and do make maps to resolve these questions, making decisions about how different elements are allocated, agreeing standards with others about how to do it. But the next step is working out how to manage the various pieces of data according to this map.

In the past, energy questions were resolved based on quite simple maps. An oil or coal mining company would seek to produce the maximum possible, provided the costs of doing so allowed a margin. Customers sought the cheapest products, the cheapest transportation costs, and then spent whatever money they had. So it led to enormous amounts of flying, big cars, big houses, masses of products. You could work out the necessary sums quite easily using box computation tools (like spreadsheets) and that's what people did.

Now we are moving to an era where the whole thing is managed much more as a system. Countries have a range of energy sources, each with their plus and minuses, for availability, CO2 and cost. Players within the sector, whether national grids, airlines, shipping companies, can develop their own system which does what they need at maximum effectiveness. So the whole thing becomes far more model based than box based.

Reducing carbon emissions is not a single challenge, and we would not wish to make a model about the whole thing, but make models about separate parts of the challenge.

For example, we can build better industrial reporting systems – how companies gather data about their CO2 emissions from different activities internally, and report the data, in their shareholder reports and regulatory filings.

We can build better tools to support carbon issues in investment decisions – where companies have to make a choice about new assets or equipment, and want to weigh up CO2 emissions their purchases will make as part of this.

We can build models to support maintenance decision making. These are important in managing carbon, because many companies have fixed maintenance 'assets', such as a crew of people able to fix gas leaks or inefficient machinery – and want to deploy these people on their biggest leaks, or leaks where it is easiest to make an impact.

We can build models for monitoring systems – where companies monitor the emissions their activities or equipment makes, and use this to learn ways to do what they need to do with less emission. Such as a warning system to a car driver which teaches the driver the emission benefits of driving smoothly.

We can build maps for data gathering systems – which track how much CO2 is emitted during an industrial process, such as producing oil and gas itself, or making steel, where there are many steps involved each with separate emission.

We can make a map for evaluating carbon emissions from family life in the whole. Transport, holidays, grocery shopping, purchases, home heating and insulation. We could develop a map which would give a family a score based on how they do everything, including answering very complex questions like whether to drive to a grocery shop or have groceries delivered by a supermarket's electric vehicle, doing multiple other deliveries at the same time.

We can make maps and models for house design – there's demand for more 'sustainable' houses, but an enormous maze of metrics. If people want to minimise energy consumption, should they just move into old houses rather than incur energy costs of building new ones? How do they play off a desire for space, or high ceilings, with heating costs, or can they have both large and low energy houses with insulation well managed? Data exists which answers all these questions based on real experience, but is not in any single model.

We can make models which support complex industrial decision making. Companies with complex industrial operations, such as oil and gas companies, have many choices to make all with different CO2 results. How they generate power, how they construct offshore platforms, how much steel is used, how much power is used, whether offshore platforms are manned or unmanned, if they are using ships or helicopters, and many other choices.

We can make models for shipping operations – do ships use less fuel for the same voyage if they have newer engines, or operate at lower speed, or have better maintenance programs, or go longer routes which avoid higher waves? What other investments can shipping companies make which will help reduce fuel costs, and where do they lead? How can customers of ships (cargo owners and shippers) better take into account the CO2 options, which may include choosing between available ships, also include rail or aviation, or not shipping at all?

How to fix IOT cyber security

In the physical (non-digital) world, security is largely about situation awareness. The most important part of security is knowing what is going on. What are your weak points, who might be looking to exploit them, have these adversaries done anything.

In the digital world, the topic of cybersecurity was developed around the PC. Having situation awareness about your PCs security is pretty hard, because it is so complex. The way computers are designed, there is no easy way for a person to know for sure if the applications it is running are legitimate, if the data communications it is sending are all for correct purposes, if the e-mail coming in has a malicious purpose. So we have this enormously complex world of digital driven security, scanning viruses, assessing e-mails, restricting software installations, making firewalls.

But in other areas of the digital world – including devices, web applications and specific software programs, the levels of complexity are not so high, at least in what the system is aiming to do. So it is not beyond people intelligence to understand the security issues.

For example, a smoke alarm may be programmed to send a message if it detects smoke, and send a “hello” message periodically, and the message destination should never be changed. It may have capability for remote access. If it starts sending different messages, or stops sending messages to the original place that is an indication it may be compromised.

