Spending time where it counts

An Artificial Intelligence Strategy for Health & Social Care in the North of Scotland 2024-2027

ANGUS Health & Social Care Partnership

Perth and Kinross Health and Social Care Partnership

Shetland Islands Council

OHAC orkney health and care

The Highland Council Comhairle na Gàidhealtachd

WI IJB & H&SCP

Shetland Health and Social Care Partnership

NHS Highland

NHS Grampian

NHS Tayside

NHS Shetland

NHS Eileanan Siar Western Isles

NHS Orkney

Aberdeenshire Health & Social Care Partnership

Aberdeen City Health & Social Care Partnership

A caring partnership

Health & Social Care Moray

Dundee Health & Social Care Partnership

Table of Contents

Welcome

The people working in health and social care across the north of Scotland are our greatest assets. They are also our greatest ambassadors.

This is true both for those working at the frontline and for those who support them. However, demands on staff continue to increase and vacancy rates have rarely been higher.

There are many things that only humans can do, such as bringing empathy and kindness to every encounter. However, we know that staff are often asked to perform tasks that are repetitive and frustrating; taking up time that could be better spent.

A professional's skills and knowledge are so precious, but could we be doing more to support some of the more laborious tasks with software or equipment?

There are already many examples of the benefits that AI has brought to different parts of health and social care; supporting patients and service users, and enabling staff to work to the best of their abilities. However, new opportunities mean that we can now go further – automating more processes and saving staff more time to deliver value-based healthcare. Some of these opportunities will be easier to grasp than others.

Recurring themes in this strategy are the needs for both skill development and closer collaboration. Getting the most out of AI technologies will require us to work together in new, imaginitive ways. Here in the north of Scotland we nurture positive collaboration; we are used to working in partnership between organisations of different sizes for the good of our patients and service users.

We also have pockets of specialist expertise in some fields of AI.

There are many questions still to be answered in this rapidly developing area. Much as with medicines, and other things that can improve health, AI contains numerous potential benefits but also some risks. We need to carefully navigate this landscape and systematically look for the opportunities that AI could bring.

We hope that Health Boards, Local Authorities and Health & Social Care Partnerships will embed assessment of AI technologies in their planning processes – both short and long-term. We also hope that these organisations, and regional planning teams, will include an assessment of AI collaboration opportunities; for example in terms of improving the efficiency of community diagnostics and other geographically distributed services.

This is not a strategy to replace humans with machines or to use algorithms to deny people treatment. This is a strategy intended to accelerate and support the adoption of inclusive and ethical technology as tools to enable staff to spend time where it brings the greatest benefits.

Mr Grant Archibald
Chief Executive
NHS Tayside

Mr Gordon Jamieson
Chief Executive
NHS Western Isles

Mr Brian Chittick
Chief Executive
NHS Shetland

Ms Laura Skaife-Knight
Chief Executive
NHS Orkney

Ms Sandra Macleod
Chief Officer
Aberdeen City H&SCP

Ms Pam Milliken
Chief Officer
Aberdeenshire H&SCP

Ms Jacquie Pepper
Chief Officer
Perth & Kinross H&SCP

Ms Vicky Irons
Chief Officer
Dundee H&SCP

Mrs Pamela Dudek
Chief Executive
NHS Highland

Mr Simon Bokor-Ingram
Chief Officer
Moray H&SCP

Prof Caroline Hiscox
Chief Executive
NHS Grampian

Ms Pamela Cremin
Chief Officer
Highland H&SCP

Ms Josepine Robinson
Interim Director
Shetland Community H&SC

Mr Stephen Brown
Chief Officer
Orkney H&SCP

Mr Nick Fayers
Chief Officer
WI IJB & H&SCP

Mrs Gail Smith
Chief Officer
Angus H&SCP

Executive Summary

How should the humans working in health and social care spend their time? What do they really need to be doing and which tasks could they hand over to technology - technology that's available now or in the near future?

Those are the questions behind Spending Time Where it Counts.

As you read this strategy you will note recurring themes.The need for collaboration and staff development. The need to act now and not prevaricate. The ethical aspects of increased adoption of AI vs the ethical issues of not adopting AI technologies that are available to us.The need to protect data and the need to use that data to improve healthcare.

Leadership is another theme. The chief executives and chief officers have signed the Welcome but we are all leaders. We all make decisions and we can all influence things more than we think we can. Wherever we are in an organisation, at every opportunity to change, at every level of the planning process we need to ask "Could AI help?"

Having asked that question, we need to be able to get an answer. We will need some knowledge ourselves and we will need someone to turn to when we get out of our depth.

And when we have the answer, that an AI technology could help, we need to rapidly adopt and scale and support that across the region.

Recommendations

Robotics

1. Commission the national Robotarium to assess the potential to use Autonomous Delivery Vehicles in NoS hospitals and Care-Assistance robots in care homes and similar settings.

2. Work with Directors of Pharmacy to assess where further adoption of pharmacy robots in the community and hospitals could bring advantages.

3. Support the Innovation-led exploration of delivery Unmanned Aerial Vehicles – supporting more proofs of concept in the north of Scotland.

4. Systematically evaluate the potential for the adoption of robotics in other areas including, but not limited to, instrument sterilisation and floor cleaning.

5. When it is deemed that robotics could bring advantages, contracts should be developed on a regional (or national) basis to gain an economy of scale and to facilitate support.

Process Automation

1. Ensure safe and comprehensive adoption of Microsoft 365 Power Automate by:

    - Working with NES, NSS and others to ensure widespread knowledge and adoption of cloud Power Automate.
    - Develop regional skills and capacity in M365-based automation, including the use of Power Virtual Agents and premium versions of Power Automate.

2. Establish regional expertise in the development of more advanced rules-based automation – beyond what Power Automate is capable of. This should include collaboration on contracts and licences to gain best value.

3. Establish a regional advisory function to help health boards and HSCPs identify opportunities for automation and then implement them at scale.

4. Departments should systematically review their processes and consider whether automation could bring value.

5. Rapidly establish proofs of concept in areas that bring value to the North of Scotland.

6. Share case studies of where process automation has brought benefits – supporting further adoption.

Machine Learning

1. Increase collaboration with the Data Lab, Universities and other relevant partners such as the SG & COSLA Data Delivery Board.

    - Promoting work done in the NoS so it can be adopted more broadly.
    - Being Fast Followers of Machine Learning techniques developed elsewhere; accelerating benefits for the citizens and staff of northern Scotland.

2. Do an options appraisal to determine the optimal Machine Learning (ML) infrastructure for the region, balancing performance and cost.

3. Create an ecosystem that supports adoption of ML from proof of concept to procurement; led by the Innovation Hubs and with industry engagement where appropriate.

4. Accelerate implementation into standard care when a ML concept has been validated. This should include increasing the publicity around successful use of ML in health and social care.

5. Develop an approach to the monitoring of algorithms, to support the regulatory requirements, that is efficient and straightforward.

6. Conduct proof of concept studies in the area of Generative AI, including Large Language Models, to see what benefits they could bring to health and social care.

Predictive Analytics

1. Develop the skills of the Health Intelligence workforce in the region to better exploit modern predictive analytics tools.

    - Enhance collaboration between teams in the region to gain benefits from this skillset.
    - Develop relationships between NHS Analysts (Territorial and National) and Academia to assist with the adoption of specialist modelling skills.

