ONE LONDON

AI CENTRE for Value Based Healthcare

A Framework for the safe, efficient and effective implementation, use and maintenance of AI in health and care in London.

North West London Integrated Care System
Working together for better health and care

North Central London Health and Care Integrated Care System

NHS North East London

South West London Integrated Care System

South East London Integrated Care System

Contents

1		**Executive summary**	3
2		**Introduction**	6
	2.1	**Purpose**	6
	2.2	**Scope**	7
3		**Partnership**	10
	3.1	**Partnership across our system**	10
	3.2	**Governance**	11
	3.3	**Roles and responsibilities**	11
	3.4	**Partnership with vendors and other external organisations**	14
4		**AI Infrastructure and data**	16
	4.1	**AI Infrastructure**	16
	4.2	**Data pipelines as reusable assets**	17
	4.3	**Information governance**	17
5		**AI use cases in health and care**	19
	5.1	**Use cases**	19
6		**AI delivery approach**	23
	6.1	**AI delivery lifecycle**	23
	6.2	**Idea pipeline**	24
	6.3	**Proof of concept**	26
	6.4	**Pilot**	29
	6.5	**Business as usual and scale**	32
	6.6	**Monitoring**	34
7		**Communication and engagement**	37
	7.1	**With our workforce**	37
	7.2	**With our community**	39
8		**Training**	41
9		**Conclusion**	42
10		**Appendix 1: Definitions**	43
11		**Appendix 2: AI Ethical Reflection Template**	46
		High level project information	46
		Common ethical considerations	46
		Impact identification and scenarios	47
		Potential harm analysis	48
12		**Appendix 3: Contributors to the development of this framework**	49

Executive summary

This document provides an overview of the agreed way of implementing and monitoring Artificial Intelligence (AI) products in the London health and care system.

The Framework covers five key areas:

Partnership
Principles for how we will work together across organisations to align approaches, share information/ expertise and reduce duplication for efficiency and to generate the greatest opportunity for value creation.

Infrastructure and data
The infrastructure, data and information governance arrangements that need to be in place to implement AI tools.

Use cases
Areas where AI products are available which may address identified challenges in the NHS.

AI delivery approach
How we take the opportunity to work together as a system to pilot, implement and monitor AI in a consistent way, as well as ensuring there are appropriate controls in place to mitigate risks associated with the use of AI.

Communication and workforce development
Our approach to communication and development (both with our community and our workforce) required to implement and use AI successfully.

This Framework does not separate the implementation and management of AI from other digital technologies. However, it is designed to provide guidance on the considerations that are specific to AI and outline a delivery approach to which any local delivery methodology can be mapped, to facilitate structured collation of data on AI projects and utilisation across London.

The Framework outlines the role of an AI Coordination and Advisory Lead (AI Lead) – with one to be identified by each Integrated Care Board (ICB) for their respective Integrated Care System (ICS). The Coordination and Advisory Lead is critical to delivery of this Framework, as they will collate information from across their system on the status of AI projects and implementations and will provide expert advice to the system on planning for the implementation of AI; focusing on the mitigation of risks.

The AI Lead's role will be to monitor the progression of AI projects as they move through the stages from idea pipeline, through proof of concept, pilot and to business as usual, and, if successful to scaling across their ICS. This includes supporting decisions about when to cease at any of these stages.

Utilisation of this Framework will facilitate collaborative working across the London region, and aims to support the safe, efficient and effective implementation, use and maintenance of AI in London. It should be noted that the regulation, legislation and guidance around the use of AI is an evolving field. This Framework will be updated from time to time, but is not intended to provide an overview of all of these requirements.

This Framework highlights some of the key aspects of regulatory processes that should be followed by providers implementing and using AI. But it remains the responsibility of the implementing organisation to ensure their approach complies with all relevant regulation, legislation and guidance.

2. Introduction

2.1 Purpose

Applications of AI are becoming increasingly available to the health and care sector. This offers opportunities to address challenges faced by our workforce created by high levels of demand for health and care services and insufficient capacity to meet that demand. However, it also brings risk with it. This includes risks associated with embedding bias into decisions, as well as risks associated with inequitable and inefficient implementation and management of AI solutions.

The UK Government AI Action Plan (AI Opportunities Action Plan - GOV.UK) states:

"Our ambition is to shape the AI revolution on principles of shared economic prosperity, improved public services and increased personal opportunities so that:

- AI drives the economic growth on which the prosperity of our people and the performance of our public services depend;
- AI directly benefits working people by improving health care and education and how citizens interact with their government; and
- the increasing of prevalence of AI in people's working lives opens up new opportunities rather than just threatens traditional patterns of work."

The speed at which this field is growing provides us with the opportunity to move to an efficient and effective strategic plan for adoption of AI projects. This is a move away from ad hoc or locally planned approaches to technology implementation that have been a feature of the past.

There are many documents prescribing and guiding the use of AI, but there is a lack of clarity about how to practically move forward with the safe and efficient use of AI. Some providers are able to build on existing expertise and governance arrangements to cover AI, but others do not have access to such expertise. This inequality risks either missing out on opportunities AI provides, or utilising AI without putting mitigations in place to address common risks and solution-specific risks.

The purpose of this Framework is to outline an approach for the implementation, use and maintenance of AI within the London health and care system. The proposed approach agrees a collaborative methodology, where providers share information about their plans and outcomes, and where projects follow a level of consistency in approach so that opportunities for spread and scale can be harnessed.

The Framework also provides an approach to engagement of patients and citizens, as well as our workforce in our plans to utilise AI in health and care pathways. This is critical to success as it is only through access to high quality data that AI tools will be able to improve health and care delivery. We need to ensure that we have the trust of patients and our workforce in the way that we are considering these new technologies, and that we have processes in place to address concerns if they are raised.

2.2 Scope

Organisations included
The intention is that the approach outlined in this Framework will be implemented wherever AI is used. However, it is acknowledged that uptake will be variable uptake. ICBs will be responsible for promoting the Framework in their systems, and for supporting alignment of their provider organisation to the Framework, focusing on ensuring sharing of information to a lead group that will facilitate sharing across London. The more organisations that agree to share information across our system, the more efficient and effective we can make AI implementation and monitoring.

Scope of solutions included
AI is an umbrella term for a range of algorithm-based technologies that solve complex tasks by carrying out functions that previously required human thinking. This includes AI being used as a medical device and AI being used in other contexts in the healthcare space. This document covers both but draws the important distinction between them, as the use of AI as a medical device has additional regulatory requirements (www.gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device/software-and-artificial-intelligence-ai-as-a-medical-device).

Types of technologies covered
The definition of AI is very loose amongst software vendors . For the purposes of this document, it is important to be discrete about products that actually contain AI and those that do not.

These software systems may use one or a combination of different AI technologies; which are of differing levels of complexity.