So a map could be drawn of what the sensor is normally doing, how security on the device is managed, and what to expect from this sensor and another 100,000 sensors which may be managed by the same person. Computer software could follow the map and be able to identify which sensors might be doing something unexpected, and pass this onto a person to make a decision of what to do.

How to regulate cybersecurity

Cybersecurity regulation is a big headache for governments today, with challenges from Chinese 5G equipment, electronic devices made in factories outside Western jurisdictions, uncertainty over how to regulate online services and so on.

An insistence on appropriate maps as a condition of doing business could be a resolution. It could be done in a similar way to how companies operating higher risk activities are required to write and share their risk assessments, and be willing for them to be scrutinised, with the level of assessment and scrutiny proportional to the risk.

Publishing a map does not guarantee security, because the actual product may not align with the map, and even if it did, may not have adequate security. But it certainly makes it much easier for regulators.

An online game for children could be considered high risk, if it provides an opening for adults with malicious intent to communicate with children, or even ask children to communicate via a different website with lower levels of control.

Operators of such games are capable of mitigating such risks with a map based approach, and regulators are capable of assessing these maps, and also asking internet companies to keep online games they are not comfortable with out of the reach of children via child safe internet controls.

How to make learning more equal

Equality in learning is often considered the pathway to equal society, if equal society means that people gain jobs based on merit alone and education is the pathway to getting there. And computer learning is a great pathway to equal learning. There will always be some teachers which are better than others, but computers bring potential to make similar level of learning available to everyone.

We should really say ‘teaching’. Anyone can learn anything they want if they have the direction and motivation, because it is all there on the internet. But most people struggle with the motivation and direction part of it – which is where a teacher comes in.

If digital technology is taking part of the teaching role, then it is weaker than a person, in that a person can use the full range of people driven motivating skills. A computer needs to compensate for the lack of this.

One way this can be done is by setting students challenges which take them somewhere along their ‘learning journey’ which are just ahead of their current competency, so set them something which has just the right amount of challenge – not too easy and not too hard for them.

This takes us into the world of map making. The computer needs to have a map of the steps someone might follow when learning something, to continuously create this challenge. This map could be similar to the student’s “learning journey” which they might be set on, under non digital teaching.

The student’s competency and development along the map can be assessed with various digital methods, and then the software can set the next appropriate challenge.

How to implement tech better in organisations

Many organisations of all types could operate better if they made better use of digital technology products which have already been created. But, as the statistics say, so many organisational digital technology projects 'fail'.

Perhaps this is because the models people have of how the organisation works, what people do, what software could do, and how it would help them to do what they do, are very weak.

Consider the analogy – if organisations implemented software the way a large supermarket chain might introduce a new range of food products. Do supermarkets just choose what to buy based on someone's supposed expertise and force customers to buy it? No. Supermarkets and suppliers have enormous expertise about what their customers like, what makes them happy, how they use things, and what might persuade them to take on something else. And when they introduce new products they would do it on a small scale, and monitor how well these products fit into their customers' lives, reflected in how much they get repeatedly bought.

People talk a lot about 'change management' – but the change management can be easier if people quickly see how well the tools help them do the job they have already been asked to do – rather than if people are given an entirely new goal.

How to have better healthcare

Healthcare service provision is about the right allocation of money, people and assets to tasks and needs. Organisations have budgets which they spend on people, equipment and systems, and they have a pipeline of people with needs, which they match together.

No healthcare organisation has unlimited resources. Nearly all organisations seem to have resources they would consider just below adequate. So there is continuous pressure everywhere to better align everything.

Human experts can build models, which might show, for example, about different times of year with more demand for different kinds of experts, such as more colds at a certain time of year. They might build models about how a healthcare system can rapidly adjust how it allocates people and tasks at a time of high demand in one area.

Another example of better ways to apply digital technology in healthcare is in body monitoring.

The human body is enormously complex, and our greatest success at working with it involves allowing people to spend their lives building up models of how the body works, starting with learning from all the people who have gone before them, and sometimes specialising in specific parts of the body. We call these people doctors. Their knowledge may draw on books, but it can never be conveyed in a book, which is the reason you can't diagnose yourself from reading a book.