2. Increase awareness of and provide secure access to health and social care data that's currently available in the region to support the development . Consider the creation of a regional data loch.

3. Determine what the key data gaps are and consider how they could be routinely collected.

4. Progress proof of concept work on predictive analytics to support logistics and whole system flow. For example predicting DNAs and length of stay.

5. Improve the visibility of the output of predictive models; making outputs available at the point of care to all staff and patients/service users who need to see them.

Introduction

It is common to start strategy documents with pronouncements about the unprecedented times in which we live, and how only radical transformation will bring us to salvation.

This document is different.

The situation we face isn't new – the increasing demands on healthcare organisations and the declining availability of workforce are trends that are long-established.

And the solutions we propose aren't new either - using technology to ease the lives of humans is a practice millenia old. Consider the humble wheelbarrow, which allows one person to do the carrying work of three.

What this strategy aims to do is to join the dots between these trends and solutions, to describe how the need for a high-quality health and social care system run with fewer staff could be facilitated by greater adoption of certain types of technology.

If implemented and supported well, these technologies should improve health outcomes and equity of access. They should also increase financial and environmental sustainability, and improve the resilience of health and social care delivery organisations.

They will also enable patients and carers to take a bigger role in managing their own wellbeing and health conditions. A key requirement is to take a systematic approach to the inclusion of AI and related technologies in service planning; involving staff, service users and members of the public as we do so.

Although the prospect of increased adoption of Artificial Intelligence fills many people with excitement and optimism, there are also those with concerns. These concerns typically cover a range of important areas, including ethics, law and employee relations. We need to address both the unrealistic expectations of some enthusiasts as well as the often unwarranted concerns of the doubters.

With Artificial Intelligence and related technologies there are many choices to make and a wide range of potential futures are open to us. This strategy will lay out a vision of how things could be in the mid-late 2020s, structured around four interlinked domains. We consider what is realistically possible in this time-frame, as well as addressing some of the concerns raised.

An essential requirement is to translate strategy into actions. While we don't include a full delivery plan, we make some key recommendations which should facilitate the development of more details plans.

What is Artificial Intelligence?

Scotland's National Artificial Intelligence Strategy defines AI as:

Technologies used to allow computers to perform tasks that would otherwise require human intelligence.

The term can seem quite futuristic, even a bit Sci-Fi, and can sometimes make people imagine dystopian scenarios of machines taking over the world and controlling our lives. The reality of AI is often a lot more humdrum. Larry Tesler, a US Computer Scientist, quipped that "AI is whatever hasn't been done yet." When we adopt AI technologies we rapidly regard them as quite normal, unremarkable and something that we can't imagine being without.

Computers have been performing intelligent tasks previously done by humans for many years – in our cars, washing machines and shops. And healthcare has being using various forms of AI for decades. When an abnormal blood test result is presented in red text, when a door opens in front of the porter pushing a wheelchair, when a machine goes ping in response to dropping blood pressure; those are all different forms of AI at work.

As we adopt AI technologies more widely, we should expect this "absorption into normality" to continue.

Adoption of AI isn't an even process; some countries and organisations are moving into the future faster than others – experiencing the challenges and gaining the benefits as they go. But even for the late adopters, things that seem unusual now such as autonomous delivery vehicles, process automation and advanced decision support, will seem entirely normal, and many people will not regard them as AI at all. In 2030 we will still be thinking of AI as something coming down the line which is technically advanced, futuristic, and perhaps a bit concerning.

The great breadth of technologies which allow computers to perform tasks that would otherwise require human intelligence can make categorisation challenging. In this strategy we have subdivided AI into the interconnected disciplines of:

- Robotics
- Software process automation
- Machine Learning
- Predictive Analytics

These are defined in the following table:

A categorisation of AI

Robotics	A robot is a device that senses its environment, processes that data and performs tasks in the physical world. They can work under the direct supervision of a human or more autonomously; using their sensors to determine the best course of action. They can be fixed to one place or be mobile on land, sea or air.
Software process automation	This is where software follows a pre-determined set of actions depending on the inputs or triggers it receives; "If this, then that". The sequence of actions could be very simple, such as filing an email received from a particular account, or highly complex. Confusingly, this is sometimes referred to as RPA – Robotic Process Automation, but no robots are involved. The trigger could be coming from another software application or could be the typed or spoken words of a human. The set of logic and actions - the algorithm, could have been specified by humans, or software might have learnt what the algorithm should be.
Machine Learning	This is the process by which software learns what the rules should be in order to achieve a desired outcome. The software could be trained using a dataset containing known examples of what "normal" and "abnormal" look like, eg images, or audio, and it then calculates which factors or rules could be used to make that distinction. Alternatively, the software could be given a dataset and asked to identify factors which are associated with the probability of a specific outcome occurring, for example delayed discharge from hospital.
Predictive Analytics	This is the use of data to predict the future. Historical data and those obtained from other sources, such as the environment, can be used to construct a software model. The model is then compared against the unfolding reality and the model is refined as needed. As more data feeds are added, the model becomes more capable and better at predicting the probability of certain outcomes

What are some current uses of AI in health & social care

Robotics	Surgical robots that minimise unwanted trauma during procedures. Autonomous vehicles that distribute laundry around a hospital. Robots that place and pick medications in a pharmacy. Unmanned Aerial Vehicles (Drones) that deliver biological samples and medication. Cutting the grass around buildings.
Software process automation	Automatically emailing a case manager when a police concern report is received about a child in their care. Co-ordinating annual leave requests and authorisations. Summarising jargon-filled test results into plain English. Recognising speech and turning it into an action.
Machine Learning	Learning what makes a chest X-ray normal Learning how to calculate the volume of a brain lesion from analysing MRI scans.
Predictive Analytics	Predicting future bed occupancy based on a range of data feeds. Predicting the best treatment for an individual based on a range of data items. Predicting which patients are at risk of deterioration

Augmented Intelligence

The current reality in healthcare is that, rather than acting in isolation, Artificial Intelligence typically works alongside human intelligence. Examples:

- Software in a laboratory recognises an abnormal blood chemistry result and highlights it in red. This draws the attention of a human clinician to the out-of-range result. The human then decides whether, in this patient's case, anything special needs to be done.

- An image analysis algorithm identifies an area of concern on a brain scan – this allows the radiologist to spend more time reviewing that area of this scan rather than the areas which have been assessed as of no concern.

This collaboration between software and human is sometimes referred to as Augmented Intelligence and is a way of supporting clinical decisions. This can reduce cognitive overhead and therefore reduce errors and stress.

In robotics, a device which closely supports the work of a human is sometimes referred to as a cobot.

As technology develops, the balance between human and software will shift – with artificial intelligence taking on an increasing proportion of the work, while humans retain a supervisory, quality assurance role.

Many solutions use developments from multiple disciplines. The example of a pharmacy robot.

The software in a pharmacy robot was trained via supervised machine learning to recognise medication packages and to distinguish them from shelving and other objects.

The robot follows a set of software instructions to pick up a list of specific products from the shelves and put them in the hopper ready for dispensing.

Data is analysed to predict when stock levels will run low and therefore which medicines need to be re-ordered.
The human pharmacist uses their intelligence to do a final check before dispensing the medication to a patient.

