with you, for you

Gloucestershire Health and Care
NHS Foundation Trust

Artificial Intelligence (AI) Usage Policy

Policy Number	AI - 01
Version:	1.1
Purpose:	*To provide the Trust with a policy position on the use of AI's to ensure consistency, standardisation, provide clear rules and responsibilities.*
Consultation:	*Digital Services, Digital Group, Clinical systems, IG, IG Group, Ops, Staff side.*
Approved by:	*The Digital Group & Director of Finance*
Date approved:	*23 March 2026*
Author(s):	[redacted], *Paul Griffith-Williams*
Date issued:	*24 Apr 2026*
Review date:	Apr 2027
Audience:	*All staff.*
Dissemination:	*The Policy will be publicised through the intranet and communications.*
Impact assessments:	*Equality Impact Assessment.*

Version History (*The version history should be updated each time the document is revised)*

Version	Date	Reason for Change
1	*Dec 25*	*Initial Draft*
1.1	Mar 26	Final version

Part 1 - Summary

SUMMARY

In the rapidly evolving field of Artificial Intelligence (AI) technologies there is the potential for it to revolutionise patient care, streamline administrative processes, and enhance overall operational efficiencies in Health and Care organisations. Gloucestershire Health and Care NHS Foundation Trust (GHC) recognises the importance of embracing AI. The Trust also recognises the current drive for embracing AI technology, by the Department of Health and Social Care, along with the moral imperative to explore such technology where there are clear benefits to patient care and outcomes. That however must be within a framework that ensures their safe, ethical, lawful and responsible use. This AI policy serves as a guiding framework to ensure the appropriate deployment, management, and oversight of AI systems across the Trust.

TABLE OF CONTENTS

PART 2 - Detail

1. INTRODUCTION

1.1. GHC recognises the potential that AI can play in enabling it to meet its core responsibilities, aims and objectives, in a more efficient and cost-effective way. GHC also recognises the drive to embrace such technologies, and the perceived moral imperative to embrace such technology where there is clear benefit to the care we provide. It is important, however, that any use is in a safe, ethical and lawful manner. It must recognise data subject rights, especially around automated processing, and manage any potential risks associated with its use.

1.2. The use of AI will need to follow an approval process, in each instance, Information Governance (IG) and Digital Services (DS) will maintain a list of approved AI technologies.

2. POURPOSE

2.1. The purpose of this policy is to establish clear guidelines for the Trust around development, implementation, and monitoring of AI systems, that will protect personal data, uphold ethical standards, and mitigate bias and other potential risks. It ensures compliance with UK data protection laws, promotes transparency, and safeguards patient rights.

2.2. GHC recognise that AI technology, including machine learning algorithms and natural language processing, can contribute significantly to research, improving healthcare outcomes and resource allocation. It is, however, essential to ensure that AI technologies are used in a manner that aligns with legal requirements, respects patients' rights, and maintains the trust and confidence of our patients, colleagues, and stakeholders.

2.3. This policy outlines key principles and procedures that must be adhered to when utilising AI technologies within the Trust. It addresses critical aspects such as data privacy, algorithm transparency, accountability, and ongoing monitoring of AI systems. By implementing these guidelines, The Trust aims to develop a culture of responsible AI technology use, where the benefits of AI are harnessed while minimising potential risks.

2.4. It is important to note that this AI policy is not exhaustive, AI technology is ever advancing and therefore may need to be adapted and updated periodically. Additionally, regulatory requirements alongside best practices in AI governance are evolving. The Trust is committed to responsible AI technology implementation.

3. SCOPE

3.1. This policy applies to all colleagues, contractors, and partners involved in the development, deployment, or use of AI technologies in health and care services for the Trust, across clinical decision support, diagnostics, administrative automation - clinical or nonclinical, and patient engagement tools.

3.2. This policy applies to any use of AI for GHC, whether it is patient facing, clinical support services, corporate services, or any work that is carried out in the interest of Trust business.

4. DUTIES

4.1. Senior Information Risk Owner (SIRO):

- Take responsibility for the overall governance and management of information risks associated with AI systems.
- Approve this AI Policy, supported by relevant technical and subject matter experts and approving groups within the Trust.
- Ensure that appropriate risk management processes, controls, and policies are in place.
- Ensure collaboration between stakeholders to address potential risks and mitigate any adverse impacts arising from AI implementation.
- Provide oversight and strategic direction to ensure the responsible use of AI technologies across the Trust.

4.2. Medical Director & Caldicott Guardian:

- Ensure data is processed in accordance with the Caldicott Principles.
- Ensure confidential patient information is processed legally, ethically and appropriately.
- Provide advice and guidance to staff on the appropriate implementation of AI with clinical data.
- Ensure that use of AI by medics complies with this policy.

4.3. Trust Directors:

- Ensure that AI use with their directorate is in line with this policy.
- Ensure that AI is only used where it has been formally approved.
- Ensure AI use is necessary, proportionate, limited and a suitable Asset Owner is identified.

4.4. Data Protection Officer:

- Oversee and ensure compliance with relevant UK data protection legislation best practice associated with AI.
- Provide guidance on data privacy related to AI systems.

- Review & approve Data Protection Impact Assessments (DPIAs), for AI projects, products, to progress to the IG Group for formal approval.
- Serve as the point of contact for data subjects and supervisory authorities regarding data protection concerns related to AI.

4.5. Head and Deputy Head of IG and Records:

- Review and support the development of DPIAs for all AI projects or products.
- Serve as a point of contact for staff with queries or concerns relating to AI.
- Provide guidance relating to data protection and AI.
- Ensure the implementation of AI in the Trust is in line with data protection legislation, and this policy.

4.6. Head of Patient Safety:

- Oversee and ensure AI use is compliant with the Trusts clinical risk practices.
- Ensure the safety risks relating to AI projects or systems are properly assessed and documented.
- Ensure that any issues, incidents or concerns relating to clinical safety of the AI projects or systems are investigated.
- Escalate any concerns, issues or incidents relating to clinical safety of the AI projects or systems as appropriate.

4.7. Clinical Safety Officer (CSO):

- Assess the safety risks associated with AI projects or systems used in clinical settings.
- Collaborate with relevant stakeholders to establish safety protocols and guidelines for AI projects or systems implementation.
- Monitor and evaluate the performance and safety of AI projects or systems.
- Investigate and address any incidents or concerns related to the clinical safety of AI projects or systems that arise and escalate where needed.

4.8. Chief Clinical Information Officer (CCIO):

- Ensure data is processed is necessary and proportionate in accordance with the Caldicott Principles.
- Ensure confidential patient information is processed ethically, safely and appropriately.

- Provide advice and guidance to staff or team on the appropriate implementation of AI in a clinical setting.

- Support the Clinical Safety Officer, with reviews of AI for clinical use.

4.9. Head of Human Resources & Organisational Development:

- Ensure support is in place for the lead manager in any change management programmes relating to AI including undertaking a people impact assessment, staff/ trade union engagement or consultation.

4.10. Associate Director of OD, Learning & Development:

- Develop the workforce to have the skills and confidence they need to make the most of digital services and improve care, through access and signposting to appropriate training programmes.

4.11. Associate Director of IT:

- Ensure there is suitable assistance for the implementation, integration, and maintenance of AI systems.

- Ensure that there is proper configuration, security, and compatibility of AI systems with existing IT infrastructure.

- Ensure IT collaborates with vendors and other stakeholders to address technical issues and provide technical support for AI systems as required.

- Ensure suitable digital guidance for staff looking to utilise or implement AI in the Trust.

- Ensure all AI for the Trust progress through the Trust Digital group, before implementation.

4.12. Head of IT Operations:

- Support the implementation, integration, and maintenance of AI systems.

- Support the proper configuration, security, and compatibility of AI systems with existing IT infrastructure.

- Collaborate with vendors and other stakeholders to address technical issues and provide technical support for AI systems as required.

- Ensure all AI have progress through the Trust's Digital group, before implementation.

4.13. Business Intelligence:

Where there are linkages to BI systems or processes (such as the Trust's data warehouse), BI services will:

- Assist in the implementation, integration, and maintenance of AI systems.
- Support in the proper configuration, and security of AI systems integration.
- Collaborate with vendors and other stakeholders to address technical issues and provide technical support for AI systems, as required.

4.14. Information Governance Group:

- Review and support approval of this AI Policy, as and when necessary, given the fast moving nature of the technology.
- Review and approve suitable guidance for staff looking to implement or utilise AI in the Trust.
- Review and approve all DPIAs relating to AI projects and systems being implemented by the Trust.

4.15. Information Asset Owners:

- Accountable for following proper purchasing process in line with Appendix 1.
- Ensure that a DPIA and, where appropriate, a DTAC is completed.
- Ensure there is a suitable implementation plan and training material in place for users.
- Link in with the Trust CSO and CCIO, where appropriate.
- Ensure there are audit processes in place for monitoring outputs.
- Ensure there are suitable controls for the access and use of the AI system.
- Report any concerns arising from the AI use.

4.16. Digital Group:

- Review and approval proposals for AI projects or systems for the Trust.
- Review and contribute to suitable guidance for staff looking to implement or utilise AI in the Trust.

4.17. All staff:

- Familiarise themselves with the Trust's AI policy and guidance (Appendix 3) ensure that any use of AI complies with this policy and any national guidelines or protocols.

- Must not use any AI, including freely available AI, including mobile devices, for Trust business that has not been formally approved by the Trust.

- Report any concerns or issues related to the AI systems in use to the Trust's IG Team.

- Be aware of what AI is in use on your own personal mobile devices and ensure they are disabled when providing care.

- Be aware of what AI patients or relatives may have in use on their own devices.

5. AI IN DETAIL

5.1. There are four types of AI, Narrow Artificial Intelligence (nAI), Generative Artificial Intelligence (GenAI) and Machine Learning Artificial Intelligence (MLAI). Agentic Artificial Intelligence (AgAI). All of which are more generically captured under the term Artificial Intelligence.

5.2. AI can use statistics and information to generate connections and patterns between words, concepts and ideas. It uses probability to produce a response based on data patterns it has seen during training and remembered from previous interactions. AI has many uses, including summarisation, creative writing, educational support, content generation, pattern recognition, enhanced search, translation and customer services.

5.3. **AI in General:** In general, AI could be used in the Trust in three ways:

- Specifically for delivery of health and care.

- For administrative functions associated with delivery of health and care or corporate services.

- For health research.

5.4. **Benefits:** AI can bring some significant benefits for the Trust by supporting it to meet its essential functions and long-term plans, as it can:

- automate routine tasks and processes, such as document processing, minutes, customer service, and fraud detection, freeing up staff to focus on higher value activities and complex issues.

- enhance creativity and innovation, by generating new ideas, insights, and solutions, and enabling collaboration across disciplines and domains.

- support decision making and policy making, analysing large and diverse data sets, identifying patterns and trends, and providing evidence-based recommendations.

- Support improvements in health care, staff education, through personalised care plans, monitoring health conditions, detecting risks and anomalies, and education support plans.

5.5. **Risks:** Alongside the benefits the use of AI can bring there are also significant risks to be aware of:

- Outputs can be, made up, false, or misleading; commonly referred to as ‘hallucinations’.
- Outputs can omit key information.
- Outputs can introduce biases and discrimination.
- Use may cross ethical boundaries.
- Information input into AI may be more widely used for developing or training the AI model.
- AI may make use of and generate biased, discriminatory, or offensive content.
- Data entered into AI applications may enter the public domain, putting at risk our confidential personal, special category or commercially sensitive information. Data released may breach regulatory requirements or data protection laws.
- Use of AIs can carry inherent risks, an assessment of risks, Appendix 3, should form part of the procurement process.

5.6. It is important that the purpose and use of AI is clearly defined and agreed, including why AI is being used and what value it will bring to the organisation. Where possible data should be anonymous so a legal basis would not be required. Where this is not possible, a legal basis for the use of the data required must be found before any data is processed, along with meeting the requirements of the common law duty of confidentiality.

5.7. It is essential, however, that data and use cases are carefully assessed to determine if individuals can be identified using the contents of the information even if common identifiers such as name, address and phone number are removed. The combined details of a local area, a rare disease and a very young age may enable a patient to be identified. In such cases you will need to treat this as personal data and identify a legal basis for the processing along with meeting the requirements of the common law duty of confidentiality.

5.8. The above requirements also apply to data used to test and develop AI systems even if there is no outcome or decision for an individual, this is because you are processing data by using it to train AI models or algorithms.

5.9. **Consent and Patient Rights:**

- **Individual Care:** Implied consent may apply when AI supports direct care, but patients must be informed that an AI is in use and in which way.
- **Secondary Use:** Explicit consent or appropriate legal basis is required for using data to train AI models for broader applications.
- **Right to Explanation:** Patients can request a more detailed explanation of how AI contributed to decisions about their care.
- **Right to Restrict:** A patient as the data subject has the right to restrict how we process their data, which may include restrict what AI can be used for in their care.
- **Right to object:** a patient as the data subject has the right to object to how we process their data and that can include the use of AI in their care.
- **Automated decision making:** A patient as a data subject has the right to opt out of automated decisions about them and their care and insist on human processing.

When considering an IA product, that will be processing person identifiable and special category data, all of the above must be able to be implemented, to remain to be lawful processing.

5.10. **Developing Artificial Intelligence Products for Healthcare:** The NHS's AI and Digital Regulations Service is an AI regulation service for people who develop or plan to use AI or a digital technology in health and social care. It brings together regulations, guidance and resources for digital healthcare technologies. The service is comprised of four partners, National Institute for Health & Care Excellence (NICE), Medicines and Healthcare products Regulatory Service (MHRA), Health Research Authority (HRA) and Care Quality Commission (CQC).

5.11. **Using AI for Research:** Health Research Authority (HRA) approval is required for research studies that take place in the NHS in England. The 'HRA AI and Digital Regulations Service' can provide guidance for NHS AI adopters, and digital health innovators. Review by an NHS Research Ethics Committee (REC) is required, as well as an assessment of regulatory compliance and related matters undertaken by dedicated HRA staff.

5.12. If you are planning to develop an AI research programme within the NHS, the Research and Development Support Services team within the Innovation, Research and Improvement System (IRIS) will be able to provide advice and guidance on how to apply for research ethics and approvals via the Health Research Authority.

5.13. **Freely Available Generative Artificial Intelligence Apps, systems and Services:** AI is a feature of many developments within commonly used applications, such as Chat GP, Grammarly, Google translate, Edge co-pilot, Google bard, Apple SIRI,

Android Gemini, that staff may have access to. Unless they have been approved for your planned type of use, by the Trust, staff must not utilise them for Trust business.

5.14. Please refer to the separate staff guide for using generative AI.

5.15. **Procuring AI:** When considering the procurement of an AI product you should consider your proposed product against the Trust's AI usage Grid, at appendix 1. This will support you in understanding the Trust's appetite and governance steps needed for your AI tool.

5.16. When moving to procuring and implementing artificial intelligence products or systems that include AI features you must complete the AI procurement and implementation check list, at Appendix 2.

5.17. **Nationally procured AI functions:** Even when considering national procured AI offerings, they must still follow the Trust governance and approval process in line with this policy. Consider:

- Is the IG and IT team aware of such implementation and system function?
- Has the national purchase procured function/offering been intended for use by a combined community Trust?
- Is there any assurance documents in place associated with the AI function(s) (e.g. a DPIA, DTAC, DCB 0160)?
- Are there any national guidance provided on how to use and implement the AI?
- A Trust DPIA covering our intended use will be needed.

5.18. **Principles to follow**:

- **Transparency:** Patients and staff must be informed when AI is used or involved.
- **Accountability:** AI must support not replace decision-making, corporate or clinical. The user remains accountable for all outputs generated by the AI and should ensure the output is accurate, compliant and fit for purpose.
- **Feedback:** Users must have feedback routes for AIs where there are concerns around the output to support accountability, fairness and bias, privacy, security, safety and effectiveness.
- **Fairness & Bias Mitigation:** We will ensure there are processes in place to identify and mitigate biases in our AI systems to ensure fairness and inclusion where AI is utilised. AI systems must be regularly audited to prevent bias and ensure equitable outcomes.

- **Privacy and Security:** All AI systems must comply with UK GDPR and NHS data protection standards. We prioritise the privacy and security of personal data and health care data we hold.

- **Safety and Effectiveness:** AI tools must undergo rigorous testing and where used for clinical care meet the NHS Digital Technology Assessment Criteria (DTAC).

- **Regulatory Compliance:** We will comply with all relevant laws, guidance and regulations governing AI technologies.

- **Training:** There must be suitable training for users, to ensure safe and effective use of the AI tool.

5.19. **Governance**: The governance for implementing an AI will follow the Trusts normal governance process for implementing an IT system or change.

- **AI & Committees:** AI purchase must follow the normal Trust committee and group processes, for approval. Depending on where the activity fits within the AI usage grid, they may need to progress through the IG team, Digital Team, the Trust's Digital Group and IG Group.

- **AI Ethics:** TBC

- **Data Protection Impact Assessments (DPIAs):** A stage 1 will be needed for all AIs in the green Yes now (Yn), Yes controls (Yc), and amber Maybe partial Governance (MpG). A full stage 2 DPIA is required for all AI products, in the Amber Maybe full Governance (MfG) and Red No not now categories.

- **Information Asset Register:** All AI systems that are processing Person Identifiable, special category or corporate confidential data must be logged with the IG team and placed on the information asset register. A tool designated as an information asset must have an assigned asset owner.

- **Permitted uses:** The AI once approved can only be used in a manner that is compatible to the approved use, any divergence will need further review and approval. The DPIA will make clear the extent of the approved use, this will also be recorded in the asset register for transparency.

- **Privacy Notice:** The Trust's privacy Notice must be updated as and when in line with the agreed use and purpose, to support transparency.

5.20. **Environmental and financial impact**:

- Better use of information and data can help optimise delivery, such as through analysis of travel routes and thus reduce the environmental impact.

- The development, maintenance, and disposal of technology, however, all come with a large carbon footprint, due to energy-intensive processing. Training AI models can produce about 626,000 pounds of carbon dioxide.

Training a single GenAI model can consume as much as 284,000 litres of water, the equivalent to an average person's consumption over the course of 27 years.

- The use of AI technology may require higher specification equipment, which may result in increased costs for more powerful laptops and desktops; alternatively, colleagues may experience delays in processing AI tasks if the current equipment is not upgraded, leading to slower performance and reduced efficiency.
- The data industry is predicted to account for more carbon emissions than the automotive, aviation and energy sectors combined. We must work to better understand the impact of digital activities on our carbon footprint in order to start reducing it.
- Decisions to enable real-time data feeds, data warehousing, the use of AI and predictive analytics, and other data heavy solutions, need to be justified in the context of the environmental impact.
- Implementing data minimisation and retention policies should be considered as mitigations, as they can help reduce unnecessary carbon emissions.

PART 3 – Explanatory Information

6. DEFINITIONS

Term	Definition
Artificial Intelligence (nAI)	Technologies that simulate human intelligence processes, including machine learning, natural language processing, and computer vision.
Generative AI (GenAI)	AI systems capable of creating new content such as text, images, or data based on training inputs, including Chat bots, Co-pilot, SIRI, Google bard.
Machine Learning (MLAI)	A subset of AI that enables systems to learn from data and improve over time without explicit programming. Including Grammarly, Duolingo, Google translate, face ID.
Optical Character Recognition (OCR)	A technology that uses automated data extraction, pattern recognition to convert images of text into a machine-readable format, more recently includes an AI element. It allows text to be extracted from scanned documents or images, rendering it editable and searchable. Including Nanonets, Acrobat pro.
Ambient Voice Technology (AVT) – sits within the generative AI arena.	An AI advanced speech technology that listens to a conversation and seamlessly interacts between humans and machines. Often used to producing transcript of a conversation, in a healthcare setting such as when engaging with a patient face to face, the spoken words are converted to text. In healthcare an AVT must be at least an MHRA class 1 medical device. Also known as AI scribe technology.

Agentic AI (AgAI)	Agentic AI refers to AI systems capable of goal-directed, autonomous behaviour, going beyond traditional or generative AI that primarily responds to prompts or generates content. Agentic AI can plan multi-step workflows, make decisions, and execute actions independently, often using external tools, APIs, and databases to achieve complex objectives. It is built on AI agents, which are specialised autonomous entities that can collaborate in multi-agent systems to solve tasks too complex for a single agent.

7. PROCESS FOR MONITORING COMPLIANCE

7.1. Adherence to this policy will be monitored through

- Staff awareness through Indi-to-go.
- Completion and review of AI project or systems documentation including DPIAs.
- ICS Cyber reviews.
- Inclusion on the Trust's registers, DPIA, Information Asset and flows.
- IT system installation & CAB process.
- Reviews by the Trust's Digital and IG Groups.

8. TRAINING

8.1. There is no specific training needed in respect of this policy.

8.2. Each AI system implemented should ensure there is suitable training in place before going live, along with training for new staff throughout the life of the product. IT 's digital instruction team should be utilised to support AI training.

9. REFERECNES

9.1. Sources of advice and guidance

- **NHS AI Lab**: Provides comprehensive guidance on understanding, developing, and adopting AI within the NHS. It includes resources on information governance, clinical validation, and practical implementation tips.
- **NHS Transformation Directorate**: Offers guidance on the implications of using AI in health and care settings, focusing on lawful and safe data use.
- **AI and Digital Regulations Service**: A one-stop-shop for advice on adopting and utilizing AI technologies in health and social care.

- **Artificial Intelligence Playbook for the UK Government**: Provides detailed guidance on AI technologies, use cases, and corporate policies for government departments.

- **Government Digital Service (GDS)**: Offers accessible explanations, sample use cases, and guidance for policy professionals on AI technologies.

- **ICO Guidance on AI and Data Protection**: Offers detailed advice on how to interpret relevant data protection laws as they apply to AI, and recommendations on good practices for mitigating risks.

- **NCSC Guidelines for Secure AI System Development**: Provides comprehensive guidelines for secure AI system development, covering design, development, deployment, and operation.

- **Principles for the Security of Machine Learning**: A set of principles aimed at ensuring the security of machine learning systems, applicable to anyone developing, deploying, or operating such systems.

- **The UK AI Safety Institute**: Focuses on minimizing risks from rapid and unexpected advances in AI, ensuring public safety and human flourishing.

- **Regulatory Sandbox for AI in Healthcare**: Provides a controlled environment for testing new healthcare technologies, ensuring both efficacy and compliance before full-scale deployment.

- **Ada Lovelace Institute Foundation Models Public Sector AI**: Examines the use of foundation models in the public sector, including their development, testing, and deployment.

- **Alan Turing Institute Ethics and Governance in AI Practice**: Offers a series of workbooks and an online platform to help the public sector apply AI ethics and safety to the design, development, and deployment of algorithmic systems.

10. ASSOCIATED DOCUMENTS

10.1. This policy should be read in conjunction with the following:

- HR Policy Manual
- IT, Information Management and Data Protection Policies, in particular;
    - Information Governance Management System Policy
    - IT Information Security Policy
    - Digital Platform Policy
    - Commissioning Policy for IT Systems

- IT Access Control Policy
- Data Quality Policy
- Portable IT Equipment and Removable Media Policy
- Accessible Information Policy

Appendix 1 AI usage Grid

Artificial Intelligence	Technologies that simulate human intelligence processes, including machine learning, natural language processing, and computer vision.

Type	No Data	Aggregated / Anonymous data	Corporate Confidential no PID	PID within tenant	Corporate Confidential with PID	Within EPRs	PID Outside tenant	Linking to EPRs	Supporting Clinical care	Supporting Diagnosis	Generating Clinical output	Diagnosis
Generative AI	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Machine learning	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Optical Character Recognition	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Ambient Voice Technology	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Agentic AI	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Combinations there of	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn
Machine Automation	Yn	Yn	Yc	Yc	MpG	MpG	MfG	MfG	MfG	Nn	Nn	Nn

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Yn	Yc	MpG	MfG	Nn
Yes now no control needed	Yes, controls needed	Maybe possible with partial governance	Maybe possible with full governance	No, not now

Term	Definition
Generative AI	AI systems capable of creating new content such as text, images, or data based on training inputs, including Chat bots, Co-pilot, SIRI, Google bard.
Machine Learning	A subset of AI that enables systems to learn from data and improve over time without explicit programming. Including Grammarly, Duolingo, Google translate, face ID.
Optical Character Recognition	A technology that uses automated data extraction, pattern recognition to convert images of text into a machine-readable format, more recently includes an AI element. It allows text to be extracted from scanned documents or images, rendering it editable and searchable. Including Nanonets, Acrobat pro.
Ambient Voice Technology (AVT) – sits within the generative AI arena	An AI advanced speech technology that listens to a conversation and seamlessly interacts between humans and machines. Often used to producing transcript of a conversation, in a healthcare setting such as when engaging with a patient face to face, the spoken words are converted to text. In healthcare an AVT must be at least an MHRA class 1 medical device. Also known as AI scribe technology.
Agentic AI	Agentic AI refers to AI systems capable of goal-directed, autonomous behaviour, going beyond traditional or generative AI that primarily responds to prompts or generates content. Agentic AI can plan multi-step workflows, make decisions, and execute actions independently, often using external tools, APIs, and databases to achieve complex objectives. It is built on AI agents, which

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

	are specialised autonomous entities that can collaborate in multi-agent systems to solve tasks too complex for a single agent.
Machine Automation	With or without an AI the machine automates a process normally carried out by a human.
Within Tenant	On Trust owned and managed servers or within the Trust's Office 365 environment.

Appendix 2: AI Procurement and Implementation Check list

Action	Yes/No	Date	Comment
Check the AI usage Grid for your product			
Engage Procurement			
Engage IG			
Engage IT			
Engage BI			
Engage the CSO and CCIO (consult throughout procurement and implementation)			
Complete a Digital Technology Assessment Criteria (DTAC)			
Have you considered and mitigated AI Biases?			
Have you considered and mitigated Ethic concerns? (including potential societal impacts)			
Is the AI associated with research? (if yes, you must obtain approval from the HRA)			
Have you considered how you will manage objections to automated decision making?			

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Have you considered a process to flag any concerns regarding the AI (including input, output, bias, ethics)			
Have you considered the risks and practical steps to			
reduce the risk associated wit the AI? (complete the ICO AI risk toolkit, available here or from IG Team)			

Other Considerations:

- You must consider Article 22 of the UK General Data Protection Regulation in relation to automated individual decision making, including profiling. – Individuals have the right not to be subject to automated decision making.
- AI outcomes or outputs must be reviewed by a human.
- You cannot rely solely on the use of AI for decision making, there must be substantial involvement from an appropriately qualified human.
- There must be an agreed process to flag any concerns regarding the output of the AI product.
- If there are concerns which have led to an incident this must be reported as per the Incident Reporting Policy.
- Use of AI must be transparent to staff and patients ensuring they understand where it is being used and how it may impact their employment, work or care.
- The logic behind it and the use must be explainable.
- Data must be collected and processed in a lawful and ethical manner, with appropriate consent and anonymisation measures in place.
- Data access and sharing must be strictly controlled, and data must be stored securely throughout its lifecycle.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

- You should conduct patient and public engagement activities that include determining if individuals support the use of data for your intended purpose, or if they have any concerns on how their data will be used.

- If the use of AI involves service change then prior to the implementation of any AI programme, formal consultation must take place with employees and their trade union representatives in accordance with the organisational change policy.

- You must be assured that any product mitigates against bias and discrimination.

- AI systems should be continuously monitored for suspicious activities, anomalies, and potential security breaches.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Appendix 3 – AI Data Protection Risk toolkit (ICO)

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Appendix 4 – AI Information Guides (developed based on the Plymouth example)

AI Information Guide 1

Computer Vision – A Basic Overview

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY ARE WE ALL TALKING ABOUT COMPUTER VISION?

Computer vision marks a significant area within artificial intelligence that provides computers with the remarkable ability to see, interpret, and understand our visual world from images and videos. A key goal of computer vision is to assist with or handle certain visual tasks, enabling new capabilities and supporting human efforts. This involves more than simply processing a digital image; it means these systems learn by recognising patterns from vast amounts of visual data, similar to how a person learns. Through this process, they learn to identify objects, analyse the context of a scene, and extract meaningful, actionable information from what they 'see'.

- **Powerful Visual Capabilities for AI:** These technologies represent a fundamental way for artificial intelligence to work with, interpret, and respond to the visual world.
- **Shaping Our Digital World:** Computer vision has already greatly influenced how we use technology, from the safety features in cars to how we unlock our phones.
- **Enabling Advanced Interactions:** It is a key reason that computers can now engage in more sophisticated visual tasks, such as understanding our surroundings or allowing for gesture controls.

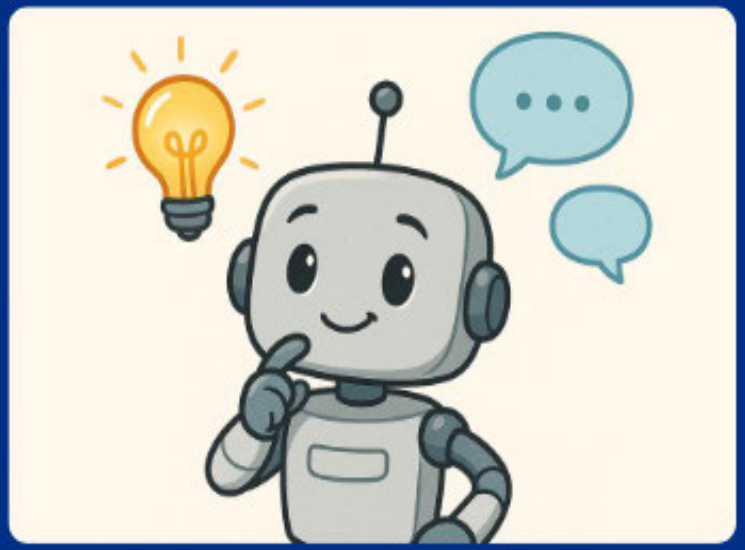

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS COMPUTER VISION?

At its heart, computer vision is a specialised area within the artificial intelligence umbrella, focused on enabling computers to 'see' and interpret the world, much like we teach a child to recognise faces, objects, and scenes. A key goal of computer vision is to assist with or handle certain visual tasks, enabling new capabilities and supporting human efforts. It is important to understand the difference between a computer simply capturing an image and actually understanding it. A standard camera or a scanner creates a digital copy. Computer vision goes a crucial step further. It does not just see a picture of a street; it works to identify the cars, read the road signs, and recognise the pedestrians. The aim is to gain meaningful information and understanding from visual data, not just to record it.

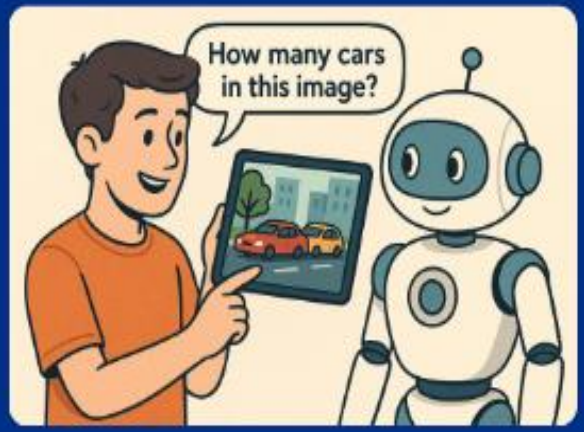

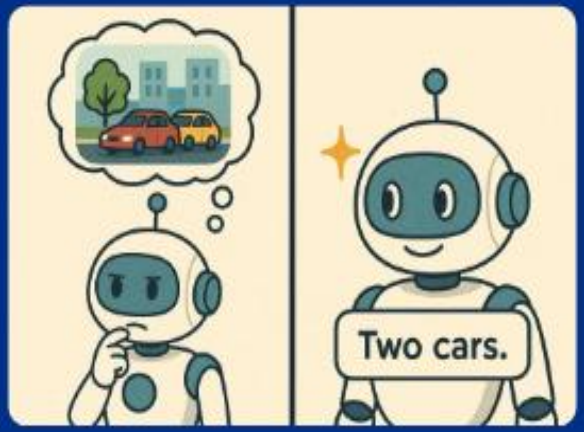

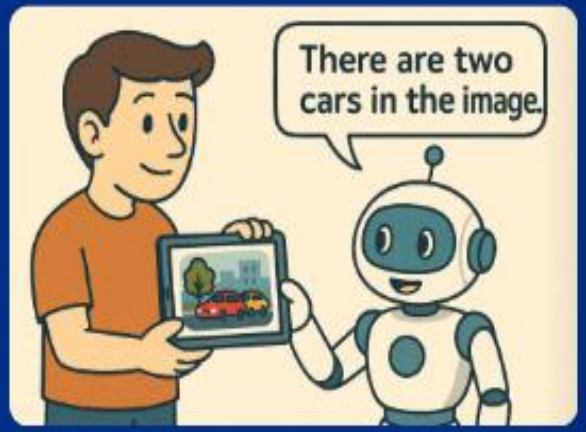

The process often starts when a person needs to understand something within an image or a video. For a computer vision system to help, its first essential step is to receive this visual input, such as a photograph or a frame from a video feed, which it sees only as a grid of numbers called pixels.

Next, the computer vision system begins to analyse the visual input to find meaningful patterns. It starts by detecting basic visual building blocks, such as edges, lines, corners, and simple textures. It then combines these simple patterns to recognise more complex shapes and parts of objects. This analytical process aims to transform the unstructured grid of pixels into a structured understanding of the image's contents.

Finally, based on its analysis of all the extracted features, the system can make a decision or provide an interpretation of the image. This output turns the visual data into useful, actionable information. This could be identifying and counting objects as shown here, classifying the entire scene, or performing another relevant task based on what it has 'seen' and understood.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION IN THE AI FAMILY

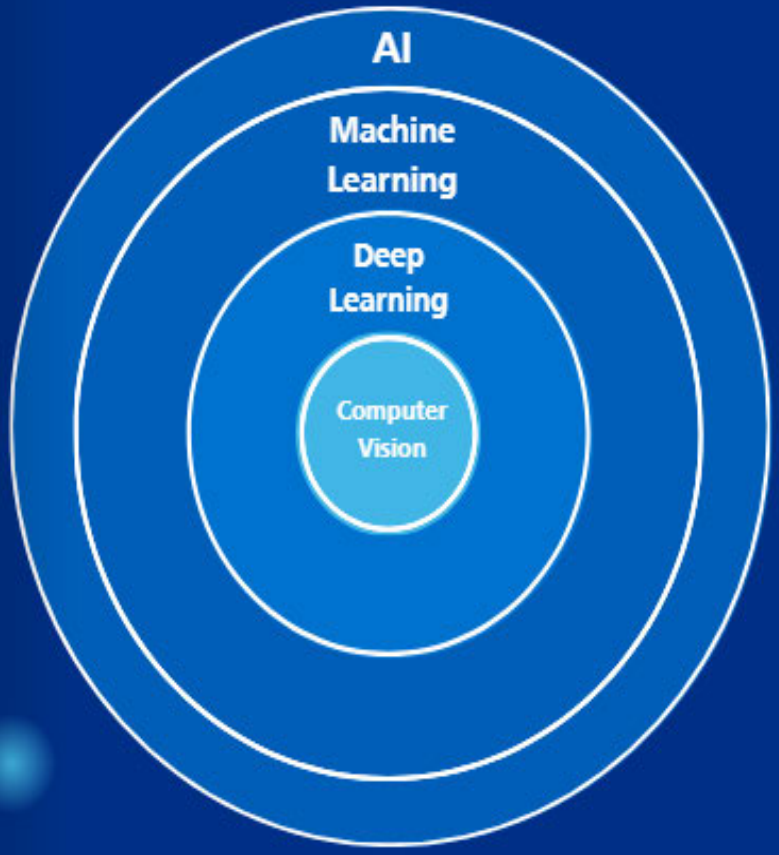

Artificial Intelligence (AI): AI is a wide ranging field focused on creating computer systems that can perform tasks normally requiring human intelligence, such as learning, reasoning, problem-solving, and decision-making. It is the overall goal of making machines intelligent.

Machine Learning (ML): ML is a core method used to achieve AI. Instead of being explicitly programmed for every task, ML systems learn patterns and make predictions or decisions by analysing large amounts of data. It is how many AI systems acquire their intelligence.

Deep Learning (DL): This is a powerful type of machine learning that uses computer systems with many layers, which is why it is called 'deep'. These layers work together to learn, starting with very simple ideas and building them up into more complex ones. This layered method is the engine behind most of the biggest recent breakthroughs in artificial intelligence.