Being people, they can work with enormously complex, rich sets of data, which no computer could ever do. A computer could work with some of this data types. It could also store more data than a person could.

Body monitoring tech is a field with a great deal of hype, but consider that a service like Fitbit is based around a device which can measure body acceleration in different directions. This is then used to estimate how much exercise someone has had. It is sophisticated technology, but the amount of exercise is just one of many factors an expert doctor takes into account when assessing someone.

Digital health monitoring might improve if technology developers saw their role as to provide useful insights to individuals (such as 'you exercised much less this week than last week and your running pace has slowed'), rather than making systems which promise to provide health in themselves.

How to recruit more black police officers

A big news story at the time of writing is protests over US police brutality to black people. Getting more black people working as police officers would be a good step to solving this problem. For example, London is 13.3 per cent black – and its police officers are 3.5 per cent black.

The pipeline of recruitment has multiple stages. Starting with people having a desire for a certain career path, developing the necessary skills or qualifications, being accepted as suitable by a recruiter, and doing well and being happy in their jobs.

What might work is to look at all the sections of the recruitment pipeline and gather as much information as we can about how black people are engaged, and that may point to ways to improve the situation. For example, we can try to understand where the drop from 13.3 to 3.5 per cent happens in London's recruitment processes and apply attention there. Are black people just less interested in being police officers, or do they have a higher dropout rate after being recruited?

How to fix more potholes in the roads, with less temporary traffic lights

Road maintenance, like healthcare service provision, is about allocating resources (people, equipment) to tasks (such as potholes to fix).

Fixing more potholes means using these resources more effectively.

It is fiddly work nobody wants to do. But can we use digital technology to better track where the biggest potholes are, when would be the most effective or quickest time to fix them? Then we can put together a plan of how we can fix the maximum potholes in the minimum time, going quickly from one to another one.

Can we make a map of the roadworks we might like to do on one street or group of streets during a certain period of time, also with roadworks to change water pipes or install cables? Can we ask all the companies involved with different types of roadworks to share their plans, so we can make an integrated maintenance plan for one street for a year?

Can we spot how different teams are better than others at certain pieces of work?

Can we map the use of temporary traffic lights and how well modelled they are to changing traffic flows during the day?

This sort of planning is not beyond capability, but it involves a mixture of computer intelligence and people intelligence, which is not often seen, particularly in planning road maintenance.

Matching supply and demand within organisations

Many organisations do a task which in itself is not complex, but they do it at a big scale. This gives them opportunities to reduce costs, but that only works if they can manage the increased complexity of bigger operations. The complexity means bigger challenges matching supply to demand, and assets to tasks.

They have to align their purchasing with what they need, and align their resources (such as people, vehicles, big equipment, beds) with where they can have the greatest impact. They have to make a schedule which can deliver what is required while having enough flexibility to be able to minimise the impact of unplanned events and minimising costs. They have to make their free cash flow work as hard as possible, while being able to make scheduled payments.

Many software tools have been developed to try to achieve this, but based on a box based data structure, like on a spreadsheet – such as a tool to connect a list of required purchases with appropriate suppliers.

The real world does not usually fit into box based data formats. Consider all the various interruptions which could occur when trying to get your children to school on time, such as a refusal to wear any of the socks which are clean this morning, or an important phone call at the wrong moment. Organisational life is the same just on a bigger scale – so an analogy is getting 30 children to school on time with 40 potential interruptions.

We can fit this into a model, such as the models we have in our minds – plenty of people manage complex challenges in their jobs like this all the time.

Similarly, organisational software might be able to handle the changes in real life much more easily if organisational software was designed around model / map based data structures, rather than box based data structures.

How to support developing countries to develop faster

All the above approaches come together to make a way for developing countries to develop faster. Providing for energy needs with less CO_2 emissions, giving people security, equal access to learning, good healthcare, trusted police forces, less road potholes, healthy organisations. We have not got into fintech but we continue the same theme for means to identify corruption.

This needs more than ideas to work

All of this outlines an approach which is different in many ways to the one currently being taken. But it takes more than explaining an approach to make anything change.