Narrow vs General AI

The examples above are of software or devices which have been developed to perform a specific task.

The software in a delivery robot might be very good at stopping it from bumping into a wall but it won't be very good at recognising an abnormal chest X-ray.

A virtual agent which logs onto a website to run a query and download a report won't be able to predict the bed-state of a hospital next weekend.

This focus on a specific task is sometimes called Narrow AI.

The situation where software can perform multiple tasks, similar to the range of a human, is sometimes called Artificial General Intelligence. This doesn't yet exist, so for the forseeable future the world will contain multiple separate Artificial Intelligences, working independently of each other.

Why does the north of Scotland need an AI strategy? The case for action

Political
The 2021 National AI Strategy was both broad and also narrowly focussed. It covered all sectors of business and industry, including health, but only in the respect of Machine Learning and downstream applications. So, while it is a useful document it doesn't fully meet the needs of the region.

The national NHS AI Hub is focussed on Health & Social Care but, like the national strategy, only focusses on Machine Learning and its applications.

The NHS AI Hub assumes that health boards will have capacity to work with them to develop and implement solutions. This will be a challenge for smaller health boards, of which the north has the majority. There is a need to provide an equal service to people in all parts of Scotland – including remote island communities – ensuring that diagnostic services are available.

Economic
Financial pressures will continue. We need to improve the efficiency of health and social care, reducing the need to unnecessarily expand the workforce.

There is a requirement to diversify the economy away from oil and gas and to support the development of more technology enterprises in the region, including university spin-offs.

Demonstrating the region's approach to innovation and encouraging new investment including research-based start-ups.

Creating consortia could produce economies of scale – eg when using virtual agents which could work across multiple health boards.

Social
Perceptions and understanding of AI are varied. Some voices make unrealistic claims that AI can solve all our problems while others are unnecessarily concerned about machines taking over the world. Being clear about our strategy will help address both.

Reducing the need for humans to perform repetitive low-value tasks could making the NoS an attractive, modern, place to work, supporting staff retention.

Automating, or otherwise improving some clinical support pathways could improve the experience for patients and clients.

The combination of a skills shortage and rurality creates particular challenges for financial sustainability. The role of autonomous vehicles and supported self-management could be more important in our region than elsewhere.

Technical
Interoperability requirements. There are unanswered questions about which systems we are going to integrate – for example electronic patient records and business support systems. Being clearer about our strategy, in particular the role of process automation, will help with planning.

There are unresolved questions about to what extent the north of Scotland should manage our own machine learning infrastructure vs outsourcing this to elsewhere.

Legal
EU AI Regulation – could become a standard to follow

Medical Device Regulations – implications for software as well as physical devices. Being clear about our strategy could support the development of the necessary specialist roles.

Environmental
Avoiding unwarranted duplication of equipment by collaborating between health boards.

Avoiding unnecessary road travel, for example by semi-autonomous delivery by Unmanned Aerial Vehicles.

Avoiding waste – eg of medicines and other consumables by better use of predictive analytics and optimised logistics.

Spend on agency and bank nursing and midwifery staff in the north of Scotland

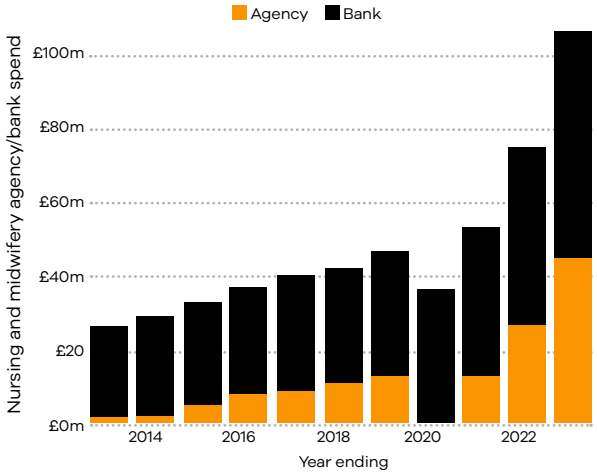

Nursing and Midwifery vacancies in the north of Scotland 2016-2023

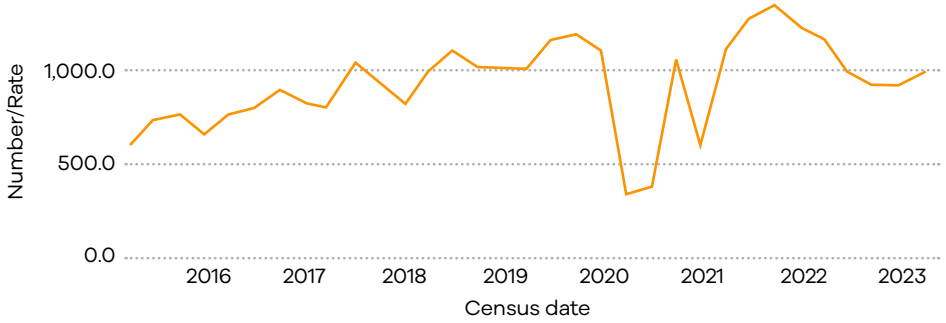

Robotics

Robotics

The word Robot conjures a range of mental images - of machines working on production lines assembling cars or androids walking around performing various tasks.

At its root a robot is simply a device that senses its environment and performs actions. A modern washing machine fits the definition as it weighs how much dirty laundry has been put in it and uses that data to plan the cycle. As well as laundering they are increasingly used for other support purposes such as delivering supplies to wards and dispensing drugs in a community or hospital pharmacy. They are also used for some direct clinical purposes such as providing "guard rails" for selected surgical procedures such as joint replacement and prostatectomies. Surgical Robotics has been covered in a separate strategy and is not covered further here.

However, even for these purposes they are far from ubiquitous, and most hospitals and pharmacies in the north of Scotland still rely on humans to push things along corridors and pick drugs off shelves.

Portering staff shortages sometimes means that clinicians have to leave wards to collect medication so that patients can be discharged on time.The potential of robotics is even greater than these examples, and spans the gamut of the domestic setting, care homes, community health facilities, hospital care and the land, sea and air between these places. There have already been trials of drone delivery systems in the north of Scotland and early explorations of autonomous delivery vehicles moving items between buildings.

The field of Social Robotics is also rapidly developing. This includes the use of devices which directly interact with people. For example to support emotional wellbeing, detect falls and provide assistance, both in the home and in care homes.

Challenges

To make progress in these areas will require new collaborations, both between and within organisations. Like many aspects of technology there is an economy of scale with robotics. Robots can be expensive and small organisations are less able to get competitive pricing than larger ones; and the north of Scotland has many small health boards. It therefore follows that health boards might want to form a consortium when negotiating contracts with robotics suppliers.

A similar principle exists for the specialist skills that are needed for conducting procurements and for supporting robots once they are in use. If we want to take advantage of the benefits that robotics could bring then we will need to share knowledge; and also learn from other health organisations and from specialist centres with robotics expertise. Regulation and legislation will have particular impacts on robotics and can therefore make planning difficult; particularly when it comes to robots moving on the open road or in the air.

We need to consider not just whether the use of robotics is possible for an area but also whether they are the right solution. We need to ensure that human-human contact is preserved.