Computer Vision (CV): This is a specialised field of artificial intelligence focused on enabling machines to 'see' and interpret visual information from the world, such as images and videos. Computer vision is largely powered by deep learning techniques, which help computers to learn and recognise patterns in visual data.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

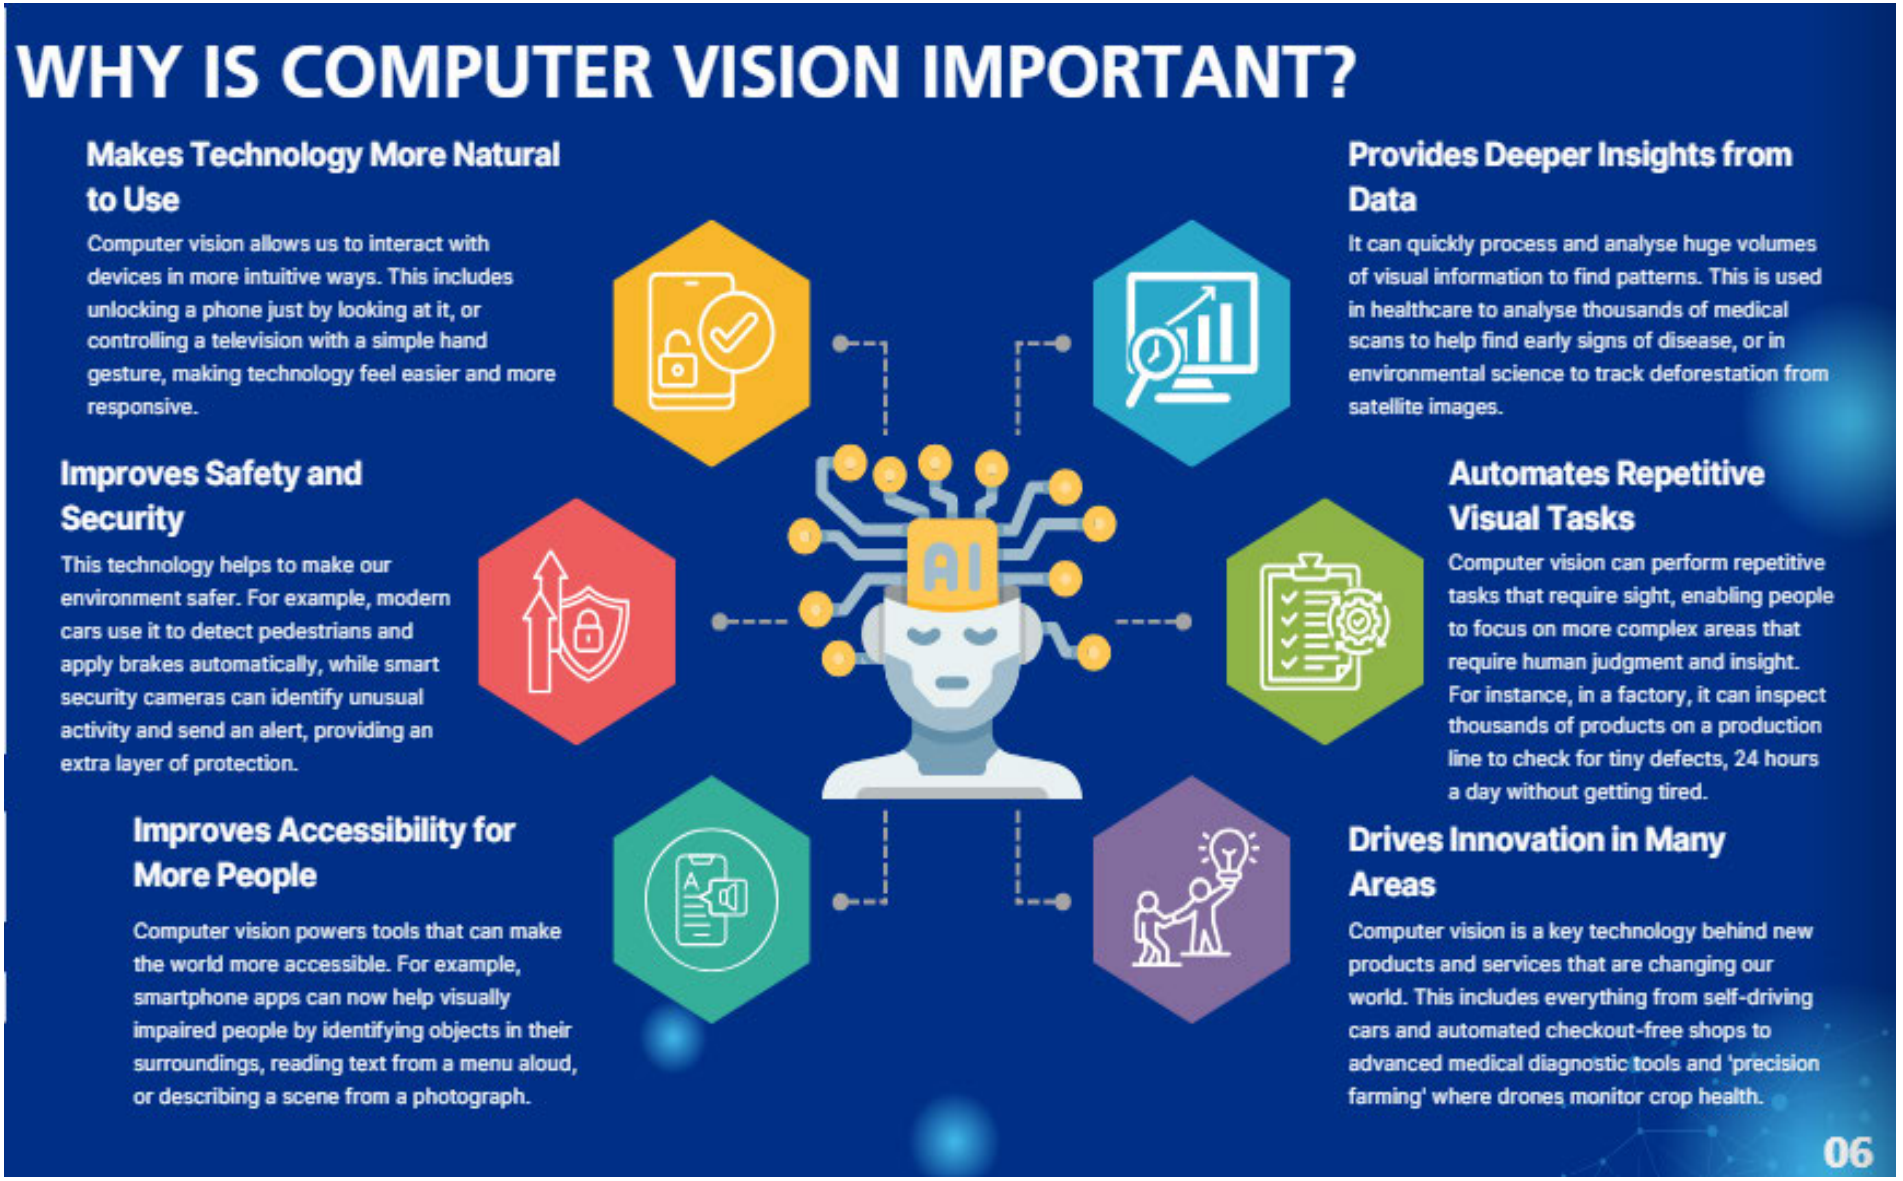

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW DOES COMPUTER VISION WORK?

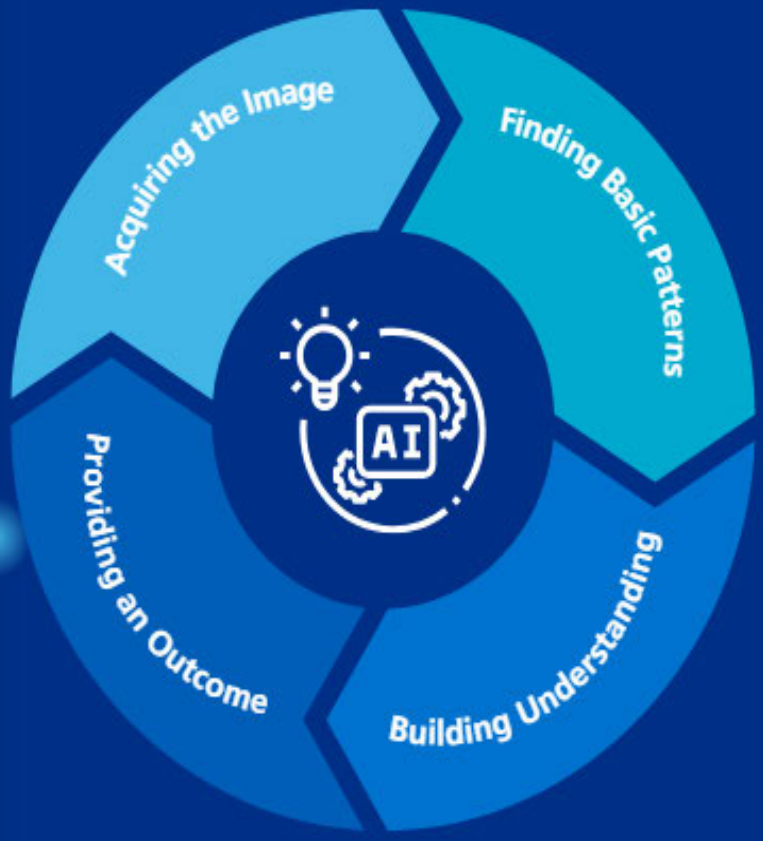

Acquiring the Image

First, the system receives the visual information. This could be a photograph you upload or a live video feed from a camera. To the computer, this image is not a picture but a large grid of numbers, where each number represents the colour and brightness of a single pixel.

Finding Basic Patterns

Next, the system begins to analyse this grid of numbers to find the most basic visual patterns. It is trained to look for simple building blocks, like straight lines, sharp corners, edges where colours change, and simple textures. This is the first step in making sense of the visual information.

Building Understanding

This is where the system starts to build a more complete understanding. It learns to combine the simple patterns it has found into more complex shapes and parts of objects. For example, it might learn that certain lines and curves often form the shape of a wheel, and that four wheels combined with other shapes often represent a car.

Providing an Outcome

Finally, based on the complete picture it has built, the system provides a useful outcome. This could be giving the image a label, like "this is a picture of a car", drawing a box around a specific object it has found, or providing a count of the items it has identified. The final outcome turns the raw image into useful information.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION WORK?

Seeing an Image as Pixels

For a computer vision system to begin understanding an image, it must first break the picture down into manageable units that it can process. A computer does not see a whole photograph in the same way a person does; instead, it sees the image as a large grid made up of many tiny, individual pixels. Each pixel has a specific colour and brightness, which the computer represents as a number. This process of seeing an image as a grid of numerical pixel values is the essential first step. Once the image is recognised as this collection of basic digital units, the system can then begin to analyse them to find patterns, which is fundamental for any further understanding.

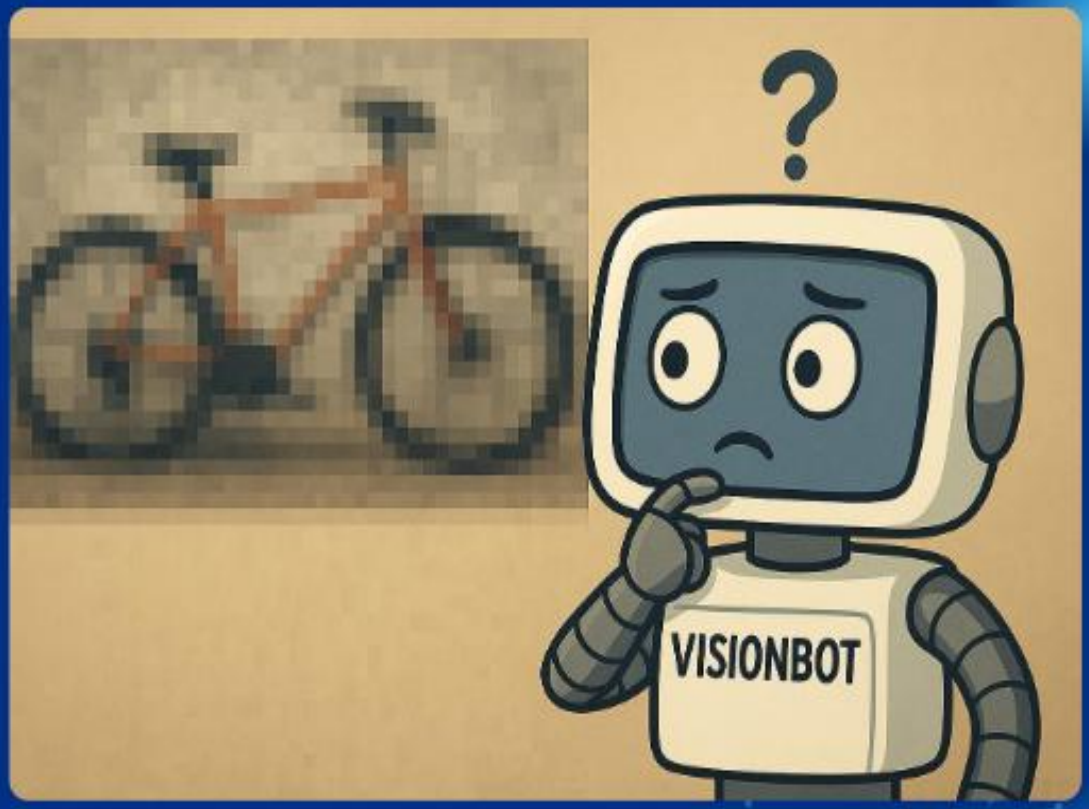

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION WORK?

Finding Basic Patterns - Identifying Visual Clues

After seeing the image as a grid of pixels, the computer vision system must begin to find meaningful patterns within that grid. This stage involves looking for the most basic visual 'clues', such as simple lines, edges where colours change, sharp corners, and basic textures. This is like a detective scanning a room for evidence using different specialised tools. The system uses many digital 'filters', where each filter is designed to find only one specific type of pattern. For example, one filter might highlight all the vertical lines, while another finds all the circular shapes. By using many of these filters, the system creates a set of 'clue maps', each one showing where a specific basic pattern appears in the image. This step is vital for being able to build up a full picture later on.

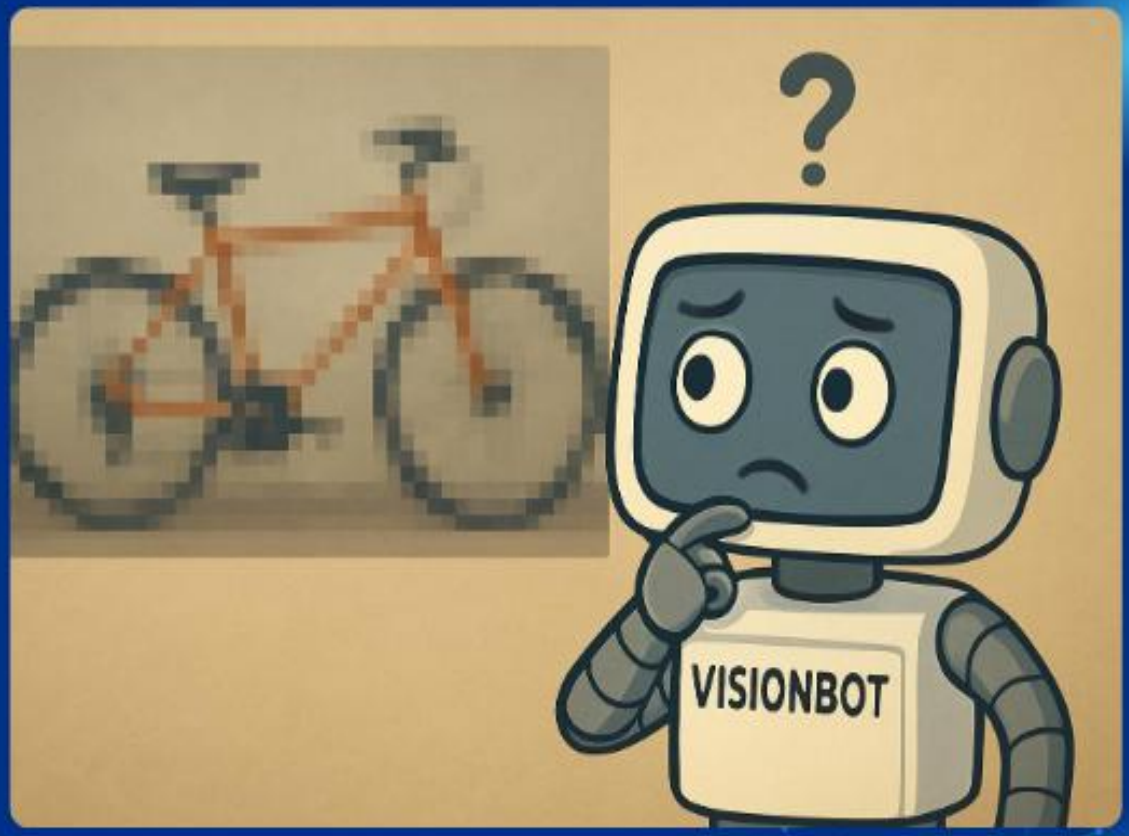

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION WORK?

Building a Full Picture - Assembling the Clues

With the basic visual clues like lines, corners, and textures identified, the computer vision system then works to understand how these pieces fit together to create a recognisable object. This stage involves learning the relationships between the simple patterns. For instance, it determines how certain lines and curves connect to form a larger shape, like a wheel, and how multiple shapes can be assembled to form a complete object, like a bicycle. It is similar to building with LEGO bricks. The basic visual patterns are like the individual bricks. By learning how to connect these simple bricks in the right order and arrangement, the system can construct a complex model. By analysing these connections, the system builds a full picture of what the image contains, moving beyond individual visual clues to a more complete and meaningful interpretation.

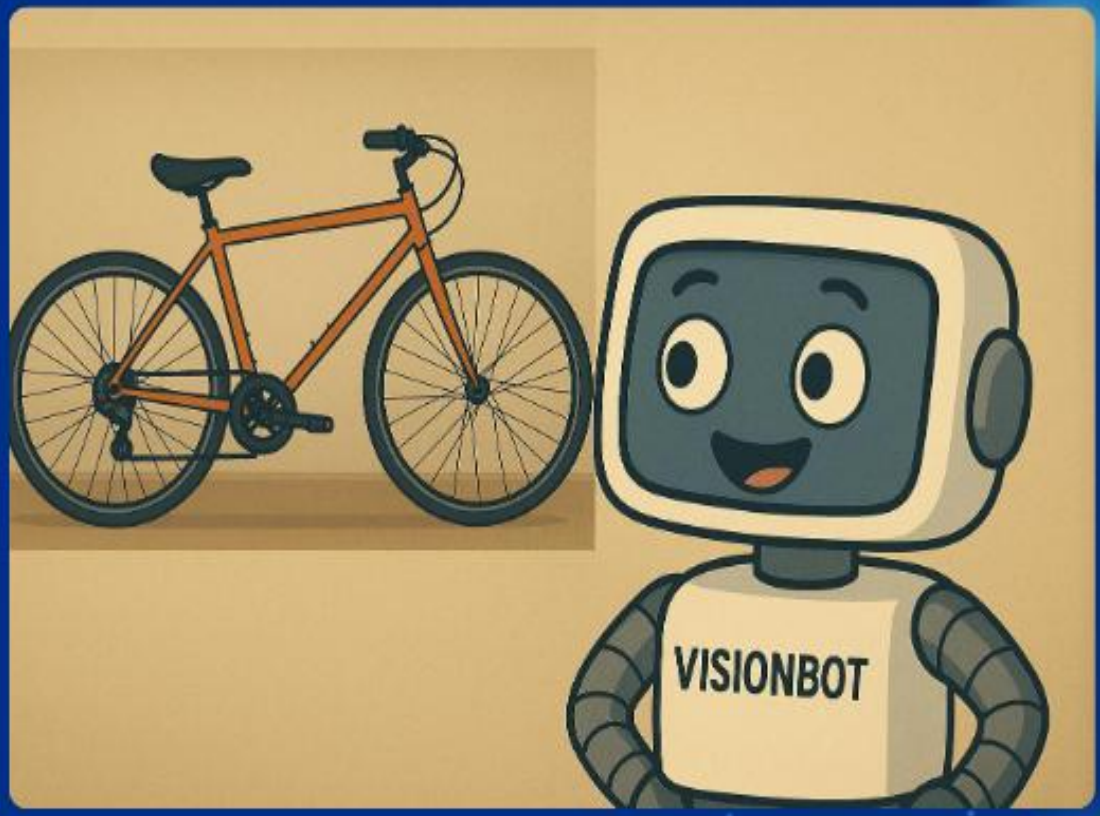

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT CAN COMPUTER VISION ACTUALLY DO?

01 Image Classification

Computer vision systems automatically sort images into categories or applies labels to them. This is similar to how a photo application might group all of your pictures into different albums such as putting all images including a cat together into one album.

02 Object Detection

These systems can find and identify specific objects within an image, usually by drawing a box around them. For example, a system could analyse a photograph of a busy street and draw a box around every car, person, and traffic light it finds.

03 Facial Recognition

This is a specialised type of computer vision that focuses on identifying or verifying a person from a digital image. This is the technology used to unlock your smartphone with your face, or to automatically suggest tagging friends in photos on social media.

04 Scene Understanding

Computer vision can analyse a whole scene and describe what is happening in a sentence. This is used in applications that help visually impaired people by reading a scene aloud, for example, "a young girl is playing with a dog in a park."

05 Video Analysis

This capability goes beyond static images to understand what is happening over time in a video. For example, it can be used to analyse sports footage to track a player's movements, or to monitor a security camera feed and send an alert when it detects someone falling over.

06 Image Segmentation

This highly detailed task goes beyond just drawing a box around an object. It works by identifying the exact outline of each item in an image by classifying every single pixel. For example, it works like a digital 'colouring-in' book, where the system knows the exact boundary of every object and can colour all the parts of the cars blue and the road grey, without going over any lines.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

REAL-WORLD APPLICATIONS OF COMPUTER VISION

Computer vision is already part of many tools and services we use every day, often working behind the scenes to improve how we interact with technology and the physical world. This field of artificial intelligence helps computers to 'see' and interpret the visual world, leading to improved safety, new interactive experiences, and innovative solutions across diverse sectors such as healthcare, automotive, and retail. The real-world applications of computer vision are continuously growing, helping organisations and individuals to automate visual tasks, gain valuable insights from images and videos, and shape a future where technology can better perceive and understand our surroundings.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION IN YOUR DAILY LIFE

Making everyday technology smarter and more helpful.

UNLOCKING YOUR PHONE WITH YOUR FACE

Many smartphones use facial recognition to provide secure access. When you set it up, the system learns the unique features of your face. Computer vision is then used to compare the person looking at the phone with the stored information to unlock it instantly.

SMARTER PHOTO ALBUMS

Photo applications on your phone or computer use computer vision to automatically identify faces of friends and family, or to recognise objects and scenes. This is how they can automatically group all your pictures of a specific person, your pet, or all your photos from your favourite places.

SOCIAL MEDIA FILTERS

Interactive filters on applications like Instagram or Snapchat rely on computer vision. The system tracks the location of your facial features, such as your eyes, nose, and mouth, in real-time. This allows it to accurately place digital items like glasses, hats, or animal ears onto your face.

VISUAL SEARCH IN SHOPPING

Many shopping applications now allow you to search with an image instead of words. You can take a picture of a piece of clothing or furniture you like, and the application will use computer vision to find that exact item, or very similar ones, for you to buy.

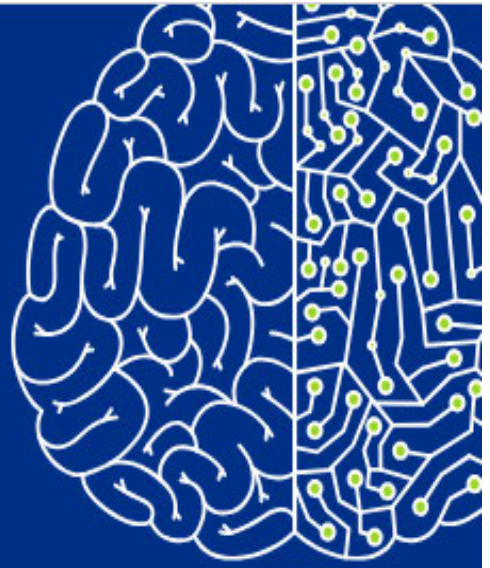

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION IN INDUSTRY AND BUSINESS

Improving efficiency, quality, and customer experiences.

AUTOMATED QUALITY CONTROL

In manufacturing, computer vision systems are used to inspect products on a fast-moving production line. These systems can spot tiny defects, working 24 hours a day to ensure that products meet quality standards.

'JUST WALK OUT' SHOPPING

Some modern shops use computer vision to create a checkout-free experience. A system of cameras and sensors tracks the items that a customer picks up from the shelves and puts in their bag. When the customer leaves, the system automatically calculates the total and charges their account.

MANAGING STOCK ON SHELVES

In retail, cameras can be used to continuously monitor shelves in a shop. Computer vision systems can automatically detect when a product is running low or is out of stock, and can alert staff to refill the shelf, ensuring products are always available for customers.

PRECISION FARMING

In agriculture, drones equipped with cameras fly over fields. Computer vision analyses the images to monitor the health of crops, identify areas affected by pests or disease, and determine which parts of the field need more water or fertiliser. This helps farmers to use resources more efficiently and increase their harvest.

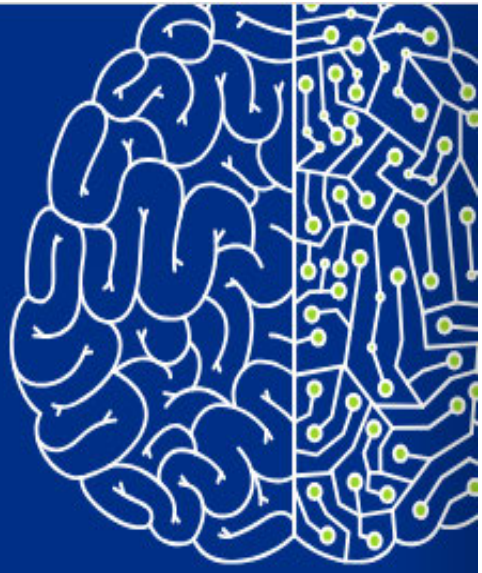

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION IN HEALTHCARE AND SCIENCE

Assisting experts and accelerating new discoveries.

ANALYSING MEDICAL SCANS

Computer vision is used to analyse medical images like X-rays, MRIs, and CT scans. It can automatically highlight areas of concern for a doctor to review, such as potential tumours, fractures, or the early signs of disease. This acts as a helpful diagnostic aid for radiologists and other specialists.

MONITORING THE ENVIRONMENT

Scientists use computer vision to analyse vast collections of satellite images taken over many years. This helps them to track large-scale environmental changes, such as the rate of deforestation in a rainforest, the melting of polar ice caps, or the impact of floods and other natural disasters.

ASSISTING IN ROBOTIC SURGERY

During some surgical procedures, computer vision provides real-time information to surgeons using robotic tools. It can help to track the position of surgical instruments, highlight nerves or blood vessels, and provide a clear, enhanced view of the area being operated on, improving precision and safety.

UNDERSTANDING ANIMAL BEHAVIOUR

Researchers use computer vision to study wildlife without disturbing the animals. Cameras placed in natural habitats can automatically track animal movements, identify different species, and analyse their behaviour over long periods, providing valuable data for conservation efforts.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION FOR SAFETY AND ACCESSIBILITY

Protecting people and making the world more inclusive.

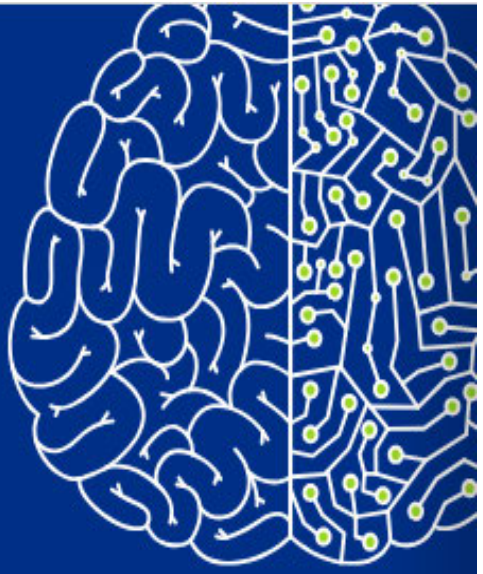

ADVANCED CAR SAFETY FEATURES

Modern cars use cameras and computer vision to power advanced driver-assistance systems. These include features like automatic emergency braking, which detects pedestrians or other obstacles, and lane departure warnings, which alert the driver if the car starts to drift out of its lane.

WORKPLACE SAFETY MONITORING

On construction sites or in factories, computer vision systems can be used to improve safety. They can automatically detect if workers are wearing the correct personal protective equipment (PPE), such as hard hats, or identify potentially hazardous situations, such as a person getting too close to dangerous machinery.

DESCRIBING THE WORLD FOR THE VISUALLY IMPAIRED

Smartphone applications use computer vision to assist people who are blind or have low vision. By using the phone's camera, the application can identify objects, read text from signs or menus aloud, and even describe the scene in front of the person, helping them to navigate their surroundings.

AUTOMATED SECURITY ALERTS

Intelligent security cameras use computer vision to monitor a location for specific events. Instead of a person having to watch all the video feeds constantly, the system can automatically detect unusual activity, such as a person in a restricted area, and send an immediate alert to security staff.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

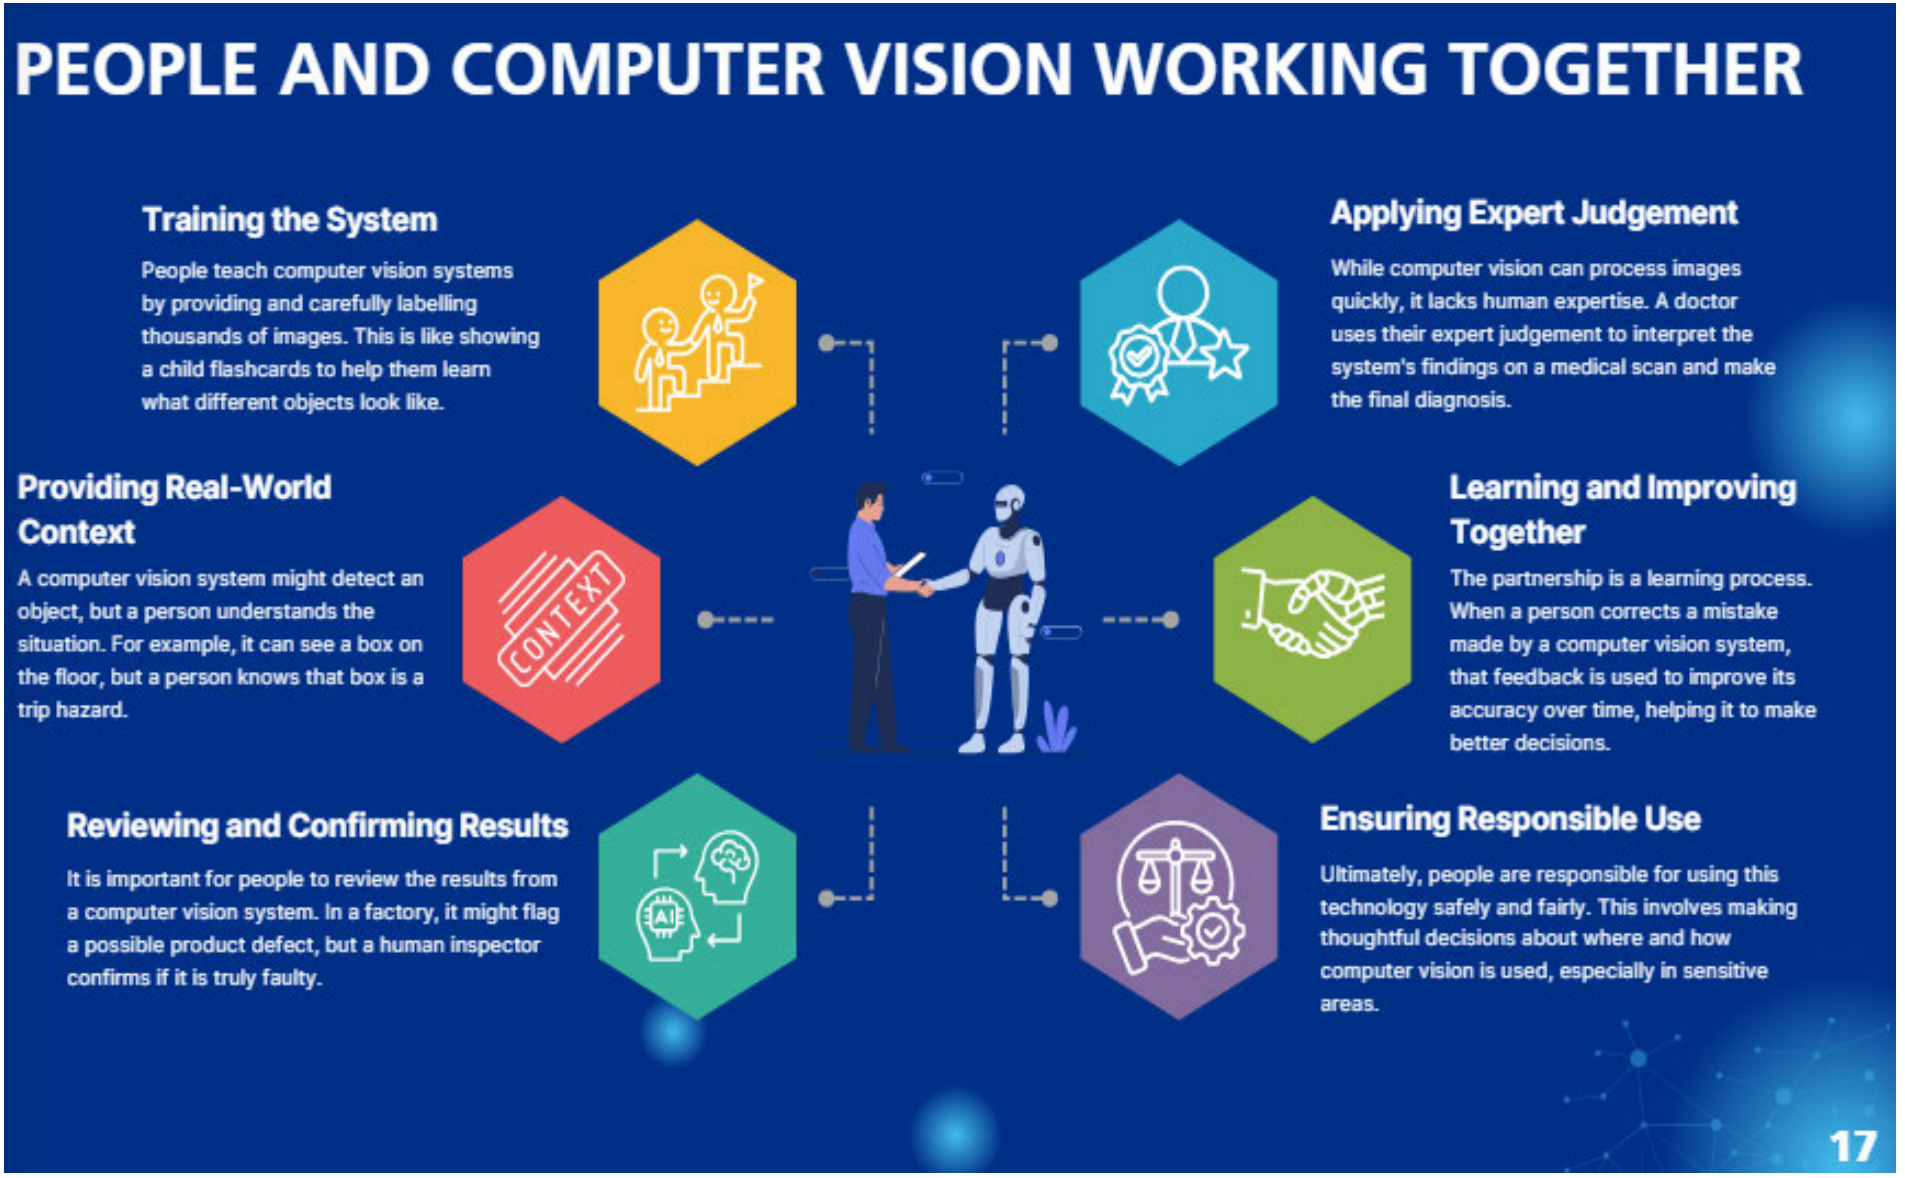

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CHALLENGES AND IMPORTANT QUESTIONS

01 Understanding Complex Scenes

Computer vision systems can struggle to understand scenes in difficult conditions, like heavy rain, poor lighting, or when an object is partially hidden.

02 Accuracy and Reliability

Results from computer vision are not always perfectly accurate. Systems can misidentify objects, so human checking is essential in critical situations.

03 Bias in Data and Systems

If the training images are not diverse, a computer vision system can become biased and may be less accurate for certain groups or environments.

04 Misinformation & Misuse

The ability to create realistic fake images and videos ('deepfakes') means the technology could be misused to spread false information.

05 Privacy and Surveillance

The use of cameras raises important privacy questions. How visual data is collected, used, and protected is a major ethical concern for society.

06 How Decisions Are Made

It can be hard to know why a complex computer vision system made a certain decision. This 'black box problem' makes it difficult to find errors or build trust.

07 Working in All Environments

Making computer vision work reliably across all the different visual and cultural environments around the globe is an ongoing challenge.