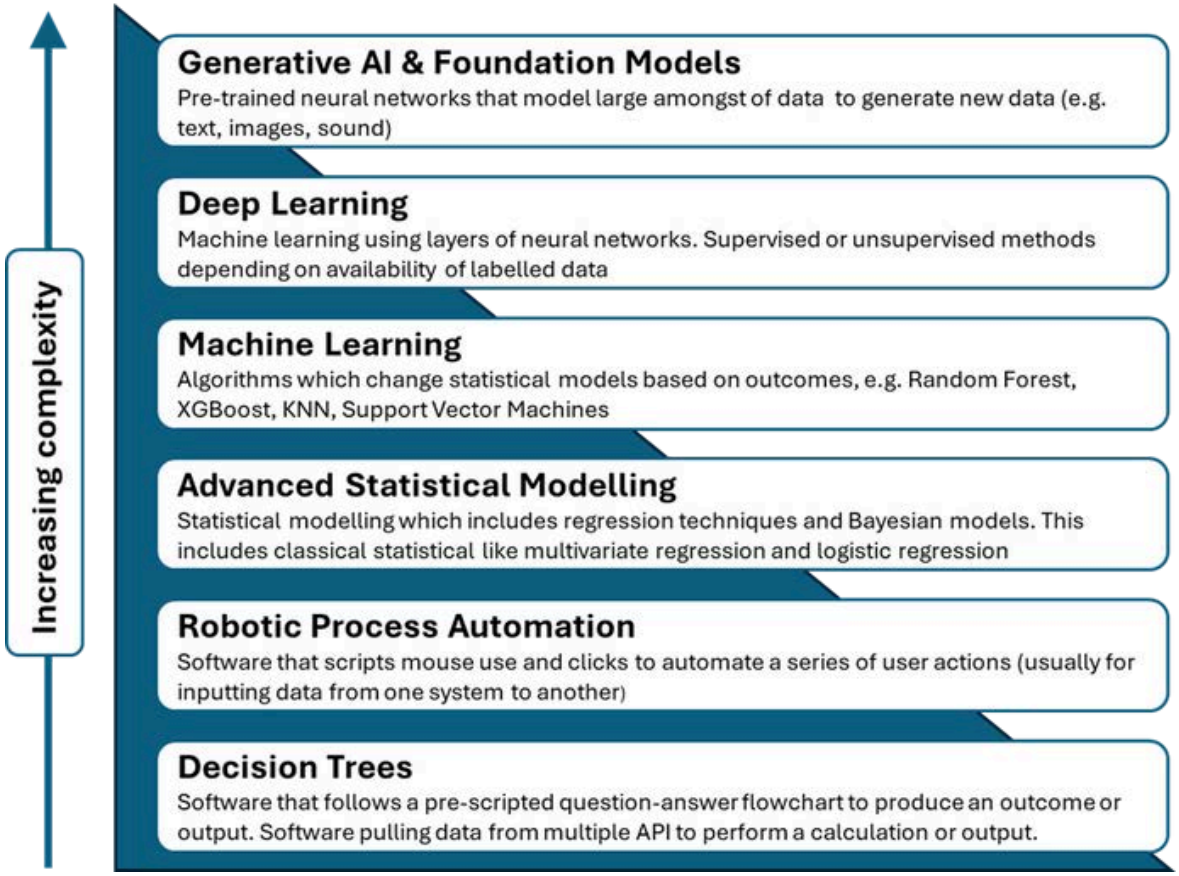

Figure 1: Definitions of AI software technologies used in this document (noting that decision trees and robotic process automation do not fall within some definitions of AI).

Decisions made using AI are either fully automated, or with a 'human in the loop'. Currently, the General Data Protection Regulation (GDPR) requires a human to be involved in the process, and this Framework supports the proposition that, as with any other form of decision-making, those impacted by an AI supported decision should be able to hold someone accountable for it (ref ICO Definitions | ICO). For this reason, this Framework only covers AI where the final decision is made by a human.

Some sections of this framework are not relevant to AI that does not encompass elements of machine learning, and some are only relevant to AI where it is used as an adjunct in clinical care (as opposed to back office functions). This is because there are different ethical considerations and monitoring requirements when the AI solution is learning from data and changing its outputs based on that learning. In all cases, however, the sharing of information and alignment with the AI delivery approach is applicable.

3. Partnership

3.1 Partnership across our system

The following principles guide how all partners work together across all initiatives:

1. Act in the best interests of the system as a whole, putting the needs of our community at the centre of our decisions.

2. Provide leadership and resources to support the system in delivering its objectives.

3. Work in partnership to ensure shared learning and efficiency in everything we do.

4. Have digital strategies and plans that support achievement of the vision including taking responsibility for delivery of projects and programmes to support system transformation, system resilience, and continuous improvement.

5. Invest appropriately in digital and data initiatives, acknowledging that digital foundations can require significant investment but are required to unlock efficiencies and transformation required for financial sustainability.

6. Consider the impact of actions of one organisation on our system, and engage broadly to ensure the impact is well-managed, including mitigation of risks and maximising opportunities.

Thinking about how this applies to the implementation and use of AI, it means that we will:

- Take extra time during some stages of the lifecycle to establish processes that will work system-wide, such as ensuring other partners are identified in our procurements.
- Agree which partners will take a lead for which use cases, testing and hosting a solution that may be used across multiple organisations.
- Take time to share expertise and experience.
- Commit to sharing progress updates, including allowing the collation of information on projects in flight.
- Agree to monitor AI products associated with a particular use case and providing advice to other partners using that product if any issues arise.
- Agree to consolidating some expertise, which will be made available across our system, noting that each partner will need to have teams who are experts in change management and implementation of AI, as they would with any other digital solution.

3.2 Governance

London has multiple sovereign NHS organisations with different priorities and constraints, and this Framework does not seek to supersede or duplicate any arrangements in place.

It is proposed instead that each ICS forms an AI Advisory and Coordination Group, to be led by the AI Advisory and Coordination Lead Organisation (see roles and responsibilities section below for more information). This would not be a governance forum as such, but would facilitate collaborative review of projects to share expertise, as well sharing of information on project and monitoring status, and sharing of information about risks and mitigations.

The frequency of, and approach to these meetings will be decisions for each ICS.

3.3 Roles and responsibilities

Coordination and sharing of information and expertise is critical to the delivery of this Framework.

Each ICB in London will identify an for their ICS. This Lead is likely to be a group or organisation, but requires a single point of contact for the roles and responsibilities of this group within the ICS. This organisation needs to not only collate and share information, but should also have access to expertise on AI implementation and utilisation that can be shared with providers across the system. The AI Lead may be shared between ICSs or there may be one per ICS. Where there are many, it is the responsibility of this AI Lead to share information in a consistent manner with other AI Leads in London.

The AI Lead is critical to delivery of this Framework, as they will collate information from across their system on the status of AI projects and implementations and will provide expert advice to the system on planning for the implementation of AI, focusing on the mitigation of risks. Their role will be to monitor the progression of AI projects through the stages from idea pipeline, through proof of concept, pilot and to business as usual, and, if successful to scaling the projects across their ICS. This includes supporting decisions about when to cease at any of these stages.

It is also expected that this organisation will lead a community of practice for their system, supporting our workforce to learn from each other.

This Framework refers frequently to the AI Lead, and each provider using AI, or considering the use of AI, should know who this is for their system. The details of the AI Lead for each ICS will be published as an Addendum to this Framework and will be updated as required.

The roles and responsibilities of the AI Lead and other organisations are identified below, with the term partners meaning all health and care organisations across our system.

Organisation	Role
Health and care providers	• Governing any AI projects or uses in their organisation. • Conducting AI projects and managing AI solutions in accordance with the Framework. • Sharing information across London on AI projects and outcomes. • Ensuring safe and responsible implementation of AI, including complying with legislation and policies, such as but not limited to clinical safety assessments and data protection impact assessments. • Leading AI projects for the system, and considering the potential for cross-organisational scale when planning and implementing projects. • Agreeing to reduce duplication in AI project implementation as much as possible by leading or learning from others prior to undertaking local projects.
The AI Advisory and Coordination Lead (AI Lead)	• Leading and supporting an ICS-wide AI Advisory and Coordination Group to oversee the programme management of AI for the system. • Collating and publishing (internally and externally as appropriate) the status of AI projects and implementations across the system (aligning with the format of other London ICSs) including the leads for each use case, the status of each use case to support sharing of information. • Providing expertise to health and care organisations on AI implementation and monitoring. • Providing expertise on identifying and managing risks and issues relating to AI implementation and monitoring including by reviewing project documentation such as the Ethical Reflection Template (Appendix 2).