Robotics: a vision

Linda is spending a couple of days in hospital for a hip replacement.
She notices that delivery robots bring medicines, laundry, food and other supplies to her ward. She was pleased though that when she was taken to theatre she was taken by a human – talking with him about her dogs helped calm her nerves.

The surgeon uses robotic assistance to increase the chances of a successful outcome.

When Linda was being treated for bowel cancer last year, the chemotherapy was flown to her local hospital by an Unmanned Aerial Vehicle. This meant she had much less distance to travel and was able to have a more normal life – helping with her successful recovery.

Less time-critical supplies such as laundry were taken to that local hospital by autonomous vans.

When Linda's dad had a stroke, a rehabilitation robot helped him maintain joint flexibility and helped him recover reasonably well. He still needs some support though, particularly as he lives alone. He has carers and his son popping in through the day to help but finds his cleaning robot entertaining as it slowly moves about vacuuming and mopping.

Robotics – Case Studies

- Some hospital and community pharmacies use robots to manage stock and assist with dispensing – placing and picking items from shelves.
- Some hospitals in Greater Glasgow and Forth Valley use Autonomous Delivery Vehicles to take food, laundry and other supplies to wards. This frees up humans for more personal tasks and has been shown to save money and reduce accidents.
- The CAELUS Project (Care & Equity – Healthcare Logistics Unmanned Aircraft System Scotland) is establishing more proof-of-concept trials of drone delivery systems. One use-case is transporting chemotherapy agents to local hospitals in order to reduce people having to travel.
- Trials of robotic phlebotomy devices are underway in some countries.
- Some care homes in England and New Zealand have successfully trialled social robots to support emotional and social interactions for isolated people. Such equipment is also under consideration for parts of northern Scotland.

Recommendations

1. Commission the national Robotarium to assess the potential to use Autonomous Delivery Vehicles in NoS hospitals and Care-Assistance robots in care homes and similar settings.

2. Work with Directors of Pharmacy to assess where further adoption of pharmacy robots in the community and hospitals could bring advantages.

3. Support the Innovation-led exploration of delivery Unmanned Aerial Vehicles – supporting more proofs of concept in the north of Scotland.

4. Systematically evaluate the potential for the adoption of robotics in other areas including, but not limited to, instrument sterilisation and floor cleaning.

5. When it is deemed that robotics could bring advantages, contracts should be developed on a regional (or national) basis to gain an economy of scale and to facilitate support.