What it really needs is many people who have the intelligence, technical capability, conceptual thinking capability and people skills, who get excited about what is possible, giving them the motivation to get there, and make this approach deliver business value, which will encourage the people controlling the budgets to pay them to do more of it.

CALLING FOR THE WEIRD PEOPLE – TO MAKE MAPS

Our proposal is that a core skill that these intelligent, technical, and conceptual people could develop could be described as the ability to make maps. Showing how the worlds of people intelligence and machines come together.

And to have the ability to make good maps, means having the motivation to really figure out how things work, and make maps of it which others can understand, you might need to be a bit weird or obsessive.

There are many people with competence to understand situations, perform well but not necessarily to make maps of them which other people can follow. This may include people currently working as artists, doctors, parents and politicians.

There are also people with technical skills but not map making skills, who we might call 'geeks'.

Weird people may face some resentment from the geeks, who have believed until now that to work with technology is to work exclusively with technology, not this odd maps for the people-digital interface which we talk about here. But you can handle that OK, if you are a bit weird.

Understanding how systems work

By making maps, we mean explaining how something works, so that it could be understood be someone else, or maybe even a machine. A map doesn't need to be drawn on paper (although it might be) – it could also be written language, like in this book.

We tried to give some rudimentary examples of maps in the previous chapter, giving an outline of what a map might look like for example to decarbonise different areas of society with the help of digital technology, described in written language.

In the cyber security example, we showed that cybersecurity is a mix of very separate domains, but we can take each one and map out the methods we are using to keep things secure, how we maintain these defences and spot a breach of them, and what to do when that happens, in the same way as you could do for security on your house.

You are not usually going to be making maps from scratch, because the map will already exist. But has probably never been written down, just held in the minds of the people who work in that domain.

Consider that, in the real world, there is already a fairly robust, rigid structure for how to do just about anything. (Cybersecurity being an exception, being such a new domain).

You could look at banking, defence, supply chains, corporate management, policing, oil and gas, education, project management, asset management, any kind of security, running all kinds of small businesses.

In any sector you can think of, government or private, there are specific roles that people do, processes which are followed, controls, means of making the right decisions, attracting revenue and avoiding problems. People have figured out how to make organisations which can function effectively enough. People have worked out what kind of situation awareness they need to make a judgement about what to do about a problem, how to communicate what is happening with other domain experts, how to find a path forward, make decisions, make things work and reduce risk.

This all happened before software was invented. But now we have software, in many cases software should follow the same old map, not try to make everything new.

Now we have software, expectations about how organisations should perform have risen, as a result of companies like Amazon who do supremely well at non exotic technology, so they can make digital technology do magic. But you can still see that Amazon does tasks which were done before e-commerce was invented (like mail order catalogues and deliveries) – but it does it supremely well.

What a map looks like

You may be wondering what a map looks like.

Well, this book itself is an example of a map. We began with an outline of an idea, and steadily added detail. Finally, we'll aim to end up with some plans detailed enough to be used as a plan for action, or as we could call it, an 'executable map'. You have not got to that part of the book yet.

Language itself is a tool for sharing maps, in the sense that we evolved language in order to explain to people how we do things, or how things work. Language is not the best sort of map for these examples, because they are hard for computers to understand, but that's a good place to begin.

A map is a conceptual idea - primarily something we hold in our heads – as in "what is our map for how we are going to do this?"

You already carry multiple maps in your head, about how to do anything, how to navigate any relationship, how something works. You could explain these maps to someone else using language, and you could write them down on paper in some form, showing how different elements connect.

A feature of maps is that if they are orientated around a goal (what you are trying to achieve), they can work at a wide range of resolution levels. An example is how you could explain how something works in enormous detail or just as an outline, depending on who you were explaining it to and what they needed to know, or were capable of absorbing. A summary takes much less effort to absorb than a full detailed map.

People are always quick to call for summaries rather than detail, but sometimes you do need the detail. An architect could plan out a building with some lines scrawled on paper, but if this plan is being given to a regulator to approve it, or a building company to build it, an enormous amount of detail would be required. The same applies if a computer is going to follow the map in any way, or software is going to be written to follow it. Or if we are writing a report with recommendations we would like people to follow.

Computer code is also a map, in that it explains to a computer how it should operate as a result of different inputs. But of course it is highly detailed, because computers cannot figure out for themselves what to do.