Organisation	Role
The AI Advisory and Coordination Lead (AI Lead) - continued	• Providing expertise to support evaluation of whether AI technologies are effective and whether they provide meaningful and measurable benefits, including through reviewing the TEST evaluation framework documentation developed by the implementing organisation. • Providing advice to support the system-wide procurement of AI solutions. • Providing advice on tools and capabilities for AI. • Providing consultancy services such as AI monitoring, implementation advice and evaluation. • Supporting efficient AI implementation through sharing of governance documentation such as data protection impact assessments templates or clinical safety assessments made available by other organisations. • Coordinating and participating in an AI community of practice • Work with the other AI Leads across London to share information and reduce duplication.
Integrated Care Board (ICB)	• Signing up to this Framework and promoting its implementation across their system. • Leading the identification and/or commissioning of the Advisory and Coordination Lead for their ICS. • Identifying system-wide priorities that may fall outside of organisational priorities and identifying lead organisations to support these projects.
NHSE London Region	• Supporting opportunities for scale across London (where pilots support this), including by facilitating funding opportunity identification and prioritisation. • Ensuring that AI training needs are considered in the development of the Digital, Data and Technology workforce plan for London as well as consideration of the plan for supporting the broader workforce to engage with AI tools and technologies.

3.4 Partnership with vendors and other external organisations

This AI Framework makes it clear that we are focusing on solving problems, rather than using AI solutions. AI is a fast-developing field, and many products are not yet fit for purpose in the UK health and care sector.

There may be value in working in collaboration with vendors to develop products, or in fact, develop products of our own. When a provider is working with a vendor or is developing a product, it is important to consider how we ensure that the product is then made available to other providers in our system in a cost-effective way, and one that recognises their role in our partnership.

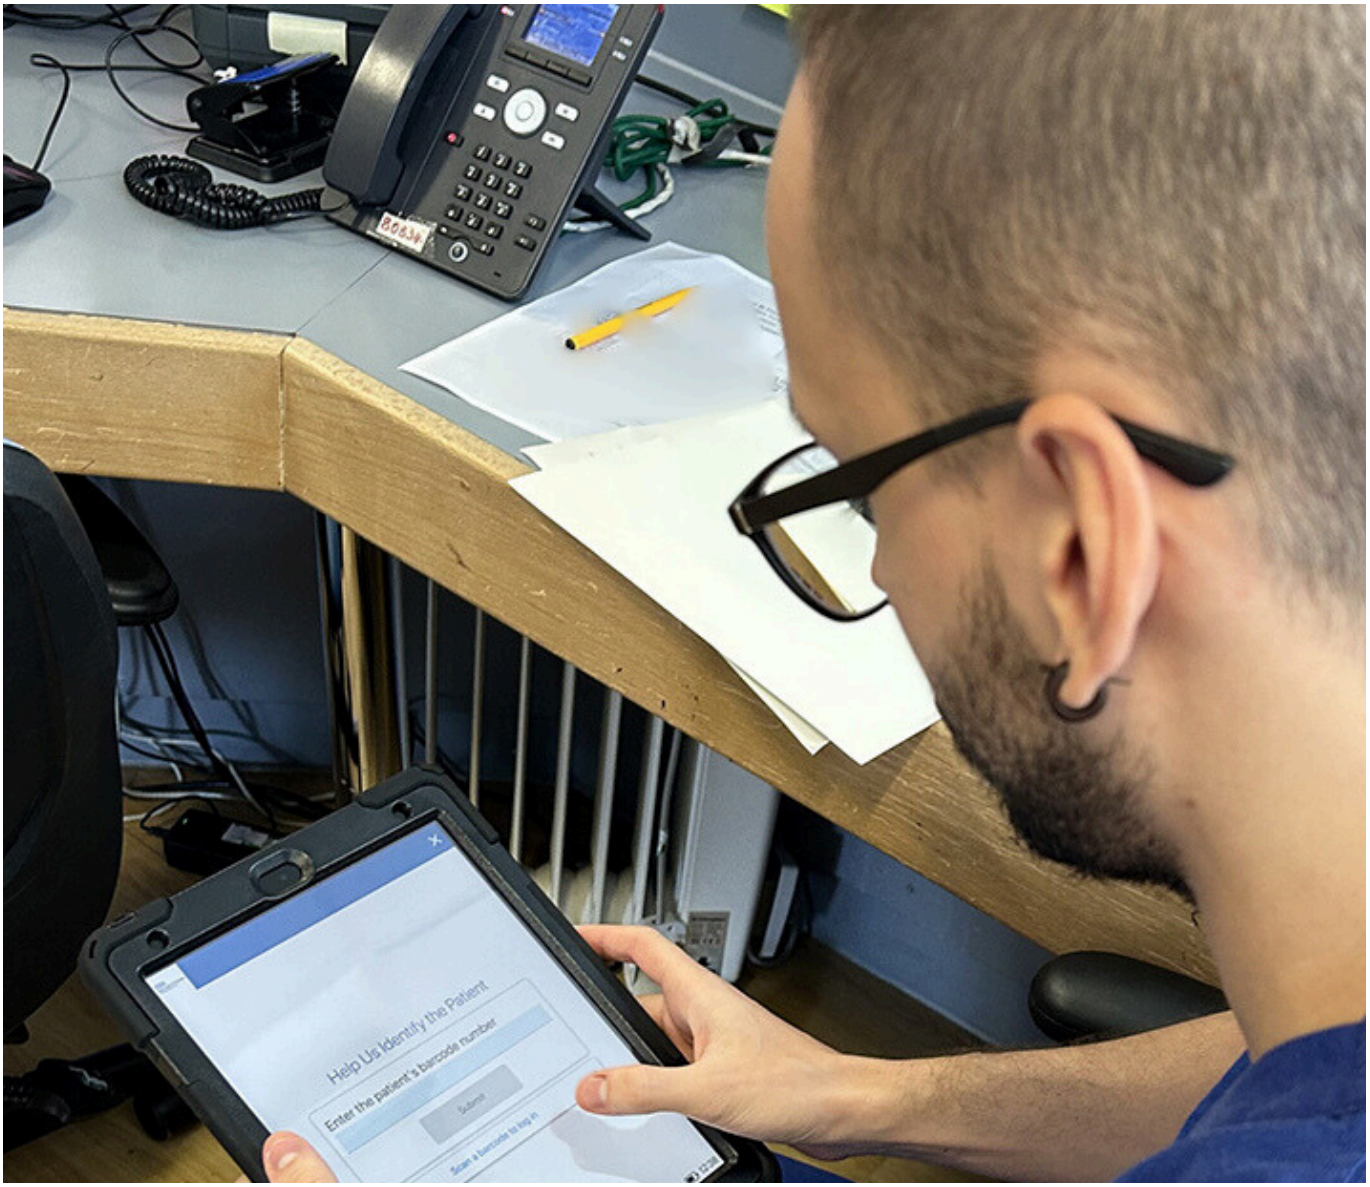

4. AI Infrastructure and Data

4.1. AI Infrastructure

To successfully implement and host AI solutions appropriate infrastructure needs to be in place, which will vary depending on the AI model. The models are outlined below:

- Vendor-hosted AI models are often the simplest (and usually cheapest) scenario. This involves using AI in a static form that is off-the-shelf with customisation limited to the user-interface. This would be the most common approach to purchasing an AI product from a vendor. Most software vendors provide their AI products within an existing vendor cloud infrastructure which may limit transferability of configuration to other use cases. It may also limit approaches to monitoring, meaning that careful consideration of the monitoring tools provided is important.

- Training AI models is substantially more complex and requires re-usable infrastructure that is not dependent on specific AI vendors. Many academic-based projects or pilots fall into this domain. The main dependencies for this high-performance compute, large-scale high-quality source data and a skilled machine learning engineer workforce. The main advantage of this approach is that a high degree of customisation is possible.