08 Ensuring Fair and Responsible Use

A key overall question is how we can develop and use these powerful technologies ethically, fairly, and in ways that benefit everyone in society.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 CV Helps Computers 'See' Our World

The main goal of computer vision is to enable computers to understand and interpret visual information from images and videos, making technology more aware of its surroundings.

02 It Is All Around Us Every Day

You use computer vision frequently, often without realising it, in things like the facial recognition that unlocks your phone, modern car safety features, and the automatic sorting of your photos.

03 A Broad Field with Many Tasks

Computer vision is a wide area of artificial intelligence with a long history. It includes many different tasks, from simply labelling an image to identifying the precise location and shape of many different objects within it.

04 CV Uses Powerful Modern Tools

Computer vision uses advanced methods that are very effective at learning from images. These tools allow the system to build its understanding in layers, learning simple lines and shapes first, before combining them to identify complex objects.

05 Human Oversight is Essential

While powerful, computer vision systems work best with human guidance. People are needed to train the systems, review their results for accuracy, and use their expert judgement to make important final decisions.

06 Challenges and Ethics Matter

It is important to be aware of the challenges of computer vision, such as potential biases and the need for accuracy, and to consider the ethical questions around topics like privacy and fairness as the technology develops.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 2

Computer Vision – A Clinical Overview

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY ARE WE ALL TALKING ABOUT COMPUTER VISION IN HEALTHCARE?

Computer vision is a significant area within artificial intelligence that is becoming increasingly important for medicine. It provides computers with a remarkable ability to analyse the complex visual information in medical images, such as X-rays and scans, in ways that can support clinicians and benefit patients. This involves more than simply viewing an image; it means these systems are learning to recognise subtle visual patterns that can be relevant for diagnosis and treatment.

- **An Essential Technology for Modern Medicine:** It is a key tool for making sense of the vast amounts of image-based information in healthcare.
- **Powerful Capabilities for Healthcare AI:** This technology represents a fundamental way for artificial intelligence to see, interpret, and analyse the visual information of health.
- **Shaping the Future of Healthcare:** Computer vision is already starting to influence how medical images are analysed and how care is delivered.
- **Enabling Advanced Clinical Tools:** It is a key reason that computer systems can now engage with more complex visual information, from identifying very subtle findings on scans to helping create detailed 3D models for surgical planning.

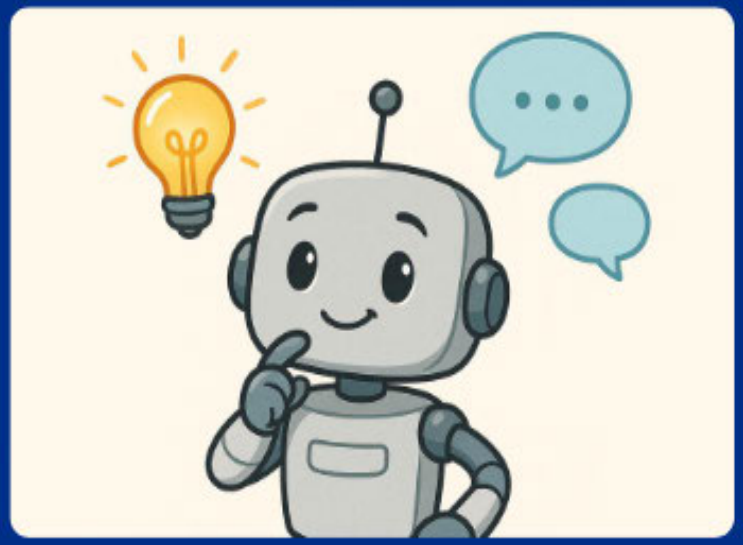

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS COMPUTER VISION IN HEALTHCARE?

Computer vision is an important area within artificial intelligence. In healthcare, its main focus is to enable computers to understand, interpret, and use the complex visual information in medical images, such as X-rays and scans. You can think of it as the process of teaching computers to become skilled at recognising features in clinical images and understanding the way anatomy and pathology appear visually. By learning from thousands of medical images, these systems identify and understand patterns in how different conditions appear. This learning allows computers to make better sense of medical images and, as a result, to provide more useful support, for example by highlighting an area of concern for a clinician to review.

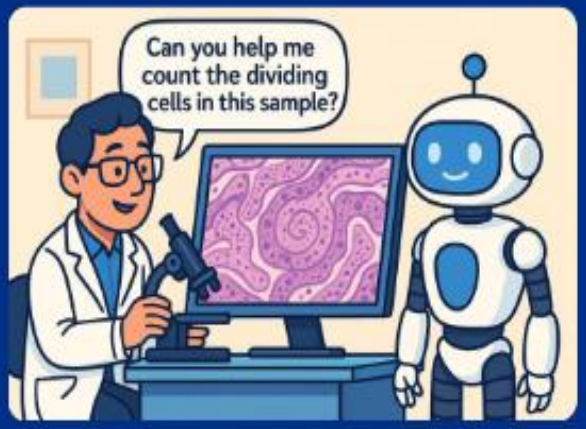

A key task in healthcare is the detailed review of medical images, when a healthcare professional needs to find specific information within a medical image. For a computer vision system to help, its first step is to receive this clinical image, such as a digital slide or a hospital scan, which it sees simply as a grid of numbers representing the visual information.

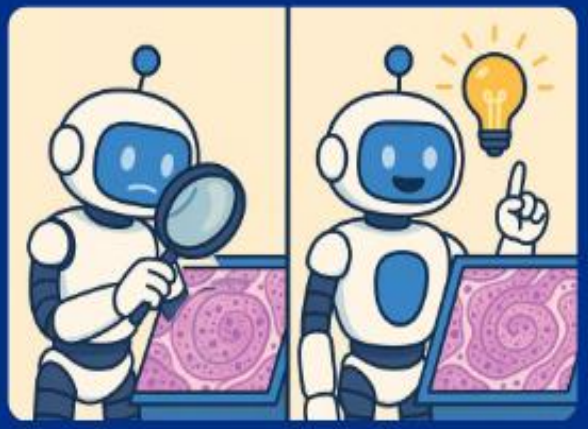

Next, the computer vision system begins to analyse the image to find meaningful patterns it has been trained to recognise. It starts by detecting basic visual building blocks, such as shapes, lines, and textures. It then combines these simple features to recognise more complex medical structures, like different types of cells, as shown in this example. This analytical process aims to transform the visual information into a structured understanding of the image's contents.

Finally, based on its analysis, the system can provide a useful result. This output turns the visual data into helpful information for the clinician to review. This could be identifying and counting specific cells as shown here, measuring a feature, or highlighting an area of interest. The clinician can use this information to support their professional assessment, always ma..ing the final decision themselves.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION IN THE AI FAMILY

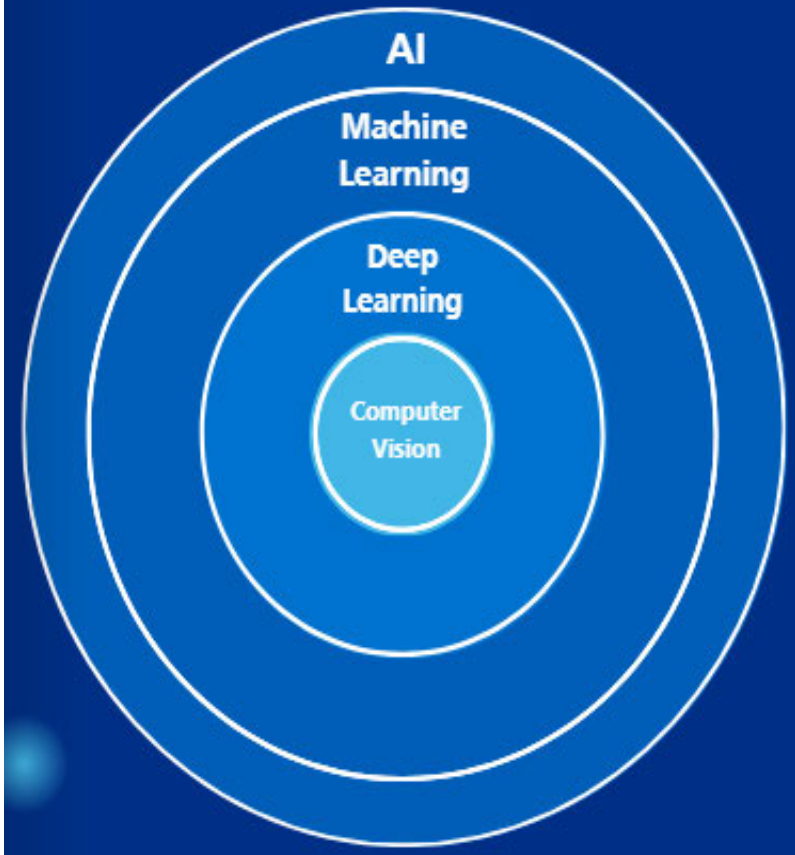

Artificial intelligence (AI): This is a wide-ranging field focused on creating computer systems that can help with tasks that normally require human thought, such as learning, problem-solving, and supporting clinical decisions. It is the overall goal of making machines helpful and intelligent.

Machine Learning (ML): This is a core method used to achieve artificial intelligence. Instead of being programmed with specific rules, these systems learn patterns by analysing large amounts of medical information, such as many examples of hospital scans. It is how many of these systems learn to be helpful.

Deep Learning (DL): This is a powerful type of machine learning. For images, its many layers work together to learn, starting with very simple patterns like lines and building them up into more complex ideas, like recognising a specific organ. This layered method is the engine behind many of the biggest recent breakthroughs.

Computer Vision (CV): This is a specialised field of artificial intelligence focused on enabling computers to 'see' and interpret medical images like X-rays, scans, and digital microscope slides. It uses the deep learning method to learn how to recognise important patterns in this clinical information.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

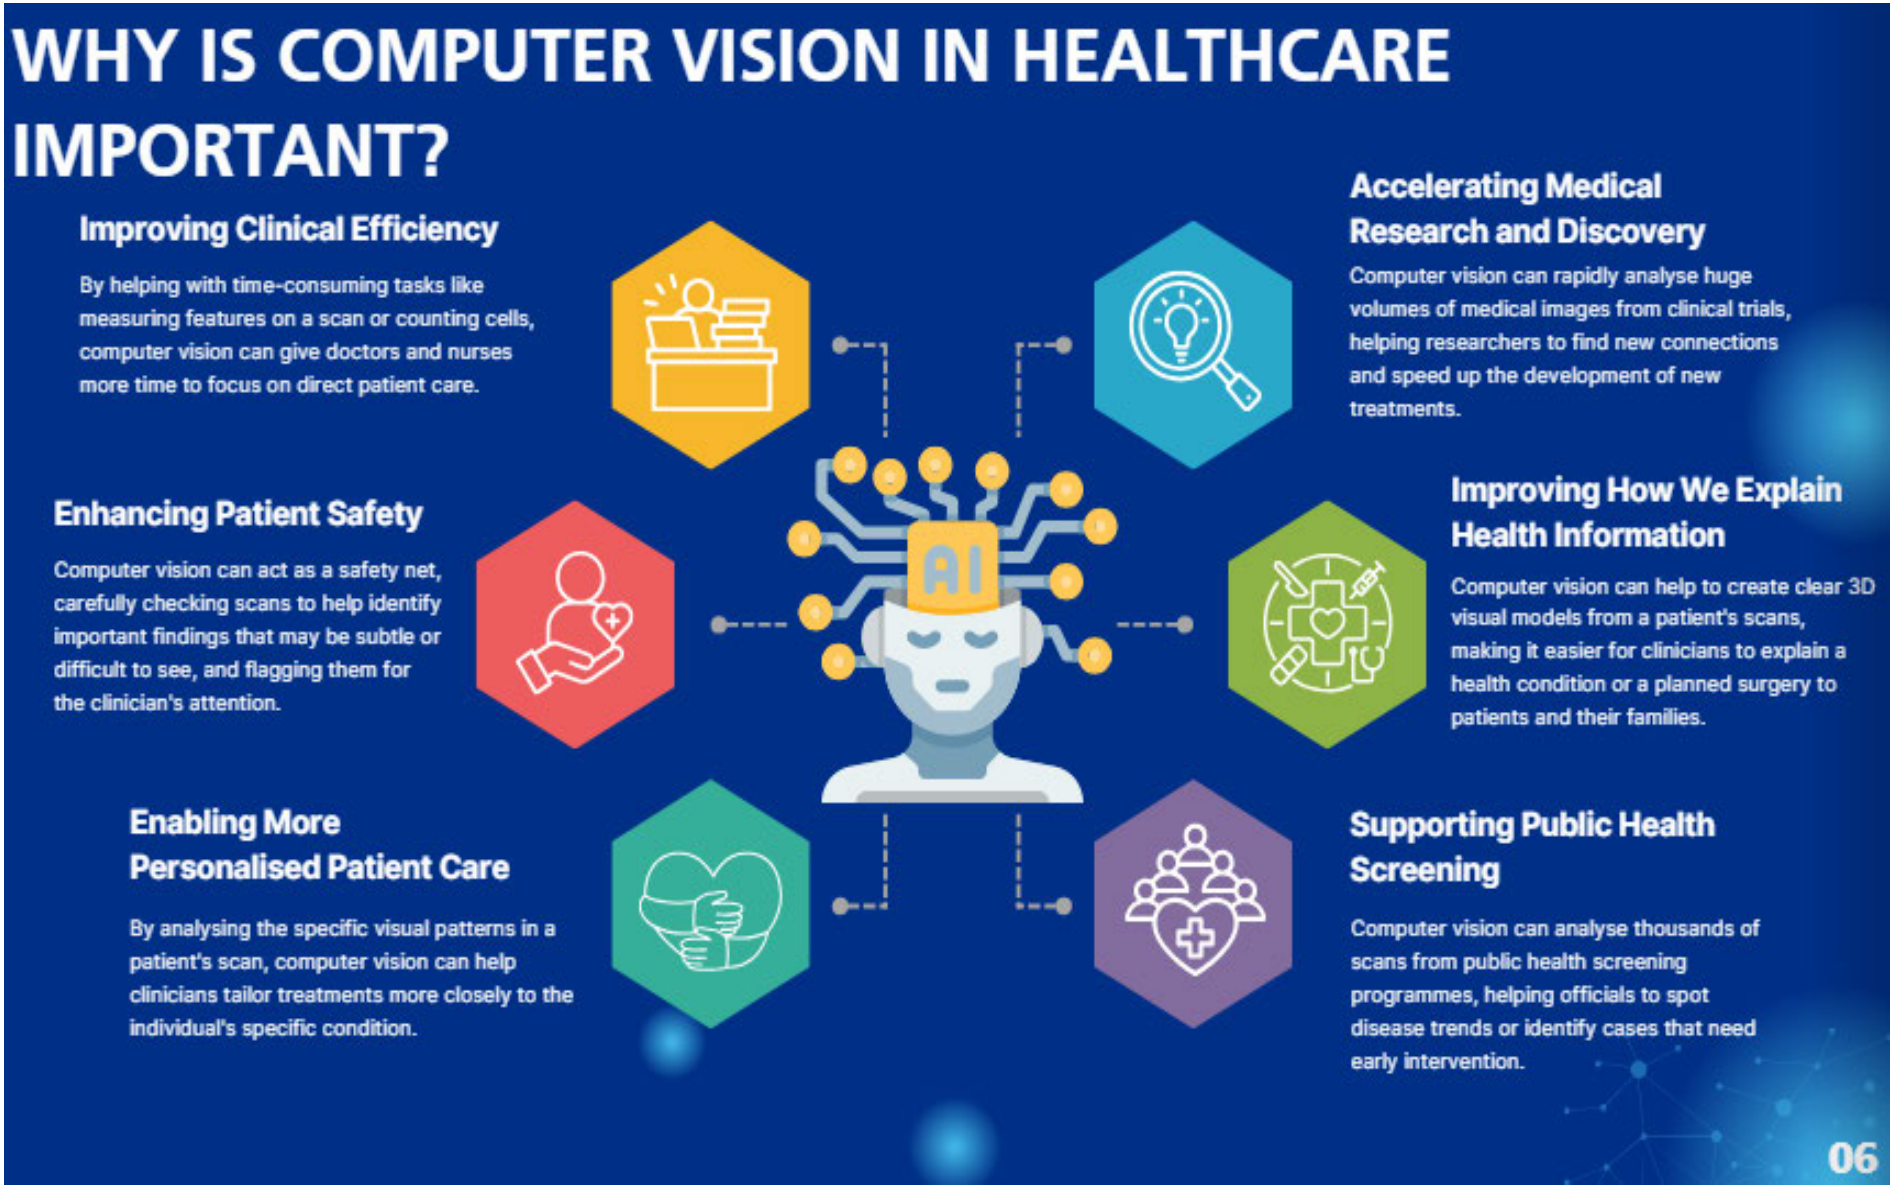

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

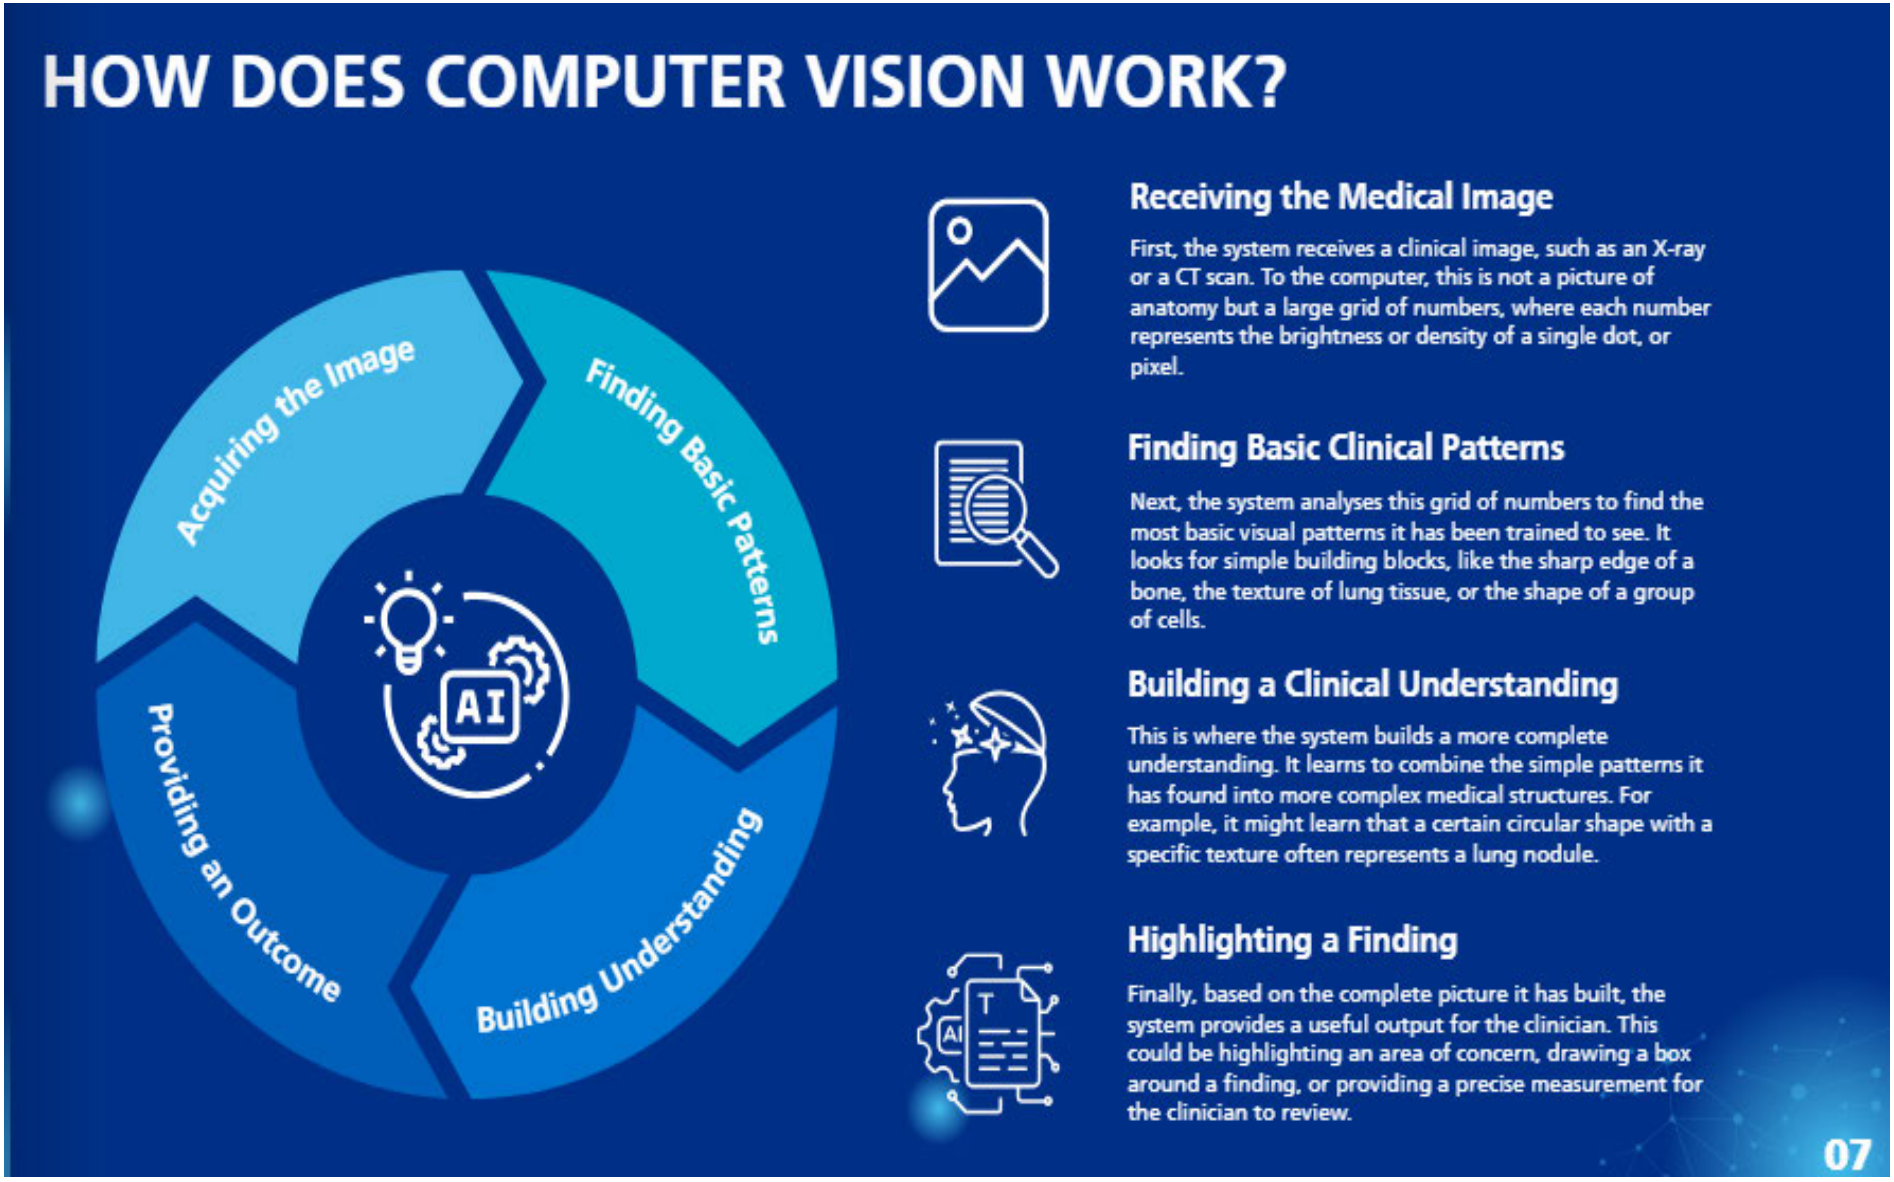

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION IN HEALTHCARE WORK?

Seeing a Medical Scan as Pixels

For a computer vision system to begin understanding a medical scan, it must first break the picture down into manageable units. A computer does not see a whole scan in the same way a person does. Instead, it sees the image as a large grid made up of many tiny, individual dots called pixels. Each pixel has a specific brightness or colour, which the computer represents as a number. This process of seeing a scan as a grid of numerical pixel values is the essential first step. Once the image is recognised as this collection of basic digital units, the system can then begin to analyse them to find patterns.

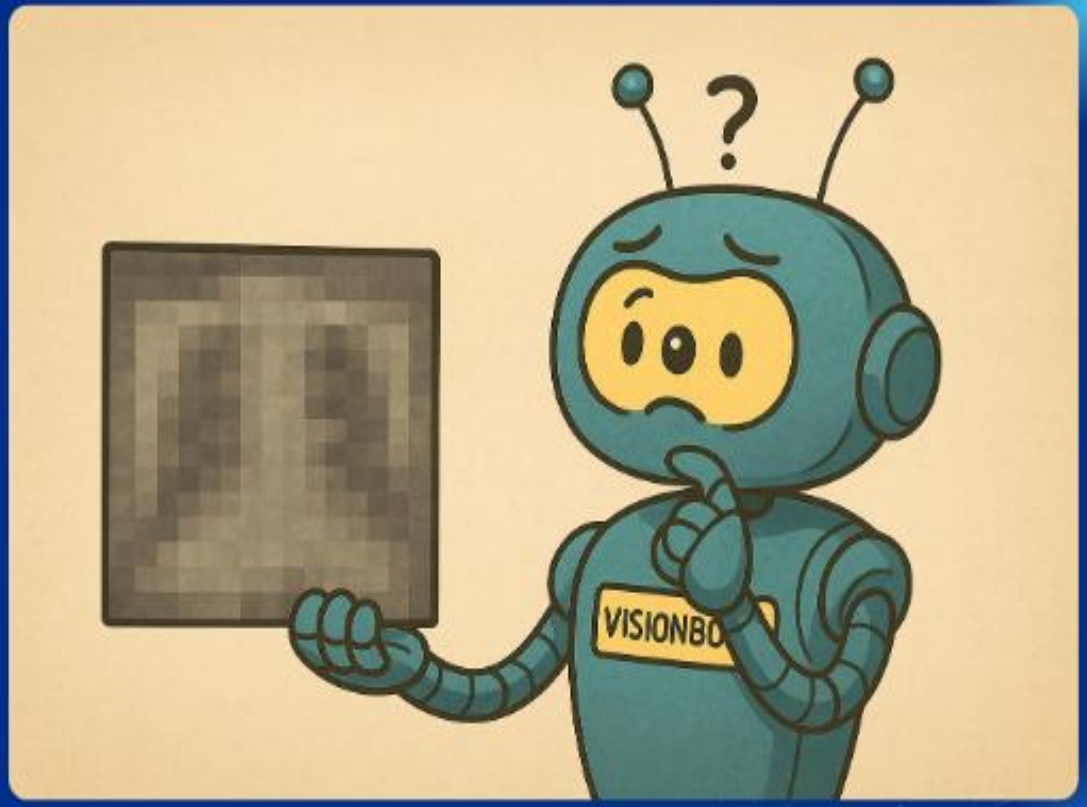

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION IN HEALTHCARE WORK?

Finding Basic Patterns - Identifying Visual Clues

After seeing the scan as a grid of pixels, the computer vision system must begin to find meaningful patterns within that grid. This stage involves looking for the most basic visual 'clues', such as the edge where a bone meets tissue, the specific shape of a cell, or the texture of an organ. This is like a detective scanning a room for evidence. The system uses many digital 'filters', where each filter is designed to find only one specific type of clinical pattern. For example, one filter might highlight all the sharp edges, while another finds all the small, circular shapes. By using many of these filters, the system creates a set of 'clue maps', showing where different basic patterns appear in the scan.

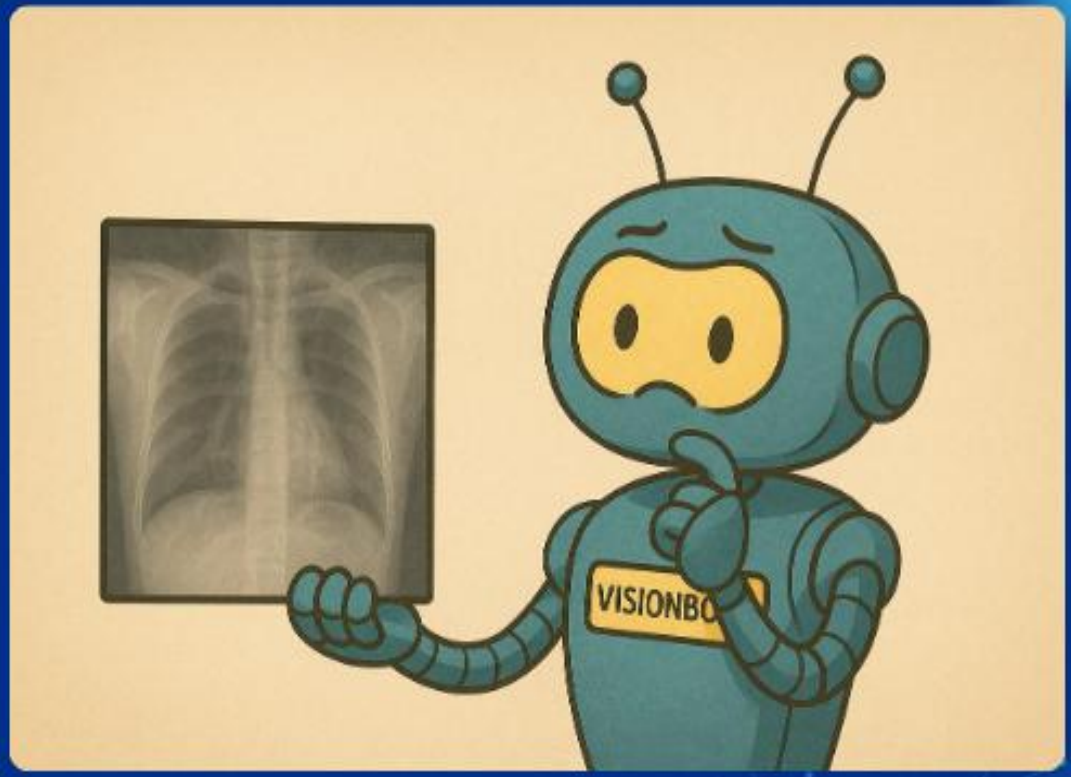

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES COMPUTER VISION IN HEALTHCARE WORK?

Building a Full Picture - Assembling the Clues

With the basic visual clues like edges, shapes, and textures identified, the computer vision system then works to understand how these pieces fit together to create a recognisable structure. This stage involves learning the relationships between the simple patterns. It is similar to building with LEGO bricks. The basic visual patterns, like the edge of a bone or the shape of a cell, are like the individual bricks. By learning how to connect these simple bricks in the right order and arrangement, the system can construct a complex model of an organ or other anatomical feature. By analysing these connections, the system builds a full picture of what the scan contains.

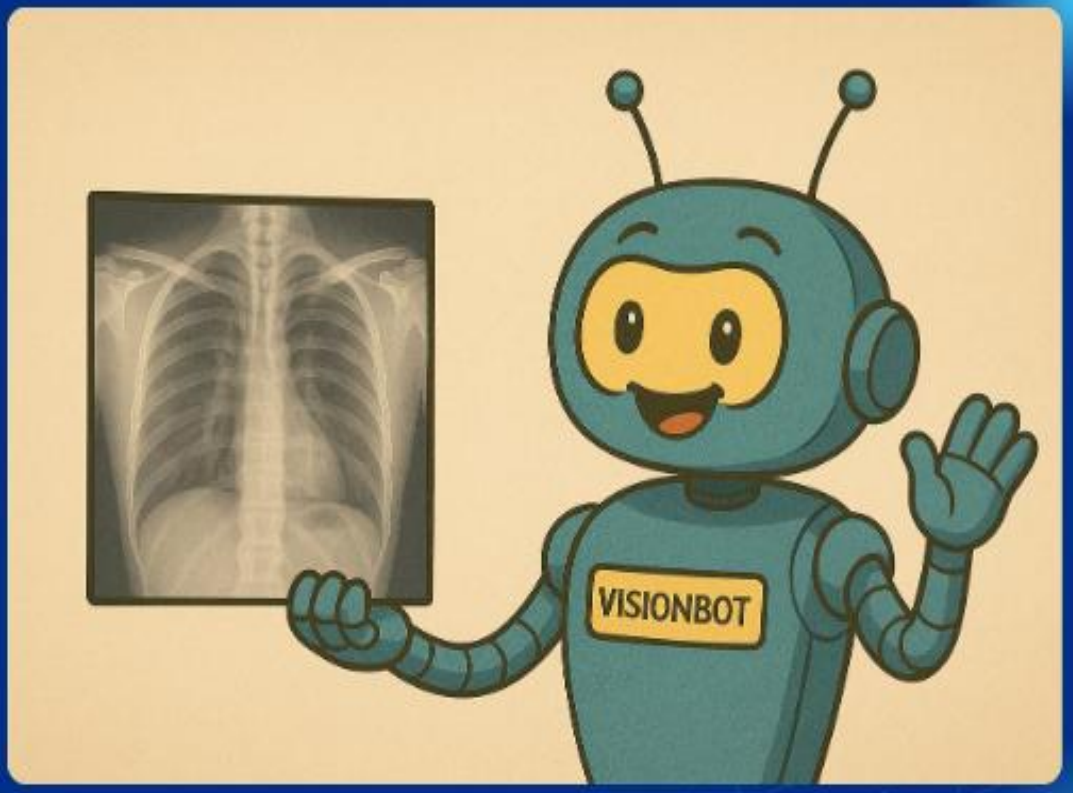

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT CAN COMPUTER VISION ACTUALLY DO IN HEALTHCARE?

01 Highlighting Clinical Findings

Computer vision can analyse medical scans and highlight areas of potential concern, such as a lung nodule on a CT scan. This acts as a helpful second check, ensuring that important findings are reviewed by a clinician.

02 Assisting Clinical Trials

Computer vision can assist with medical research by precisely measuring features in thousands of clinical trial scans. This helps researchers to assess if a new treatment is effective much more quickly and consistently.

03 Tracking Treatment Response

It is important to know if a treatment is working. Computer vision can compare a patient's scans over time to accurately measure very small changes, such as whether a tumour is shrinking, providing supporting evidence of a treatment's impact.

04 Providing Key Measurements

A complex scan contains lots of information. Computer vision can automatically provide key measurements, like the precise volume of a brain bleed. This automates a time-consuming task and gives clinicians important data quickly.

05 Creating 3D Visual Aids

Computer vision can turn a series of flat scan images into an intuitive 3D model. Surgeons can then use these clear visual aids to help explain a planned operation to patients and their families, improving understanding.

06 Prioritising Urgent Scans

In a busy hospital, computer vision can help by reviewing scans as they arrive and flagging those with potentially urgent findings. This helps clinicians to prioritise their worklists, ensuring the most critical cases are seen first.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CLINICIANS AND COMPUTER VISION WORKING TOGETHER

An Assistant, Not a Replacement

Computer vision should be seen as an advanced assistant for clinicians. It can handle routine visual tasks, but it does not replace the knowledge and judgement of a healthcare professional.

Clinician Oversight is Essential *(Human-in-the-Loop)*

It is vital that a clinician always reviews and checks any information provided by a computer vision system. The healthcare professional makes the final decisions about patient care.

Using the Tool Correctly

Clinicians play an active role by ensuring the system is used for the right task on the right type of medical scan. This ensures the technology provides the most useful and relevant support.

Training for Safe and Effective Use

For this partnership to be successful, healthcare staff need training. This helps them understand how the tools work, what their benefits are, and, just as importantly, what their limitations are.

Refining the Technology with Clinical Feedback

The partnership works both ways. When clinicians provide feedback on the performance of these tools, it helps developers to make them more accurate, safe, and useful in real clinical situations.

Focusing on Patient-Centred Care

This partnership aims to reduce the time clinicians spend on repetitive screen-based tasks. This frees them up to focus on direct patient interaction and making complex care decisions.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT CONSIDERATIONS FOR USING COMPUTER VISION IN HEALTHCARE

01 Patient Safety and Accuracy

In healthcare, accuracy is vital. Computer vision tools must be reliable, and their outputs must always be carefully checked by a clinician to ensure patient safety is the highest priority.

02 Data Privacy & Cyber Security

Protecting sensitive patient health information is essential. When computer vision systems use patient data, strict rules like the UK GDPR and Caldicott Principles must be followed to ensure data is secure and used appropriately. Robust cybersecurity defences, including encryption and secure networks, are vital to prevent breaches or attacks that could compromise data.

03 Fairness and Avoiding Bias

These systems learn from the scans they are shown. If the data is not diverse, they may be less accurate for some patient groups. Efforts are needed to ensure they work fairly for everyone.

04 Transparency and Explainability

It is important to understand how a system arrives at a conclusion. This "explainability" helps to build trust and allows clinicians to assess the reasoning behind a tool's suggestion.

05 Accountability and Responsibility

Clear lines of responsibility are essential. While the technology can assist, the healthcare professional remains accountable for all clinical decisions and patient outcomes.

06 The Complexity of Medical Images

Medical scans can have subtle variations and imperfections. A constant challenge is building systems that can handle this visual complexity with a high degree of reliability.