Ultimately we can foresee a world where the map has enough detail that it gives the computer all the information that it needs to run. This is the philosophy behind "low code" – that you can create a model (which is the same as a map) for a computer, and the code can be generated automatically. The human input is only in creating the model, not coding. We call this an 'executable map'.

Many types of maps

We can make a map for anything we can explain, or might want a computer to do, so that's a broad spectrum of things. But here's some examples of different maps which may help explain the concept.

We could make maps of how our organisations operate and run. How organisations allocate funds, people and equipment to purchases, tasks and customers.

We could make maps of how people learn, or could learn something specific – what are the stages of how someone could develop their understanding or competency of something – like a stage of swimming badges.

We could make maps of how a personal relationship works – how we respond when certain things happen, so as to keep the relationship on track rather than descending into anger and upset.

We can map different sources of CO2 emissions, and how we gather them together to assess total emission for something.

We can map how people use digital technology to do their work – what situation awareness they need and how the computer provides that.

We can map how a complex digitally enabled system works, such as a water plant with lots of digital sensors. How people should understand it and spot where something is going wrong.

We can map how cybersecurity in a system is managed, such that an external regulator can see and approve it.

We can map how a computer makes its own understanding of what is going on, based on some kind of algorithm or statistical analysis. We can have maps of artificial intelligence which someone else can understand.

Conceptual thinking

Conceptual thinking is a label we can give to the process which people do, which computers cannot yet do, as we develop an understanding of a situation and pathways how to achieve our goals. It may involve understanding what elements impact other elements, understanding patterns, and how to make something work.

Not every person is good at conceptual thinking, but conceptual thinking is needed just about any time a person does something good. Conceptual thinkers can play football, design buildings, manage organisations, teach a class, create art.

Conceptual thinking is not just about the power of the brain (although that helps a lot), it is also about the training and experience. This is what enables a doctor to put together different pieces of knowledge and experience to diagnose a patient with something not seen before.

Conceptual thinking could be another name for map making. Although not all conceptual maps will be shared with others, or need to be, and not all shared conceptual maps need to be understood by machines.

If we are going to use machines in a domain which has conceptual or expert thinkers – and that is just about any organisation – then the machines need to have a map which aligns with these conceptual thinkers.

A mistake surely made too often is that people have tried to introduce software to an organisation based around a completely different map to the one its experts use, and then found that the conceptual thinkers did not like the software or find it helpful, and did not wish to rethink their understanding of how the organisation works based on a different map used by the software.

An example is a book publishing company which introduces software to manage book printing and supplies to bookshops, but the software was originally developed for use in a manufacturing operation, and proves entirely impractical for book publishing, where the ways of working, and conceptual models helped by the experts are different, and a poor match with the software.

Instead, organisations would benefit from weird people who can sit down with the conceptual thinkers, understand their conceptual maps, and write down a version of them which can also be used as a basis for making the software.

Going deeper than user interface design

If we are having discussions about a people-machine interface, some people may assume that we mean the user interface. This is not what we mean. The user interface is just how information from the software is displayed on screen.

People use the term 'user experience', which can technically mean all aspects of how the user feels when using the software, but in practise this term is often used for how someone feels when using the screen.

We are going deeper here – by the people-machine interface we mean the underlying map which people use to understand the domain, and the map which the machine uses, and how well they match.

Consider an airline pilot who has a mental model of how the aeroplane flies, what impacts the flight, and what may go wrong. The pilot is interacting with a screen while flying, but is the screen giving information which fits with the pilot's mental model? When we see a problem such as a faulty sensor or misleading information on the screen, as with a number of recent cases, can the pilot quickly identify the problem, because the information being provided by the screen does not fit with the pilots mental model of what is going on? Or is the pilot expected to only interact with information on the screen, so is unable to fly the plane if the information on screen is incorrect?

An analogy for superficial integration with people and machines could be a company which looks initially good but the problems lie just below the surface. It has a good looking CEO and website but is a really difficult place to work in, and not very helpful to customers who have problems.

People learn slowly

People who have experience and skill in a domain have rich mental maps about how it works – but these did not come about overnight. We start learning something in baby steps.