The full AI life cycle includes both inference and training, with a layer of monitoring. This is a very complex and sophisticated use-case and is gold-standard in most industries, as the performance of AI models drift or degrade over time. The full AI lifecycle requires model hubs for model management, bias detection and dataset drift.

To avoid limits on scalability to provider or ICS-level, it is important to consider if the AI vendors provide limited model hubs serving only their own products.

4.2. Data pipelines as reusable assets

Describing the full scale of data depth, quality and complexity is not within the scope of this Framework. However, it is important to note that the implementation and impact of AI are highly dependent on the quality of data at both the point of training and at the point of use, with quality encompassing the following:

- interoperability
- missingness
- granularity or fidelity
- semantic depth
- dimensionality
- standardisation
- automation/ fluidity.

Limits to any of the above affect the use, performance and impact of AI implementation projects, and this should be considered when evaluating products. Defining the cause of the failure of a product to deliver the intended benefits will be useful to other organisations, as they may be able to mitigate those circumstances to deliver a successful implementation (for example, if their data quality was higher than that of the organisation conducting the pilot).

4.3 Information Governance

All AI products require access to data. In some cases, this may be data to train the models. In all cases, local data is required so that the AI model can apply the algorithms to generate the outputs. This means that for all use cases there is a need to consider both information governance and data sharing.

The implementing organisation is responsible for ensuring that a Data Protection Impact Assessment (DPIA) is undertaken, and in some cases, a data sharing agreement. However, there is value in having consistency in decisions around the use of data for AI solutions, and pooling expertise to make decisions may be valuable to support our IG experts make what are sometimes difficult decisions about the balance of risk against benefits to patient care.

The AI Lead in each ICS will host a database of products and associated documentation, such as DPIAs, including those that were developed by organisations leading the implementation of a specific use case (and will share these across London as appropriate). If other organisations do not agree with the decisions that were made (and where there is no clear rationale for the inconsistency), there should be a collaborative discussion between the organisations to achieve consistency. Where necessary, consideration could be sought by the AI Coordination and Oversight Group.


Oxleas
NHS
Seema

5. AI use cases in health and care

5.1 Use Cases

Health Innovation Networks in each area often have information available on the use cases for digital products including AI, and this may be the most up-to-date source of information.

A study conducted by the Health Innovation Network (Home - The Health Innovation Network) found that at present, utilisation is greatest in the areas of automation. This is to be expected given the lower complexity of this type of functionality, which has been around for some time in various forms, but is becoming increasingly sophisticated. The other area where there has been a sharp rise in utilisation is in supporting staff efficiency, such as ambient voice technology which uses AI to listen to and summarise a consultation between a clinician and a patient.

The tables below outlines the areas where AI products have been identified in the HIN report as available for use (with diagnostic capabilities added as this was out of scope of the HIN work).

Category	**Area where AI products are available to support care provision and support functions**
Corporate	• Communications • Corporate governance and legal • Estates and facilities • Finance • Human resources and training • Procurement
Administrative Operations	• Booking and scheduling • Logistics and supply management • Performance management and reporting • Quality, safety and clinical governance • Staff rostering and deployment • Smart automation

Use Cases

Category	Area where AI products are available to support care provision and support functions
Clinical Operations	• Clinical interactions (automated or simulated) • Medicines management • Clinical decision support • Staff efficiency, including ambient voice technology
Diagnostics	• Medical imaging • Digital pathology • Other imaging modalities

Readiness for adoption of AI products is highly varied, and the following factors should be considered when assessing where to focus initial efforts when implementing AI:

1 Start with the problem. Is the AI product really the best answer to the problem you are trying to solve or are you just listening to vendor hype?

2 How complex is the AI model? Consider where the type of AI sits on the complexity diagram in Figure 1. The more complex the model, the greater the consideration required for ethical, legal and privacy considerations, as well as the approach to ensuring model outputs initially and on an ongoing basis.

For models that are more complex, there is also a greater risk of error or bias. This can occur initially and is often picked up in piloting phases, but can also occur over time with models degenerating or collapsing. While the response to this is often that the final decision on the output rests with the user of the system (the clinician or administrative staff member), this is not that simple as confirmation bias could impact on the information available to the staff member..

3 How much integration is required with other clinical and / or administrative systems? Implementing a stand alone product could work in the proof of concept phase, but it is critical that products are integrated into the work of your organisation to improve consistency of uptake and reduce risk of loss of structured data collection. If a highly integrated solution is required, ensure that you account for this in your pilot planning.

4 How readily available is the infrastructure required and the data source? Many products are vendor-specific, which could mean it is faster to adopt but does not then allow for expansion of the model to other settings or easy utilisation of other data sources into the model. It also has impacts from a commercial point of view, as many vendors offer cheap trials and the ongoing costs may not be clear. When balancing the decision between vendor-specific and locally-supported models, try and consider the total cost of ownership over a period of years, not months.

5 Are there products in use already in the NHS? One of the best ways to determine if an AI solution is valuable is to gather information from peers as to the benefits they have achieved when using a particular solution (or product category). It is important that each provider using an AI product or solution has a clearly defined benefits plan and captures information on benefits. This information can then be shared with other providers to support scaling, or support a decision to focus on other areas as a priority.

The intention will be that health and care organisations prioritise use cases based on their organisational need. This would be discussed at the ICS AI Advisory and Coordination Group to reduce effort between organisations by seeking to identify a lead organisation for each priority use case.

6. AI delivery approach

6.1 Delivery lifecycle

There is value in partners across our system adopting a similar approach to AI implementation so that we can have a consistent approach to communication and reporting. It will also assist other organisations to understand if, when and how they can use the information from pilot projects to use successful tools in their own organisation. However, it is acknowledged that there will be a need to align with local organisational processes for other digital projects and products. Therefore, where another approach is taken, mapping to the approach outlined in this Framework is encouraged.

It is proposed that there are four implementation phases, followed by ongoing monitoring and evaluation. These are outlined below.

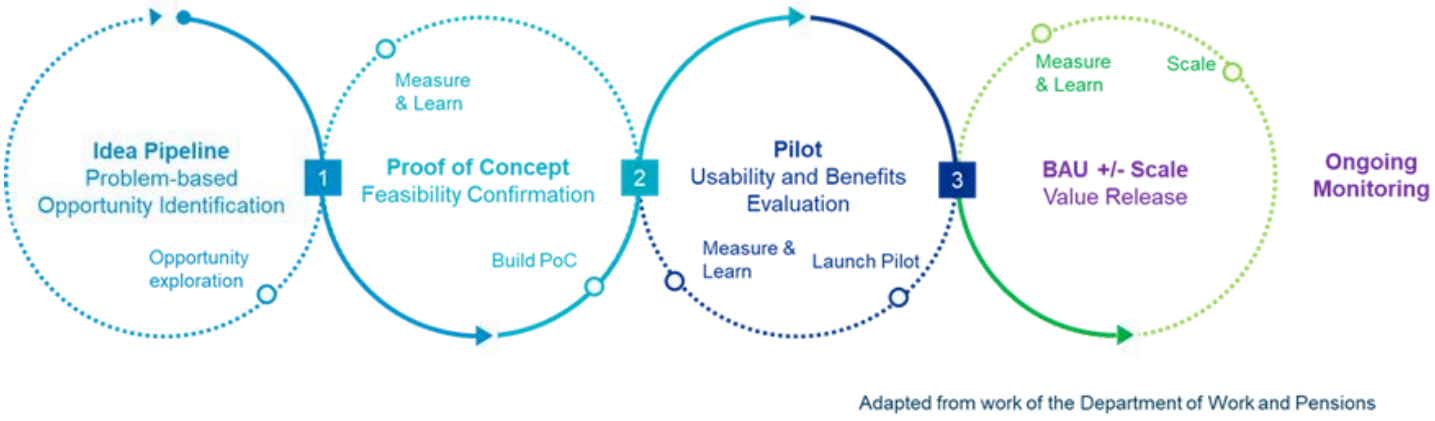

Adapted from work of the Department of Work and Pensions

The purpose of, and approach to each of the stages is outlined in the following sections.