Software process automation

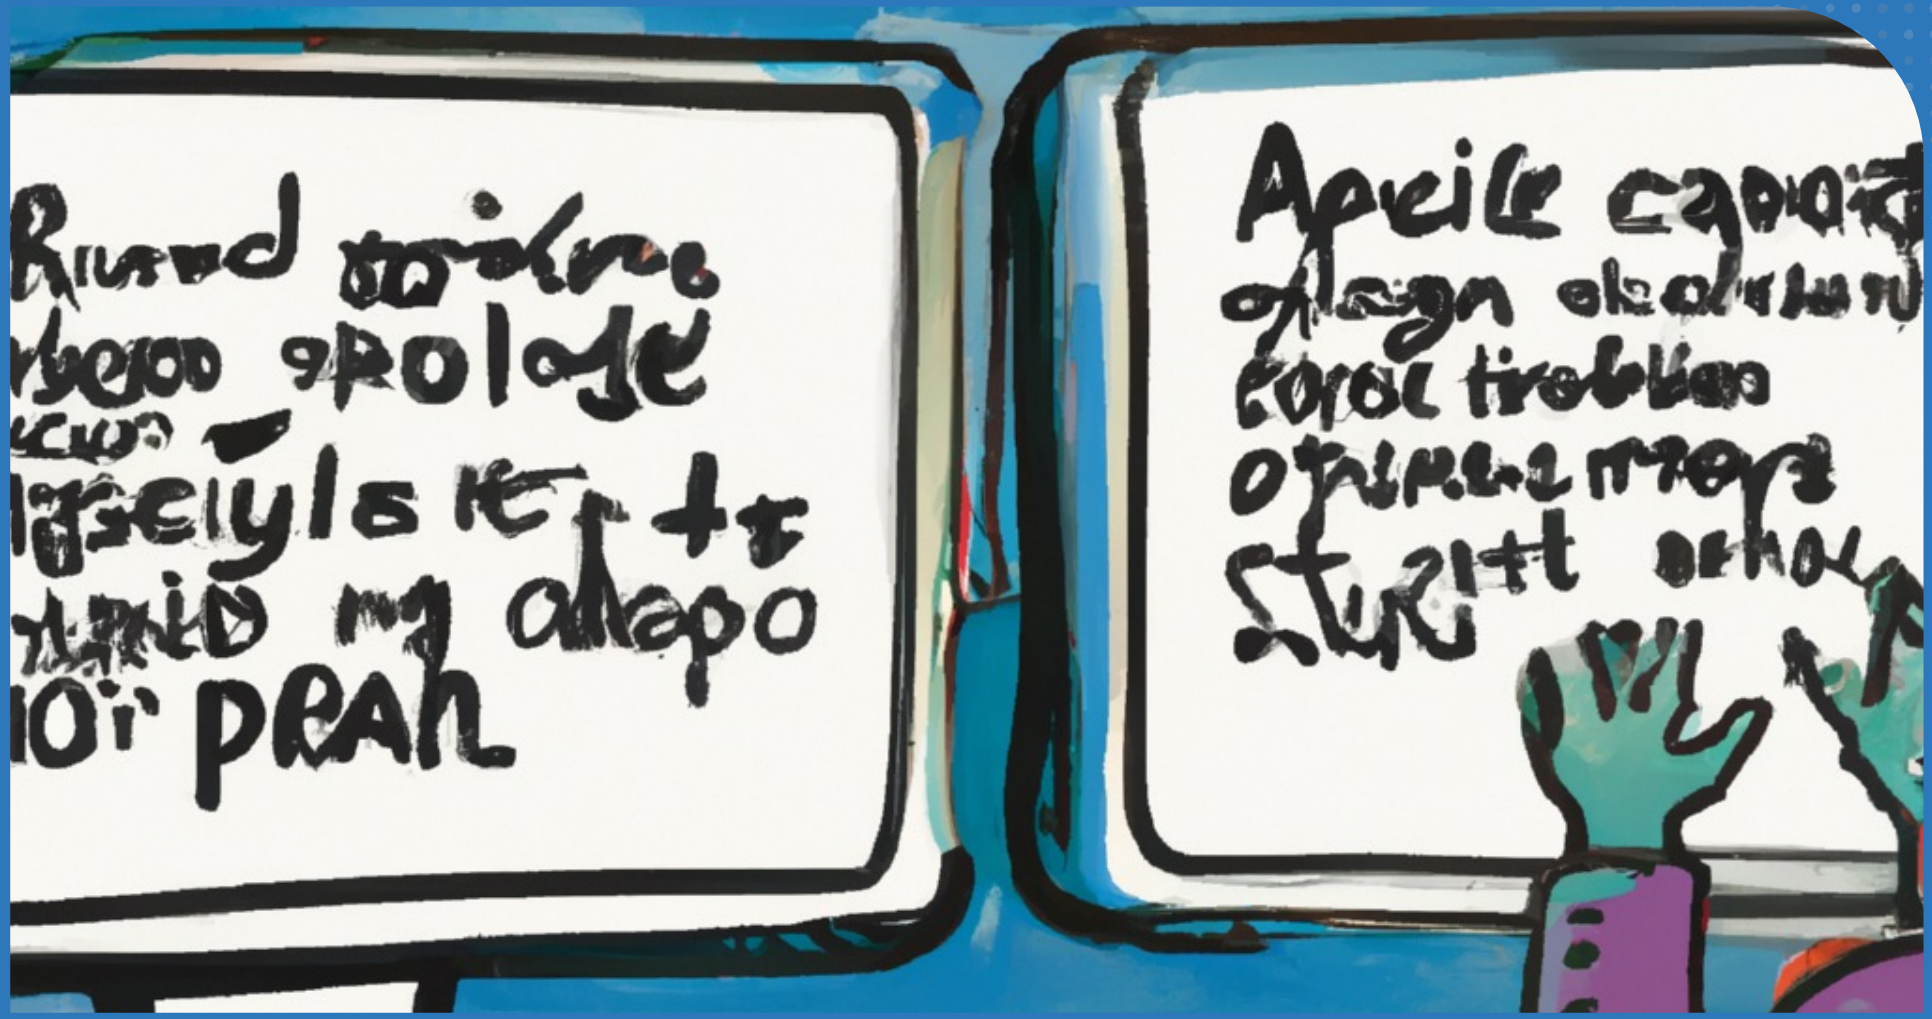

Automating processes using software

Software is increasingly used to support clinical, social care and administrative processes

- To record findings and document care plans.
- To request investigations and record the administration of drugs.
- To access best-practice guidance.
- To give members of the public access to their own data and to support them with managing their own health and care.
- To produce reports of activity.

By the late 2020s it will be very rare to find a clinical or administrative process that doesn't rely on software. Currently, working with clinical or administrative software typically involves humans clicking through a series of screens, choosing options and entering text. Sometimes it's necessary to copy and paste between different applications, and admin staff may have to repeatedly download files for reporting or quality assurance purposes.

Clinical management often involves responding to phone calls and other notifications and then working through software forms to perform tasks, and then documenting that that task has been done.

Increasingly though it's becoming possible for some of this notifying, interpreting and action-taking to be at least partially automated – whether in the back-office or at the front-line.

The Microsoft365 suite, which the NHS and Local Authorities have licences for, includes many automation tools out of the box. Some features require premium add-ons but others are available to those with the more widely available E5 licence.

This section covers the scenarios where humans have decided in advance what steps should be followed in a particular scenario. Depending on the techniques used, this is referred to as RPA – Robotic Process Automation, or Intelligent Automation.

Challenges
Adoption of M365-based automation has been slow due to a skills deficit and a delay in implementing the necessary security controls.

Identifying the best automation opportunities requires greater awareness and specialist skills; which are in short supply.

More advanced automations have a revenue cost and there is often no budget to cover this. Savings may come to staffing budgets and rebalancing may be needed.

Clinical Decision Support

A range of approaches can be used to increase the probability that a clinician makes the right decision. The right decision based on evidence, as well as the characteristics of the person and the situation.

Some of these approaches are quite basic but easy to introduce while others offer potentially more benefits but are much more challenging to implement.

Many, but not all, approaches to clinical decision support (CDS) include aspects of artificial intelligence.

Collectively they augment the intelligence of the clinical workforce while still involving "real" human intelligence at crucial steps.

Similar approaches can be used to support patients to decide between different treatment options, or none.

Different levels of Clinical Decision Support

Level	Examples	Types of AI that might be included
0 – Providing information [not on list in cited article]	Providing quick links to information such as BNF or guidelines. The Right Decision Service is mainly operating at this level.	Generally none. But AI could introduce some contextual awareness that proposes information resources that might be useful.
1 – Alerting	A blood result outwith the normal range is highlighted in red. Prescribing system alerts to drug interactions.	Human-derived rules-based automation.
2 – Interpreting	An ECG includes a suggestion that there is ST elevation in some leads. An X-ray is reported as normal, or needing a radiologist review.	Human-derived rules based (ECG example) Machine-learnt rules (X-ray example)
3 – Assisting	Patient admitted with upper abdominal pain – recorded in EPR. An admission order set of bloods proposes a request for serum amylase, in case it's forgotten. A new client's voluminous care and medical history is summarised for a new care practitioner	Human-derived rules based. Large Language Model
4 – Critiquing	Scriptswitch suggests an alternative antibiotic to the GP who's prescribed one that's not currently recommended.	Human-derived rules based.
5 – Diagnosing	IBM's Watson Image analysis suggests a differential diagnosis for a skin lesion.	Human-derived rules based, with automated information retrieval. Natural language processing. Machine-learnt rules – eg via image analysis.
6 – Managing	"Pathway" features in some EPRs which could order investigations, prescribe meds and refer a patient when a clinical diagnosis of, say, pancreatitis is made.	Human-derived rules based. For this area, as with others, the more "phenotype" data gets added to the EPR, the more that machine learning could be used to define the best pathways.

Adapted from: Pryor, T. A. (1990). Development of decision support systems. Int J Clin Monit Comput 7(3):137-146.

Process automation – a vision

- Muhamad is a service manager in a clinical department. His line manager needs twice-weekly reports on the department's activities. This used to require quite a bit of manual effort but is now almost entirely automated thanks to some training and direct help from the automation support team.

- Muhamad manages annual leave requests for his admin team and schedules their time on reception and back-office duties. Much of this is also automated now.

- When new members of staff join, often at short notice in the case of locums, Mohamed is under pressure to make sure that everything is set up for them – door passes, IT system access and so on. Thanks to the new process linking the joiner, the department, HR and the Digital Directorate this generally works like a dream.

- The department's secretaries have to manage internal meetings and get approvals for minutes and reports. This sign-off used to be part of the agenda but the secretaries received training in using the M365 Approvals app and that saves everyone time both before and during meetings; entirely avoiding some meetings.

- The department has been carrying a vacancy for a deputy service manager. With the efficiency gains from automation, Muhamad and the nursing manager are now considering diverting funds towards a new healthcare support worker instead; to help increase capacity and reduce waiting times.

- Automation, in the form of clinical decision support, is built into the electronic records that the department's clinicians use. This removes basic tasks and reduces "cognitive overhead", making it easier to do the right investigations and get the right treatments for people.

Process automation – Case Studies

- Aberdeen City Council has an automation which files a police Concern Report in the social care record of a person and then notifies the relevant care manager.
- The council uses a rules-based chat-bot to respond to questions – whether from members of the public or staff. This frees up specialist employees to focus on the more challenging questions.
- The council has a team which actively seeks out otehr automation opportunities across the organisation.
- Some Scottish GP practices have automated the process of filing eConsult summaries in the person's electronic record.
- Most electronic patient record systems use rules to make abnormal test results more prominent to clinicians, such as by presenting them in red text.
- Account creation and closure for joiners and leavers is becoming a semi-automated process in many organisations.
- NHS Lothian is piloting semi-automated theatre scheduling software which uses prioritisation and procedure duration data to build theatre lists.
- In the English NHS the Transformation Directorate promotes RPA and Intelligent Automation as ways for Trusts to improve their efficiency and contain costs.