This applies particularly for someone who already has complex mental models – because it means learning something new takes even more effort, because this person is fitting what they are hearing into the mental models they have already.

If we are making maps to help people understand something new, then we might draw them as simply as possible, as we tried to do in the short first chapter of this book.

Given time, our minds can develop internal models of enormous complexity - think about the knowledge which a doctor has, or a business leader, about their domains, or the ability for a movie producer or writer to hold many details of their work in their mind at once.

But there is no point in making models which are beyond human capability to understand them, and you may not have the sort of people available to you who would become doctors or senior business leaders.

If it is not possible to explain what you are doing in a way that a person can understand, then perhaps can only manage it with computer models. For example, nobody fully understands what a PC is doing to the level of knowing what the processor will do next or what data it will send next.

Box and model based knowledge

The way machines and people manage knowledge is fundamentally different. It could be thought of as box based and model based knowledge.

For box based knowledge, think of a database or spreadsheet, which can have unlimited rows, but each one is given the same fields or attributes. For example, a list of vehicles with their owners and ages, a list of sensor readings over time, a list of sales. Computers can do amazing stuff with a large volumes of structured information like this.

As an example of model based knowledge, think of something you carry in your mind which has multiple facets which you could never fit into a box. For example, different aspects of your job, how some equipment works. You could explain this to someone else, perhaps using language, or maybe lines on paper connecting to other lines, but not with numbers in boxes. People can do amazing stuff with large complex models like this.

This illustrates what kind of situations are better for computers and what are better for people.

There is crossover. There is box based information which people may carry in their minds and not in a computer, for example a list of 5 things you plan to buy today from a supermarket.

There are also ways that computers can do model based knowledge – but usually by converting it into boxes. Consider how a chess computer can do a model-type task – choose the right next move – but does it by making millions of options for moves and evaluating them, so creating a box of data.

Another example is a search engine indexing system, which is able to reduce the mess of the internet into models and lists of pages and what sort of searches they may be relevant to, so it is able to provide the pages it thinks are most relevant to your search.

Computers are making big advances in their ability to think in models rather than lists – and developments in machine learning, statistics modelling and graph databases are part of this. But in terms of their ability to understand a complex situation (which they have not already been programmed to reduce to boxes) and determine what to do, they are usually nowhere near as good as people are.

And programming a computer to reduce a complex situation to data boxes can be pretty hard, particularly in domains where there is no score, and many different events may happen, which is why we do not see much AI in use today.

The idea of computers figuring things out for themselves, as people can do, remains a distant dream (now labelled “Artificial General Intelligence’.

Data and information

Software programmers often see the technology world in terms of simple layers of data, contextualised / integrated / analysed data, logic and user interfaces. If it is all built well, the data in the database is processed so it can serve up useful ‘information’ to the ‘user’.

The problem with this approach is that it assumes that the information people need can be generated from the data which the system stores. For example, in the COVID-19 crisis, policy makers needed to know how fast the virus was spreading. But only data to create this information was the positive results of tests, which did not provide a full enough picture, because it depended on everybody potentially having the virus getting regular testing.

Software design can go further in this error, and act on the assumption that the information machines are capable of serving is the information which people need, which may even force people to adjust their way of working to be able to use it. So we see managers making decisions based on "key performance indicators" which a machine is capable of serving them, although they may not convey a picture of what is actually going on.

The solution to this problem, once again, is to build systems which are more capable of working with people intelligence, and which are designed to support the ways which people have already figured out.

In any domain which functions effectively, people have worked out what information they need to run that domain. For example police departments everywhere have learned how to figure out if they are doing an effective job without relying on any KPIs.

There is nothing wrong with all the extensive infrastructure which companies are building to gather, contextualise, integrate, analyse and present data, they just need to recognise that it may not be what people need, or may only give them a part picture.

For example, if you have a purchasing system, you can build a super sophisticated system which can take a list of required items, send it to multiple suppliers and let them bid to provide it. But the picture you may miss is that other goods are available which can do what is needed but have a different part number so do not show up in the computer system – and also that many suppliers do not see it is worth spending time explaining this to you, because they do not expect to win the order anyway.

If you need to predict cashflow, you can build a system which will make predictions based on your invoices sent out and payment terms, but which does not accommodate that suppliers may pay you any time from now to 120 days after they were invoiced, and there is no way to find out when they will pay before the funds arrive.