07 Regulatory Oversight

Computer vision models used in clinical settings, especially those influencing diagnosis or treatment, must meet the stringent safety, efficacy, and quality standards set by official medical regulatory bodies like the MHRA to ensure they are fit for purpose.

08 Integration with Clinical Workflows

For these tools to be helpful, they must fit smoothly into the day-to-day work of clinicians and be easy to use with existing hospital computer systems, like the image viewing software.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

REAL-WORLD APPLICATIONS OF COMPUTER VISION IN HEALTHCARE

Computer vision is already being used in a variety of practical ways within healthcare, often working behind the scenes to help clinicians analyse and interpret medical images. This field of artificial intelligence helps computers to 'see' and find patterns in scans, leading to better-supported clinical decisions, new ways to monitor patient health, and innovative solutions for patient care and medical research.

The real-world applications of this technology in medicine are continuously growing. They are helping healthcare organisations and professionals to analyse visual information more effectively, support clinical judgements, and shape a future where technology can better assist with the complexities of medical imaging.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION FOR CLINICAL DOCUMENTATION

Improving accuracy and saving time for clinicians.

AUTOMATED MEASUREMENTS

A computer vision tool can be used to precisely measure features on a scan, such as the size of a tumour. This provides specific, recordable data that can be added to a patient's notes to help track their condition.

CELL COUNTING

For pathologists, a tool can automatically count thousands of cells on a digital slide. This helps to provide a score for certain biomarkers that are important for cancer grading and treatment planning.

ORGAN VOLUME CALCULATION

Computer vision can calculate the exact volume of an organ, like the liver, from a CT or MRI scan This information is useful for monitoring diseases that cause organs to change in size over time.

CARDIAC FUNCTION ANALYSIS

In cardiology, these tools can automatically calculate the heart's ejection fraction from an echocardiogram. This is a key measure of heart health that can otherwise rely on a visual estimation.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION FOR CLINICAL DECISION SUPPORT

Providing helpful prompts to support clinical judgement.

FRACTURE DETECTION

A tool can analyse an X-ray, such as one of a wrist or ankle, and highlight a faint line that may be a subtle fracture. This prompt ensures the area is carefully checked by the clinician.

EMERGENCY STROKE TRIAGE

In emergencies, some tools can analyse a head scan and immediately alert the stroke team if a major blood vessel blockage is suspected. This allows for faster treatment when every minute counts.

LUNG NODULE DETECTION

During lung cancer screening, a computer vision system can highlight small, suspicious nodules on a CT scan. This acts as a second check for the radiologist, helping to find cancers at an earlier stage.

POLYP DETECTION

During a colonoscopy, a computer vision tool can analyse the live video feed and draw a box around polyps in real time. This has been shown to help doctors find more polyps and prevent future cancers.

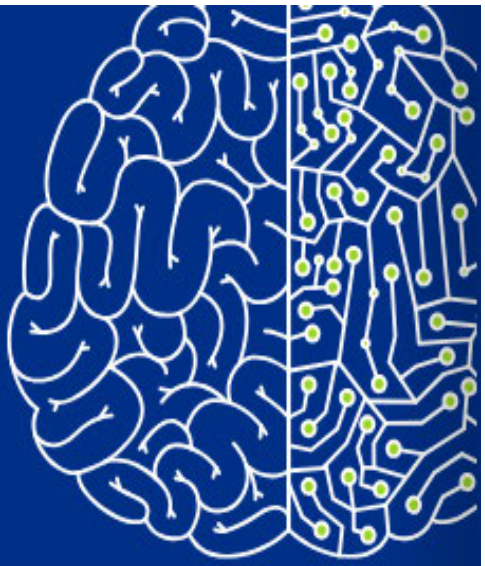

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION FOR PATIENT COMMUNICATION AND ENGAGEMENT

Using visual information to improve patient understanding.

CREATING 3D SURGICAL MODELS

Computer vision can turn a series of flat 2D scans into a detailed 3D model of a patient's own anatomy. Surgeons can use these clear, intuitive models to help explain a planned operation.

VISUALISING TREATMENT PLANS

A 3D model of a patient can be used to show exactly where treatment, such as radiation therapy, will be targeted. This helps patients to better understand what their treatment will involve.

SHOWING PROGRESS VISUALLY

It can be powerful to show a patient how their condition is improving. Computer vision can create before-and-after 3D models to, for example, visualise how a tumour has shrunk in response to treatment.

EXPLAINING ANATOMICAL FINDINGS

For something like a complex joint injury, a 3D model from an MRI allows a doctor to show a patient exactly what is wrong and how it will be treated, making it much easier to understand.

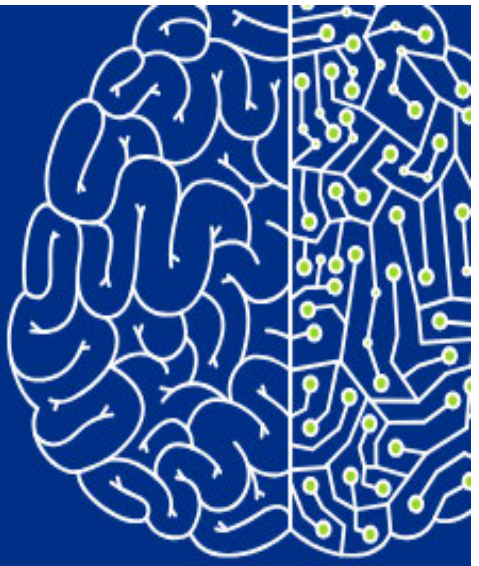

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

COMPUTER VISION FOR MEDICAL RESEARCH AND DISCOVERY

Speeding up scientific progress and the development of new treatments.

ANALYSING CLINICAL TRIAL IMAGES

Computer vision can analyse thousands of scans from a clinical trial to measure treatment effects. This allows researchers to compile a wealth of detailed visual data from the trial, supporting them in analysing the results more thoroughly.

FINDING NEW IMAGING 'BIOMARKERS'

These tools can sometimes detect very subtle patterns in scans that are invisible to the human eye. Researchers analyse these patterns to see if they can predict how a patient might respond to a new drug.

UNDERSTANDING DISEASE AT SCALE

Researchers can use computer vision to analyse medical images from thousands of patients. This helps them to better understand how a disease progresses or how it might appear differently across a large population.

VALIDATING NEW IMAGING METHODS

When a hospital gets a new type of scanner, computer vision can help to rapidly compare the new images with old ones. This helps to ensure that the new equipment is working correctly and producing high-quality scans.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CHALLENGES AND IMPORTANT QUESTIONS

01 Patient Privacy

How do we ensure sensitive patient scans are always kept completely private and secure when using these powerful tools?

02 Accuracy and Reliability

Computer vision systems are not perfect and can make mistakes. How do we ensure all findings are carefully checked by a clinician to guarantee patient safety?

03 Fairness and Bias

If the scan data used for training contains biases, the system can learn and repeat them. How do we prevent this from creating unfair disadvantages in healthcare?

04 The Complexity of Medical Images

Medical scans vary and can have imperfections. A major challenge is building systems that can understand this visual complexity with a high degree of reliability.

05 Clear Roles and Responsibilities

How do we ensure there are clear guidelines for how these tools are used as part of the clinical team, and that everyone understands their specific role in the process?

06 Understanding How Decisions Are Made

It can be difficult to know exactly how a system arrived at a suggestion. How can we make its reasoning more transparent to build trust with clinicians?

07 Regulatory Standards

How should these advanced clinical tools be regulated to ensure they meet the high standards for safety and effectiveness required in medicine?

08 Integration with Clinical Workflows

How can we ensure that new tools are designed to fit smoothly into the day-to-day work of busy clinicians and are easy to use?

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

THE FUTURE OF COMPUTER VISION IN HEALTHCARE

Looking ahead, we can expect computer vision to become even more seamlessly integrated into our clinical imaging tools and digital health services. It will continue to refine how we analyse medical scans and visualise patient information, making technology feel more like a helpful partner in delivering care. Ongoing advancements are consistently enabling these systems to see more subtle clinical signs and create richer, more detailed 3D models. The long-term journey of teaching computers to truly understand the depth and nuance of complex medical imaging is very much still in progress. As these technologies become more powerful, there is a strong and growing focus on guiding their development responsibly, ensuring that progress aligns with human values and the core principles of ethical medical practice.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 Unlocks Visual Medical Information

It is a key tool for finding subtle patterns and precise measurements that are hidden within the complex data of medical scans.

02 A Tool to Support, Not Replace, Clinicians

This technology is best seen as a helpful assistant that can handle time-consuming visual tasks, giving clinicians more time to focus on direct patient care.

03 The Clinician is Always in Charge

A qualified healthcare professional must always use their expert judgement to review any information highlighted on a scan and make the final clinical decision.

04 Responsible Use is Essential

Patient safety, data privacy, and fairness must be the highest priorities when developing and using these visual analysis tools in a healthcare setting.

05 Partnership Between People and Technology

For these tools to be successful and safe, they require ongoing collaboration between the technology developers and the clinical staff who use them.

06 The Goal is Always to Improve Patient Outcomes

Ultimately, this technology is used to enhance patient safety, accelerate medical research, and help make healthcare more effective and personalised for everyone.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 3

Generative AI – A basic overview

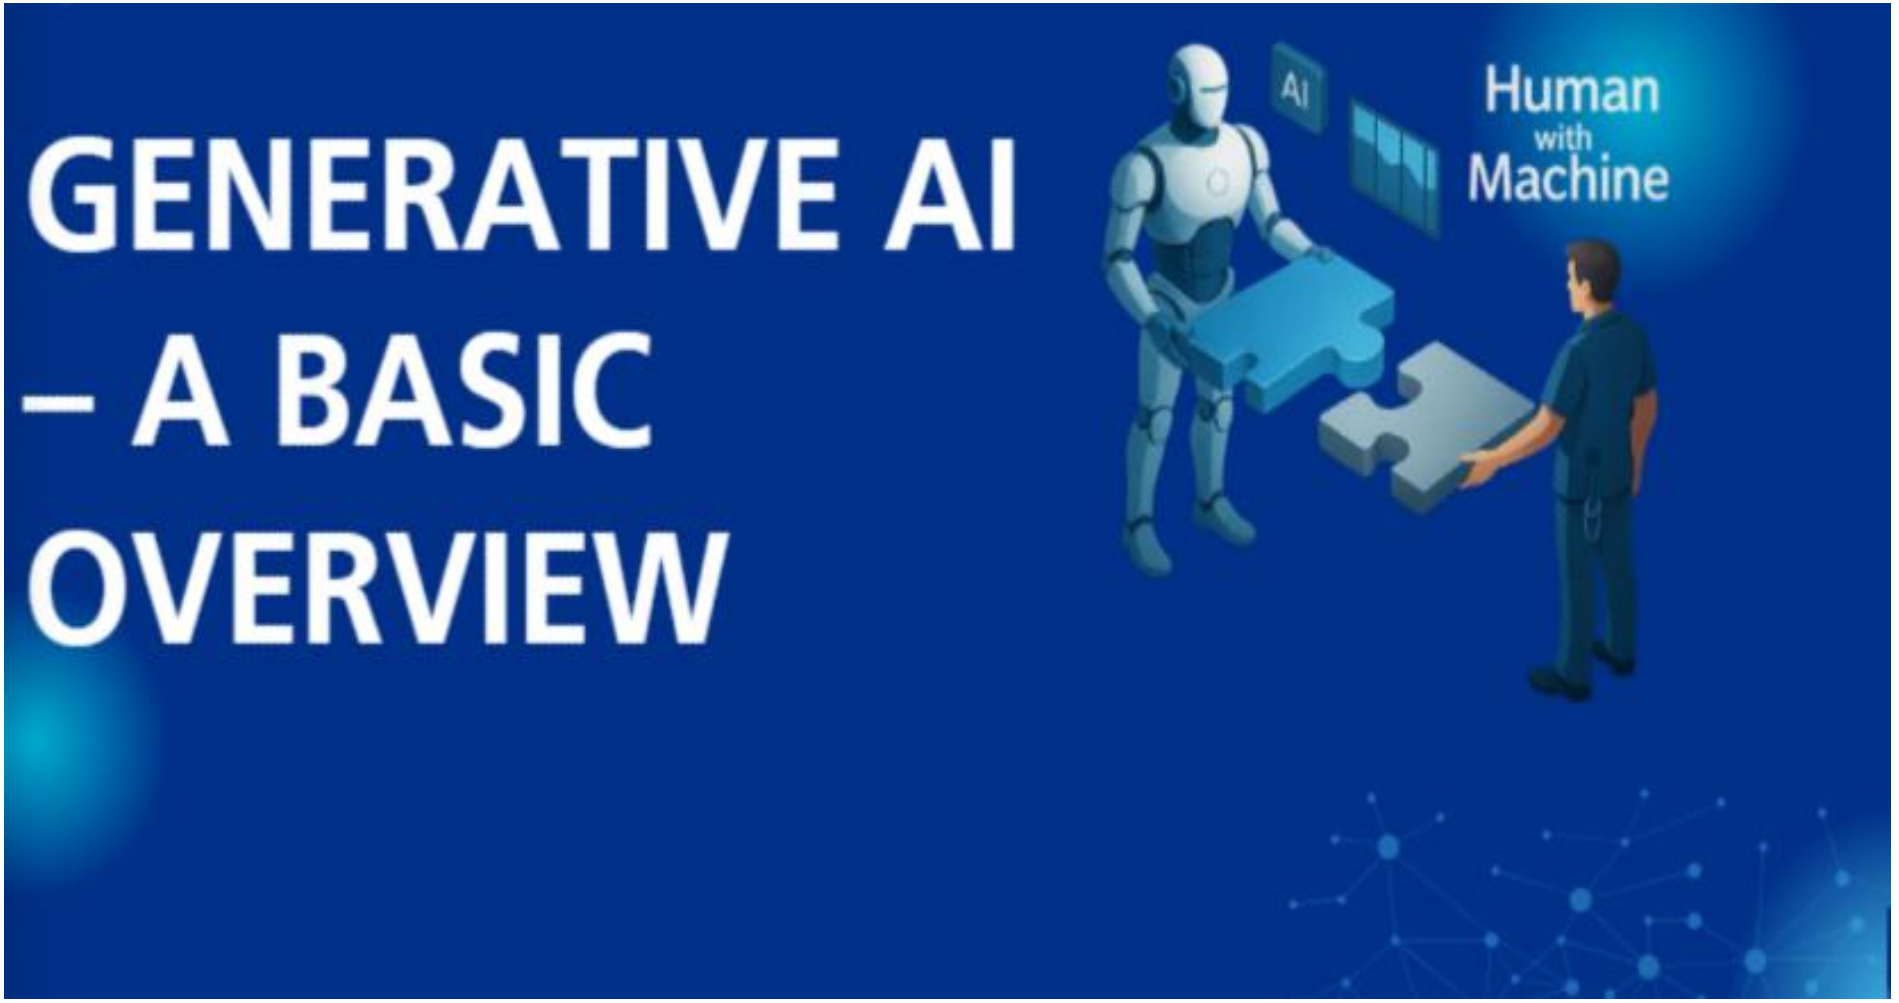

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS GENERATIVE AI?

Generative AI is about enabling computers not just to understand information, but to create entirely new content on their own. This allows generative AI to assist us in novel ways such as generating text, images, music and code by learning patterns and styles from vast amounts of existing data. Unlike traditional AI, which might analyse or classify, generative AI acts like a creative partner, synthesising new outputs that resemble human creativity. By learning from examples, these systems can produce diverse and useful content, from drafting emails and reports to designing artwork and composing music.

First, the generative AI learns by studying lots of existing examples, like looking at different paintings of a lighthouse to understand its features and common styles.

Then, using the patterns learned, the generative AI can generate something completely new and original, like painting a unique version of the lighthouse.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

GENERATIVE AI IN THE AI FAMILY

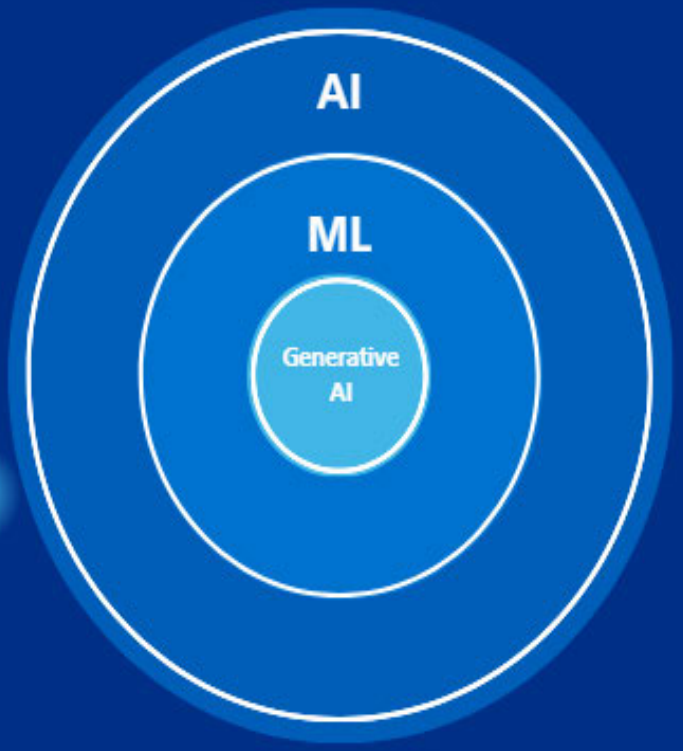

Artificial Intelligence (AI): AI is a wide ranging field focused on creating computer systems that can perform tasks normally requiring human intelligence, such as learning, reasoning, problem-solving, and decision-making. It is the overall goal of making machines intelligent.

Machine Learning (ML): ML is a core method used to achieve AI. Instead of being explicitly programmed for every task, ML systems learn patterns and make predictions or decisions by analysing large amounts of data. It is how many AI systems acquire their intelligence.

Generative AI: Generative AI is a specific type of AI, often built using ML techniques like deep learning. Its unique focus is on creating new, original content, like text, images, or code rather than just analysing existing data or making predictions based on it. It learns the underlying patterns of data to generate novel examples.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

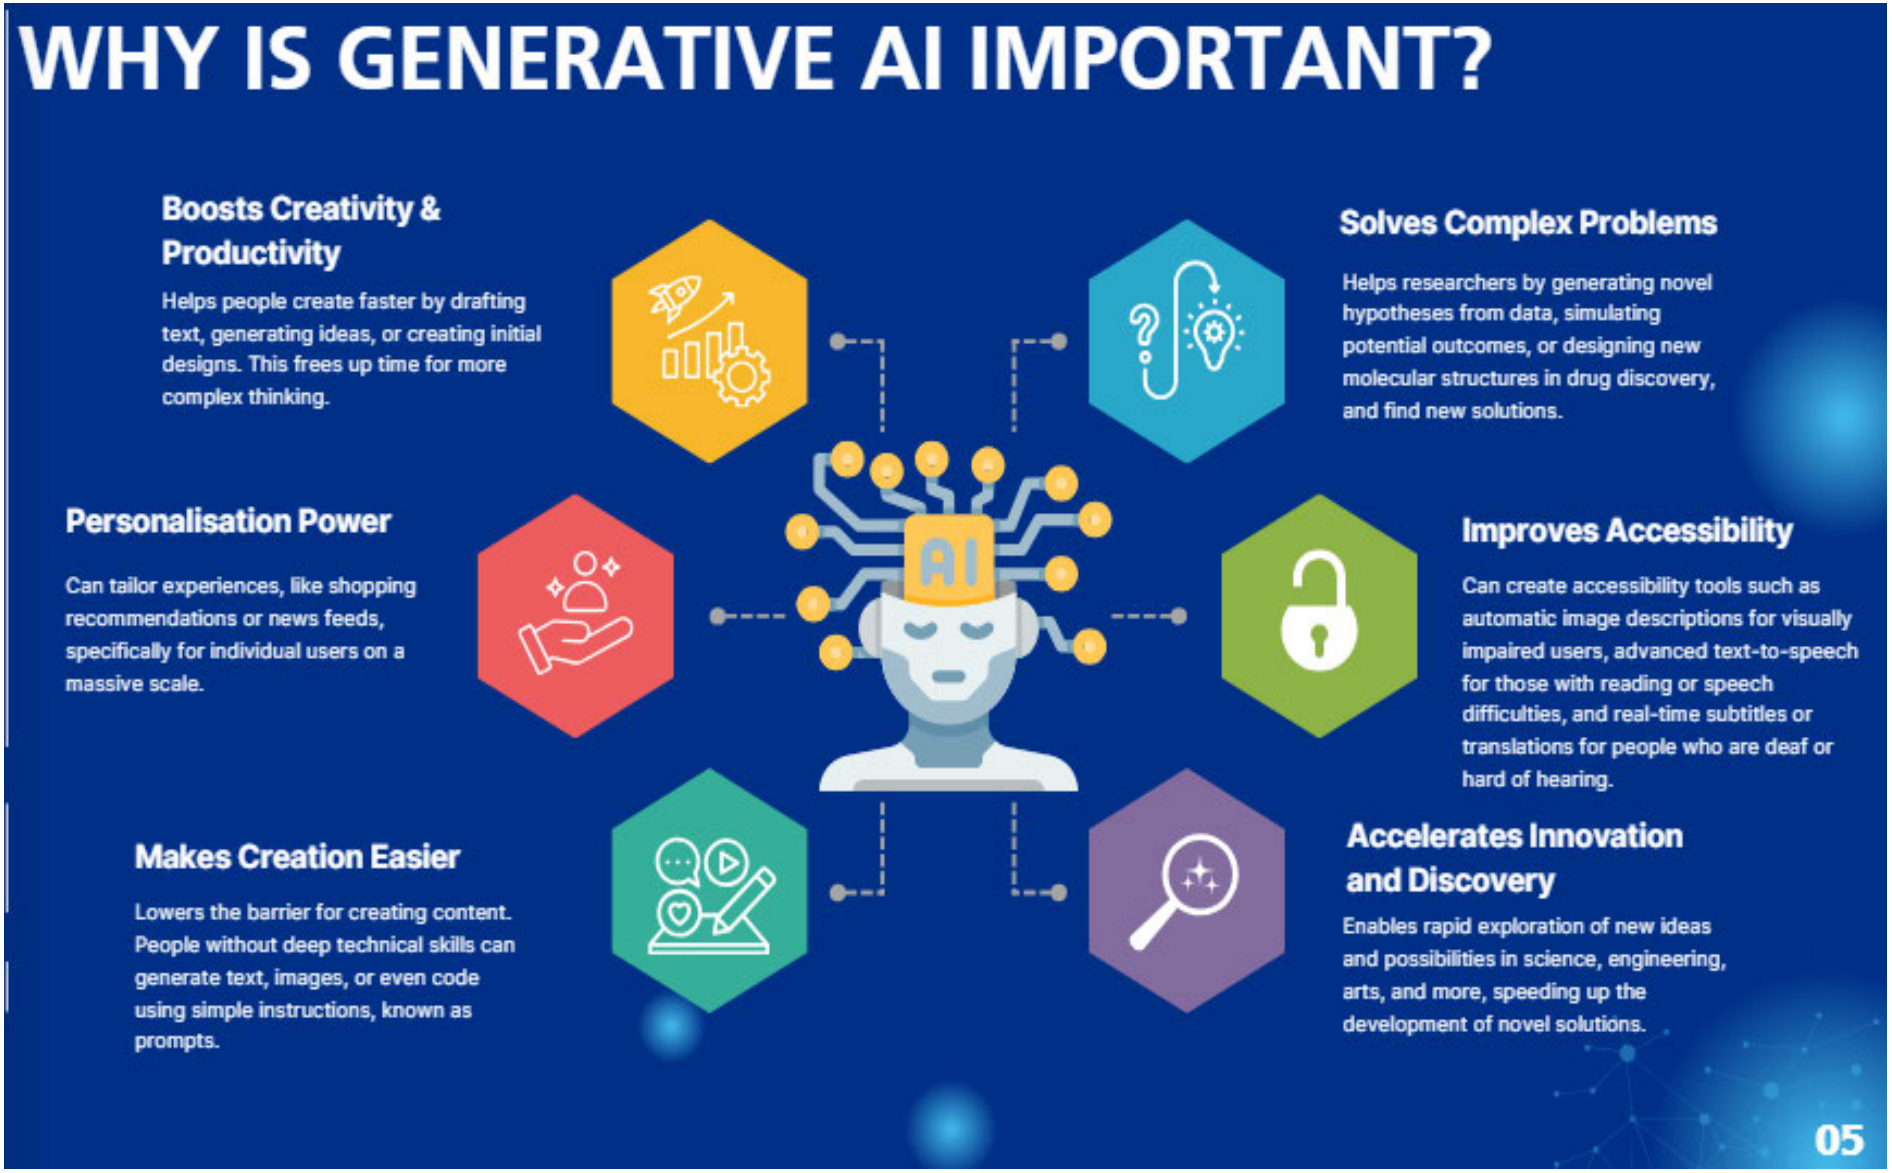

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

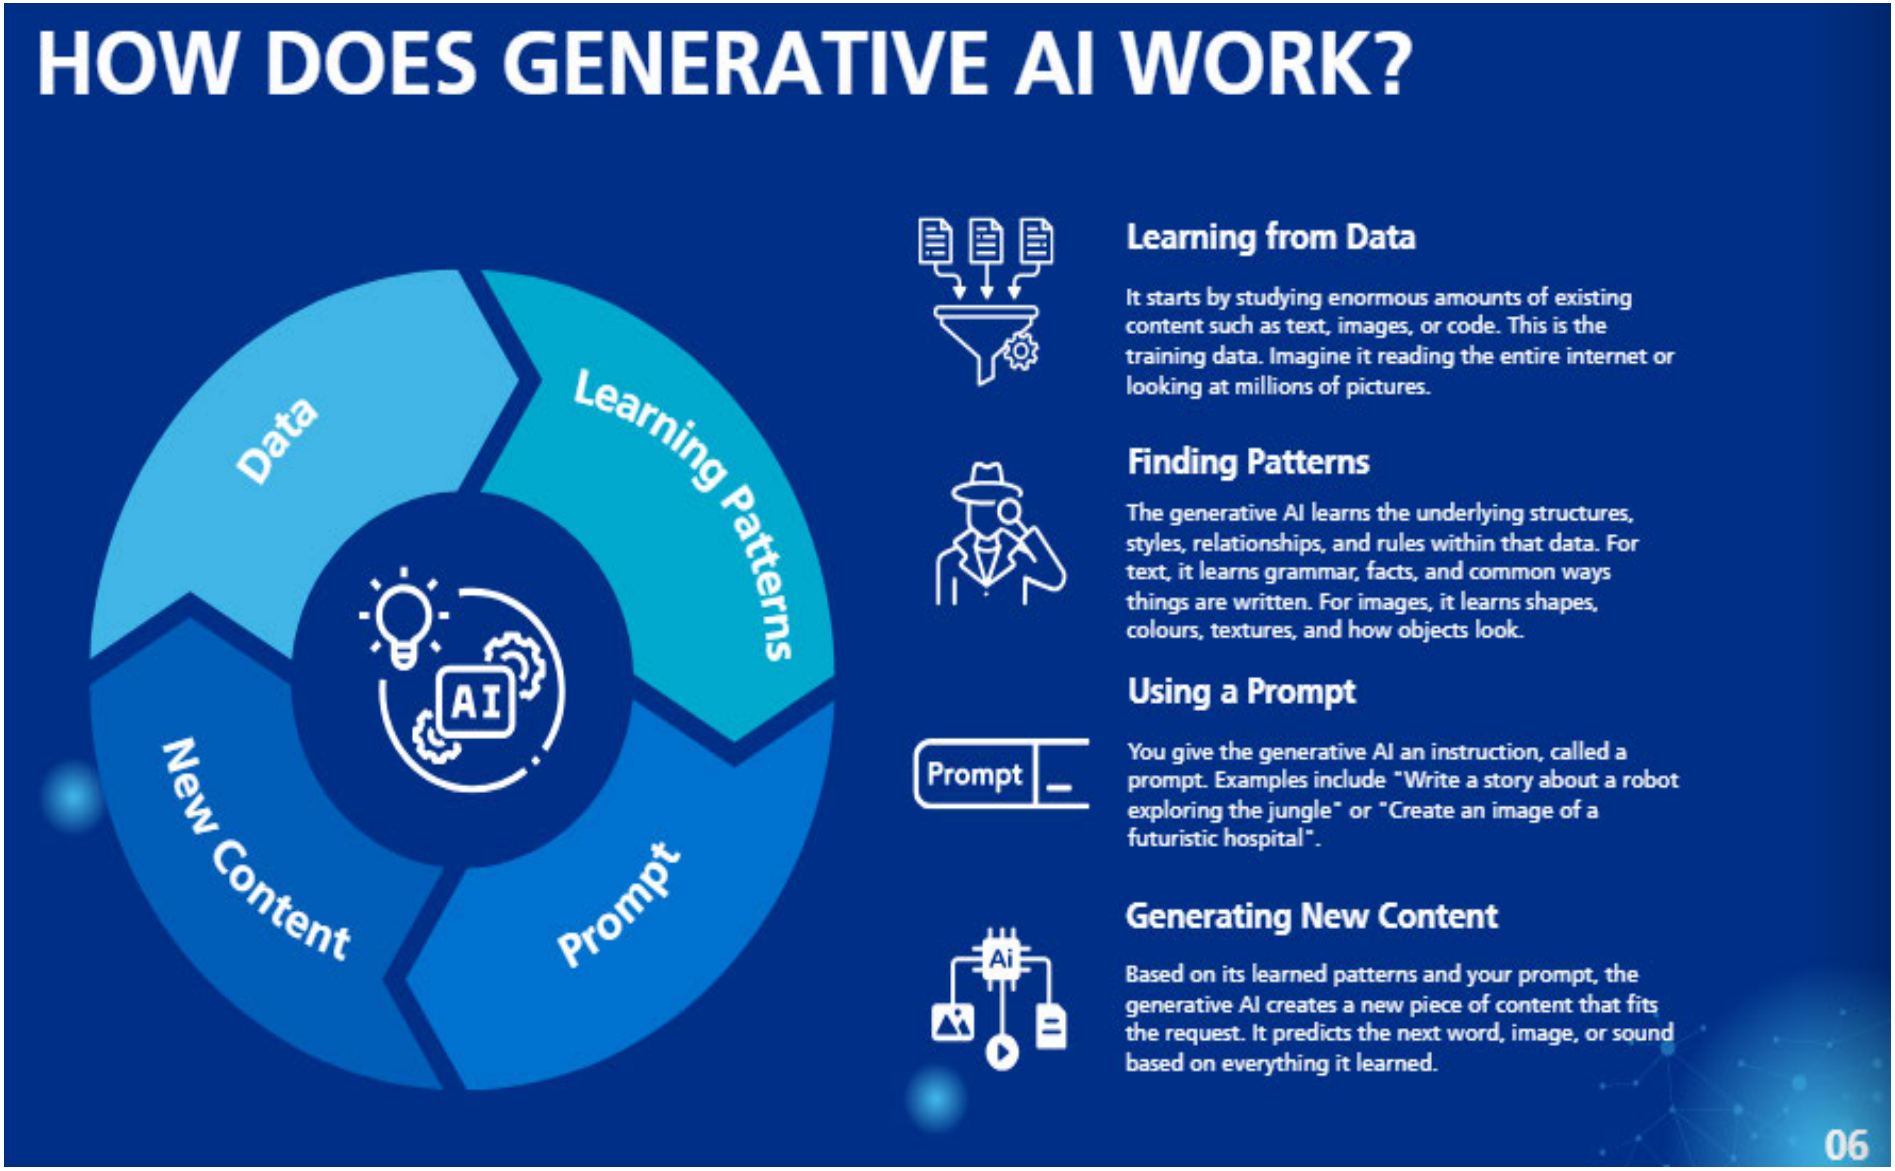

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES

The Smart Story Writer

One way generative AI creates is by being a super-smart predictor. For text, it looks at what has been written and predicts the best word to add next, then the next, and so on, building up sentences and stories piece by piece.

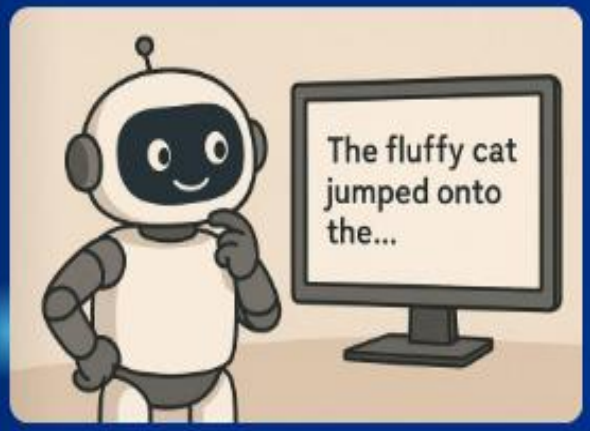

The generative AI looks at the current text to understand the context.

It thinks of many possible words that could fit next.

It chooses the best word and continues, building the story piece by piece.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES

Sculpting from Digital Noise

Another way generative AI creates, especially for images, is by starting with something random or very basic, like a fuzzy pattern. Then, step-by-step, it carefully shapes and refines this pattern until it becomes a clear and detailed new picture.

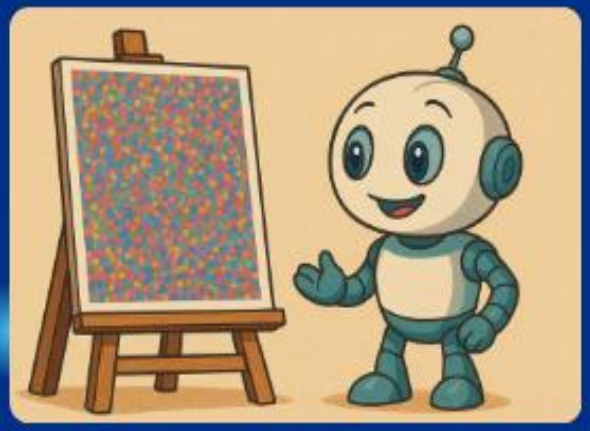

The generative AI starts with a canvas of random digital "noise"

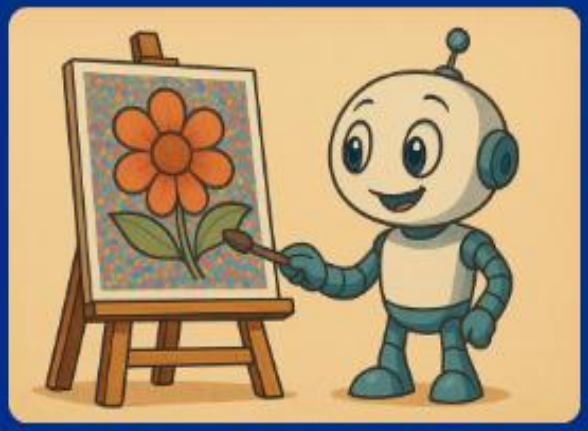

Guided by the user's request, it begins to shape the noise.

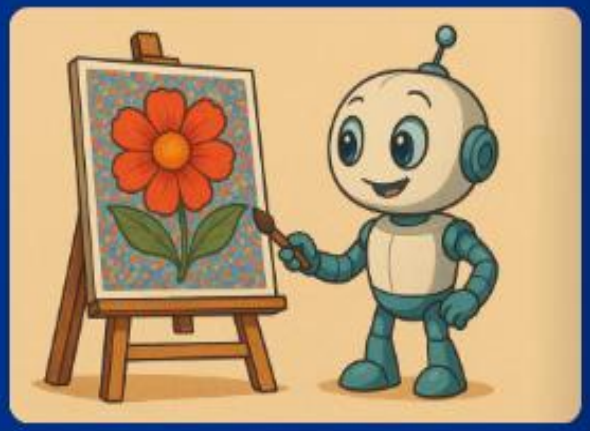

Step-by-step, it refines the noise into a detailed, original image.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES

The Creator and the Critic

Sometimes, two generative AIs work as a team. One AI, the 'Creator', tries to make new things. The second AI, the 'Critic', judges how good or realistic the creation is. By constantly getting feedback from the Critic, the Creator learns to make better and more believable new content.

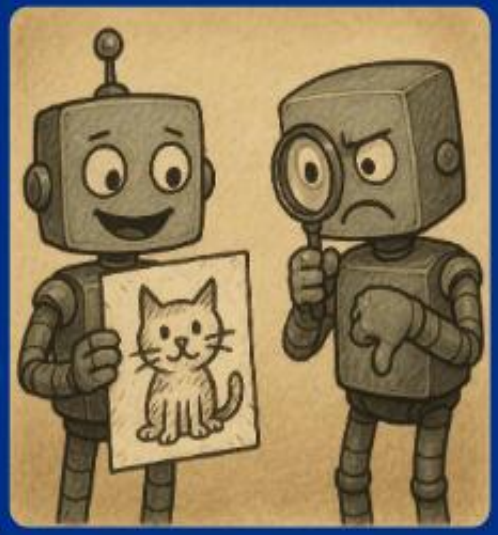

The "Creator" AI makes a new picture. The "Critic" AI judges if the picture looks real enough.