Recommendations

1. Ensure safe and comprehensive adoption of Microsoft 365 Power Automate by:
   - Working with NES, NSS and others to ensure widespread knowledge and adoption of cloud Power Automate.
   - Develop regional skills and capacity in M365-based automation, including the use of Power Virtual Agents and premium versions of Power Automate.
2. Establish regional expertise in the development of more advanced rules-based automation – beyond what Power Automate is capable of. This should include collaboration on contracts and licences to gain best value.
   - An RPA virtual agent can work 24hrs per day supporting multiple processes across more than one organisation.
3. Establish a regional advisory function to help health boards and HSCPs identify opportunities for automation and then implement them at scale.
4. Departments should systematically review their processes and consider whether automation could bring value.
5. Rapidly establish some proofs of concept in areas that bring value to the North of Scotland.
6. Share case studies of where process automation has brought benefits – supporting further adoption.

Machine Learning

Machine Learning

Machine learning is the process by which software learns how to interpret data. For example:

- Which features in a photo of a mole best predict that it's not a melanoma.
- The patterns of white areas on mammograms that best predict a cancer.
- How word order affects meaning; as used by ChatGPT and other Large Language Models.
- Which sounds within a person's speech are most typically associated with a diagnosis of Parkinson's Disease.

The software, or "machine" is usually trained by setting it to analyse thousands of annotated files on the topic in question, such as "mammograms with cancer" and "mammograms without cancer". This is called Supervised Learning.

Sometimes the software is left to work things out for itself, Unsupervised Learning, for example to learn what factors may lead to a person being discharged late from hospital.

Once learnt, these rules can then be used to automate a range of processes in support of social care and health – for example to support diagnosis or prioritisation.

Challenges

For software to learn in this way it needs access to large volumes of high quality, accurate data which are relevant for the population in which it will be put to work. For some scenarios this exists – for example radiology images. But for other scenarios that's much harder to come by. Health and Social Care data typically means details of symptoms, problems, abilities, clinical signs, diagnoses, test results, treatments and outcomes. This remains a challenge as most of our current Electronic Record systems don't yet capture this information comprehensively; it may be in free-text, meaning may vary between record systems or it might even just be written on paper.

Machine learning itself, including the ability to extract meaning from free text, and the increased adoption of citizen-facing digital services, where the patient or client may record data about themselves, gives an opportunity to capture more of this information in the necessary detail.

As with other areas in AI, there are pockets of expertise scattered across the region but we don't yet have a systematic approach to getting these technologies adopted more widely, including in smaller organisations. For imaging-based ML this adoption process could require re-evaluation each time an X-ray machine in a rural hospital has a software upgrade.

As with other areas, along with the benefits come some risks; of algorithms getting it wrong. And, while humans also make mistakes, algorithms need to be compliant with the increasingly stringent safety requirements of the Medical Devices Regulation.

Generative AI

Generative AI is a type of machine learning that can learn from data and then generate new data that is similar to the data it was trained on. Examples include OpenAI's ChatGPT and Google's Bard which use Large Language Models for text analysis. OpenAI's Dall-E, and Midjourney do a similar thing for image generation.

The images in this strategy were produced using Dall-E, with the prompt used to generate the image at the start of this section being: "A photorealistic image of a robot looking at a chest x-ray". Bard provided the definition of Generative AI at the top of this page, the prompt being "Define Generative AI in a sentence"

Although a new field, a number of potential uses of Generative AI in Health & Social Care are emerging:

- Summarising information for service users and staff.
- Translating material, including into Doric.
- Answering general questions about health and wellbeing issues.
- Drafting reports and other documents.
- Providing novel images without copyright infringement.

Generative AI could also be combined with other AI technologies such as Process Automation to increase the productivity of staff. For example to read a referral and propose some actions for the clinician to then approve or modify.

As a new field, for clinical practice the focus should currently be on understanding the capabilities and limitations of Generative AI, publishing or otherwise sharing findings as we make them. However, the technology is already sufficiently mature to be used for drafting non-clinical and non-confidential material such as the headings of reports.

Machine Learning – a vision

- Maggie is a community nurse. For several years she's used her electronic record to gather information about her patients and make notes. More recently those notes have started to give other benefits. Information about the people in her caseload is used to plan and prioritise her day – recommending a route plan and highlighting people at increased risk of hospital admission.

- Maggie has a new patient – Andrew. Before she first sees him Maggie reads an automated summary of his health and social care situation.

- Andrew has an unusual patch of skin on his leg – next to an ulcer Maggie's been dressing. She's not sure what it is so takes a photo using an app on her phone. Within a few seconds the app tells her that the patch of skin has a very low probability of being cancerous and gives Maggie a short list of some possible diagnoses. She reassures Andrew and decides to discuss these possibilities with a colleague later.

- Andrew was in hospital recently with a bad chest infection. He had a couple of chest X-rays and, as a smoker, was a bit worried about what might be found. He was pleasantly surprised that he got a result back within a couple of minutes, delivered by the radiographer, that nothing serious had been found.

- Andrew is going to be having his gallbladder removed soon but his wife is quite concerned about complications. Together they looked through an individualised risk assessment that had been prepared to help Andrew decide if he wanted to go ahead. He was pleased to see that the risk of complications was less than his wife had thought.

- When he's discharged from hospital he'll be able to access an easy-reading summary of all his notes. His wife will be able to get this automatically translated into her native Ukrainian if she prefers.

Machine Learning – Case Studies

- An NHS Scotland AI skin cancer consortium, led by NHS Tayside, has been set up with the goal of being able to use Computer Vision and Machine Learning to diagnose skin cancer within 25 minutes by 2025.
- An 80,000 woman study of breast screening and breast cancer, led by NHS Grampian and the University of Aberdeen, demonstrated AI's ability in cancer detection and future cancer prediction on a North East cohort
- The Brainomix eStroke system has been adopted in many hospitals. It uses image analysis of CT scans of people who've had a stroke to rapidly guide treatment decisions.
- SELFBACK is an EU funded Horizon 2020 project that Robert Gordon University has been involved in. The goal is to develop a monitoring system to assist patients in deciding and reinforcing appropriate physical activity and exercise plans in order to self-mange Low-Back Pain. SelfBACK uses case-based reasoning as the primary AI methodology.
- Cervical screening in Scotland uses automated image analysis of liquid based cytology samples to more rapidly identify potential abnormal cells.
- Image analysis has also been used for many years to identify retinal signs of diabetes complications.
- The UK Joint Intelligence Committee uses Generative AI to summarise source data for analysts; to support better decision making in the field of national security.

Recommendations

1. Increase collaboration with the Data Lab and other relevant partners.
    - Promoting work done in the NoS so it can be adopted more broadly.
    - Being Fast Followers of Machine Learning techniques developed elsewhere; accelerating benefits for the citizens and staff of northern Scotland.
2. Do an options appraisal to determine the optimal Machine Learning (ML) infrastructure for the region, balancing performance and cost.
3. Create an ecosystem that supports adoption of ML from proof of concept to procurement; led by the Innovation Hubs.
4. Accelerate implementation into standard care when a ML concept has been validated. This should include increasing the publicity around successful use of ML in health and social care.
5. Develop an approach to the monitoring of algorithms, to support the regulatory requirements, that is efficient and straightforward.
6. Conduct proof of concept studies in the area of Generative AI, including Large Language Models, to see what benefits they could bring to health and social care.