In our mental understanding of the world, we do not use 'data', we think more in terms of what we see is changing, what our priorities are, what problems may be coming up. It is quite a challenge for digital data to meet that need – and perhaps makes too high an expectation on digital data that it always can.

Breaking models down

What do you do when your model is too complex to hold in your mind, or even draw?

There are two ways to reduce it. One is to reduce the granularity, and another is to chop it up.

Reducing the granularity means taking the detail out of it, but keeping the model around the same goal, you do not show so much information about the routes of getting there. Think of it like a road map which just shows the motorways. Knowledge about the smaller roads can be made available as required, but doesn't complicate the map you are looking at.

Chopping up the map would be like having separate maps of all the different countries you are going to drive through.

When you need to reduce a map to make it manageable, we suggest you always keep the goal in mind. Without a goal, the purpose of map can be lost – and conversely, with a goal, a map can still work whether it has a tiny amount of detail or an enormous amount.

When an architect plans a design, you start with a rough concept for it, and then gradually build in resolution, and also break it up into pieces of information so it is manageable, as the design gets more complex. The pieces could represent physical things (i.e. different parts of a building), or different parts of the buildings' experience (inside and outside), or elements of the structure (steel, cladding), or work different kinds of contractors.

So you could chop up an architectural design to give different sub models for different people which are all geared to the same goal (creating the building). You could do models for different parts of the building, because it is easy to see how the different parts of the building come together to make a goal. But if you chop up the map in a way which loses relevance to the goal, you should be wary.

If we keep chopping up models without caring about how we do it, we might ultimately get to something which keeps us away from the overall goal completely – such as a sub-model built around the data units we are working with. Then we are now in a project to work with data, not to achieve our goal.

– for example if you have a map of purchases, it may push someone to think that you now have a new goal, to reduce the purchasing cost.

This book may damage your career

That is because they are invested in the concept that technology needs to be exotic to get interest - and they have got very good at it.

Big tech companies make money from making people believe that they need these advanced technologies, their companies will be in a mess if they don't have them, and conveniently big tech companies are in the best position to provide them.

Big consultancies make money from telling company boards how important these technologies are, and how they need big consultancy companies to explain how to use them.

Small start-ups figure out that the best way to win investment - and sales - is to talk about how they have artificial intelligence embedded in what they do.

There are tech people everywhere who are more excited by technology than the real world issues behind the technology, and have learned how to see the world in this way. When they see a problem related to technology, such as challenges planning out a software implementation project, or cyber security challenges, they automatically think that the solution must be technical too.

Software companies want to convince clients their products are unique. Obfuscation sometimes suits their interests.

Investors are often only looking for start-ups which have some promising 'artificial intelligence' - not because they necessarily believe in it (although they might), but because they want to sell their shares to smaller investors one day, and think an artificial intelligence story makes that easiest to do.

Companies which use technologies will often show interest in buying something which promises artificial intelligence over something which will actually help them run their company better. Because the sales system extends within their companies as well, as you have IT professionals who earn their promotions from getting people excited about new technologies.

Politicians and government agencies will also often advocate exotic technology over technologies which work, in the projects they choose to fund. It makes them look good, and makes the political leaders look good.

The cogs of the system are turned by the short term profiteers, and there is no short term profiteering available here.

The ideas in this book will never create new multi million pound software products, make venture capitalists rich, or fundamentally change organisations, they will just drive small improvements.

So if you value your career, and see that it should be somewhere related to technology, working in the companies described above, the ideas in this book could be very dangerous indeed.

WHY WE NEED YOU TO LOVE NON EXOTIC TECHNOLOGY

Non exotic technology needs you to love it because it is very difficult, requiring a lot of focussed, difficult work. To be motivated to do it, you need to love where it leads to.

Society also needs you to do this, because society has many difficult problems which non exotic technology can help solve.

So you need to decide whether you care more about your career, your income, politicians, big tech, investors and maybe impressing your friends, or doing something useful for the world around you, even if the real world does not tell you or give you any awards.

There may be no awards for work you do to help organisations to be healthier, doing more with less. Or work to help online learning function more effectively, reduce CO2 emissions, provide healthcare where needed, or help developing countries improve. The awards may go to the people who made the software, or who run the organisations which achieved the results.