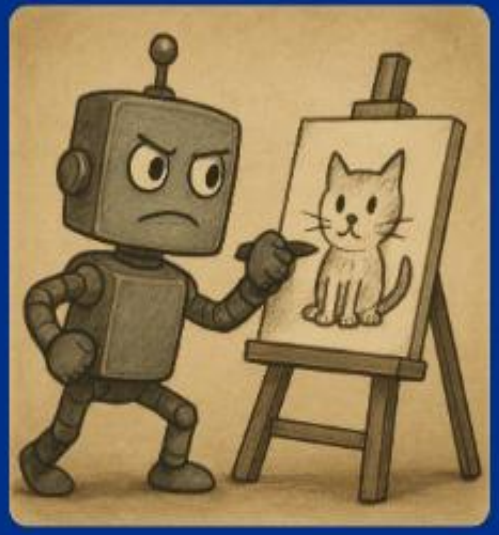

The "Creator" tries again, learning from the feedback.

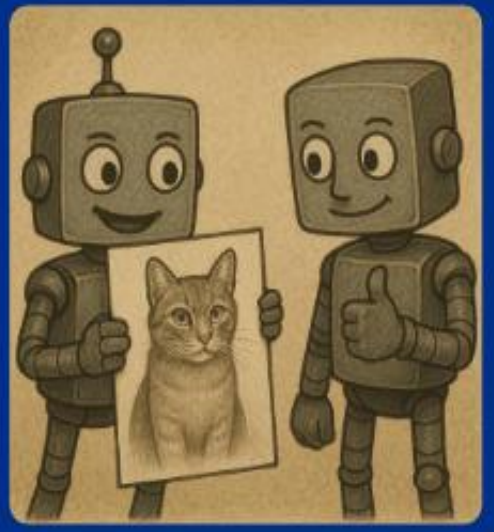

Through this teamwork, the "Creator" learns to make very believable new images.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

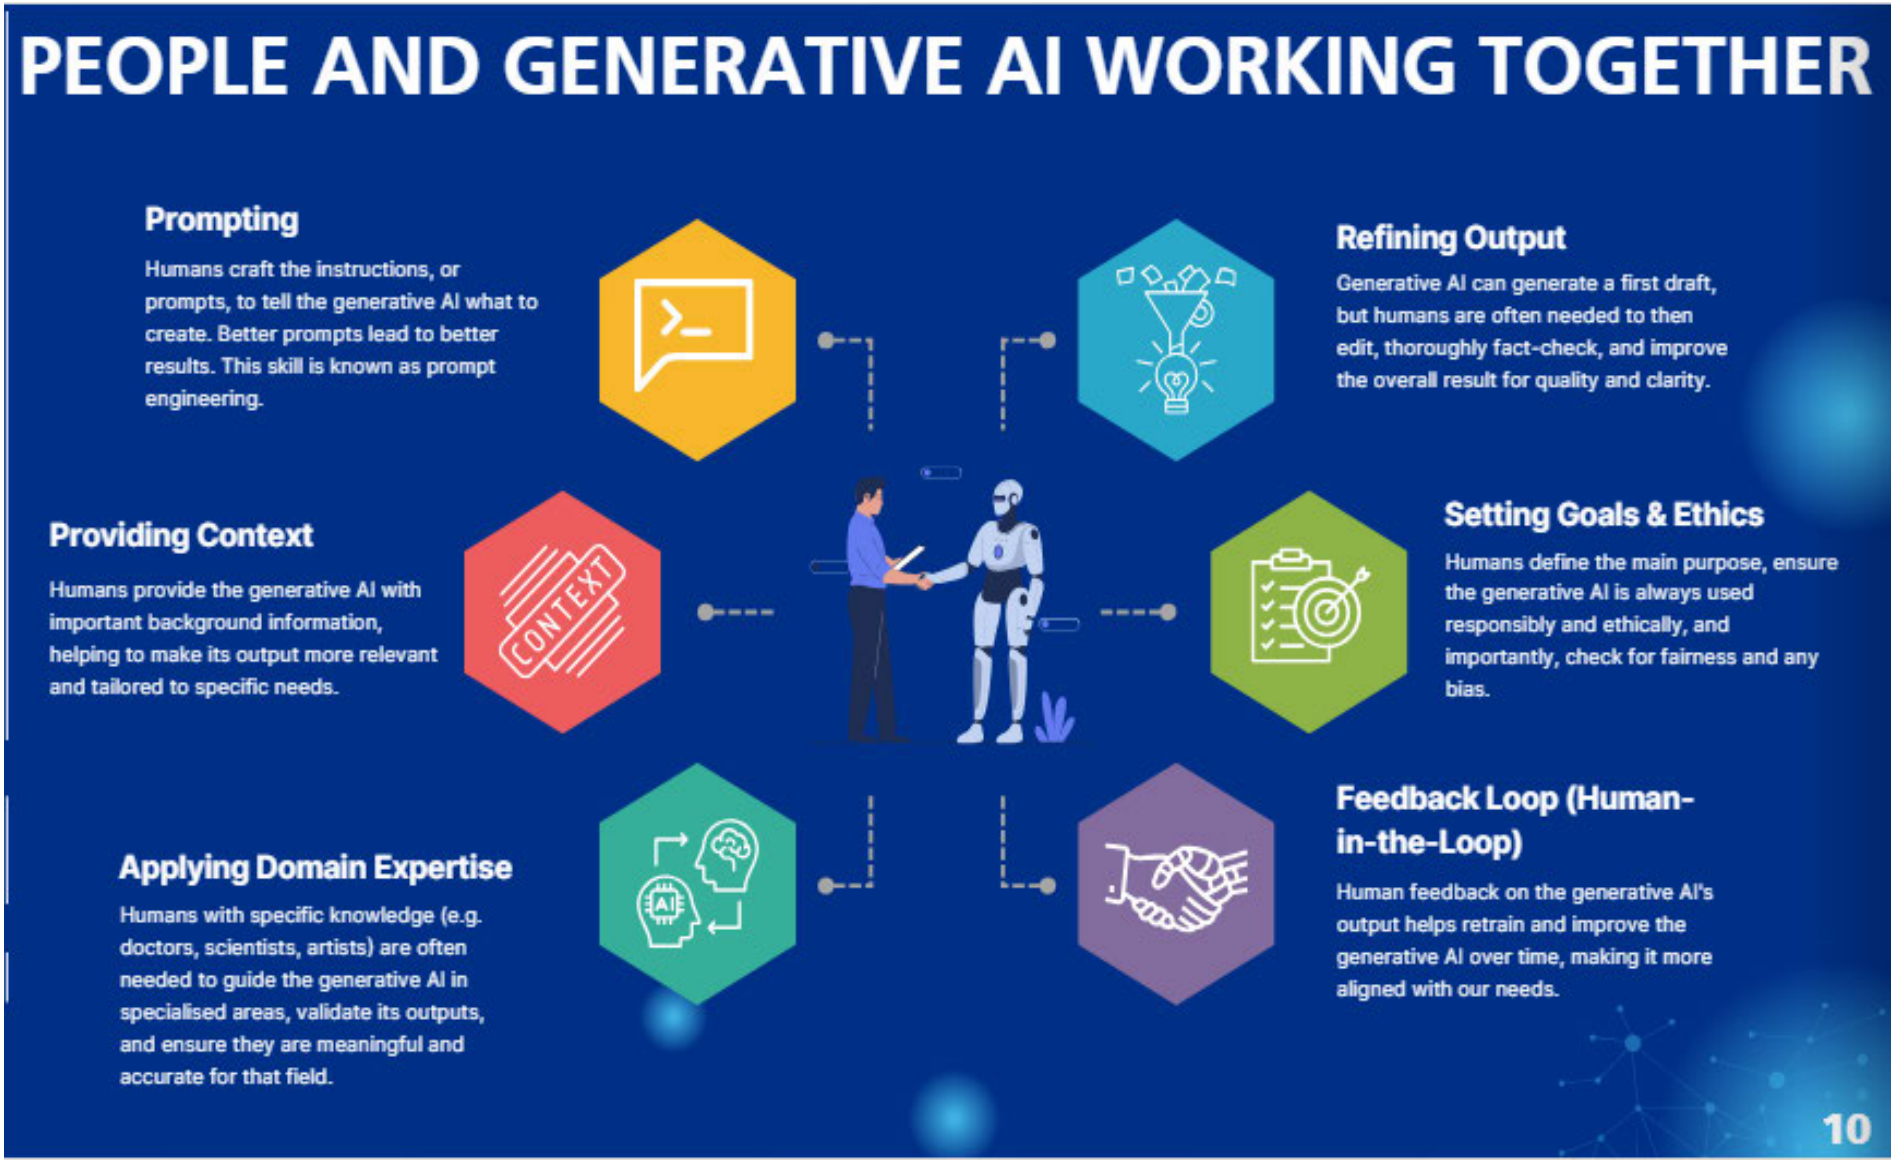

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

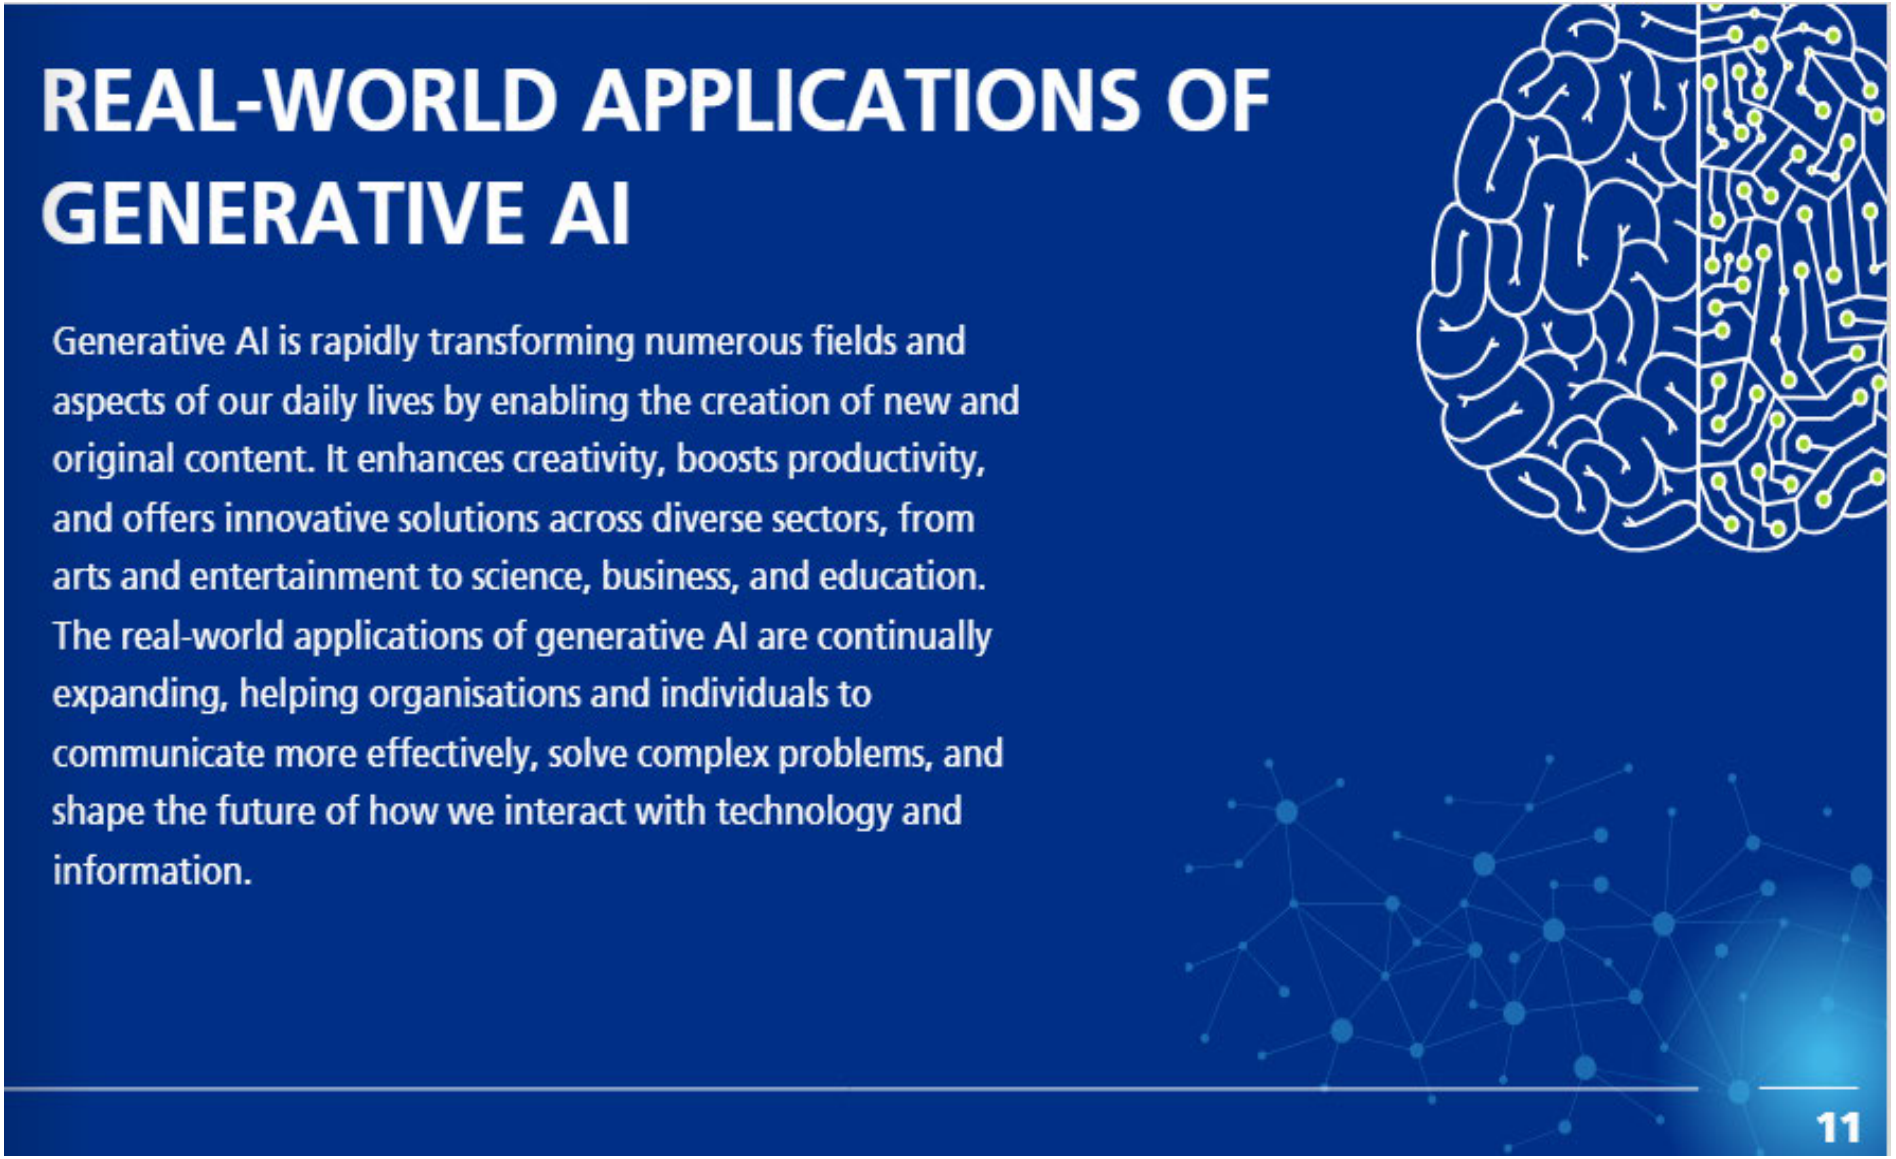

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

TEXT GENERATION

Empowering communication, creativity, and development.

AUTOMATED CONTENT DRAFTING

Generative AI produces initial drafts for emails, reports, articles, and marketing copy, significantly reducing the time spent on writing routine communications and correspondence.

CREATIVE WRITING ASSISTANCE

Generative AI acts as a co-writer, helping authors and scriptwriters develop plotlines, character dialogues, poetry, and song lyrics by offering suggestions and overcoming creative blocks.

CODE GENERATION & ASSISTANCE

Generative AI tools write and complete code snippets in various programming languages, explain complex code, identify bugs, and even generate documentation, boosting developer productivity.

SUMMARISATION & INFORMATION EXTRACTION

Generative AI can read lengthy documents or articles and produce concise summaries, extracting key information and insights quickly and efficiently for easier understanding.

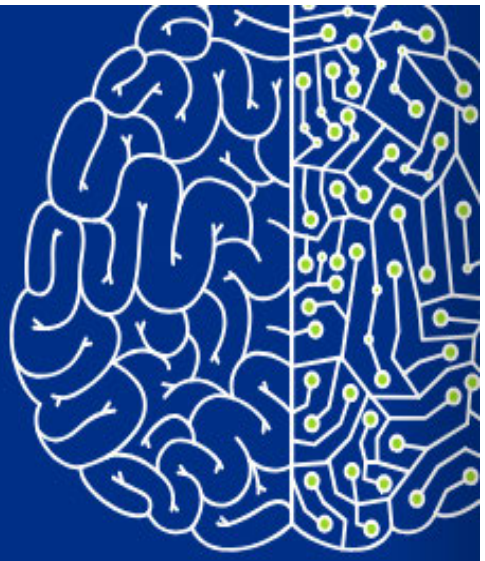

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMAGE GENERATION

Bringing concepts and artistic expressions to visual life.

ORIGINAL ARTWORK & ILLUSTRATIONS

Generative AI creates entirely new images, illustrations, and digital art from text descriptions, allowing users to bring imaginative concepts to visual life with diverse styles.

PRODUCT DESIGN & PROTOTYPING

Generative AI produces realistic visualisations of products and concepts, enabling designers to iterate quickly on ideas and present them effectively before physical creation.

MARKETING & ADVERTISING VISUALS

Generative AI generates custom images and graphics for advertising campaigns, social media content, and presentations, tailored to specific brand needs and target audiences.

CONCEPT ART FOR MEDIA & ENTERTAINMENT

Generative AI tools assist in creating concept art for films, video games, and animations, helping to visualise characters, environments, and scenes rapidly for production pipelines.

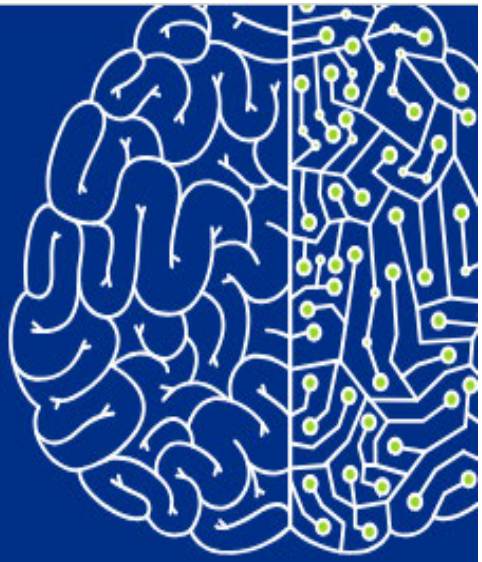

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

ENTERTAINMENT & GAMING

Creating immersive worlds and dynamic content.

PROCEDURAL CONTENT GENERATION

Generative AI generates vast and varied game environments, levels, characters, and items, reducing manual creation time and enabling more dynamic and unique game worlds.

AI-POWERED NON-PLAYER CHARACTERS (NPCS)

Generative AI can give NPCs more realistic and dynamic personalities, dialogues, and behaviours that respond intelligently and convincingly to player actions and the game environment.

VISUAL EFFECTS & ANIMATION ASSISTANCE

Generative AI tools assist in creating stunning visual effects, generating character animations, and even producing short animated sequences from scripts or storyboards for film and games.

INTERACTIVE & EVOLVING NARRATIVES

Generative AI can be used to create storylines that adapt and change based on player choices and actions, leading to more personalised, replayable, and emergent gaming experiences.

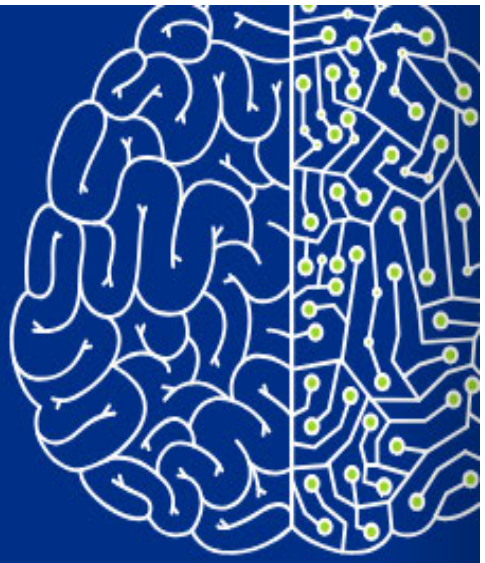

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

ACCESSIBILITY: ENHANCING INCLUSION

Breaking down barriers and promoting equal access.

AUTOMATED IMAGE DESCRIPTIONS (ALT TEXT)

Generative AI generates descriptive text for images, making visual content on websites, documents, and social media accessible to visually impaired users via screen readers.

ADVANCED SPEECH-TO-TEXT & TEXT-TO-SPEECH

Generative AI powers highly accurate and natural-sounding speech synthesis and recognition, benefiting users with hearing, speech, or motor impairments in daily communication.

PERSONALISED COMMUNICATION AIDS

Generative AI can enhance augmentative and alternative communication (AAC) devices by predicting words or phrases, making communication faster, easier, and more fluent for non-verbal individuals.

CONTENT SIMPLIFICATION & ADAPTATION

Generative AI can simplify complex texts or adapt content formats to suit users with cognitive disabilities or different learning preferences, making information more understandable and usable.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CHALLENGES AND IMPORTANT QUESTIONS

01 Misinformation & Deepfakes

Generative AI can create realistic fake text, images, or videos used to spread false information or scam people. How can we detect this?

02 Bias

If the training data contains biases, such as stereotypes, generative AI can learn and repeat them. How can we ensure fairness?

03 Copyright

Models are trained on vast data, sometimes including copyrighted work without permission. Who owns AI-generated content? How should creators be credited or compensated?

04 Security

Generative AI can be used to create malware or phishing scams. The AI systems themselves can also be attacked.

05 Authenticity & Trust

Generative AI can hallucinate, meaning it makes things up confidently. How can we know if the output is accurate?

06 Environmental Cost

Training huge generative AI models uses significant amounts of electricity and water. How can we make AI more sustainable?

07 Job Impacts

How will generative AI impact jobs currently performed by humans? How might roles change?

08 The Alignment Problem

How can we ensure generative AI systems consistently behave according to human values and intentions?

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

THE FUTURE OF GENERATIVE AI

We can expect generative AI to become even more deeply woven into our digital experiences and professional tools, often working to enhance our creative and productive capabilities. Ongoing advancements are enabling these generative AI systems to understand more complex requests, generate richer and more varied content across different formats like text, image, and video, and interact in more intuitive ways. Exciting progress is being made in areas such as personalised content creation, scientific discovery, software development, and immersive entertainment, with these capabilities evolving at a remarkable pace. Crucially, there is a significant and growing emphasis on ensuring that generative AI develops responsibly, with safety, fairness, transparency, and alignment with human values as core principles guiding its future innovation and deployment.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 **Creates New Things**

Unlike AI that just analyses existing information, generative AI uniquely makes completely new and original content, such as text or images, based on the vast knowledge learned by generative AI.

02 **Learns from Data**

Generative AI has impressive abilities to generate new content derived entirely from the complex patterns, styles, and relationships discovered within the massive and diverse datasets the model is trained on.

03 **Prompt is Key**

The way you ask generative AI for something, known as the prompt, is critically important because it greatly influences the relevance, quality, and specific nature of the output you receive.

04 **Powerful Assistant**

Generative AI can significantly boost productivity by automating tasks and enhance creativity by offering new ideas and perspectives across many different fields and industries.

05 **Not Perfect**

It is important to remember that generative AI can make factual mistakes, sometimes called hallucinations, may reflect biases present in the training data, and its use raises significant ethical questions that need careful consideration.

06 **Human Guidance is Crucial**

Human guidance and oversight are absolutely crucial. People are needed to direct its use, critically evaluate its outputs, ensure it is applied responsibly and ethically, and ultimately confirm that it aligns with our broader societal goals and values.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 4

Generative AI – A Clinical Overview

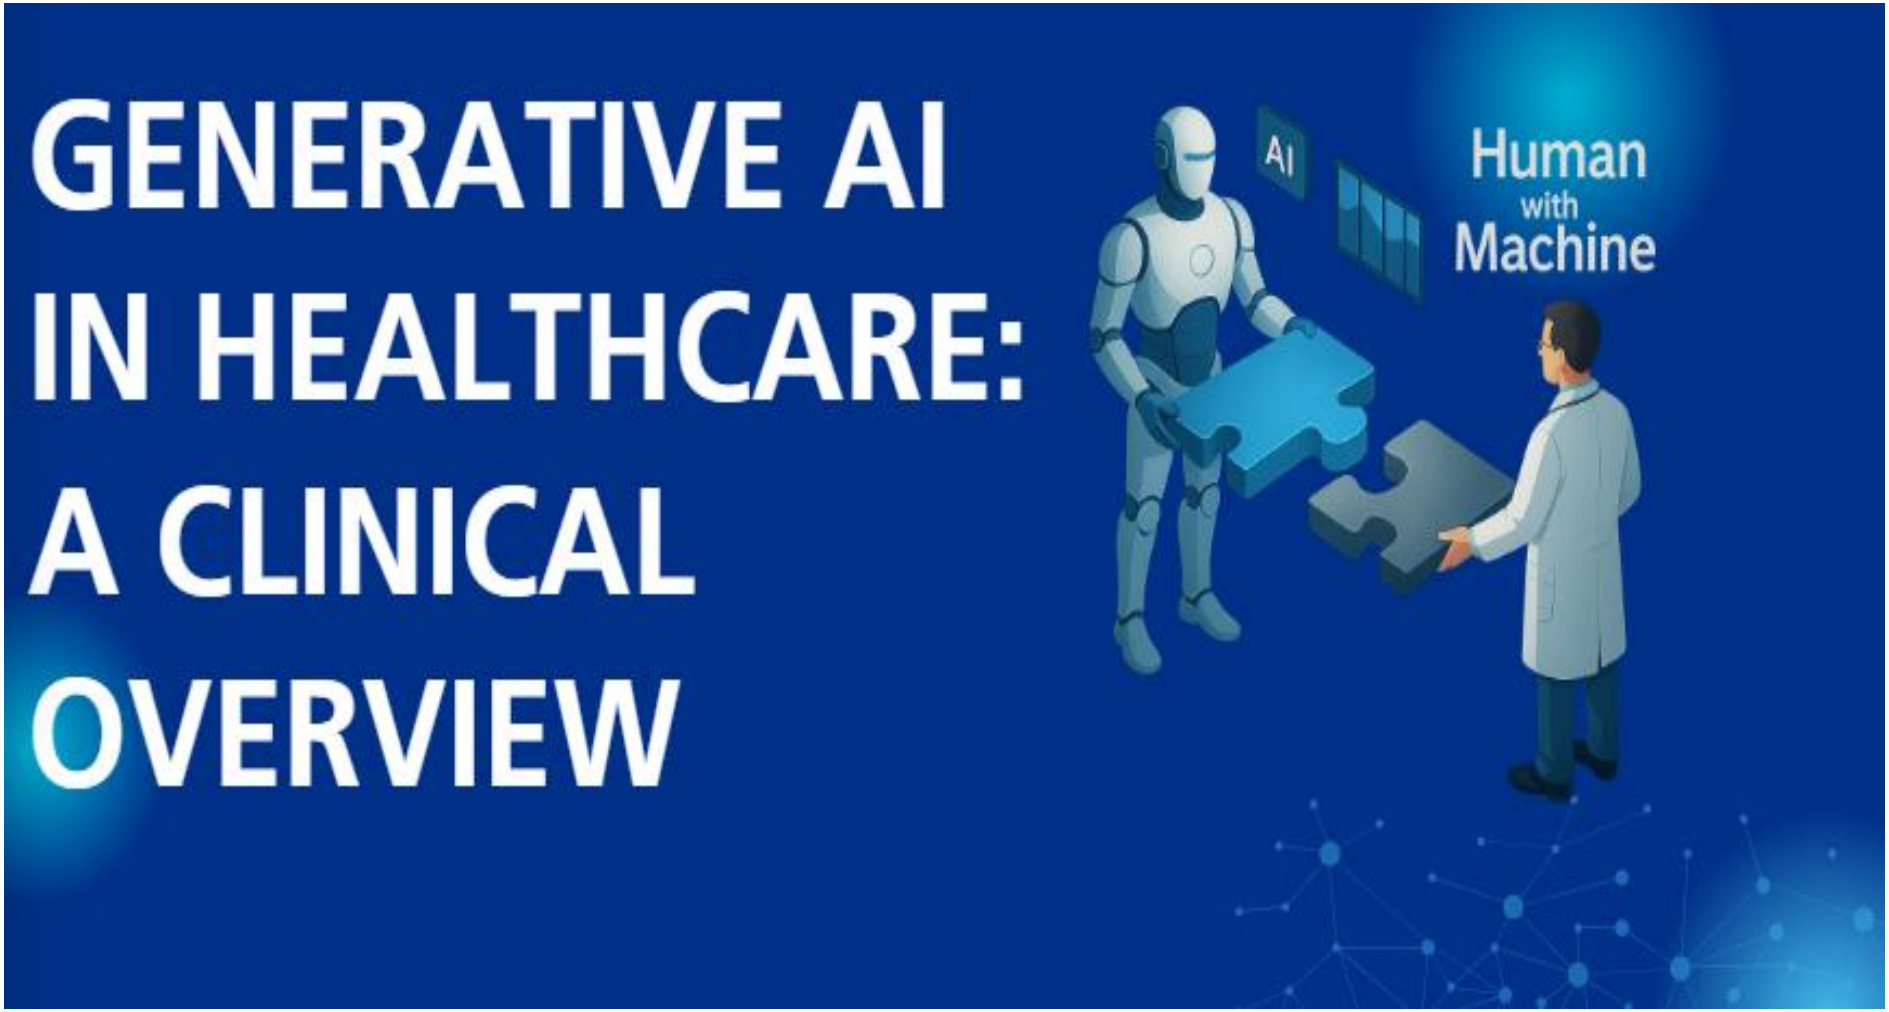

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS GENERATIVE AI IN HEALTHCARE?

Generative AI in healthcare refers to artificial intelligence systems that can learn from vast amounts of medical data to create entirely new, original content relevant to patient care, research, and operations. This new content can range from synthetic patient records that look real but are entirely fictional for training material, to novel drug compounds, draft clinical notes, or even personalised treatment simulations. Unlike traditional AI that primarily analyses or predicts from existing data, generative AI synthesises new information, mimicking human-like creativity and reasoning, but applied to the complexities of medicine. The critical difference in healthcare is the responsibility and the demand for accuracy and ethical consideration. An error in an AI-generated artwork is trivial; an error in a generated drug molecule or a diagnostic suggestion can have life-altering consequences. Therefore, generating in healthcare requires thorough validation and a deep understanding of its potential impact from the outset.

Generative AI reads and learns from real patient data.

Generative AI synthesises knowledge to draft a clear, accurate discharge summary.

Clinicians receive ready-to-use summaries, saving time and reducing admin burden.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI, MACHINE LEARNING, AND GENERATIVE AI: HOW DO THEY FIT TOGETHER IN MEDICINE?

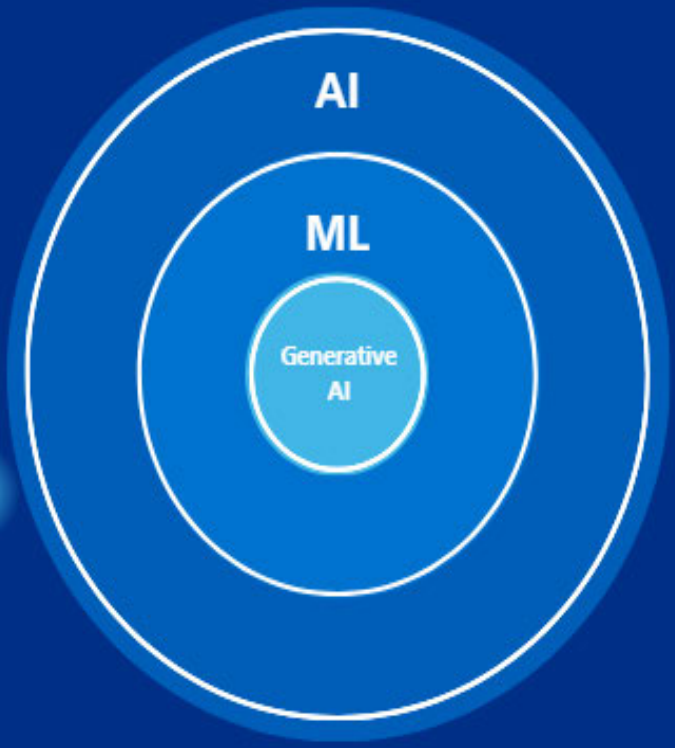

Artificial Intelligence (AI): This is the broadest concept. It refers to the overall ambition of making computers perform tasks within a medical setting that typically require human intelligence. This could be anything from helping to interpret medical scans to assisting with clinical decision support or managing hospital logistics. Think of it as: The entire field of creating 'smart' computer systems for medical applications.

Machine Learning (ML): A way to achieve AI in healthcare. Instead of being explicitly programmed for every single medical scenario, machines are trained using large amounts of medical data, like patient records or imaging results. They learn patterns from this data to make predictions, e.g. patient risk of a certain disease or decisions, e.g. flagging an abnormal scan. Think of it as: Teaching a computer about medicine by showing it many examples from past patient cases, rather than writing detailed instructions for every single possible illness or presentation.