Note that this approach does not outline an intention to have a central fund for AI projects and products at the London or ICS level. Projects would be taken forward by each organisation as they would with any digital project, but with an intention to work collaboratively across the region to reduce duplication, consider opportunities for scale and share learnings. Where there is value for all (or a large number of) systems across London to implement a product or address a use case, this may be considered in the prioritisation processes that are already in place when considering allocation of or proposals for national funding.

6.2 Idea pipeline

The objectives of this stage are to:

- collate a list of problems/demands for each provider; and
- prioritise problems at the provider and system level for action based around impact on operational and strategic objectives,
- identify potential product(s) for inclusion in a proof of concept,
- make a decision about whether to progress any ideas thorough to proof of concept,
- check with the AI Lead if any other projects have been undertaken in this use case,
- advise the AI Lead that your organisation will be the lead organisation for this use case.

The activities at this stage are likely to align with activities that your organisation has in place to respond to requests for digital tools or the escalation of problems that may have a digital solution (the problem definition stage of a user-centred design approach). There should not be a separate process for AI-driven problems at this stage, even though the requestor may come with a request to implement a specific AI solution.

When considering if AI is a potential solution, consider the following:

- Are AI products used for this problem already in NHS in London or nationally? – check the register of projects and products with your AI Lead.
- Are there vendors offering an AI solution to this problem for the health and care sector?
- Are AI products used for this problem in other sectors or internationally in health and care?

If the answer is no to the above, and you still think that AI may be an appropriate solution, consider whether an innovation partner approach may be appropriate. This is likely to require a longer lead time, but partnership with a vendor to develop a product may have benefits for both organisations. If this approach is being considered, it is suggested that seeking expert advice from your ICS AI Lead may be valuable.

When selecting a product to take to the proof of concept stage, ensure that you are considering the impact on potential future procurements. Often vendors offer test products free of charge. Make sure you are selecting the most appropriate product to meet your need rather than the most available product. Reach out to the AI Lead for support in product selection. In addition, your Health Innovation Network team may be able to support you with a market scan or a pre-market analysis if appropriate.

The key activities associated with this stage are outlined in the table below.

Clinical safety	Privacy	Ethical considerations	Explainability
Confirm that products being proposed for Proof of Concept have a DCB129 that can be provided. Check if the product is a medical device.	Consider the potential level of risk based around the AI models being considered.	Consider any ethical concerns specific to the AI models being considered.	Consider whether available products allow for explainability in the model; is there a human making the final decision and does the product give them the information required to do that?
Security	**Procurement**	**Communication**	**Implementation and Monitoring**
Consider any security concerns specific to the AI models being considered. Consider how access would be controlled if the product were scaled.	N/A Market analysis should be undertaken but this does not need to be comprehensive at this stage.	Ensure the digital team in your organisation are involved. Advise your AI Lead that you are going to be taking a product through to proof of concept and that you will be lead for that use case for the ICS.	N/A

6.3 Proof of concept

The primary objectives of this stage are to:

- define the requirements and benefits case.
- ensure that required documentation is approved prior to use including clinical safety assessments, data sharing agreements and data protection impact assessments as applicable.
- test whether solution(s) available on the market are sufficiently mature to address the problem, or whether solutions developed locally or in partnership with vendors are ready for piloting.

Note: When proposing the use of technology that is already in use in other NHS organisations, this stage may be based only on a desktop analysis of the results of others.

The intention is for this to be a fail fast and fail early stage, which is as simple as possible. Where possible, it is suggested that using test data (anonymised), and not implementing full integrations with other systems, but ensuring that the standards are met that would allow some integration to be implemented during a pilot. The outcome should be focused around whether the product actually delivers the expected outcome in a way that is high enough quality to progress to pilot. Using Ambient AI as an example, the proof of concept would check that the quality of the summary was valuable to the clinician.

Where there is a co-creation approach to the AI model development, this stage may be much more complex than when an off-the-shelf model is being considered. Prior to implementation, T.E.S.T Section A should be completed [NB: requires adaptation for usage beyond Ambient AI].

Proof of concept - key activities

Clinical safety	Privacy	Ethical considerations	Explainability
Identify a clinical safety lead and conduct the DCB160 – note this is required by legislation where the product is used in clinical care. If the product is a medical device, ensure that it is registered with the MHRA. Consider whether the product has been certified against British Standard 30440 – Validation Framework for the Use of AI in Healthcare.	Complete the DPIA and have it approved. This is required in all cases. Determine if a data sharing agreement is required, and ensure it is in place when it is required. Establish, document and implement the appropriate data consent controls (if not included in current consent agreements). Consider whether anonymised, psudonymised or test data can be used.	If the product uses machine learning, deep learning or generative AI, conduct the ethical reflection exercise at Appendix 2 of this Framework. This may be submitted to the AI Advisory and Coordination Group for a review meeting if the lead organisation chooses to do so.	Consider whether available product allows for explainability in the model. I.E is there a human making the final decision and does the product give them the information required to do that.

Proof of concept - key activities

Security	Procurement	Communication	Implementation and Monitoring
If you are integrating the product with another system or sharing clinical information or personal information with a third party, engage your Chief Information Security Officer (or system administrator) to discuss any cyber security risks. Actions may include: • Conduct preliminary vulnerability assessments. • Implement foundational security measures (e.g., data encryption monitoring, logging). • Ensure compliance with relevant security standards and regulations e.g. ISO 27001, Data Security and Protection Toolkit (DSPT).	Where possible, use test licences or products. Conduct the feasibility assessment on the basis of the product type, rather than the specific vendor's product. Consider the impact of vendor selection on future procurements. Consider inclusion of requirements to comply with BS30440. Include requirements relating to the supplier's approach to mitigating environmental impacts of their AI products.	Check with your AI Advisory and Coordination Organisation if this project would also be a priority of other organisations (suitable to scale). Advise the AI Advisory and Coordination Organisation of your project and the products being used so that it can be included on the register. Provide the AI Advisory and Coordination Organisation with the results of your evaluation.	Develop the evaluation plan. After use, evaluate whether the product provides the required functionality to an acceptable level of quality.

6.4 Pilot

If the product passes the evaluation at the Proof of Concept stage, you may choose to move to the Pilot stage. This will likely require additional investment. If you are unable to lead progress through to this stage, please advise the AI Lead in your ICS. They will then seek advice as to whether there is another organisation that is willing to lead a pilot, or whether this product/use case will be put on hold at this time.

The primary objectives of this stage are to:

- procure a solution to pilot, taking into account potential to scale
- ensure that required documentation is approved prior to use including clinical safety assessments, data sharing agreements and data protection impact assessments as applicable.
- mitigate risks associated with the product usage including through completing the ethical reflection exercise (Appendix 2).
- evaluate the pilot outcomes to determine if success criteria were met, or whether the pilot should be ceased and the product removed from use.
- identify any risks and issues associated with implementation.
- identify the change management activities required for successful implementation.

Note: If a product is already in use in the health and care sector in the UK, this stage may involve a desktop review of the evaluation results of other organisations using the product. If sufficient information is available to inform development of a business case for scale, there may be no need for a project at this stage, although confirming benefits realisation locally may be appropriate.