But life is like this, relying on people who do hard, focussed work, for no immediate reward. You may as well get used to it.

Technology will play a big part of our lives for the rest of your lives – but we can all show how to do it in a better way.

How we move forward with this

It is not time to talk about our executable plans – how these ideas could be taken forward to something which actually happens.

Perhaps a starting point is if we encourage the organisations we work with to recognise that the current approach to digital technology might be described as very amateur – technology plays a central role in our companies, yet we make big experiments as though it is fine to not have a clear idea of what we are trying to do. Documenting how things should happen using maps could also be described as professionalising.

We can promote a picture of the future where the digital world integrates with the people world, rather than looking endlessly for perfect digital technology which does everything by itself.

The domains this approach can add the most value is where activities could run more effectively if the people in them could understand the digital systems better.

We are looking for areas in society where there is too much data, and it is too complex, for people to work with properly. People see a lot of data, and think the hassle of getting value out of it is too high. But these domains do have experts who do not necessarily work with the data.

Some domains we have looked at in detail include maritime operations, oil and gas exploration and production, cyber security of the internet of things. Other domains we have not looked at in so much detail, but where these ideas could have strong relevance, include electricity grid management, financial / banking, emergency response, marketing, and anything involving scheduling, including healthcare, and broad areas of education.

We have looked at maritime operations, oil and gas operations, cyber security of internet of things devices. Other domains we've looked at less, but where it could work, include financial / banking, emergency response, marketing, and anything involving scheduling, such as healthcare.

Four paths

We see four ways to make more action into these ideas, which all overlap.

The first way is to incorporate these ideas into work which is happening already. If you are a digital project manager, cybersecurity planner, software architect, or have any role related to digital transformation, we hope these ideas about maps will be useful in developing better ways to bridge the digital and people world, which you would be doing already. If you would like to work in these roles, but do not already, we hope these ideas might help you build your skills and present to a potential employer. These are roles in - we could say - the sharp end of tech, trying to get tools in use in organisations, where there is not much room for getting excited about novelty.

Perhaps the way to have the biggest impact, using these ideas, could be to call for more map based thinking as part of all our projects, where we need people intelligence and computer intelligence to come together.

No project should be too complex for people to understand it (and if it is too complex, the complexity needs to be reduced). A map can enable people to use their intelligence to have a shared view.

If you cannot understand a situation, or it is too hard for a person to understand, we should not accept a response that we could understand it if we were more technical, or we need a computer to understand it. We need to make everything in organisational and technical life easier for people to understand.

The second way to build these ideas is to form groups of people working in these roles, perhaps in specific domains, to discuss in more detail what the people-digital interfaces and maps look like. Can we build models of, for example, the biggest priorities for cybersecurity managers working in the shipping industry, or how people typically work with subsurface data in the oil and gas industry?

The third way would be to establish a new businesses developing models, which would be open source (not proprietary), but where software products could be built on top of them and sold. For example, you have a map of how a complex sensor system works and how sensors manage their own cybersecurity, you could build a software product running on top of it, using the map to determine whether something is happening which does not fit with the map, or something unusual is happening, and warn somebody appropriately.

A fourth pathway is to build the skills of making maps into a discipline which people can learn and develop, and which companies could recruit for.

Perhaps it could help attract people who are more comfortable in conceptual domains, such as art students, to work in more technical domains, such as cybersecurity.

These technical domains offer unlimited scope to add value to organisations – and so, in theory, unlimited employment potential for individuals who develop the skills and apply them to add value. The same goes to organisations which develop the capability to do it.

Making maps is not a known discipline as of 2020. Perhaps the closest is journalism. A journalist's role is to go into a new situation, figure out how it is working, and write it down so that someone else could understand it, and be attracted to understanding it. Teachers and leaders of all kinds also have similar skills. But it also draws on the conceptual analytical skill base found in many arts subject degree courses, from history to psychology. Technical skills are important too of course but most graduates today probably have these as a base level.

If you are interested in the ideas in this book and exploring ways to build them in your organisation or personally, perhaps along with one of these pathways, the authors would be very pleased to hear from you. Please contact Karl Jeffery on jeffery@d-e-j.com