Generative AI: A type of artificial intelligence that focuses on creating new medical content. Generative AI systems often use machine learning and deep learning techniques to learn from existing medical data and produce outputs like drafting clinical summaries, synthetic medical images for training, or even ideas for new therapeutic approaches. Think of it as: A specific application of AI that uses its learned medical knowledge to generate something original and useful for healthcare, rather than just analysing or predicting from existing patient data.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

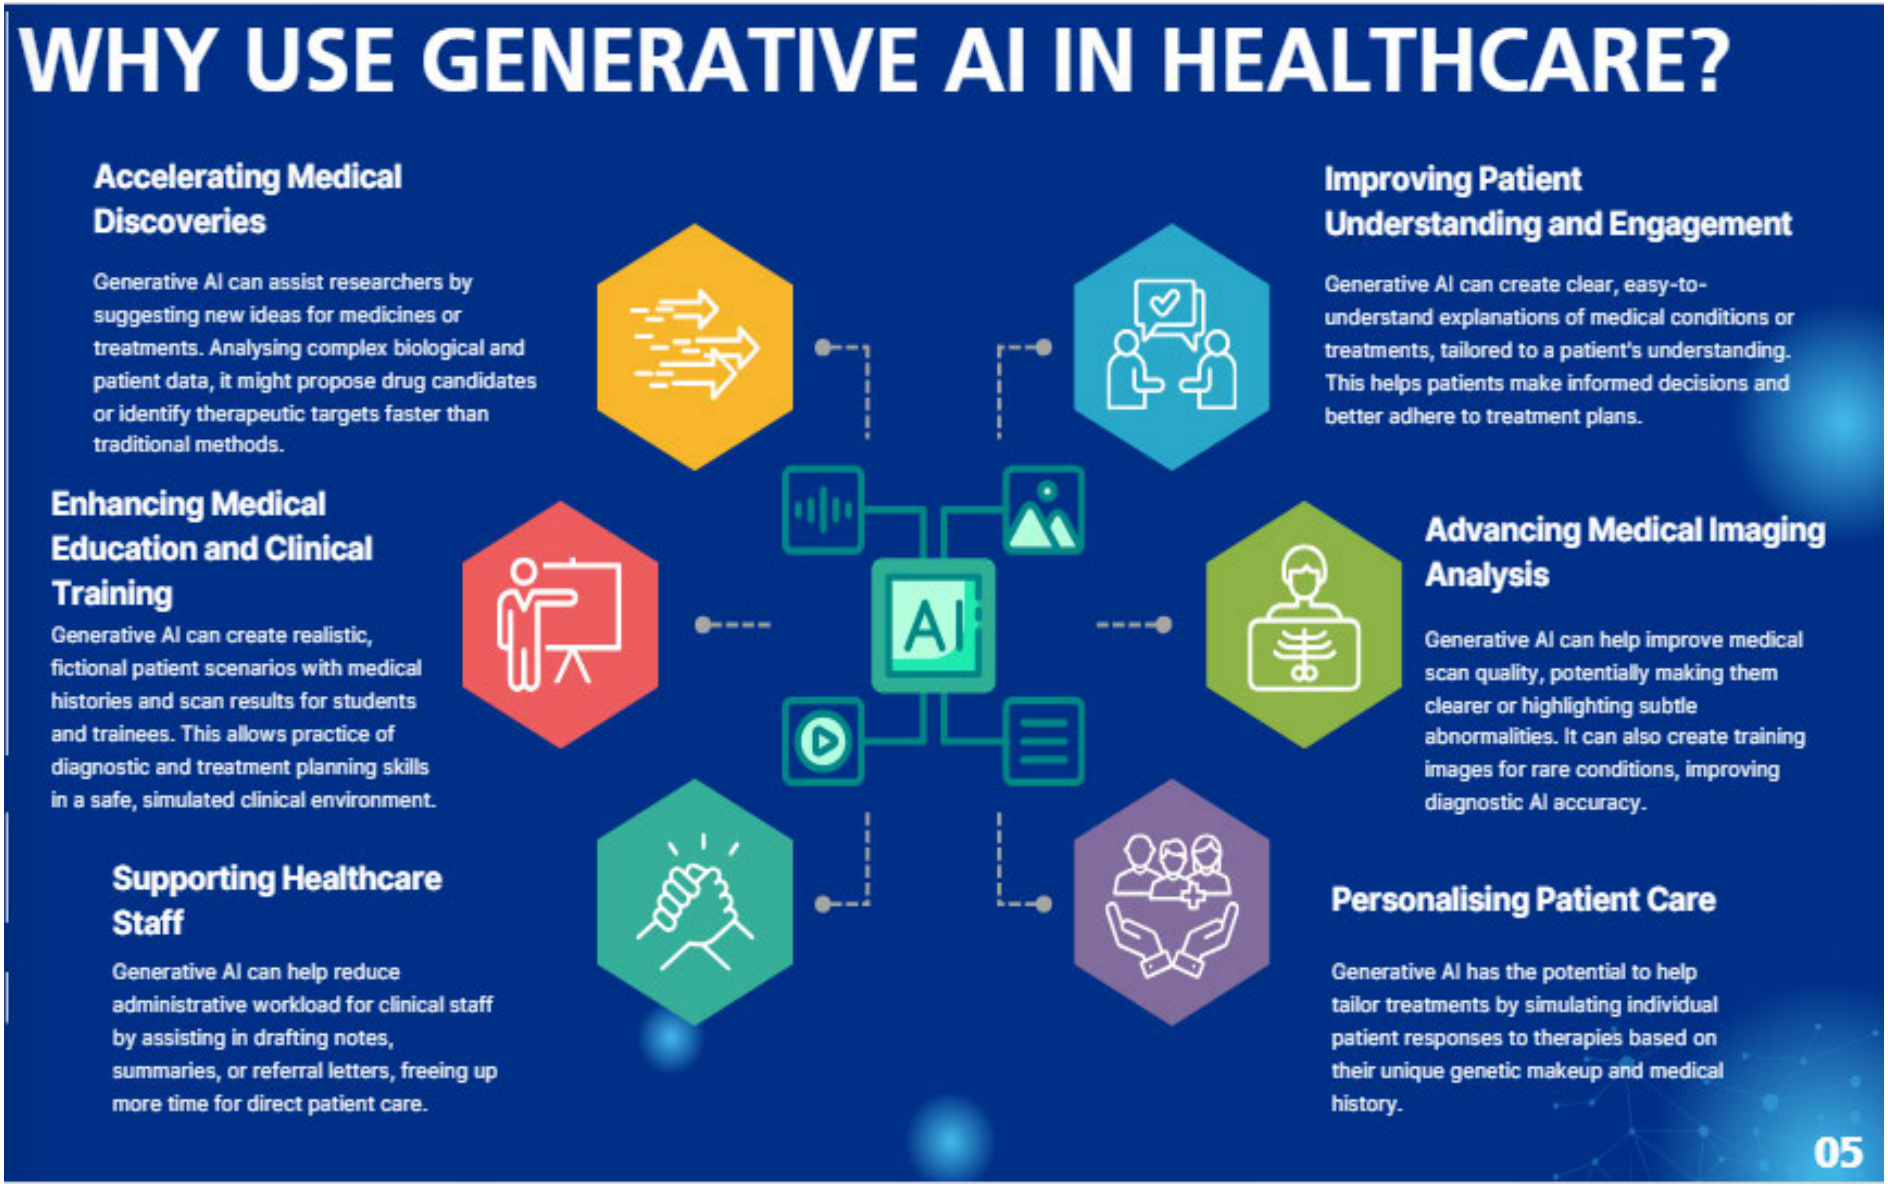

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW DOES GENERATIVE AI WORK IN HEALTHCARE?

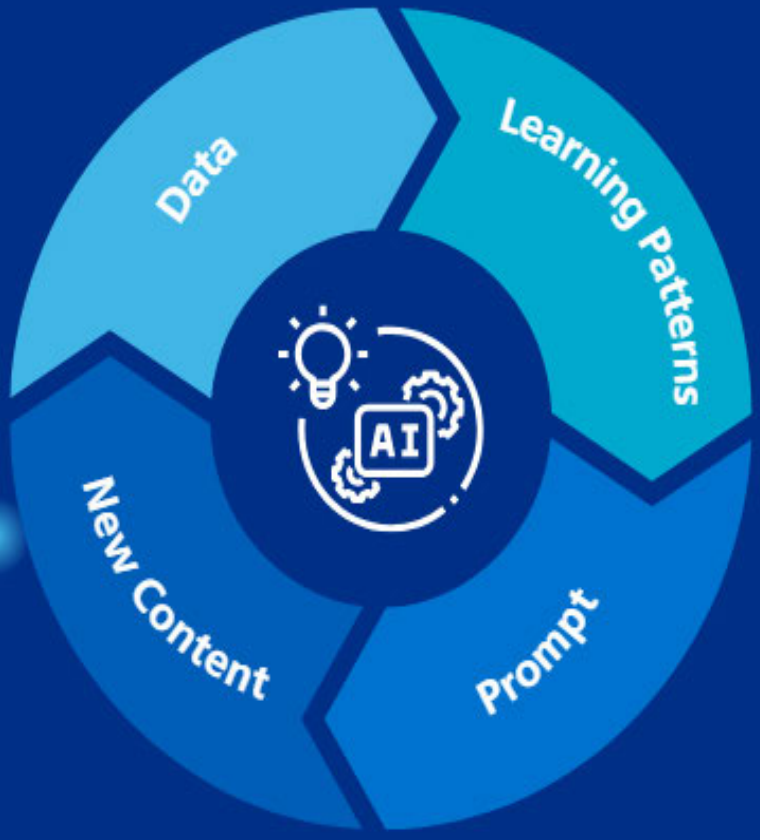

Learning from Data

Generative AI begins by studying enormous amounts of existing medical information. This training data can include anonymised electronic health records, medical images (like X-rays, CT scans, MRIs), genomic sequences, medical research papers, clinical trial data, and established treatment guidelines.

Finding Patterns

The generative AI learns underlying structures, styles, and rules within this medical information. For clinical text, it learns medical terms, common descriptions of symptoms/diagnoses, and note structures. For images, it learns normal anatomy and visual signs of diseases. For genetic data, it learns sequence patterns linked to health or illness.

Using a Prompt

A healthcare professional or researcher then gives the generative AI an instruction, known as a prompt. This tells the AI what kind of medical content to create.

Generating New Content

Based on its learned medical patterns and the clinical prompt, generative AI creates new content. It predicts the next piece of information (word, image detail, data point) to fulfil the request consistently with its medical learning. The output is new, yet grounded in its training.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES IN HEALTHCARE

The Smart Clinical Note Drafter

One way generative AI creates medical text, like clinical notes or patient summaries, is by acting as a highly intelligent predictor. By examining the existing text and, based on training on vast amounts of medical literature and clinical documentation, predicts the most suitable medical term or phrase to add next.

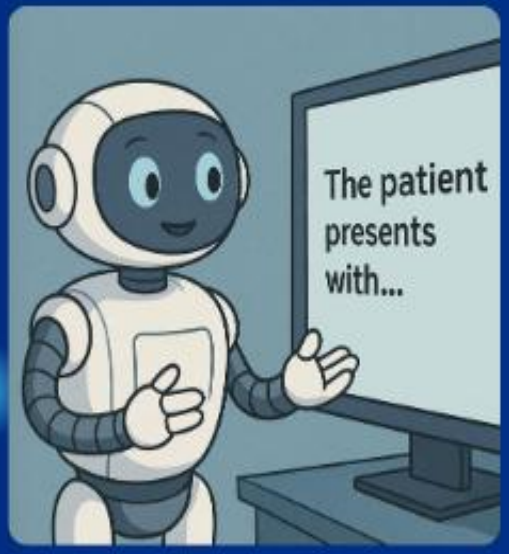

The generative AI looks at the current text to understand the context.

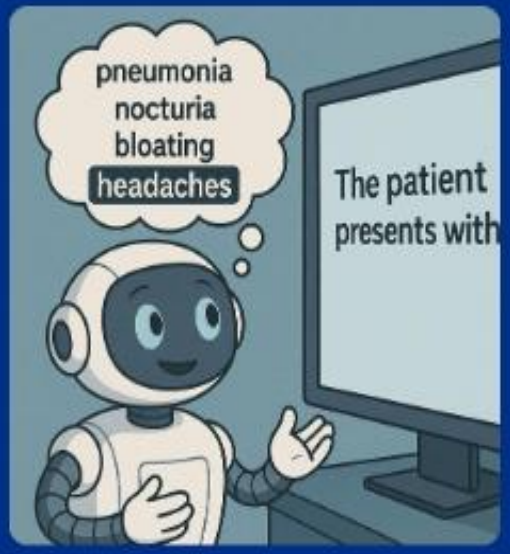

It thinks of many possible words that could fit next.

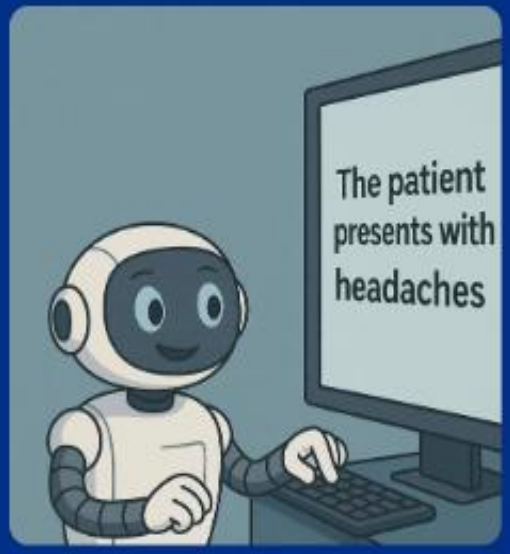

It chooses the best word and continues, building the summary piece by piece.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES IN HEALTHCARE

Shaping Medical Images from Digital Noise

Another way generative AI creates, particularly for medical images like X-rays or pathology slides, is by starting with a very basic or even random digital pattern. Then, guided by the desired medical features e.g., 'a healthy lung' or 'cells showing inflammation', it carefully shapes and refines this pattern, step-by-step, until it becomes a clear, detailed, and medically plausible new image.

The generative AI starts to examine a blurry, low-detail X-ray.

Generative AI begins enhancing the image, highlighting a clearer, more defined X-ray.

Step-by-step, generative AI refines the image into a detailed, fully enhanced X-ray, with lung nodules clearly marked for interpretation.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW GENERATIVE AI CREATES IN HEALTHCARE

The Medical Content Creator and the Clinical Quality Checker

Sometimes, two generative AI systems work together, especially for creating highly realistic medical content. One AI, the 'Creator', attempts to generate new medical data e.g., a synthetic patient record or a medical image. A second AI, the 'Clinical Quality Checker' (or 'Critic'), evaluates how realistic or clinically plausible this creation is, comparing it to real medical data it has learned from. By constantly receiving feedback from the 'Checker', the 'Creator' AI learns to produce better, more accurate, and more believable new medical content.

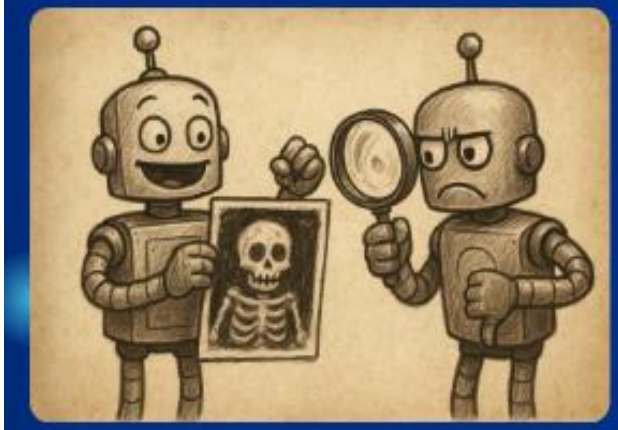

The "Creator" AI makes a new picture. The "Critic" AI judges if the picture looks real enough.

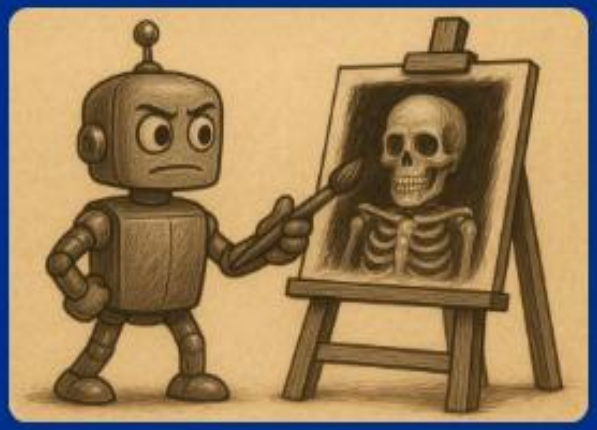

The "Creator" tries again, learning from the feedback.

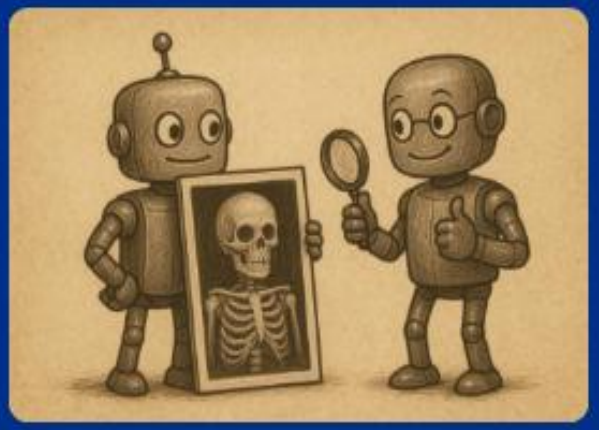

Through this teamwork, the "Creator" learns to make very believable new images.

09

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

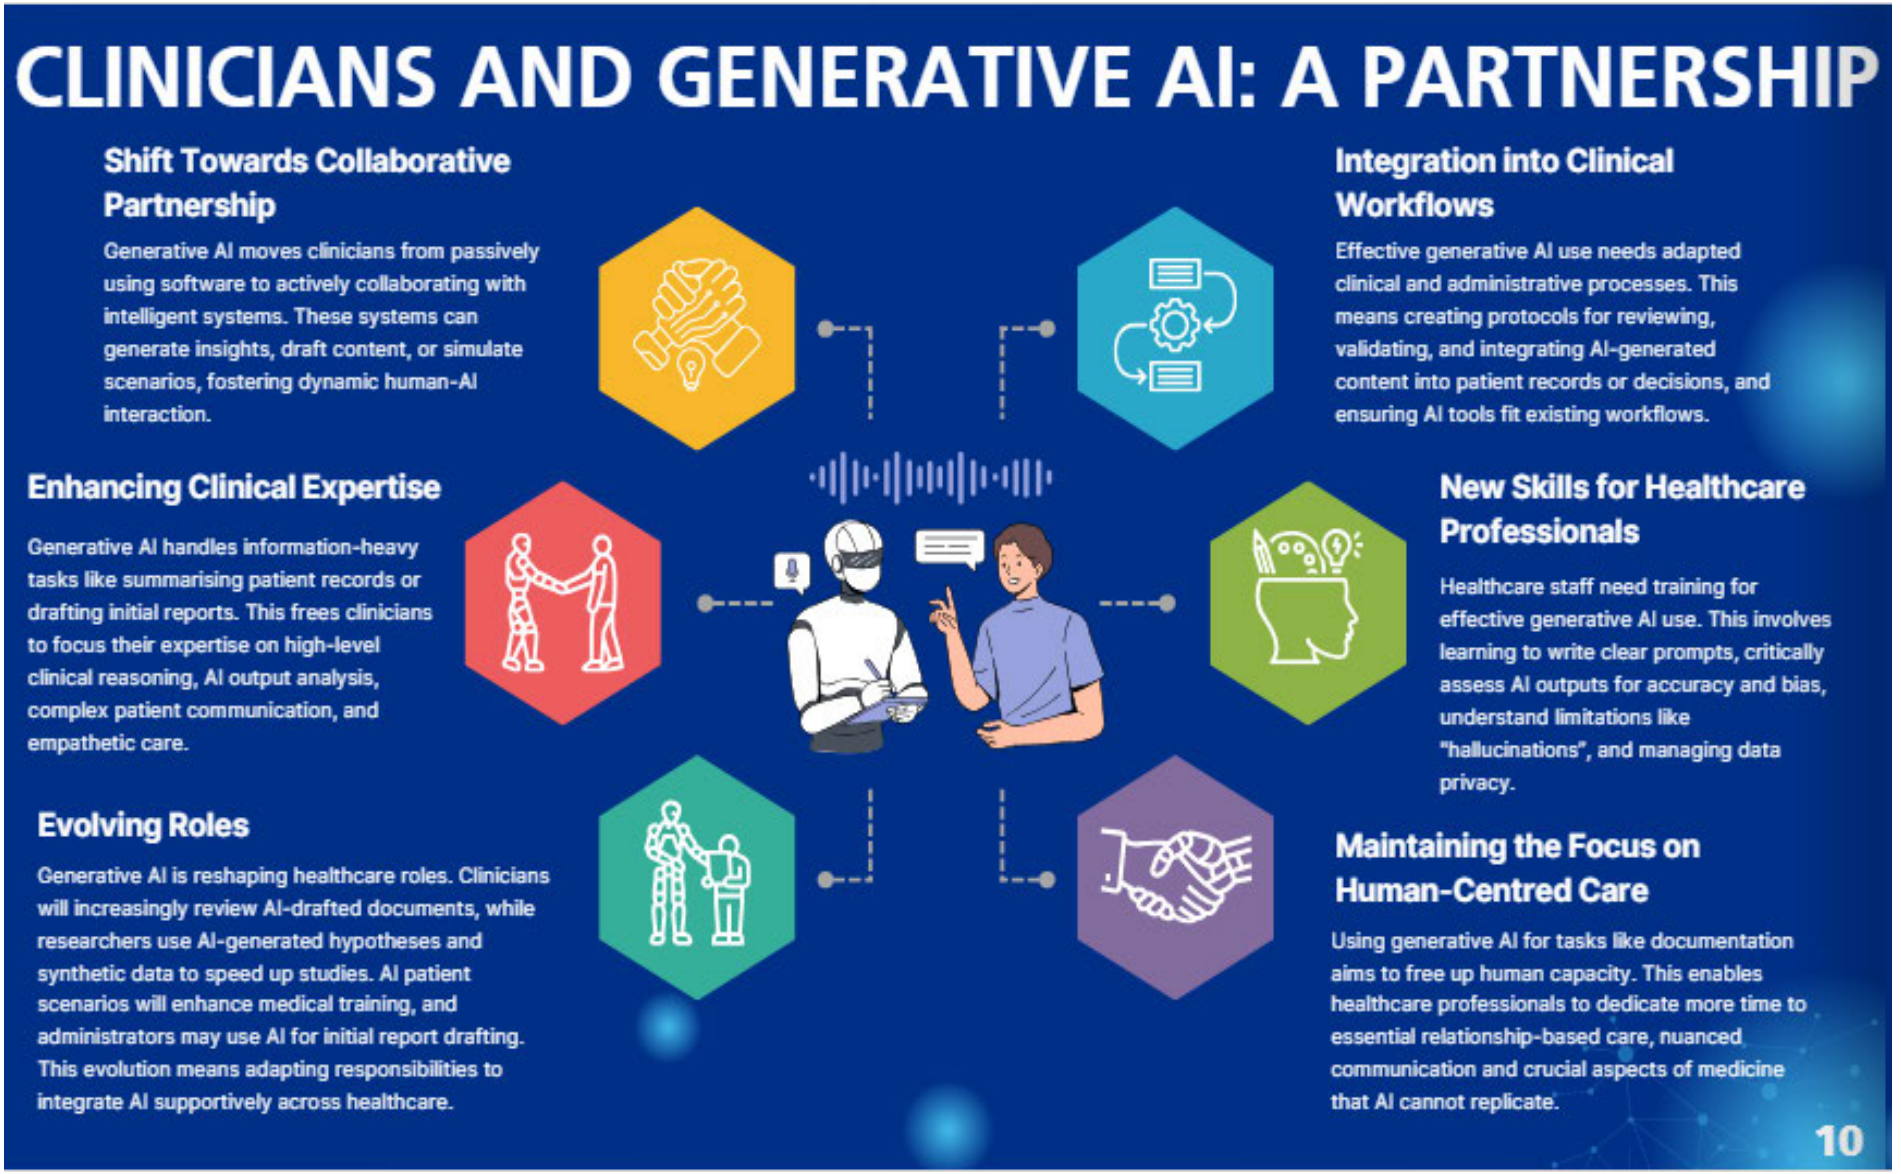

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

ETHICAL CONSIDERATIONS & GOVERNANCE IN CLINICAL GENERATIVE AI

01 Patient Safety and Clinical Validation

AI-generated medical information must be rigorously validated by qualified clinicians to prevent patient harm from potential inaccuracies or "hallucinations", ensuring outputs are clinically sound.

02 Data Privacy & Cyber Security

Strict adherence to data protection laws like GDPR and Caldicott Principles is vital when training generative AI with sensitive patient data. Robust cybersecurity defences, including encryption and secure networks, are vital to prevent breaches or attacks that could compromise data.

03 Fairness and Mitigation of Bias

Generative AI models require training on diverse and representative medical datasets, and must be carefully tested to prevent biases that could lead to or worsen health inequalities among different patient groups.

04 Transparency and Explainability

Ongoing efforts are needed to make AI's decision-making processes as understandable as possible to clinicians, which helps to build trust, identify errors, and support accountability.

05 Accountability and Responsibility

Clear lines of professional responsibility must always be established for any AI-assisted clinical decisions and their subsequent impact on patient care and outcomes.

06 Human Clinical Oversight

Healthcare professionals must always retain final clinical judgement and ultimate responsibility for patient care, using generative AI as a supportive tool rather than an autonomous decision-maker.

07 Regulatory Oversight

Generative AI tools intended for use in clinical settings, especially those influencing diagnosis or treatment, must meet the stringent safety, efficacy, and quality standards set by official medical regulatory bodies like the MHRA to ensure they are fit for purpose.

08 Informed Consent and Patient Autonomy

Where appropriate, patients should be clearly and understandably informed about the role generative AI may play in their care, ensuring their autonomy and choices are respected.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

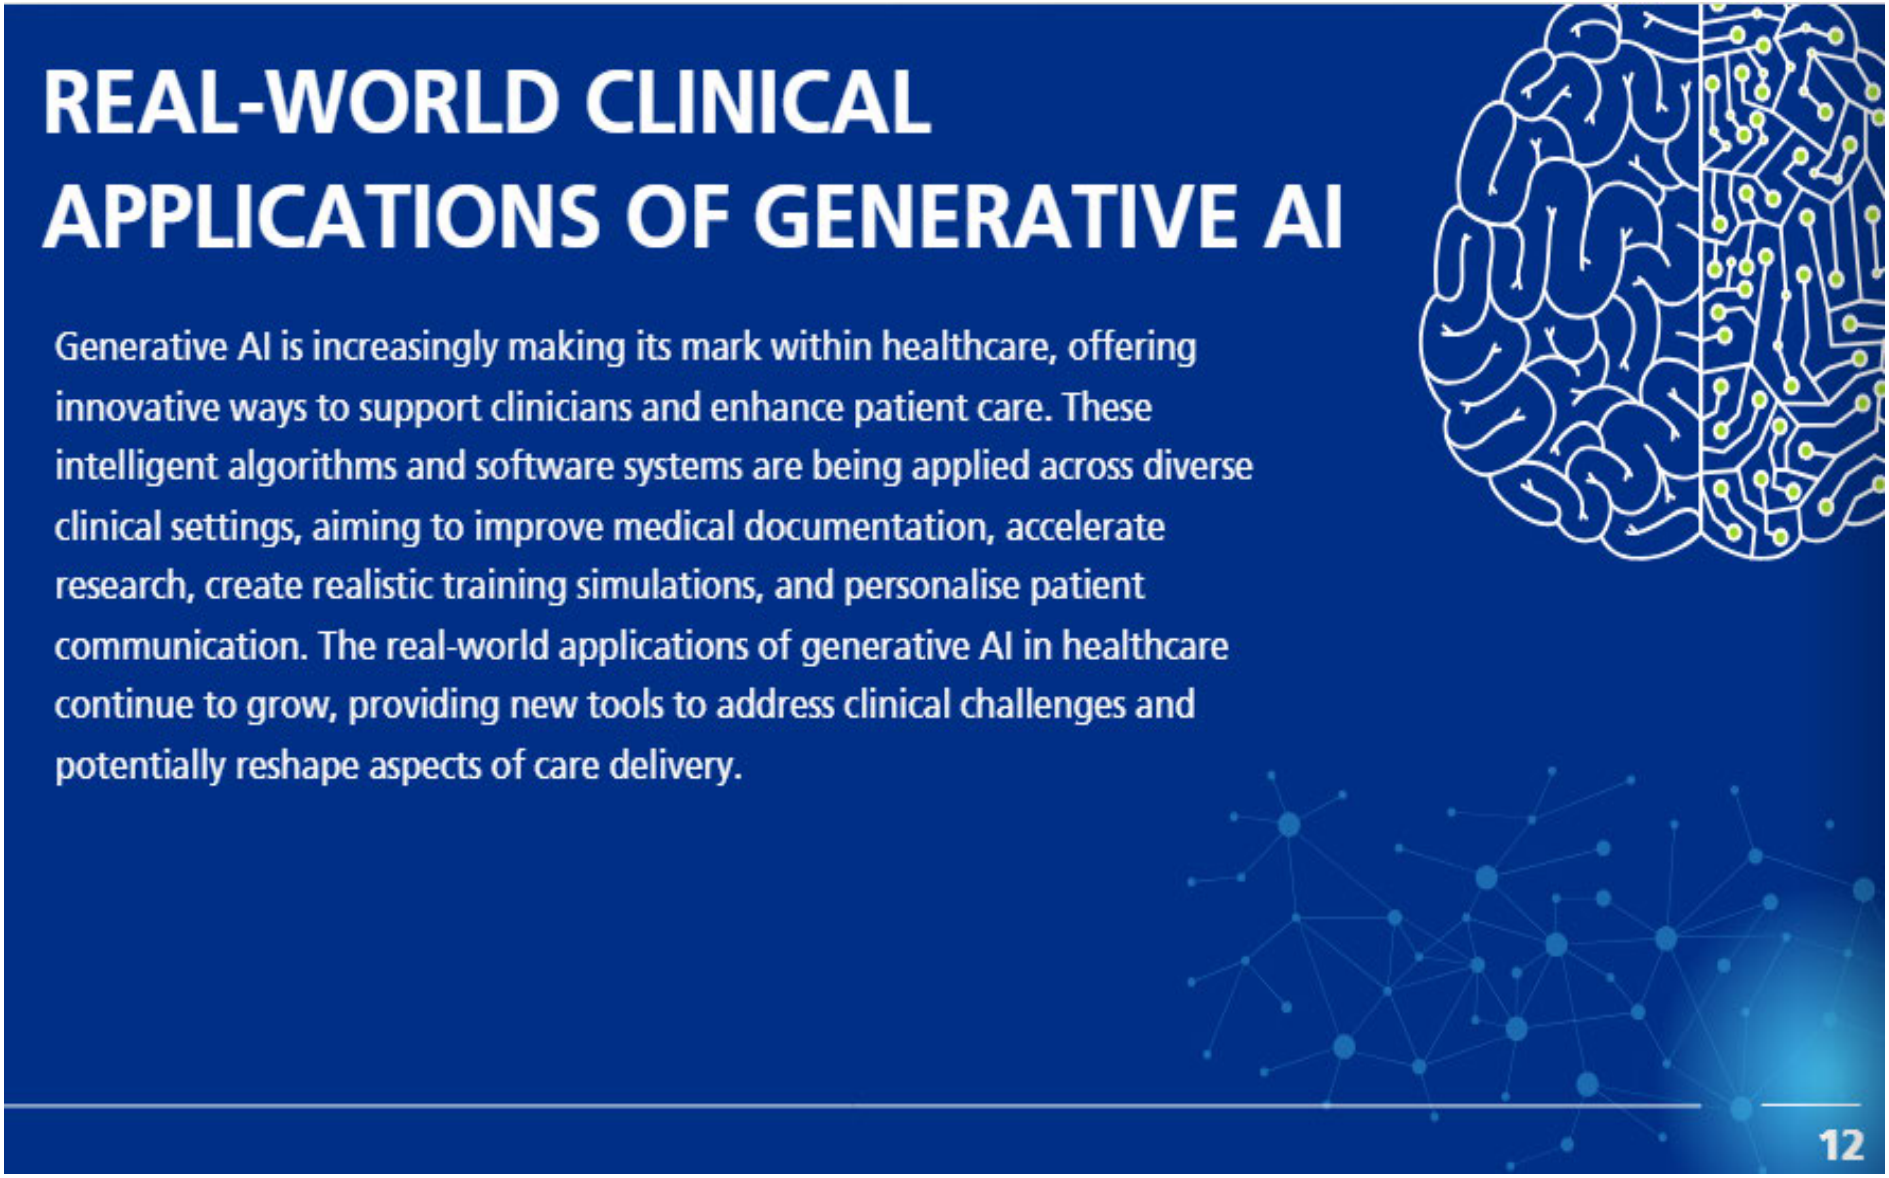

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

DRUG DISCOVERY AND DEVELOPMENT

Accelerating the journey from laboratory to new medicines.

GENERATING NOVEL MOLECULAR STRUCTURES

AI algorithms can analyse vast databases of known chemical compounds and biological targets to design entirely new molecules with the potential to become effective drugs for specific diseases.

PREDICTING DRUG EFFICACY AND SAFETY

Generative models can simulate how potential new drugs might interact within the human body, predicting their likely effectiveness and identifying potential side effects early in the development process.

OPTIMISING CLINICAL TRIAL DESIGN

Generative AI can help design more efficient clinical trials by, for example, generating synthetic patient data to model trial outcomes or identifying optimal patient cohorts for inclusion.

REPURPOSING EXISTING MEDICATIONS

Generative AI can analyse data to find new uses for existing, approved drugs by identifying new biological pathways they might affect, speeding up the availability of treatments.

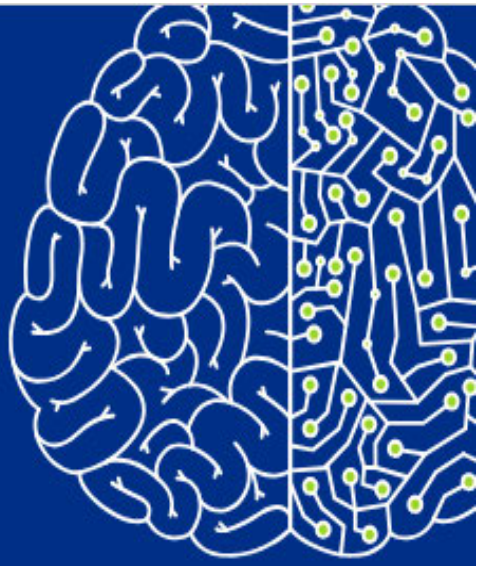

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

MEDICAL DIAGNOSTICS AND IMAGING SUPPORT

Enhancing the way we see and understand medical images and diagnostic data.

GENERATING SYNTHETIC MEDICAL IMAGES FOR TRAINING

Generative AI can create realistic but artificial medical images (X-rays, CT scans, pathology slides) to train other diagnostic AI systems or medical students, especially for rare conditions where real data is scarce.

ENHANCING IMAGE QUALITY AND RESOLUTION

Generative techniques can improve the clarity of medical scans by reducing noise or increasing resolution, potentially leading to more accurate diagnoses from existing imaging equipment.

ASSISTING IN DIAGNOSTIC REPORT GENERATION

Generative AI can analyse medical images or other diagnostic data and generate initial draft reports or summaries for clinicians to review, edit, and finalise, streamlining the reporting process.

SUPPORTING ANOMALY DETECTION

While not its primary function, generative models can learn what "normal" looks like in medical data e.g., healthy tissue and thus help highlight unusual patterns or anomalies that might warrant further investigation.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

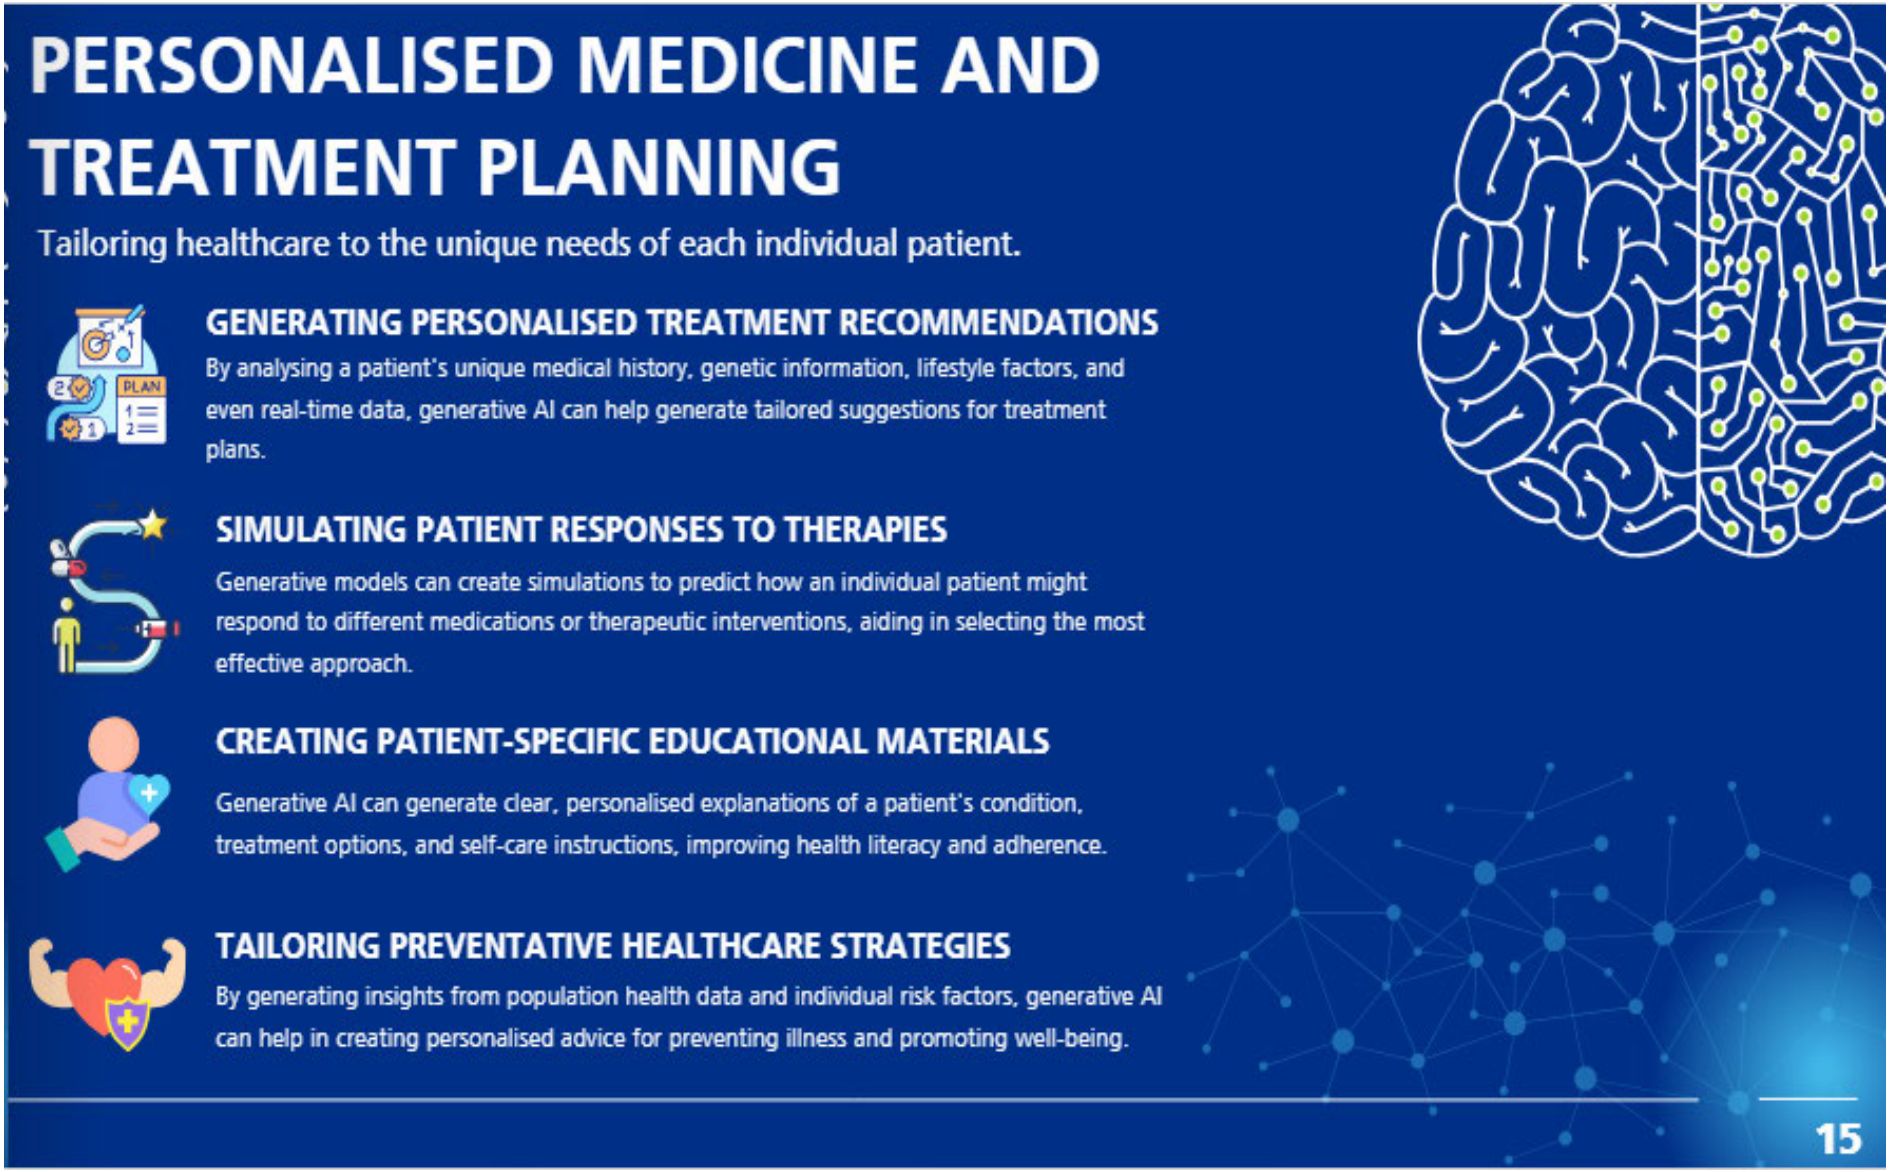

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CLINICAL DOCUMENTATION AND ADMINISTRATIVE EFFICIENCY

Reducing the burden of paperwork and streamlining healthcare operations.