It is critical in this stage that the pilot is undertaken with a clear objective of identifying whether the solution is fit for purpose and will progress to retention as business as usual and potentially scale (subject to business case definition), or whether it did not achieve the intended outcomes and the pilot should be ceased. One of the challenges faced in many organisations is that pilots do not have clearly defined end points and benefits criteria to allow for decisions about whether to continue with the use of the product/solution or whether to cease the pilot. This leads to retention of solutions that may not offer best value for money, and actions at this stage should address this challenge.

The actions at this stage should be conducted to facilitate long-term usage and scale of the product if appropriate. This means that there is a need to conduct procurement of the product in a way that supports this. Please seek advice from your AI Lead about what to include in procurement documentation to allow for use by other providers if appropriate.

Prior to implementation, T.E.S.T Section B should be completed following implementation and prior to the end of the pilot [NB: requires adaptation for usage beyond Ambient AI].

The key activities associated with this stage are outlined in the table below.

Pilot

Clinical safety	Privacy	Ethical considerations	Explainability
Update the DCB160 and ensure it is approved by the trained clinical safety lead – note this is required by legislation where the product is used in clinical care. If the product is a medical device, ensure that it is registered with the MHRA. Consider whether the product has been certified against British Standard 30440 – Validation Framework for the Use of AI in Healthcare. You may consider requiring this in your procurement.	Update the DPIA. Update the data sharing agreement as applicable.	Update the ethical reflection exercise at Appendix 2 of this Framework.	Ensure as part of the evaluation that the decision-makers are comfortable that the rationale for recommendations is appropriate, and informs their final decisions.

Pilot

Security	Procurement	Communication	Implementation and Monitoring
Update or undertake the security activities listed in the PoC stage.	Procure the solution in a way that will enable scale which may include to other work areas in your own organisation or to other organisations. At this stage, it is likely you will need to approach the market, either as an open procurement or through a framework agreement. Consider inclusion of requirements to comply with BS30440. Include requirements relating to the supplier's approach to mitigating environmental impacts of their AI products.	Advise your AI Advisory and Coordination Lead Organisation of your project so that it can be included on the register. Report on the outcomes of the Proof of Concept to the AI Coordination and Advisory Organisation, even if the decision was to withdraw the product.	Identify the intended benefits and pilot plan, including pilot end date. Assess benefits realisation and determine whether to proceed to implementation as BAU and/or scale, or whether to withdraw the product. Develop the evaluation plan and conduct the evaluation. Ensure that this includes sufficient information on costs and benefits to inform a business case for scale. If the solution uses machine learning, deep learning or generative AI, monitor the use of the AI solution, ensuring that the outputs are appropriate.

6.5 Business as usual and scale

Note: This stage will require significant tailoring depending on the circumstances. However, the below outlines a typical approach that may be taken.

Following successful completion of a pilot, and where the pilot demonstrates that use of an AI product provides value for money, decisions need to be made about whether to implement the product into business-as-usual workflows, and also to consider whether and how to scale implementation.

In relation to implementation as business as usual, there may be a need for development of a business case and approvals. In this case, it may be appropriate to continue the pilot until this process is finalised. If the business case to retain the product/solution as business as usual is not approved, the pilot should be ceased. Organisations may also want to consider scaling within their organisation to areas not included in the pilot. This is appropriate and should be included in the business case, although it may be necessary to undertake proof of concept steps and undertake the appropriate change management and benefits realisation activities.

To support decisions about scale to other organisations, the AI Advisory and Coordination Lead Organisation will circulate information about the use cases and the cost/benefit case to the London ICB Chief Information Officers (CIOs) and Chief Clinical Information Officers (CCIOs) that the product may be relevant to. CIOs will coordinate consultation with relevant health and care providers in their system to confirm the problem being solved by the AI solution is a priority problem for their organisation. Where the answer is yes, they will also seek to identify potential funding sources for any investment required to implement the solution.

A London-wide business case and procurement may be considered if all ICBs are interested in pursuing implementation of the AI product/solution, and advice may be sought from the London Digital Transformation Portfolio Board.

The key activities associated with this stage are outlined in the table below.

Business as usual and scale

Clinical safety	Privacy	Ethical considerations	Explainability
All implementing institutions to identify a clinical safety office, and to review the DCB160 provided by the piloting organisation, make any adjustments required to reflect local circumstances and approve – note this is required by legislation where the product is used in clinical care. If there are any concerns with the analysis conducted by the piloting organisation, there should be a discussion with this organisation and it should be taken to the AI Advisory and Coordination Group as required (noting that unwarranted variations in assessment of risk can generate risk).	All implementing institutions to review the DPIA provided by the piloting organisation, make any adjustments required to reflect local circumstances and approve. If there are any concerns with the analysis conducted by the piloting organisation, there should be a discussion with this organisation and it should be taken to the AI Advisory and Coordination Group as required (noting that unwarranted variations in assessment of risk can generate risk).	All implementing organisations to review the ethical reflection exercise at Appendix 2 of this Framework conducted by the lead organisation.	Ensure as part of the evaluation that the decision-makers are comfortable that the rationale for recommendations is appropriate, and informs their final decisions.

Business as usual and scale

Security	Procurement	Communication	Implementation and Monitoring
Update or undertake the security activities listed in the Proof of Concept stage.	Where possible, use the procurement undertaken by a lead organisation on behalf of London. Consistency in the use of products is important for reducing the burden of monitoring and of integrations.	Advise your AI Advisory and Coordination Lead Organisation of your use of the product.	If the solution uses machine learning, deep learning or generative AI, monitor the use of the AI solution, ensuring that the outputs are appropriate. Review the benefits realisation and whether expected benefits have been achieved.

6.6 Monitoring

AI products that use machine learning, deep learning and generative AI require ongoing monitoring to ensure continued appropriateness of their outputs. This is required as performance of AI models drift or degrade over time. The full AI lifecycle requires model hubs for model management, bias detection and dataset drift. Most healthcare AI vendors provide limited model hubs serving only their own products, which limits their scalability at a use-case level, provider level or ICS level. It is expected that medical device regulation will eventually require such model hubs. Each project should consider the monitoring requirements for the products being proposed, and should seek expert advice through the ICS Advisory and Coordination Group as required.

For each AI product in use in London, a monitoring plan must be developed. This plan may have a lead organisation who is responsible for the majority of monitoring of that product for our system, but each organisation will also need to confirm the analysis wherever there is local data that the algorithms are trained on, or local implementation considerations.

Monitoring must consider:

- Data monitoring
- Model monitoring
- Infrastructure modelling

Modelling should consider both the ongoing quality of model outputs and cyber security threats.

Each project should consider the monitoring requirements for the products being proposed, and should seek expert advice through the ICS Advisory and Coordination Group as required.

Organisations should advise the AI Coordination and Advisory Lead Organisation of any new risks or issues identified so that they can track them and advise others that may be using that product.

7. Communication and Engagement

7.1 With our workforce

Engaging our wider workforce is critical to ensuring the successful adoption of AI in NHS. Our staff, spanning diverse roles and responsibilities, bring invaluable perspectives on how these technologies can enhance communication, collaboration, and service delivery. By involving staff from various disciplines, we can foster a sense of ownership and trust, which is essential for the sustainable and ethical integration of AI into NHS practices.

Feedback from our workforce will be instrumental in refining and improving our AI offering. By listening to the experiences, concerns, and suggestions of staff across all areas, we can ensure that our approach addresses real-world challenges, maximises opportunities, and mitigates potential risks. This iterative process reflects our commitment to a collaborative and responsive strategy that supports all NHS staff and aligns with the values and operational priorities of the organisation.

Workforce engagement approach

A timetable and plan for workforce engagement will be developed to ensure staff across the NHS are informed, empowered, and able to shape the use of AI in their work. This may include a mix of webinars, interactive surveys, and hands-on workshops to maximise accessibility and impact.