Predictive analytics

Estimating the future

- Although we don't know for sure what will happen in the future, we commonly use historical data to make estimates of what it could be like. The weather this morning can give us a fairly good idea of what to expect tomorrow; even in northern Scotland.

- The branch of statistics and data science called Predictive Analytics takes this a few steps further. By gathering and analysing large amounts of historical and current data from multiple sources it becomes possible to estimate, or model, the probability of specific events or outcomes further into the future.

- While traditional statistical approaches and known influencing factors are often used, Machine Learning techniques can also be used when the number of potentially influencing factors is particularly large.

- The Goldilocks zone to find is the topic which has important questions and also the available data with which to answer them.

Challenges

- As with all data science, when the data is missing, of low quality or only available infrequently then this will impact on the ability to make meaningful predictions. We need to mature our electronic records, and use additional sources of data if we're to get the best out of the opportunities that Predictive Analytics could bring.

- Experience in Predictive Analytics has increased, spurred by the needs of the pandemic response. However, it is still predominantly traditional statistical methods that are used, and the skills to adopt the more modern Machine Learning approaches are in short supply. This skills gap is likely to particularly affect smaller organisations.

- Public Health Scotland are a centre of expertise for Predictive Analytics. However, their own strategy is in early development and it is unclear what advanced predictive analytics support or services they will be able to provide to health boards and HSCPs.

Predictive Analytics – a vision

Andrea is a Care Manager who liaises with the site & capacity and discharge planning teams in an acute hospital.

Her role is to make sure that patients get the care they need, whether to prevent an admission or to ensure a safe discharge to a homely environment.

She uses data on the predicted capacity in the acute hospital, step-down beds, care homes and home care teams to guide her decisions.

The site and capacity team can see an analysis of whole system flow including admissions and discharges. Admissions data includes people who have an imminent scheduled admission, emergency admissions expected in the next few hours and a probability-based model looking ahead 24 and 48hrs. The model predicts unscheduled admissions for various presentations based on historical data as well as a range of other parameters such as the weather and rates of infectious diseases.

For admitted patients, the likely date and time of discharge is predicted to enable more efficient planning. Andrew (see Machine Learning), has arrived in the hospital for his operation and he and his wife are hoping for a swift getaway when he's medically fit.

Consultants and hospital managers also have access to this data and can use it to refine their schedule, and staffing requirements.

The current and predicted data are accessible on the staff member's phones as well as in the hospital control room and via ward PCs.

Analysts from across the region collaborate between teams and with Public Health Scotland to provide this data. The models are continuously refined.

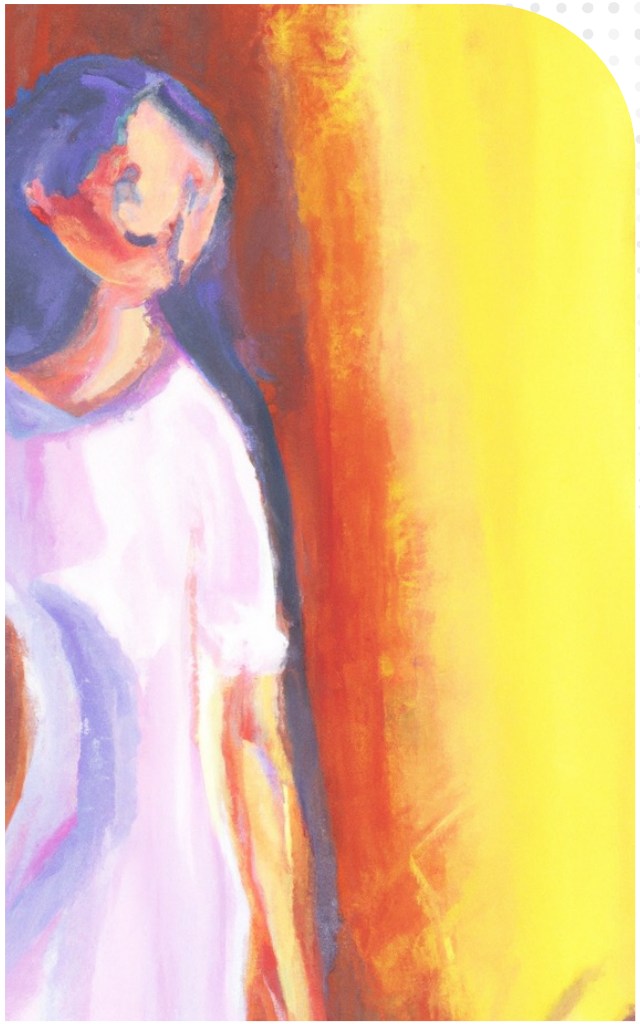

Predictive Analytics - Case studies

- System Watch is a tool created by Public Health Scotland that allows health boards to predict unscheduled admissions over the next 1-6 weeks. It relies on data submitted by health boards, other organisations and flu statistics, some data items are only provided on a weekly basis.Individual health boards have similar systems in place.

- A trust in SE England, uses predictive analytics to identify reasons why patients might not be able to attend – this could include the weather, traffic or people's jobs.Once potential non-attenders are found, automated software processes rearrange dates; for example, if somebody is less able to take time off during the day, they will be given an evening or weekend slot.

NHS set to save billions as innovative artificial intelligence tool launches | UK Healthcare News (nationalhealthexecutive.com)

Recommendations

- Develop the skills of the Health Intelligence workforce in the region to better exploit modern predictive analytics tools.
    - Enhance collaboration between teams in the region to gain benefits from this skillset.
    - Develop relationships between NHS Analysts (Territorial and National) and Academia to assist with the adoption of specialist modelling skills.
- Increase awareness of and provide secure access to health and social care data that's currently available in the region to support the development . Consider the creation of a regional data loch.
- Determine what the key data gaps are and consider how they could be routinely collected.
- Progress proof of concept work on predictive analytics to support logistics and whole system flow. For example predicting DNAs and length of stay.
- Improve the visibility of the output of predictive models; making outputs available at the point of care to all staff and patients/service users who need to see them.

Ethics and AI

As we think about the increasing role that AI may play in health and social care in the years ahead, it is important to consider the ethical implications. Questions have been asked across the core domains of ethics - Autonomy, Justice, Beneficence and Non-Maleficence.

Autonomy is about the right of an individual to direct their own future. Stories about algorithms making diagnoses and deciding what treatment people should have, remind us of "Computer says no"; the idea that people may be denied treatment or conversely railroaded into treatment that they don't want. Current technology is not at the stage of removing humans from treatment decisions, and nor is this on the horizon. All current technologies are probability based, categorising X-rays, for example, into probably normal or probably abnormal. The step of initiating treatment is left to clinicians.

Some AI technologies have the potential to increase the autonomy of individuals, for example by providing a person with tailored information that could help them give better-informed consent for a procedure.

Justice is essentially about equality and being fair and appropriate. Concerns in this area have centred around the bias sometimes found in machine-learnt algorithms. For example, if the data the algorithm was based on came from a drug trial that had excluded women of reproductive age, then the decision-support provided to clinicians would be less reliable for this group.

AI technologies that can influence practice need to have an Equality Impact Assessment performed. This also needs to happen for any citizen-facing products, where digital inequality might limit access.