AUTOMATING CLINICAL NOTE DRAFTING

Generative AI systems can listen to, with consent, clinician-patient conversations or process dictated notes to automatically generate draft clinical documentation, such as consultation notes or SOAP notes.

SUMMARISING COMPLEX PATIENT RECORDS

Generative AI can quickly read through lengthy electronic patient records and produce concise summaries of a patient's medical history, current conditions, and key findings for efficient review.

GENERATING REFERRAL LETTERS AND DISCHARGE SUMMARIES

Generative AI can assist in drafting standardised documents like referral letters to specialists or discharge summaries for GPs, ensuring all necessary information is included consistently.

IMPROVING PATIENT COMMUNICATION WORKFLOWS

Generative AI can help draft responses to routine patient enquiries or generate patient-friendly appointment reminders and follow-up instructions, enhancing communication efficiency.

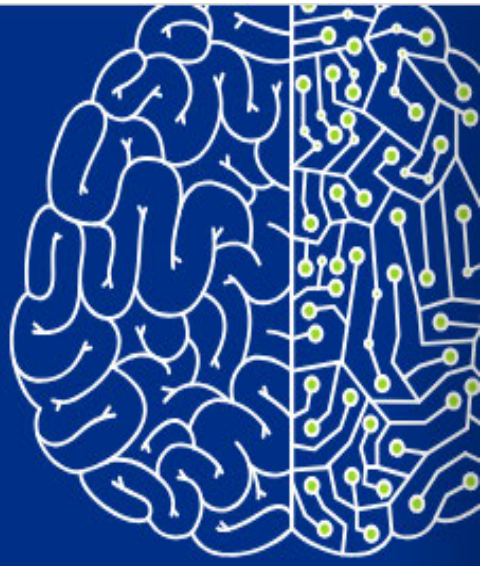

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

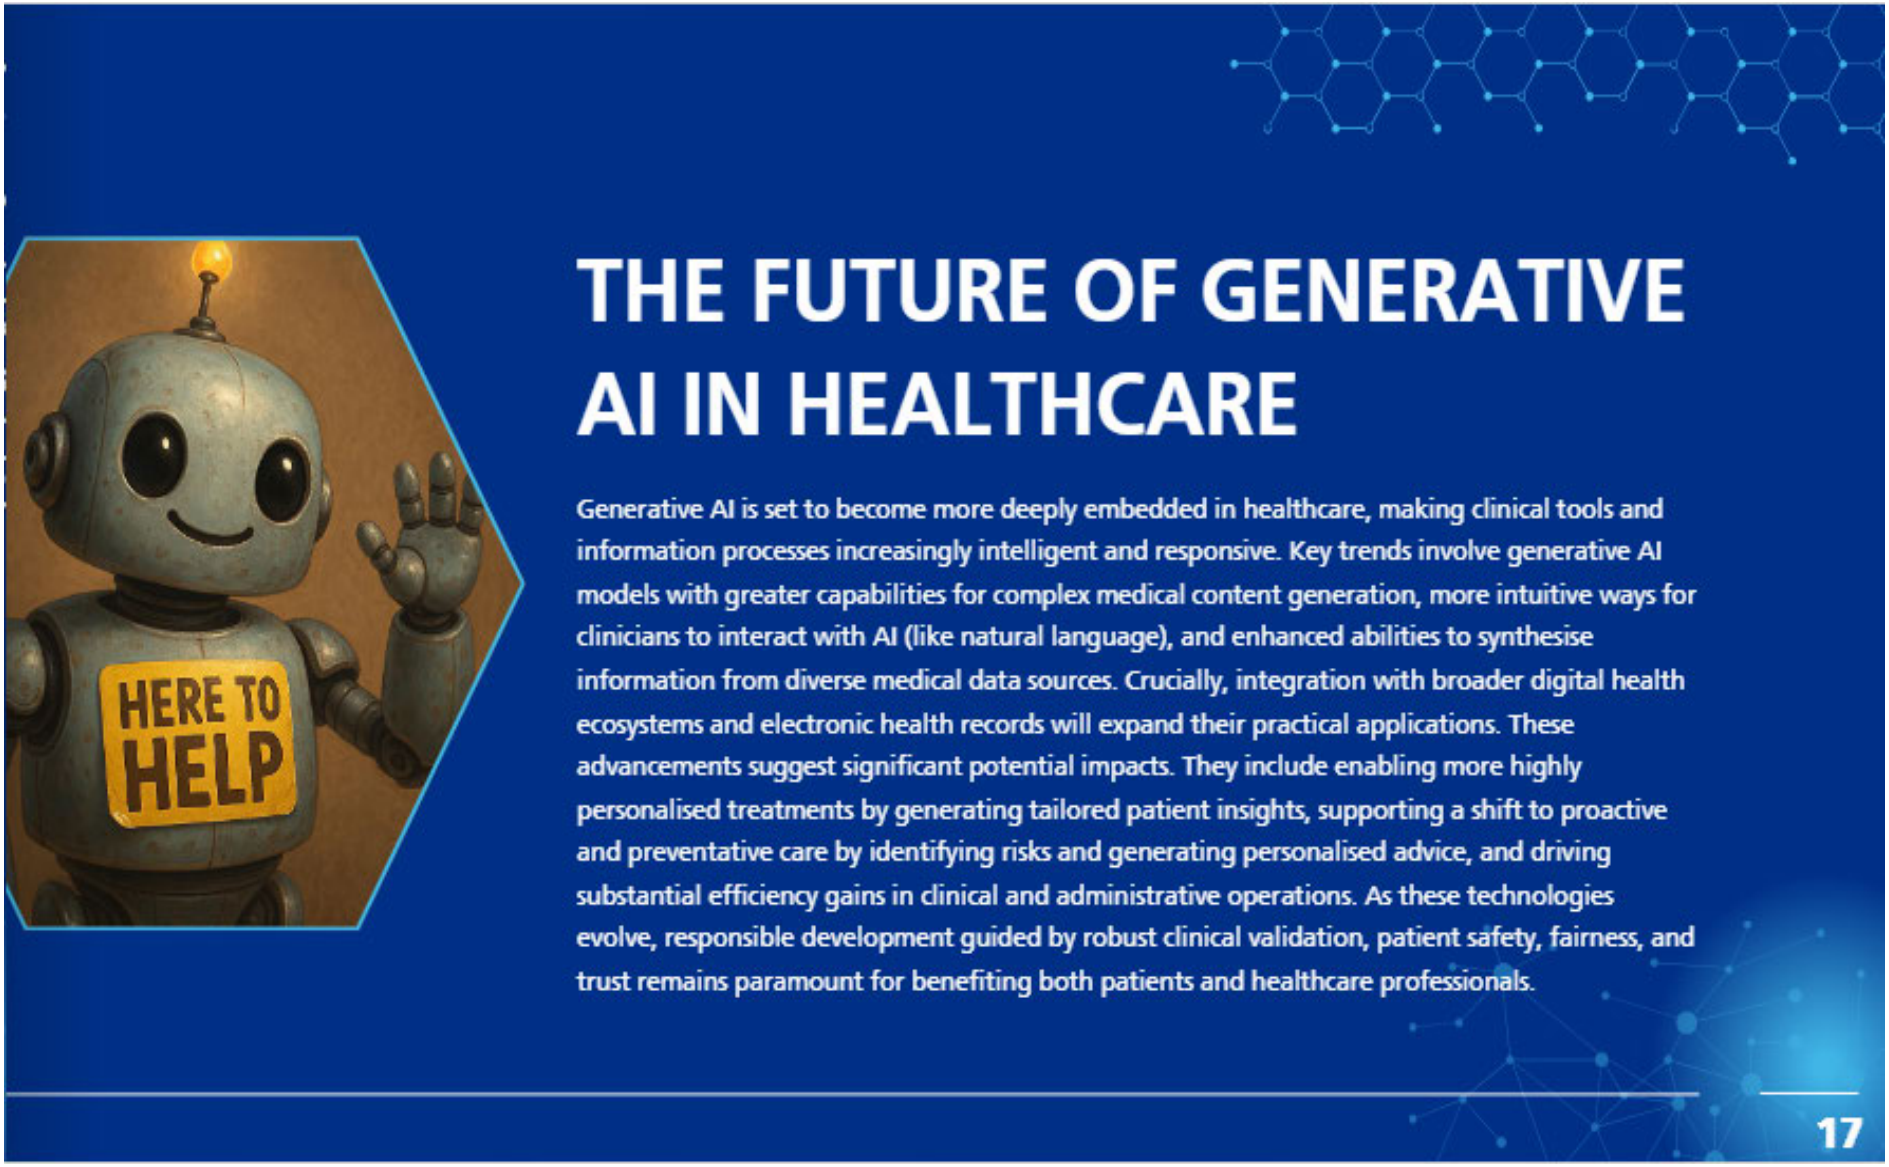

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 Generative AI Creates New Medically-Relevant Content

The core function is generating medical information; draft notes, synthetic images, drug formulations, from learned data, differing from AI that only analyses or predicts.

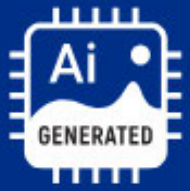

02 It is a Tool to Assist Clinicians, Not Replace Them

Generative AI supports healthcare professionals by automating tasks and providing information, but does not replace their essential expertise, empathy, or judgment.

03 Human Clinical Oversight is Absolutely Essential

All AI-generated outputs impacting patient care must be carefully reviewed, validated, and approved by qualified clinicians, who remain responsible for decisions.

04 Ethical Medical Use and Patient Data Privacy are Paramount

Responsible generative AI use in healthcare means prioritising patient safety, ensuring fairness, avoiding bias, protecting patient data, and adhering to ethical and regulatory guidelines.

05 Potential for Significant Benefits to Medicine

Thoughtfully developed and ethically implemented generative AI can accelerate medical research, improve clinical efficiency, support staff, and enhance personalised patient care.

06 Continuous Evolution and Lifelong Learning

Generative AI models are not static; they evolve as they are exposed to new data and techniques. This requires ongoing monitoring, re-validation, and adaptation of how they are used in healthcare to ensure continued safety, accuracy, and effectiveness over time.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 5

Large Language Models – A basic overview

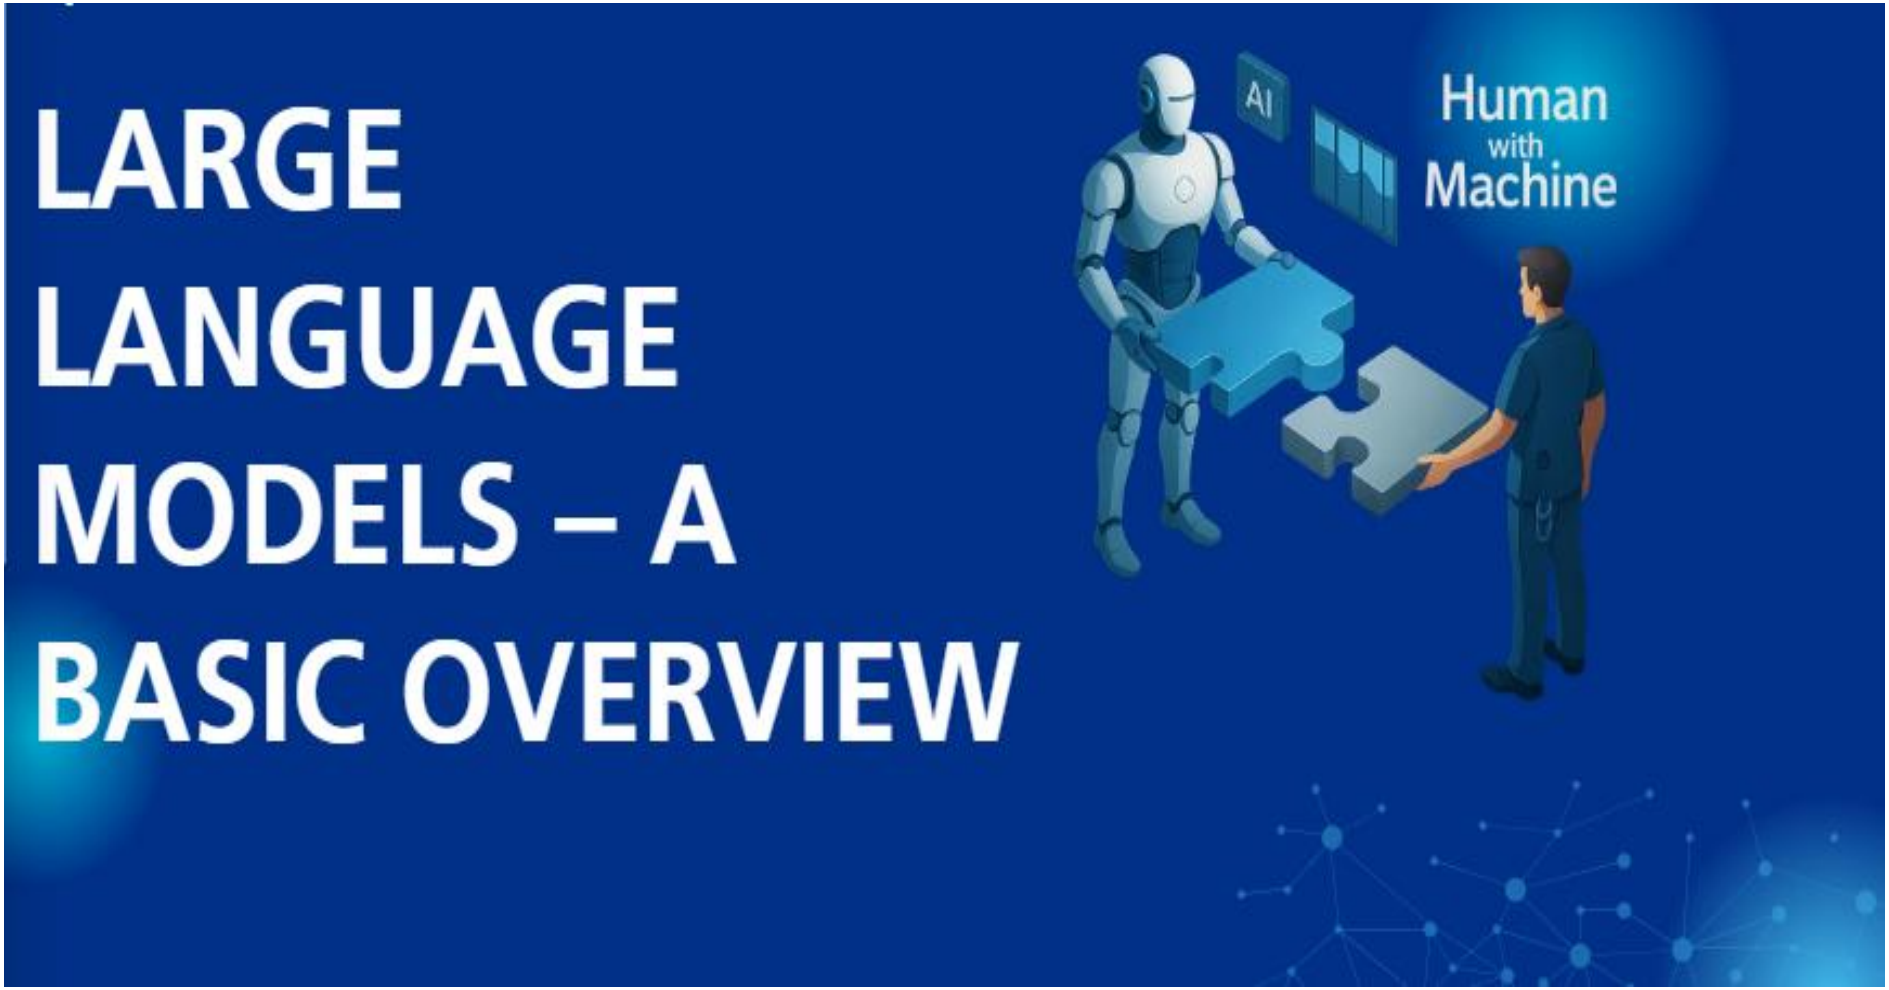

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY ARE WE ALL TALKING ABOUT LARGE LANGUAGE MODELS?

Large Language Models (LLMs) mark a significant advancement in artificial intelligence. They demonstrate a remarkable ability to understand the complexities of human language and to generate text that feels very natural. This goes beyond simply recognising individual words; it involves these systems understanding context, subtle differences in meaning, and even creative writing approaches.

- **A Popular Topic:** You may have heard about LLMs; they are often discussed in technology news.
- **Powerful Language AI:** These systems show a significant advancement in how AI can work with and produce text.
- **Changing Our Digital Experiences:** LLMs are influencing how we use technology, create written content, and find information.
- **Driving Generative AI:** LLMs are a key reason that generative AI can now create new and original written material so effectively.

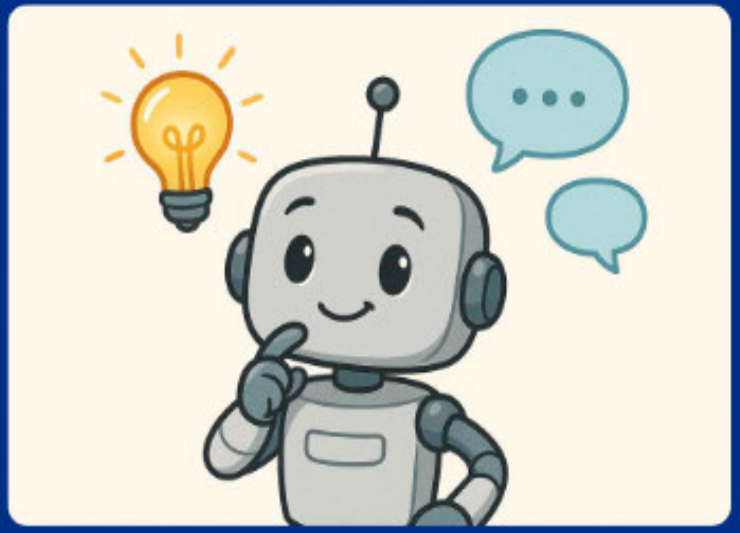

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT ARE LARGE LANGUAGE MODELS?

A large language model (LLM) is essentially a very advanced artificial intelligence programme. These sophisticated tools are specifically created to understand, interpret, and generate text that closely resembles human language. You can think of an LLM as an AI that has become exceptionally skilled with words. By learning from vast amounts of written material like books, articles, and websites, LLMs identify complex patterns in grammar and communication styles. This learning allows them to understand text you provide and also to create entirely new written content. They often work by predicting the most likely next word or part of a word in a sentence. This process enables them to construct coherent sentences and paragraphs that sound remarkably natural, assisting with many language tasks such as answering questions, summarising documents, or drafting emails.

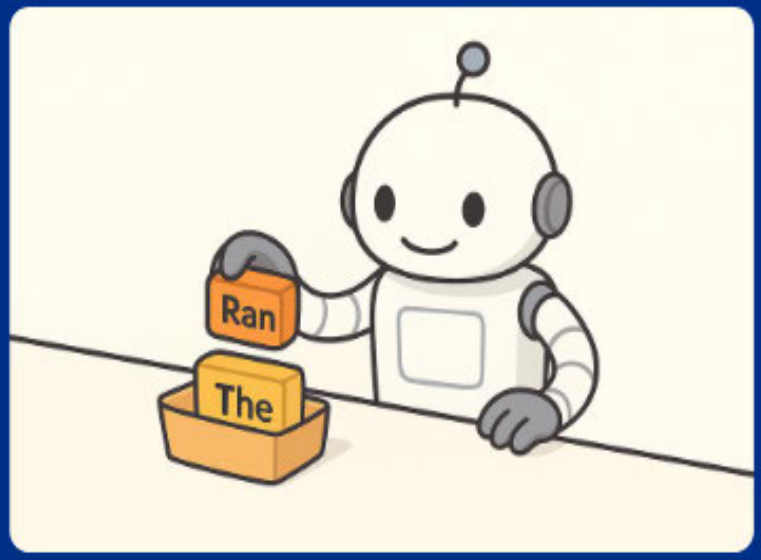

Large language models begin their process by learning to identify and select individual words and common phrases. Each 'word brick' is chosen for its potential role in forming a meaningful idea.

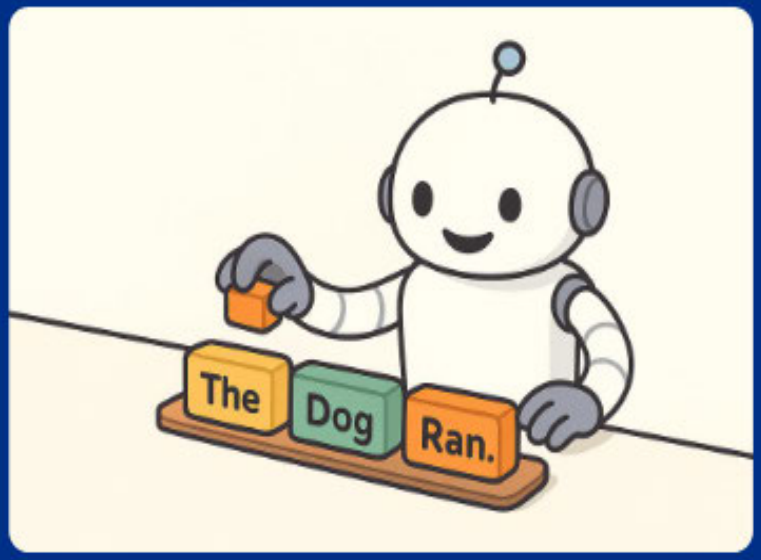

Once the necessary 'word bricks' are selected, the large language model then carefully arranges them in a logical and grammatically correct order, ensuring that the final text is complete, understandable, and makes sense.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

LARGE LANGUAGE MODELS IN THE AI FAMILY

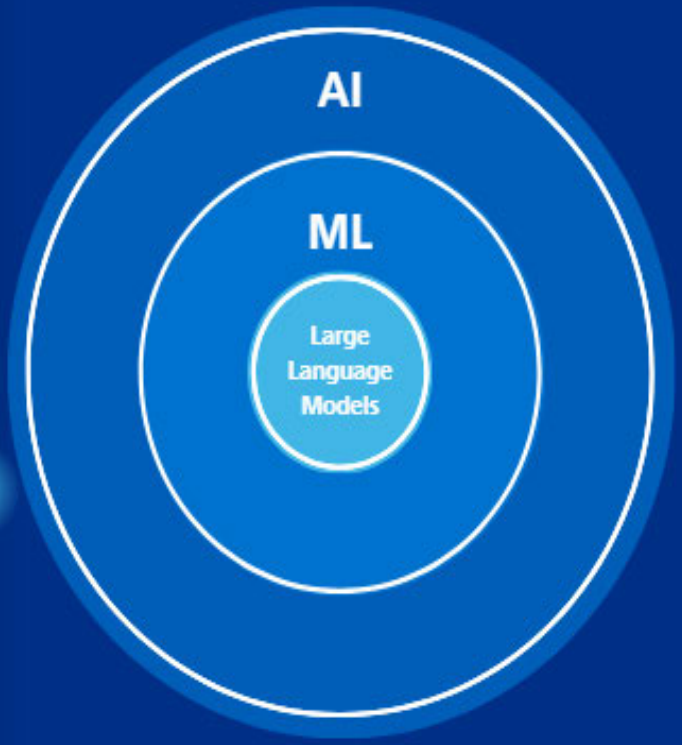

Artificial Intelligence (AI): AI is a wide ranging field focused on creating computer systems that can perform tasks normally requiring human intelligence, such as learning, reasoning, problem-solving, and decision-making. It is the overall goal of making machines intelligent.

Machine Learning (ML): ML is a core method used to achieve AI. Instead of being explicitly programmed for every task, ML systems learn patterns and make predictions or decisions by analysing large amounts of data. It is how many AI systems acquire their intelligence.

Large Language Models: These are a specialised type of machine learning, built using advanced methods often called 'deep learning'. Deep learning involves systems, inspired by the way brains work, learning from enormous quantities of data. LLMs use these powerful techniques to become experts in understanding and generating human-like language, making sense of text and creating new written content.

05

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

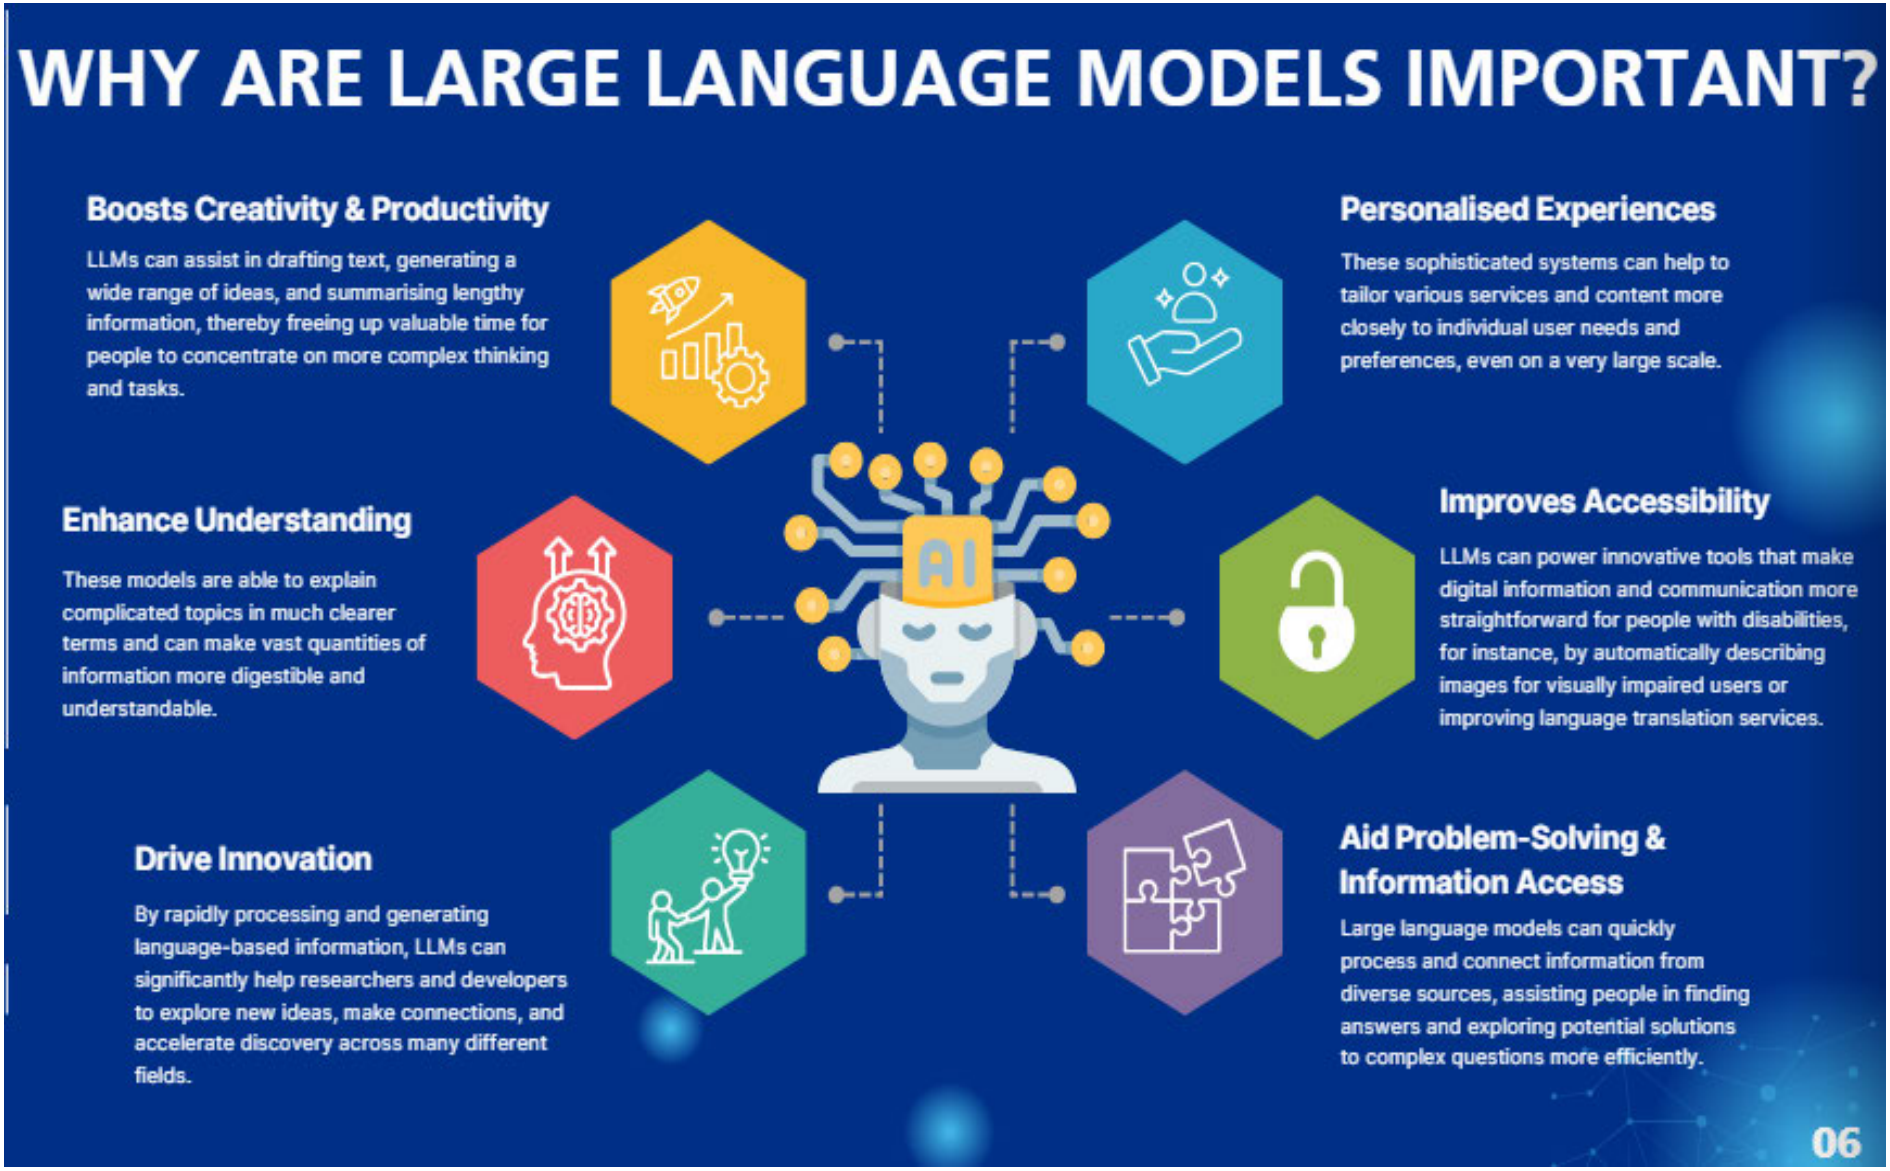

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW DO LARGE LANGUAGE MODELS WORK?

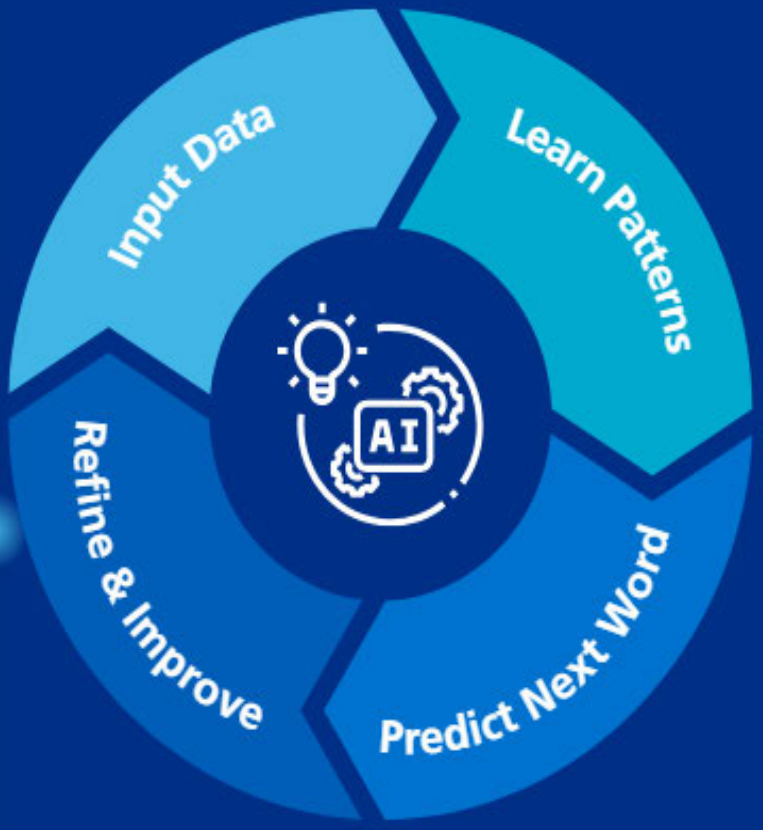

Input Vast Text Data

LLMs are initially fed billions of sentences from diverse range of written sources like books, articles, and websites. This huge collection of text is called training data.

Analyse & Identify Patterns

By analysing this massive amount of text, the models learn common language structures, grammar, facts (as presented in the training data), typical phrases, and different writing styles.

Learn to Predict Next Word

A core part of this learning is becoming very skilled at predicting which word or part of a word is most likely to appear next in any given sentence or phrase, based on the context.

Refine & Improve Predictions

Through this continuous process of prediction and adjustment, often self-guided, without direct human labelling for each step, the models get better and better at understanding and generating natural-sounding language.

07

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES LARGE LANGUAGE MODELS WORK?

The "Transformer" - A Breakthrough Design

Older ways of understanding text were like trying to cook by seeing only one instruction of a recipe at a time, without knowing what comes next or what the final dish is. The 'Transformer' design is like being able to see the entire recipe at once. This allows the LLM to understand how all the steps and ingredients relate to create the final meal, helping the LLM understand the context and how words relate more effectively.

Older systems were like trying to follow a recipe by seeing only one instruction at a time. It is hard to know what to do next.

The 'Transformer' design lets the LLM see the whole recipe at once, understanding how everything connects.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES LARGE LANGUAGE MODELS WORK?

"Attention" - Focusing on What Matters

A key part of the transformer design is something called an "attention mechanism." In simple terms, this allows the model, when looking at a particular word, to also "pay attention" to other relevant words in the sentence, no matter where these words are located. This helps the model to understand relationships between words and how the words influence each other's meaning to build a better understanding of the overall context.