The approach will follow three phases to build ownership and trust across our system. These phases will be designed to build awareness, interest, and desire, with the aim of enabling staff to actively engage with AI projects and products if and when they are introduced.

Phase 1 - Awareness

This first phase will provide a starting point for staff to engage with AI.

Aims:

- to increase awareness of the use of AI in the NHS and the need for robust governance to manage risk for our staff, people, and communities.
- to gather staff views on AI, benefits, risks, opportunities, ethics and training and development needs etc.
- to capture the outputs of this engagement phase to help shape next stages of engagement.

Phase 2 - Interest

The second phase will focus on supporting our workforce to develop an understanding of how AI could operate with their organisations.

Aims:

- to increase willingness to engage with AI and to demonstrate the value of using AI tools.
- to raise awareness of this AI Framework, and to communicate their responsibility in following process in relation to AI implementation, to ensure that risks are mitigated, and that implementation is appropriate.
- to gather information about potential uses of AI tools identified by colleagues across the system and to gather feedback on how these could be applied to each organisation.
- to begin to create an appetite for change.

Phase 3 - Desire

The third phase will build a deeper understanding in our workforce of how AI could operate within their organisations to help their teams become more efficient. It will look specifically at overcoming barriers to change around adopting tools from other parts of the NHS.

Aims:

- to co-produce an action plan to address challenges raised in previous phases.
- to provide advice to the AI Advisory and Coordination Lead Organisation about key areas where AI and automation could have the biggest and most meaningful impact in a safe and considered way.
- to identify the key areas of communications and engagement support needed to help organisations succeed in the successful adoption of AI and automation ideas.
- to identify the best way of building trust with the people we serve in our use of AI and automation in supporting their needs more effectively.
- to co-produce a short guide for managing AI and automation requests in line with the AI Framework.

Details of the approach will be developed at London level and led by ICS communication and engagement teams, working collaboratively with digital, clinical and operational stakeholders to ensure appropriateness of content and approach.

7.2 With our community

Alongside our communication and engagement with our workforce about the use of AI in healthcare, it is important that we talk with the people and communities we serve across London.

A plan to support engagement with people and communities will be developed to ensure they feel informed, empowered, and able to be involved in plans in relation to the use of AI in their health and care.

It is likely that there will be differing perspectives within our communities about the use of AI. For example, some may see the use of AI as a simple extension of other tools available to their health and care team to support them in making decisions about the best care for the individual. Others may be more concerned about the impact on their privacy and the security of their personal information.

It is also likely that the level of engagement and awareness within communities about the use of AI in their health and care will differ depending on the type of product being used and how it is being used. For example, if the AI is being used in the background to assist with diagnostic tools, some people may feel that they need less engagement than if they are interacting more directly with the AI (for example, in ambient voice technology or chatbots).

The community engagement approach will seek to draw out these perspectives and will enable the development and implementation of targeted communication and engagement approaches. Given the complexity of the topic, one of the tools that may be adopted is a public deliberative event, which allows our community to come together to learn about a topic, discuss it, and consider evidence to come to a conclusion. This enables involving the public in a complex issue in a meaningful way.

8. Training

The implementation of AI has the potential to transform work practices. Both the specialist Digital, Data and Technology (DDaT) workforce and the wider workforce will need skills to implement and adopt AI safety and effectively.

For the DDaT workforce and others interested in the technical side of AI implementation, there is a range of opportunities on offer. These include the below (but each ICB may have a link to other opportunities on their websites):

- The Kings Health Partners (KHP) Learning Hub provides a variety of free modules training in Data Science and AI. These are suitable to healthcare professionals who wish to know more and start their journey to learn a variety of data science technologies (https://learninghub.kingshealthpartners.org/catalog?pagename=Innovation_Scholars)

- NHS Fellowship in Clinical AI is, a national programme where participants are embedded in NHS Trusts undertaking projects as part of their learning (https://gstt-csc.github.io/fellowship.html).

- The EPSRC Centre for Doctoral Training in Data-Driven Health (DRIVE-Health) offers training to develop the next generation of health data scientists, computer scientists, software developers, and AI and machine learning researchers – there are around 100 funded projects per year (https://www.drive-health.org.uk/).

In addition, the establishment of the communities of practice by the Lead AI Coordination and Advisory Organisation aims to provide support and mentoring to our workforce in the implementation of AI.

For the broader workforce:

- Many trusts are establishing apprenticeship programmes which are available to staff who want to be trained in health analytics – from basic levels to AI.

- Additional training needs and opportunities need to be considered when developing the DDaT workforce plan for London.

9. Conclusion

This AI Framework seeks to provide a practical guide for AI implementation and monitoring, as well as to outline an approach to facilitating the sharing of information and best practice across London. The intention is to reduce duplication of effort and share expertise so that we are well-placed to make optimal use of AI tools and technologies in a safe, efficient and effective way.

Due to the rapidly-changing nature of the AI landscape, this Framework will be a live document that is updated as we learn from its implementation and as new regulations and processes emerge.

Appendix 1: definitions

Note: Not all of these terms are used in this Framework, but they may be used by people in their evaluations of the products. They have been included here to ensure that all people using this Framework are talking the same language.

AI Term	Definition
Algorithm	A set of rules a computer follows to make decisions or solve problems. In healthcare, algorithms analyse data to help make diagnoses or predict outcomes.
Anomaly Detection	Identifying unusual patterns, like spotting rare side effects in treatment data or abnormal lab results.
Bias	When AI makes unfair predictions due to either inappropriate training data or inappropriate interpretation of data, potentially impacting patient care or diagnosis accuracy.
Computer Vision	AI focused on interpreting visual data, such as medical imaging, to help identify diseases or abnormalities.
Data Mining	Extracting useful patterns or knowledge from large data sets, such as finding trends in patient symptoms or treatment responses.
Deep Learning	An advanced type of Machine Learning that uses neural networks with many layers, often used for analyzing medical images or detecting patterns in complex data.

AI Term	Definition
Data Privacy	Protecting patients' personal information in AI applications, following regulations like Data Protection Act to ensure data security.
Electronic Patient Record (EPR)	Digital version of a patient's medical history that can be analysed with AI to improve patient care.
Explainability	How easily humans can understand AI's decision-making process, critical in healthcare for trust and transparency.
Interoperability	The ability of different healthcare systems to share and understand each other's data, crucial for AI to access comprehensive patient information.
Integration	The connection between one platform (eg the EPR) with other digital systems (eg AI solutions) to share data. This is important for surfacing the AI outputs to the system used by the health and care team.
Machine Learning (ML)	A type of AI that learns from data to find patterns and make predictions or decisions. In healthcare, ML is used for risk assessments, imaging analysis, and disease prediction.

AI Term	Definition
Natural Language Processing (NLP)	AI that helps computers understand and process human language, like interpreting doctors' notes or patient feedback.
Neural Network	A model inspired by the human brain, made of "neurons" that help AI learn patterns in data, like spotting anomalies in x-rays or MRIs.
Predictive Analytics	Uses data to predict future health outcomes, like hospital readmissions or disease progression.
Reinforcement Learning	An AI learns through trial and error to make decisions, often applied in complex treatment planning, like adjusting radiation doses.
Supervised Learning	Training AI using labeled data (with known outcomes) to make predictions, such as identifying diseases from annotated images.
Training Data	The data used to teach AI models; in healthcare, this might be images, patient records, or test results.
Unsupervised Learning	Training AI on unlabeled data to find hidden patterns, like grouping patients with similar symptoms but no clear diagnosis.