Ethics continued

Beneficence is about the need to do good. The biggest risk here is that through inaction and excess caution we fail to gain the benefits offered to service-users and tax-payers by the increased adoption of AI technologies.

If AI technologies can improve access to health and social care, increase its quality and reduce harm to staff at lower cost than existing approaches then failure to choose to adopt them could be widely seen as unethical.

Non-maleficence is the need to avoid causing harm. Any new technology, whether a novel design of bed or diagnostic software, can introduce new risks. As noted elsewhere, any AI technology that could impact on patient care will need to have a clinical safety assessment prior to its adoption. Many will need to be registered with the Medicines and Healthcare Regulatory Agency.

We need to have capacity to perform these tasks – local or regional capacity with support from national service teams.

The law and AI

The Information Commissioner's Office rightly takes an interest in AI due to the potential risks to the rights of data subjects. The ICO has published a range of guidance documents for organisations considering increasing their use of AI, including an AI Toolkit.

In practice, the data protection risks will vary considerably between different AI technologies. Some, such as better washing machines or autonomous delivery vehicles will pose little threat to personal data, while others, such as citizen-facing personalised decision support, could be seen as higher risk. The potential hacking of autonomous delivery vehicles or decision support systems raises particular concerns and we will need a robust approach to cybersecurity.

In order to gain benefits from expanded adoption of AI across the region in a reasonable time-frame we will need a rapid process for gaining Information Governance advice and data protection assessments. Subsequent adoption decisions need to take into account both the assessment of data protection risk and of the potential care quality impact that the initiative might offer.

Trust

Essentially, these legal and ethical aspects are about trust. If we want AI to form an increasing part of health and social care delivery then the public and members of staff need to be able to trust it. Scotland's AI Strategy covers this area well, referencing the Organisation for Economic Cooperation and Development's five complementary values-based principles for the responsible stewardship of trustworthy AI. It is worth reiterating these here:

- AI should benefit people and the planet by driving inclusive growth, sustainable development and well-being.
- AI systems should be designed in a way that respects the rule of law, human rights, democratic values and diversity, and they should include appropriate safeguards – for example, enabling human intervention where necessary – to ensure a fair and just society.
- There should be transparency and responsible disclosure around AI systems to ensure that people understand AI-based outcomes and can challenge them.
- AI systems must function in a robust, secure and safe way throughout their life cycles and potential risks should be continually assessed and managed.
- Organisations and individuals developing, deploying or operating AI systems should be held accountable for their proper functioning in line with the above principles.

Other challenges and risks to be considered

- For all new technologies there is a "change hump" to get over; initially there are increased costs and no improvement in productivity. Those benefits may only come after a significant interval when the changes are fully embedded. To get over that hump requires strong leadership, and investment.

- The economy of scale favours larger organisations whereas the north of Scotland includes many small organisations. There are few current incentives to collaborate but without that collaboration it is likely that people living in smaller health boards will experience a less modern, and possibly less safe, healthcare environment.

- Public and Staff understanding of AI is variable – partly down to poor communication. Media coverage can flip between stories that over promise and those that fill the viewer with dread. Recent research in Aberdeen has shown that the more people understand about AI technologies, the more they are likely to accept it as part of their care.*

- There is already a general technical skills shortage in the region, particularly in rural areas. This is particularly true of the highly specialised skills required for increasing the adoption of AI. These factors further highlight the need to collaborate and share that specialist knowledge.

* de Vries CF, Morrissey BE, Duggan D, Staff RT, Lip G. Screening participants' attitudes to the introduction of artificial intelligence in breast screening. Journal of Medical Screening. 2021;28(3):221-222. doi:10.1177/09691413211001405

Conclusions

It's never been more important for clinicians and practitioners to focus their time where it will bring the greatest value to their patients and service users. The workforce:demand imbalance doesn't leave us any other option.

Increasing the adoption of AI technologies is one way to help our humans spend more of their time where it counts – on high value activities where only the human touch will do.

While some AI can seem like science fiction, other aspects are already widely adopted elsewhere. We have a good understanding of what would be needed to gain these benefits for the north of Scotland.

In the north of Scotland, the most achievable way to achieve these benefits will be through collaboration and targetted investment.

Membership of writing group

This document has been the product of a series of meetings from late 2021 to mid 2023.

Membership and attendance has varied over that time but the full list is:

Name	Role/Job title	Organisation
Steve Baguley	Clinical Director – eHealth (chair)	NHS Grampian
Carol Mayo	PA/Administrator (minutes)	NHS Scotland North Regional Team
Jon Eilbeck	Head of Programmes & Development	NHS Grampian
Andy Keen	Innovation Lead	NHS Grampian
Lorna Cameron	Project Manager, Innovation	NHS Grampian
Iain Ross	Head of eHealth	NHS Highland
Norry Thom	Head of Technical Services	NHS Highland
Ally Gentles	Head of Operations - Digital	NHS Tayside
Richard Rae	IT Manager	NHS Orkney
Sandra Stankus	eHealth & Infrastructure Team Lead	NHS Orkney
Carrie Heddle	Programmes Lead (PIP)	NHS Grampian
Jim Cannon	Director of Regional Planning	NHS Scotland North Regional Team

Name	Role/Job title	Organisation
Kate Cheeseborough	Innovation Project Manager	NHS Tayside
Colin Fleming	Lead for Innovation	NHS Tayside
Gerald Lip	Director of Breast Screening	NHS Grampian
Sandra Murdoch	Public Representative	
Jennifer Nicholls	Centre for Health Science	NHS Highland
Lesley Anderson	Prof of Health Data Sciences	University of Aberdeen
Kathleen Carolan	Director of Nursing	NHS Shetland
Lee Heaney	Innovation Manager	NHS Highland
Martin Malcolm	Head of Health Intelligence/Research and Innovation Lead	NHS Western Isles
Gordon McFarlane	Clinical lead for eHealth	NHS Shetland
Jeff Shaw	Partnership Manager	Aberdeenshire HSCP

Summary of the consultation process

1. **First phase**
   - Planning leads within the north
   - Regional Innovation teams
   - NoS Directors and Clinical Leads for Digital
   - NoS Workforce analysts and planners
   - Relevant University departments of Institutions already engaged
   - National AI – Data Lab and Scottish Government AI
   - SG Digital Health & Care

2. **Chief Executives and Chief Officers Review**

3. **Second phase**
   - Management and advisory structures within health boards
   - Health & Social Care Alliance
   - Scottish Care
   - Digital Health & Care Institute
   - Centre for Sustainable Delivery
   - SG Planning - Policy
   - SG Planning - Operational
   - Remaining academic institutions within the region
   - Plus others as advised by first phase of consultation

Glossary

Direct readers to the glossary of the National AI Strategy Scotland's AI Strategy - Glossary — Scotland's AI Strategy (scotlandaistrategy.com)

https://www.scotlandaistrategy.com/scotlands-ai-strategy-glossary

Acknowledgments

Many thanks to the members of the writing team and to those who gave feedback as part of the consultation exercise.

Dall-E2 for their assistance in creating the art work.

Spending time where it counts

An Artificial Intelligence strategy for health & social care in the north of Scotland 2023-2027

Final version, January 2024

DESIGN: WLS0723