Appendix 2: AI Ethical Reflection Template

This template has been adapted from the Ada Lovelace Institute - Algorithmic Impact Assessment

High-level project information

1 Your organisation name

2 The ICS(s) that your organisation works in

3 If you are not an NHS organisation, please describe your organisation purpose, size and role in the health and care sector

4 Describe the purpose of your project. What problems is it trying to solve and how. This should be a concise summary, of no more than 250 words. You can write this in the form of an abstract for a paper. Assume your audience doesn't have much technical knowledge – perhaps you are explaining this to a stranger.

5 Describe the intended outputs of your project.

6 Who are the stakeholders who will be affected by this system? For instance, who are the intended users and who will it serve? Try to be as specific as possible when listing these out e.g. clinicians, nurses, hospital administration staff, patients of a particular kind, etc.? Make sure to explicitly articulate the affected population and differentiate, if necessary, from the set of users of the system.

7 List the project outcomes.

8 List the intended benefits (at a high level).

Common ethical considerations
This section guides you through specific ethical considerations that are common in the context of healthcare, AI and the algorithmic literature.

9 Could this project lead to the creation or exacerbation of inequalities or unlawful discrimination against particular communities? For example, through worsening differential access to care? What might your current plans overlook for evaluating or monitoring bias and fairness?

10 How does your project consider consent and autonomy? Are there risks related to increased surveillance? For example, how will the intended beneficiaries of the system be informed about its use? Could the system be interpreted as increasing surveillance?

11 What kinds of environmental impacts will this project have? What amount of compute and energy will be required to train and run this system? Are there other environmental impacts of software, hardware or equipment that would occur as a result of the system?

12 How might the use of this system impact relationships between patients and health and care professionals? Could this system make some patients or service users less likely to seek care, or be less frank in discussion with health and care professionals?

13 How will individuals using, or affected by, this system be able to contest its findings? How can they appeal a result? Are there options to refuse use of the system? How and to whom will the system's outputs be made explainable and interpretable?

14 How could this system be unintentionally or intentionally misused? What are the ways in which this system could lead to an accident or error? Could the outputs of the use of health and social care data be manipulated to serve purposes that are not in the public interest?

Impact identification and scenarios
This section is for reflecting on broader potential impacts of your system in implementation.

With the above questions in mind, provide a plain-language summary of the best and worst-case scenarios that could arise from use of this system after it has been deployed, what socio-environmental conditions are required for success and what are the expected hurdles or challenges to overcome?

If technical concepts must be used here, ensure they are explained carefully.

When thinking about affected stakeholders, consider both direct stakeholders (such as clinicians, patients, other users of the technology in its intended setting) and indirect stakeholders (such as certain identity groups, regulators, civil society, the general public).

15 What is the best-case scenario that could arise from use of this system? Discuss when the system works as designed/intended, but also how failures, errors, mistakes or unexpected behaviours would be handled.

16 What kind of socio-environmental requirements are necessary for the success of this system in operation? E.g. stable connection to the internet, training for doctors and nurses, collaboration between particular clinical and administration staff etc.
In answering this question, consider which stakeholders will use this system, how they would optimally interact or work together for the system to succeed, how information would be shared (and with whom), and what social, technical and workflow dependencies may need to exist. You might also consider what kinds of infrastructure stakeholders will need to use this system successfully.

17 What are likely challenges/hurdles to achieving the best-case scenario?

18 What is the worst-case scenario that could arise from use of this system?
18a When the system works as designed/intended.
18b When the system fails or doesn't work as designed/intended in some way.

Potential harm analysis
This section asks project teams to list the potential harms and benefits for different stakeholders that arise across all scenarios described in point three above.

It then asks teams to make an assessment of what, in your view, is the perceived importance, urgency, difficulty and detectability of each harm.

Finally, this section asks project teams to consider potential mitigations for these harms, such as specific design decisions that will reduce the potential for this impact to cause harm. This helps teams consider how harms might be distributed and ascertain which harms should be prioritised, based on their potential immediacy or irremediability.

19 Based on the scenarios identified above, what are the potential harms resulting from the implementation of this system that your team will need to actively design for? Who is most at risk of being harmed, and how?

When thinking about harms, consider this from the perspective of a clinician, patient or someone else who would be affected by use of this technology.

For each identified harm, please make note of the following considerations:

- Importance – how consequential is this harm for the wellbeing of stakeholders? Which are irremediable and serious?
- Urgency – how immediate is this threat?
- Difficulty – how difficult will it be to mitigate this harm?
- Detectability – how perceivable is this harm given the current design?

20 Based on the scenarios identified above, what mitigations could you put in place to minimise these harms?

Appendix 3: Contributors to the development of this Framework

This Framework started as a framework for the South East London Integrated Care System and was expanded in scope following consultation with the other four London ICSs.

The Framework development was informed by a series of three workshops hosted by the South East London ICB and the AI Centre for Value Based Healthcare.

The following people were involved in the workshops and contributed to the development of this Framework. Many others subsequently reviewed the Framework and provided input to enable its finalisation and it is not possible to list them all, but all were valuable in development of this Framework.

- Philippa Kirkpatrick, Chief Digital Information Officer, NHS South East London (co-lead)
- Sigal Hachlili, Director of AI, Data & Digital Innovation, The Artificial Intelligence Centre for Value Based Healthcare (co-lead)
- Kirsty Ayton, Project Officer, NHS South East London
- Jack Barker, Clinical Lead for Informatics, NHS South East London
- Alex Drake, Associate Director of Performance, Lewisham and Greenwich, NHS Trust
- Fiona Howgego, System sustainability lead, NHS South East London
- Ranjeet Kaile, Director of Communications and Engagement, NHS South East London
- Zoe Keddie, Chief Information Officer, South London and Maudsley NHS Foundation Trust
- Denis Lafitte, Chief Digital Information Officer, King's College Hospital NHS Foundation Trust, Guy's and St Thomas' NHS Foundation Trust
- Meera Nair, Chief People Officer, Lewisham and Greenwich, NHS Trust
- Seb Ourselin, Director, London Institute for Healthcare Engineering, The Artificial Intelligence Centre for Value Based Healthcare
- Vijay Sivipalan, Primary Care Chief Clinical Information Officer, NHS South East London
- Jean Straus, Patient representative, Member of the KHP Digital Health Hub Patient and Public Involvement and Engagement Group
- Lawrance Tallon, Deputy CEO, Guy's and St Thomas' NHS Foundation Trust
- James Teo, Clinical Director of AI and Data, Kings College Hospital NHS Foundation Trust
- James Woollard, Chief Clinical Information Officer, Oxleas NHS Foundation Trust

We would also like to thank the insights provided by those that presented to our workshops and engaged in conversations including:

- Hatim Abdulhussein, Chief Executive Officer, Health Innovation Kent Surrey Sussex
- Brhmie Balaram, Head of Responsible AI Adoption, NHS England
- Jacob Cronin, Deputy Director, Transformation Portfolio, Department of Work and Pensions
- Aleksandra Foy, NLP Workstream Lead Cogstack, Guy's and St Thomas' NHS Foundation Trust
- Alice Morrisey, Head of Digital Transformation, Centre for Innovation and Transformation
- Donald Taylor, Chief Data Scientist, Department of Work and Pensions
- George Verghese, Primary Care Lead, NHS South East London
- Sumit Wadia, Group Head of Public Affairs and Building Your Future Hospitals, St George's, Epsom and St Helier University Hospitals and Health Group
- Rodney Young, Director of Business and Artificial Intelligence, South West London ICB
- Joe Zhang, Head of Data Science, AI Centre for Value Based Healthcare

ONE LONDON

AI CENTRE for Value Based Healthcare

For any questions, please contact: digital@selondonics.nhs.uk

North West London Integrated Care System
Working together for better health and care

North Central London Health and Care Integrated Care System

NHS North East London

South West London Integrated Care System

South East London Integrated Care System