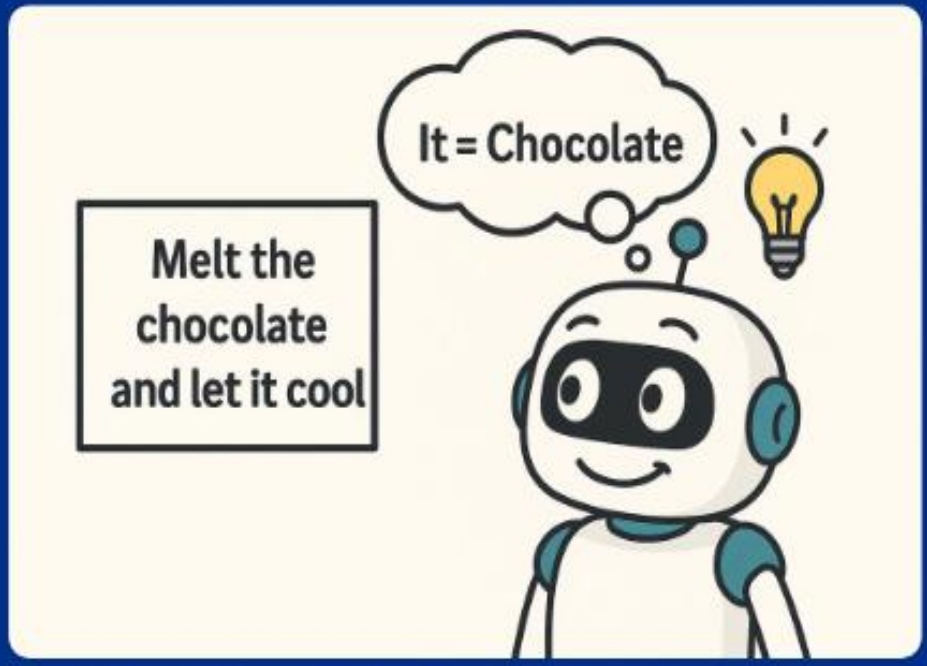

Think about reading the sentence: 'Melt the chocolate and let it cool.' To understand what 'it' refers to, your brain quickly links 'it' back to 'the chocolate'. The 'attention mechanism' helps the large language model do something similar. The system is very good at finding these links between words in a sentence, helping the system understand the full meaning correctly by connecting the right pieces of information.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT CAN LLMS ACTUALLY DO?

01 Text Generation & Writing

These models can create new written content from a simple instruction or starting point. This includes drafting articles, composing stories or poems, and even helping to write emails or marketing material.

02 Summarisation

Large language models can process long documents, research papers, or articles and then provide concise, easy-to-understand summaries that highlight the main points and key information.

03 Translation

These systems are able to convert text from one language to another. Modern LLMs can do this with increasing accuracy, often capturing subtle meanings and nuances between many different languages.

04 Answering Questions

LLMs can understand questions asked in everyday language and provide relevant and informative answers. This capability is based on the vast knowledge the models learned during their training.

05 Coding Assistance

For those who work with computer code, LLMs can be helpful assistants. These models can generate new code snippets based on a description, explain what existing pieces of code do, or even help to find and fix errors (debugging).

06 Conversation & Chatbots

Large language models are excellent at engaging in dialogues. This makes these models ideal for powering advanced chatbots for customer service, acting as creative partners for brainstorming ideas, or simply holding a helpful conversation.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

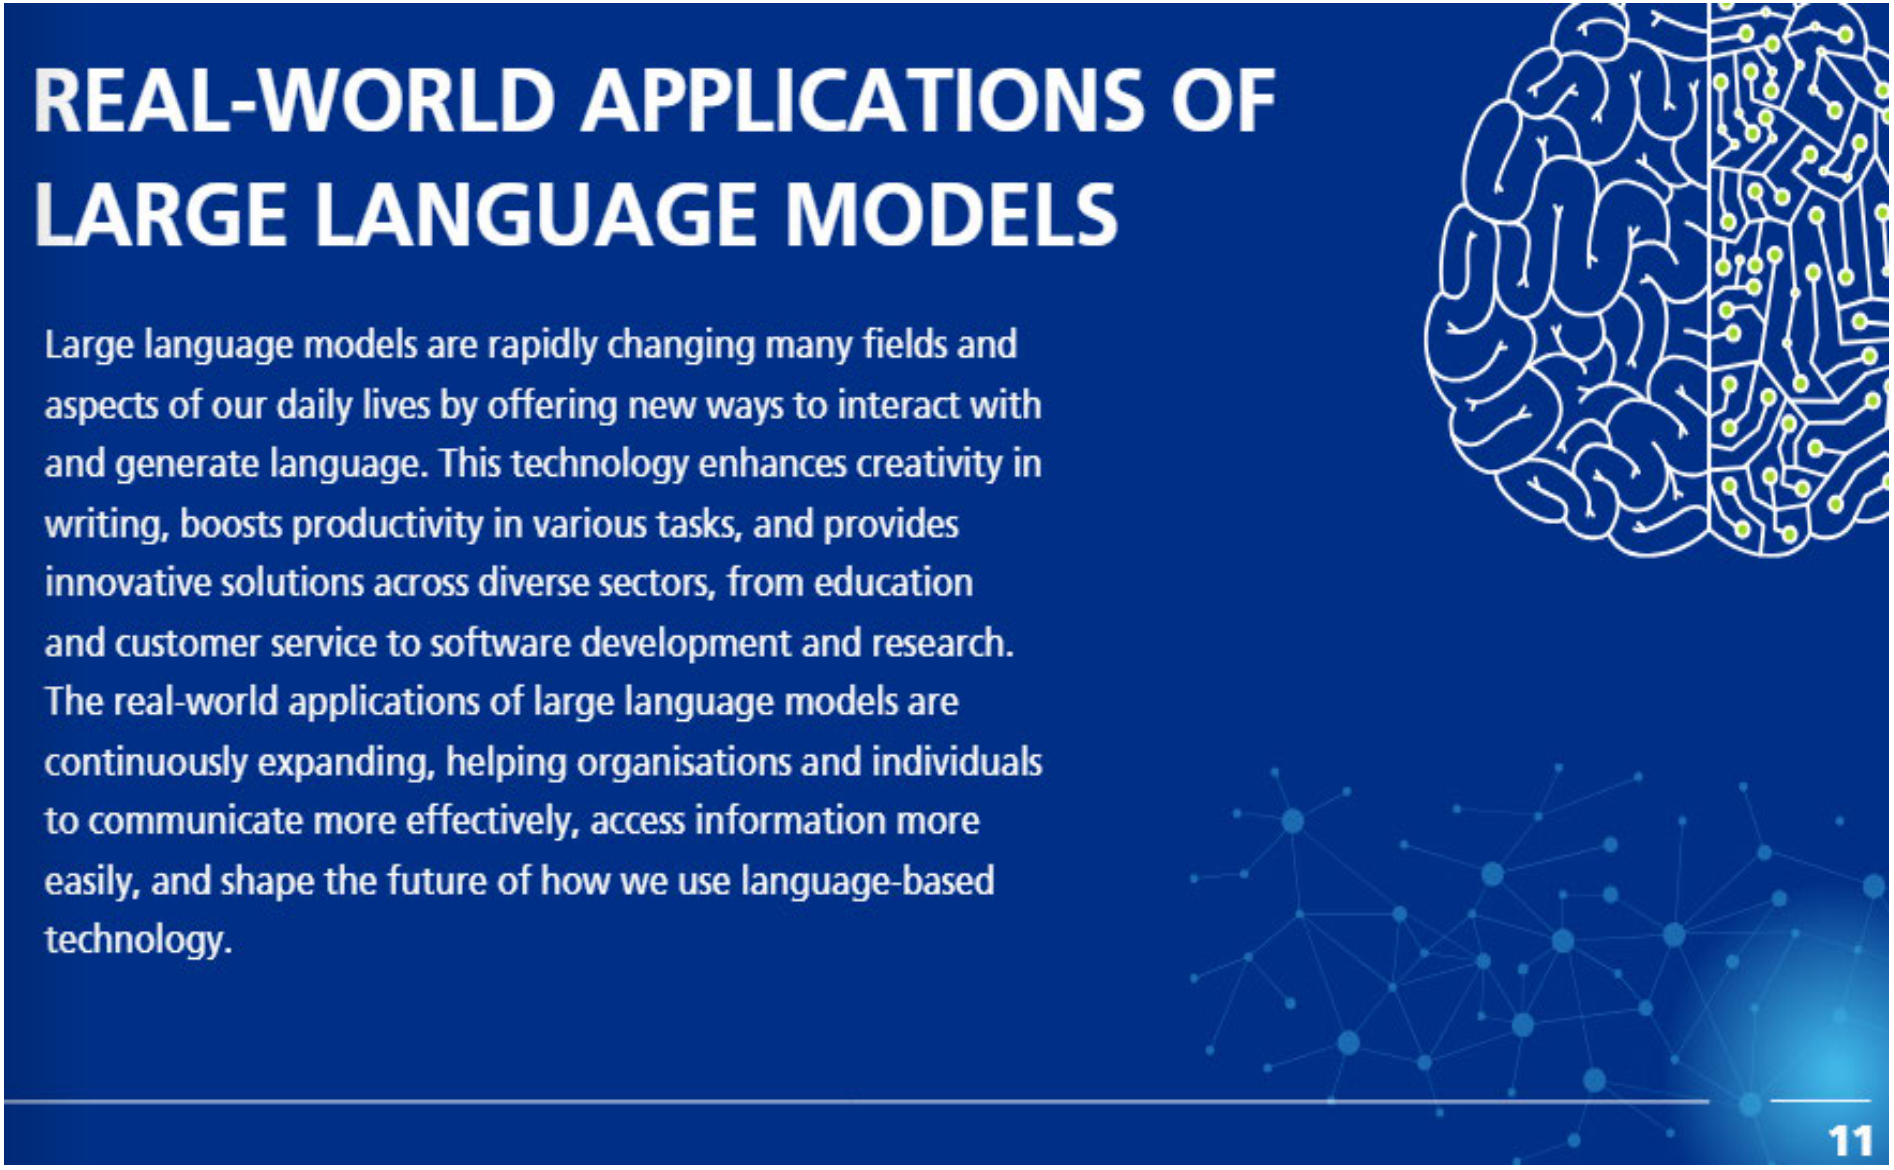

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

TEXT GENERATION & WRITING ASSISTANCE

Empowering communication, creativity, and development.

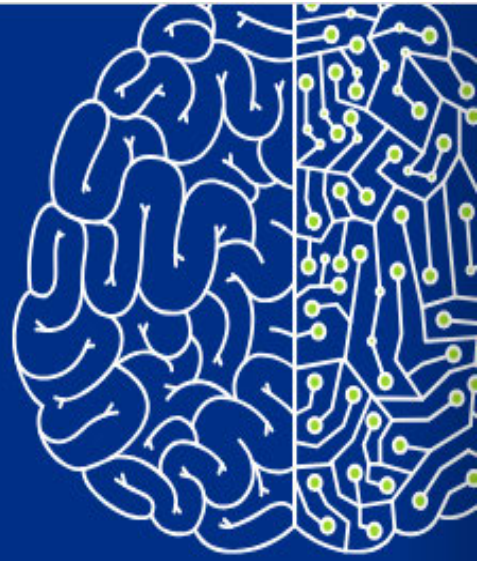

DRAFTING INITIAL CONTENT

LLMs can quickly produce first drafts for emails, reports, articles, blog posts, or marketing copy.

CREATIVE WRITING PARTNER

For authors, poets, or scriptwriters, these models can act as a creative partner. LLMs can suggest plot ideas, help develop character dialogues, write poetry in various styles, or offer alternative phrasing to enhance creative work.

IMPROVING WRITTEN WORK

Beyond just generating text, LLMs can assist in refining existing writing. This includes helping with grammar, style, tone, and clarity, acting like an advanced editing assistant.

PERSONALISED COMMUNICATION

These models can help tailor messages for specific audiences or purposes, for example, by adjusting the formality of an email or simplifying complex information for a wider readership.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CONVERSATION & CHATBOTS

Powering natural dialogues and helpful AI assistants.

SMARTER CUSTOMER SERVICE

Many companies now use LLM-powered chatbots on their websites or in apps to provide instant customer support. These chatbots can understand a wide range of questions and provide helpful answers or guidance 24/7.

HELPFUL VIRTUAL ASSISTANTS

Personal virtual assistants, like those on phones or smart speakers, are becoming more conversational and capable thanks to LLMs. These assistants can understand more complex requests and engage in more natural dialogues.

INFORMATION & LEARNING COMPANIONS

LLMs can act as interactive learning tools or information companions. People can ask questions on various topics and receive detailed explanations.

BRAINSTORMING & IDEA EXPLORATION

Engaging in a dialogue with an LLM can be a great way to brainstorm ideas, explore different perspectives on a topic, or even get feedback on your own thoughts in a conversational way.

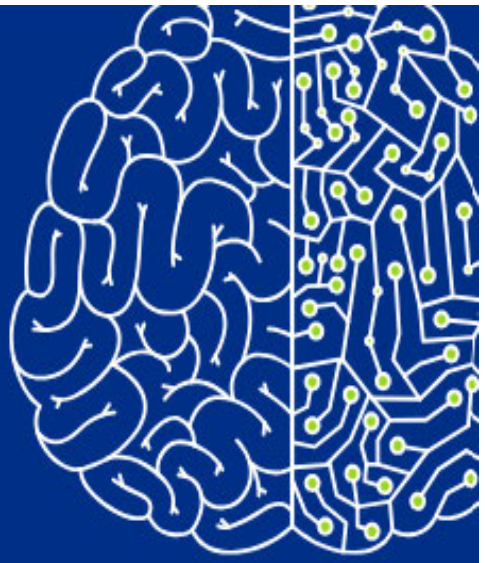

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

SUMMARISATION & INFORMATION PROCESSING

Making sense of large amounts of text quickly.

CONDENSING LONG DOCUMENTS

LLMs can take lengthy reports, research papers, or articles and quickly produce a short summary. This helps people to understand the main ideas without needing to read every single word.

EXTRACTING KEY INFORMATION

These models can scan through texts to find and pull out specific pieces of information, such as important facts, dates, names, or topics.

UNDERSTANDING COMPLEX TEXTS

For dense or technical material, LLMs can help by breaking down complex information into simpler terms or by highlighting the most important arguments and conclusions.

KEEPING UP-TO-DATE EFFICIENTLY

In fields where a lot of new information is published regularly, like news or scientific research, LLMs can help people stay informed by summarising new content quickly.

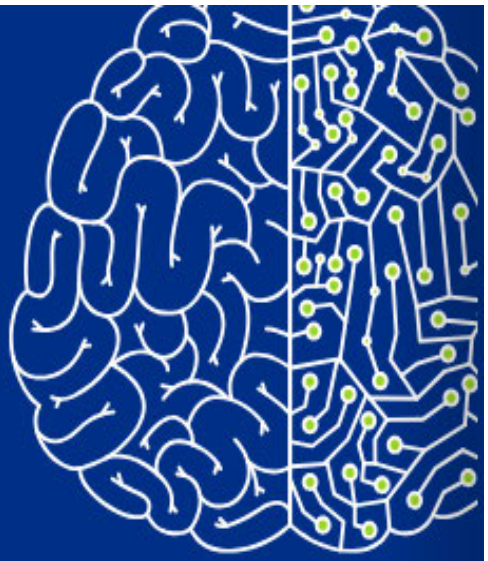

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

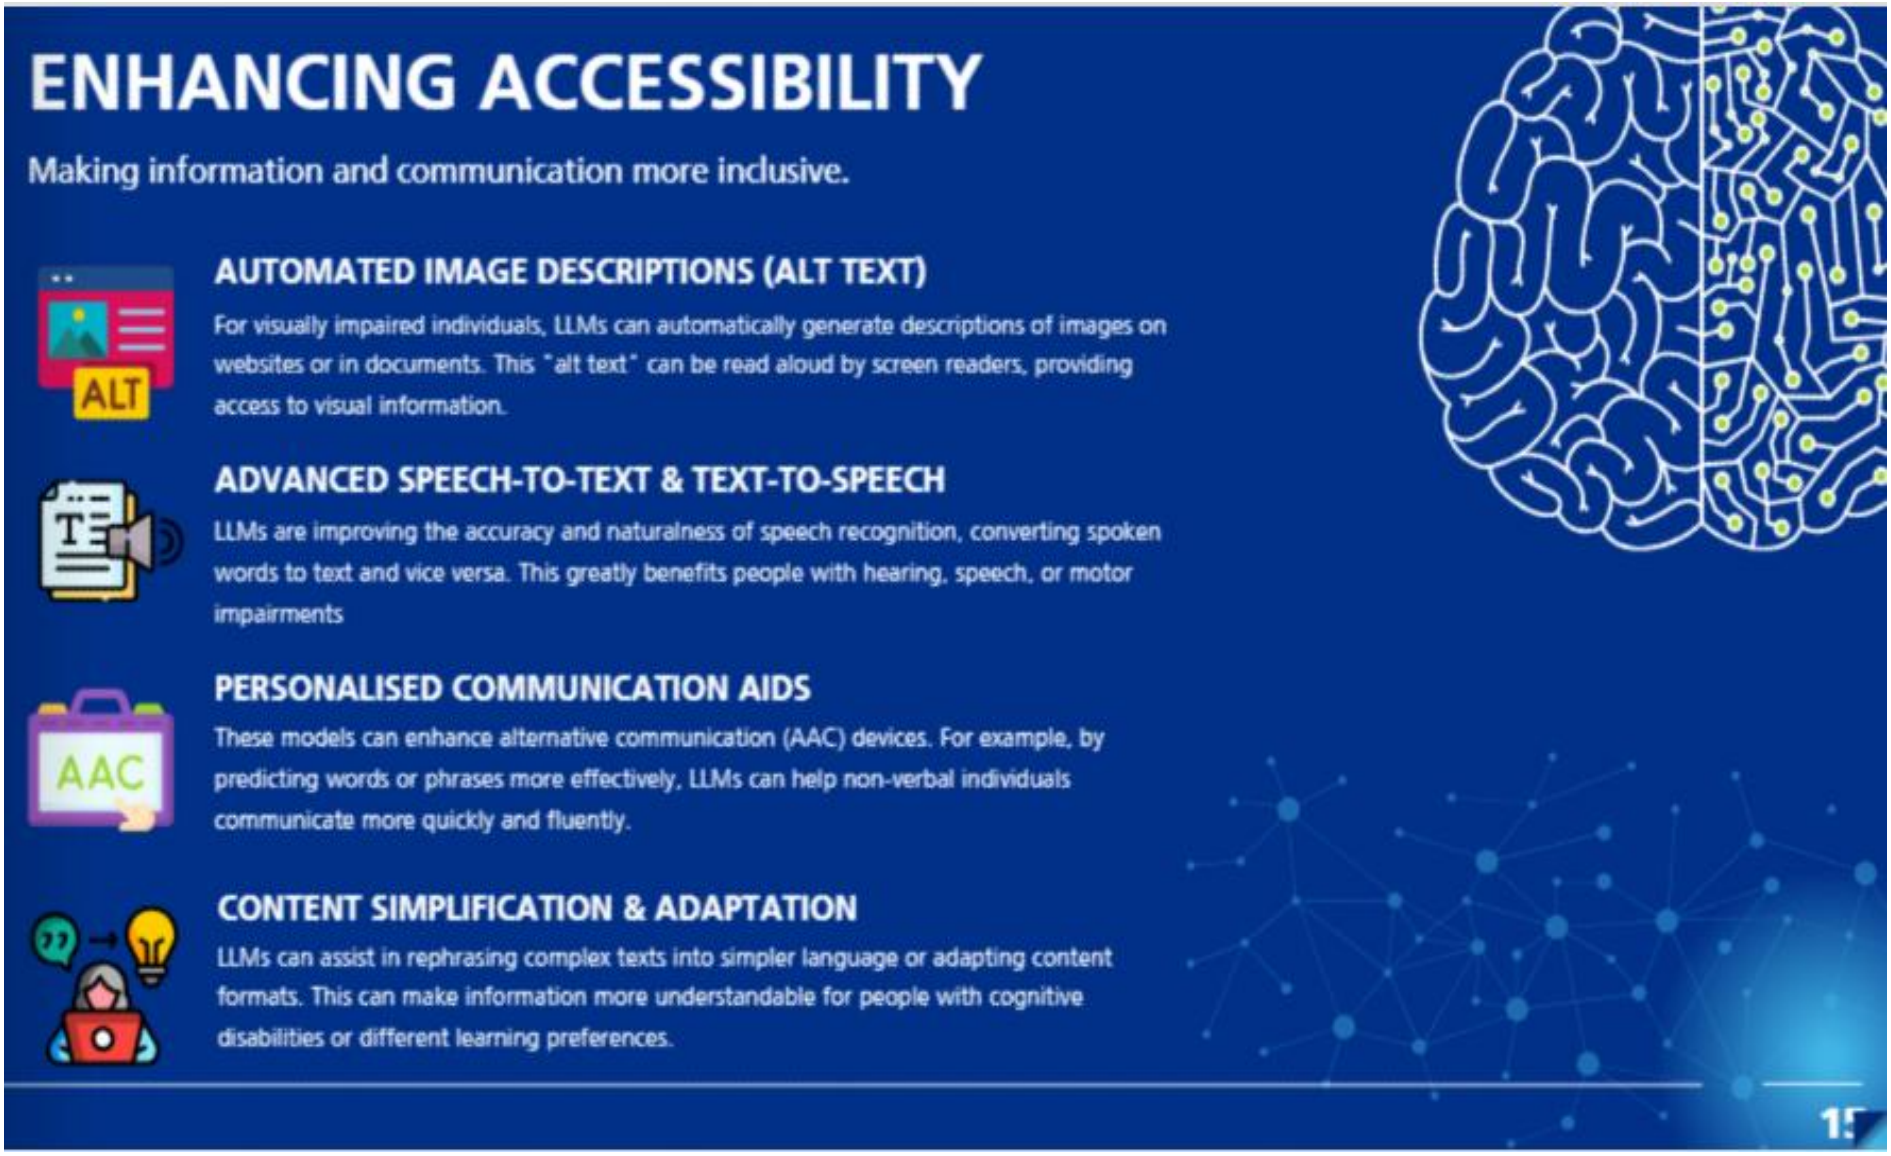

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CHALLENGES AND IMPORTANT QUESTIONS

01 **Accuracy & "Hallucinations"**

LLMs can sometimes generate information that may be inaccurate or unclear, often called "hallucinations", even if the information sounds convincing. Outputs always need careful checking.

02 **Bias in Data and Outputs**

Because LLMs learn from vast amounts of internet text, these models can unintentionally learn and repeat societal biases related to things like gender or race, if such biases are present in the training data.

03 **Lack of True Understanding**

LLMs are very good at recognising patterns in language, but these systems do not "understand" concepts or the real world in the way humans do.

04 **Misinformation & Misuse**

The ability of LLMs to generate realistic text means these tools could be misused to create and spread false information like "fake news" or for other harmful purposes like scams.

05 **Copyright and Ownership**

Questions exist about the use of copyrighted material in training data and who owns the content created by LLMs. These are complex legal and ethical areas.

06 **Environmental Impact**

Training very large language models requires significant computing power, which in turn uses a lot of electricity and can have an environmental cost.

07 **Impact on Jobs**

Like with any new technology, there are conversations about how LLMs could be used in tasks like writing or processing information

08 **Data Privacy**

When people interact with LLMs, questions arise about how the data shared in prompts and conversations is used and protected.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 LLMs Learn from Vast Text Data

Large language models develop their abilities by processing and finding patterns in enormous collections of text. The quality and nature of this data heavily influence what the model learns, including potential biases.

02 These are Pattern-Matching Systems

At their core, LLMs are very sophisticated at recognising and reproducing patterns in language; these systems do not "understand" or "think" like humans.This explains why LLMs can be fluent but sometimes incorrect, and why human oversight is vital.

03 "Prompts" are Key to Guidance

The instructions (prompts) people provide to a large language model significantly shape the output the model generates. Clear, specific prompting helps to get more useful and relevant results.

04 Versatile Language Tools

LLMs can perform many different language tasks, from writing and summarising to translation and answering questions. This versatility makes these models useful in many different areas and for many different people.

05 Human Oversight is Essential

Outputs from LLMs should always be reviewed by people for accuracy, appropriateness, and potential bias. Critical evaluation and human judgment are necessary for responsible use of this technology.

06 A Rapidly Evolving Field

The capabilities of large language models are improving quickly, with new developments happening all the time. Staying informed helps us to understand both the new opportunities and the new questions that arise. It is also important to be aware of the challenges and potential risks discussed. A balanced perspective allows for responsible exploration and use of this powerful technology

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 6

Machine Learning – A basic overview

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS MACHINE LEARNING?

Machine Learning is a type of Artificial Intelligence (AI), which aims to make computers perform tasks that typically require human intelligence, think of machine learning as the 'learning' part of AI. While powerful at finding patterns, machine learning doesn't 'think' or 'understand' context the way humans do.

01

Imagine teaching a child to recognise a cat. You show them lots of pictures of cats and point out their features like pointy ears and whiskers.

02

Machine Learning does something similar by examining huge amounts of data. It automatically learns the key patterns and features, such as the ears or whiskers, enabling it to recognise a cat itself.

03

Just as the child starts to spot the 'cat patterns' (ears, whiskers) on their own, machine learning searches through the examples to find patterns automatically, without needing step-by-step instructions.

04

Once machine learning has learned the patterns, it can look at new data it has not encountered before and make a useful prediction or decision.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

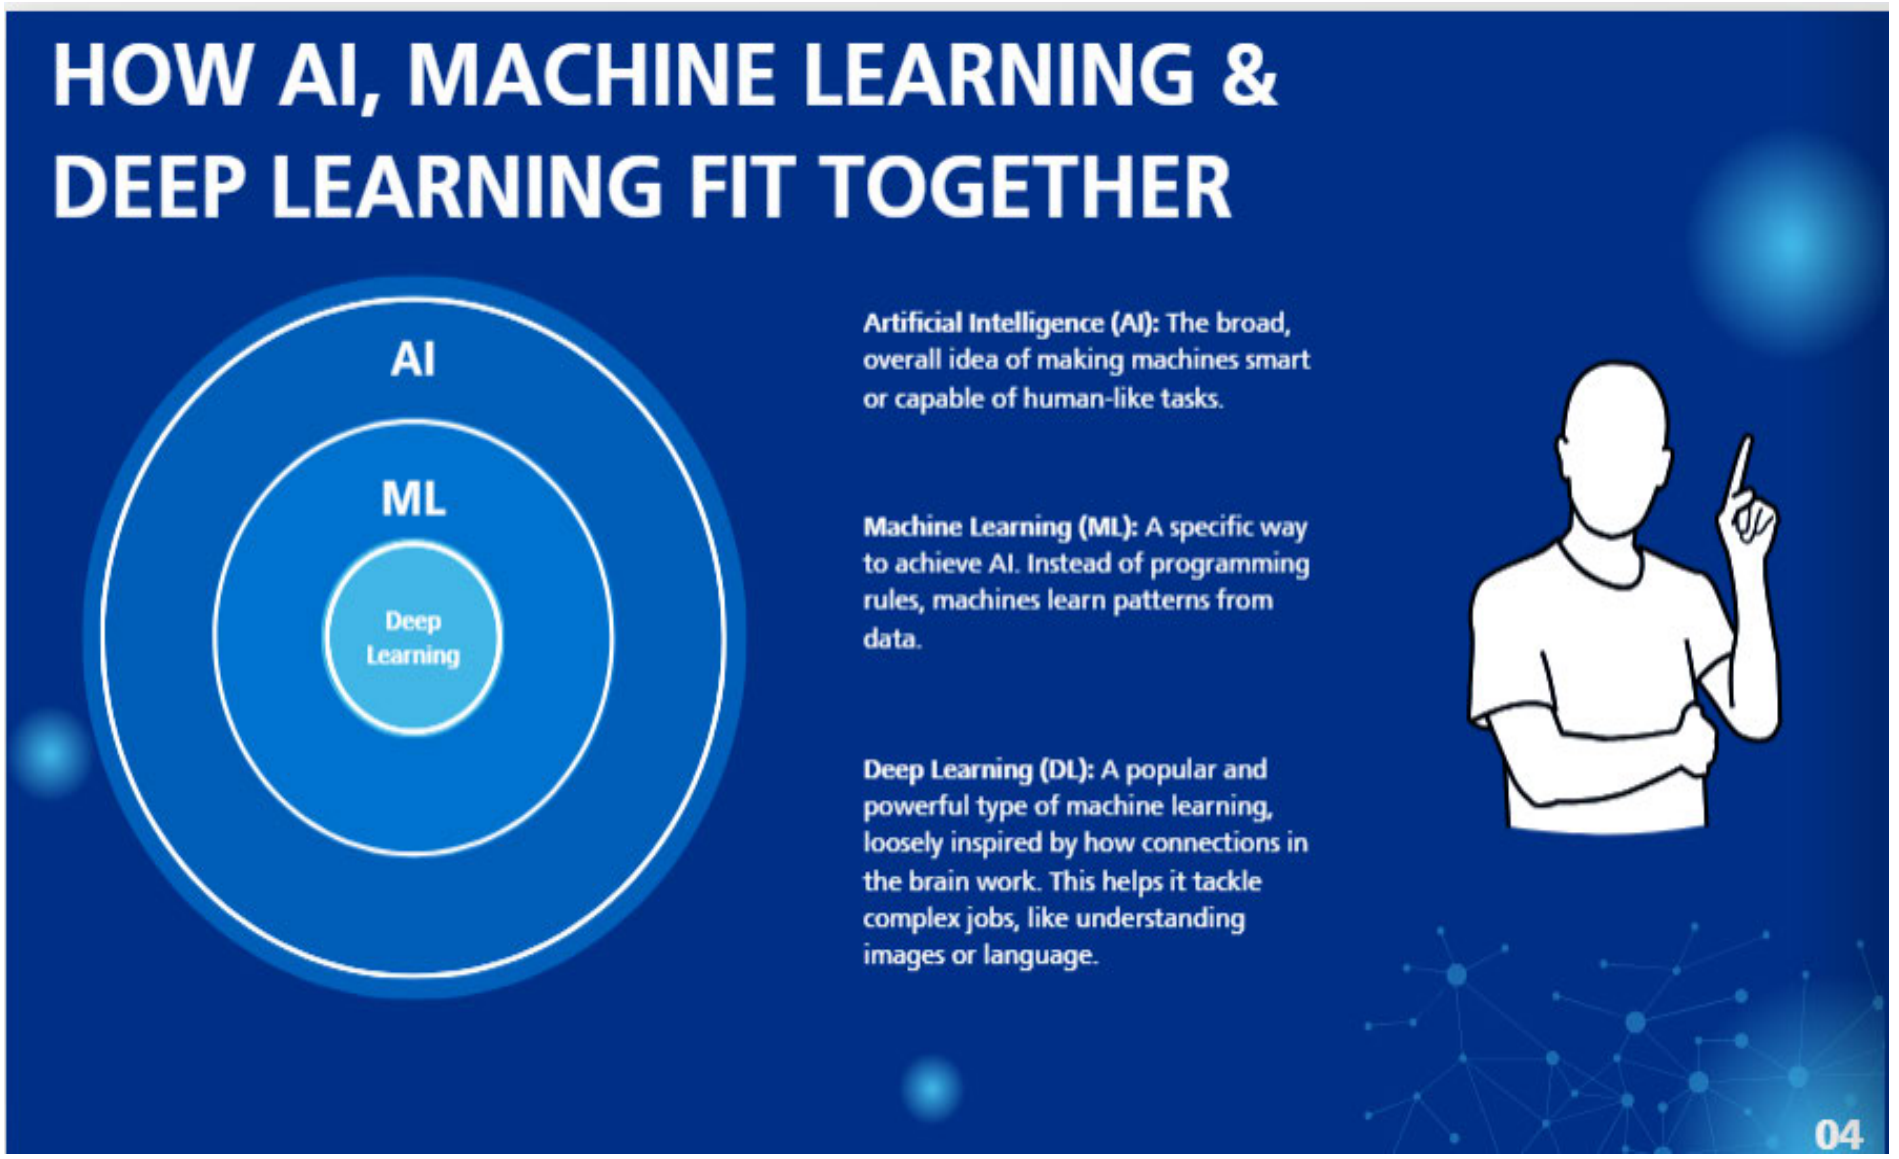

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY USE MACHINE LEARNING?

01 Making Sense of Huge Amounts of Information

Machine Learning makes it possible to search through enormous amounts of data, far more than any person could manage, to find valuable information that would otherwise be missed. Think of finding a needle in a giant haystack.

02 Finding Hidden Patterns and Insights

Machine Learning is great at spotting connections and trends in data that are not obvious to humans, helping us understand complex situations or human behaviour better.

03 Automating Tasks and Making Decisions Faster

It can learn to perform repetitive jobs automatically, like sorting emails into spam, or help make quick decisions based on lots of information, such as suggesting the best route in traffic.

04 Making Experiences More Personal

Companies use it to tailor services just for you. Think of Netflix recommending shows you might like, or Amazon suggesting products based on what you have viewed.

05 Predicting Future Outcomes

By learning from past data, machine learning can help make accurate predictions, like forecasting weather, predicting sales trends, or identifying potential problems before they happen.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW DOES MACHINE LEARNING WORK?

01

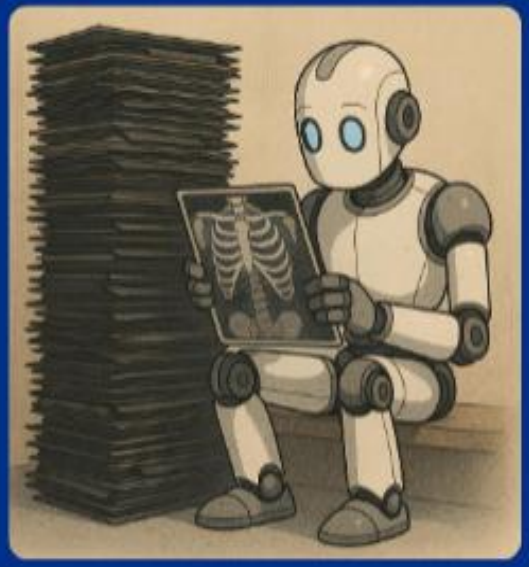

The machine learning model is first given lots of information where the outcome is already known. For example, it might be shown thousands of X-ray images already labelled by experts as showing pneumonia or not. Crucially, the quality and accuracy of this initial information directly impacts how well it learns.

02

The machine learning model then analyses this information. It searches for patterns that connect specific features, like certain shapes in an image, to the known outcomes, like the pneumonia diagnosis.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

03

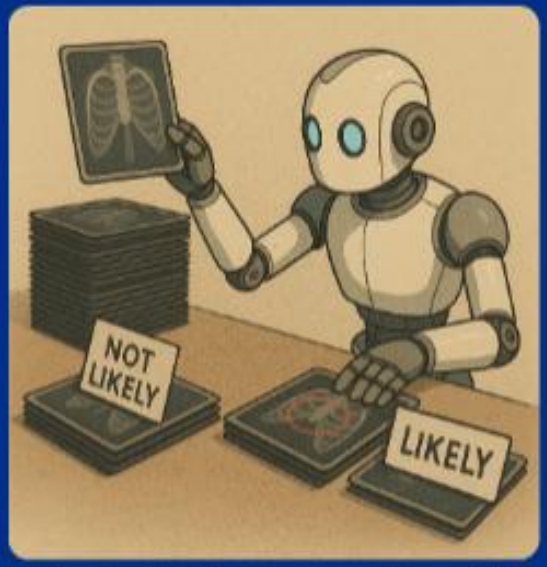

Having learned these patterns, the machine learning model can then look at new, unseen information, like a new X-ray, and make a prediction, for example, estimating the likelihood of pneumonia.

04

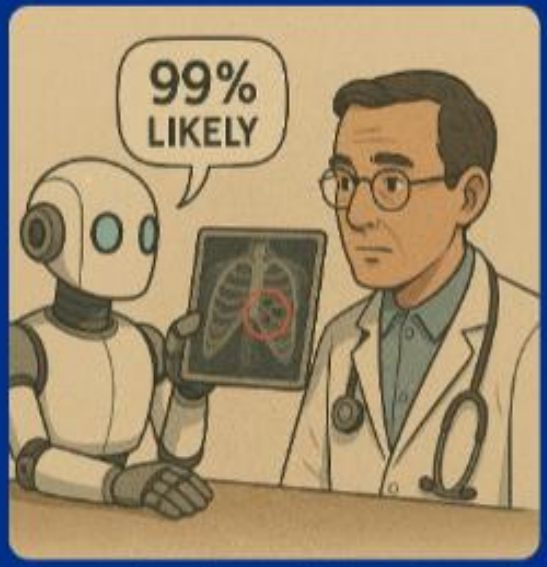

The more data the machine learning model sees, the better it gets at making accurate predictions.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

DIFFERENT WAYS MACHINES LEARN

Learning with a Teacher (Supervised Learning)

Machine Learning learns by using examples where the correct answers are provided. It is like studying with a teacher who tells you if you are right or wrong.

Finding Patterns on Its Own (Unsupervised Learning)

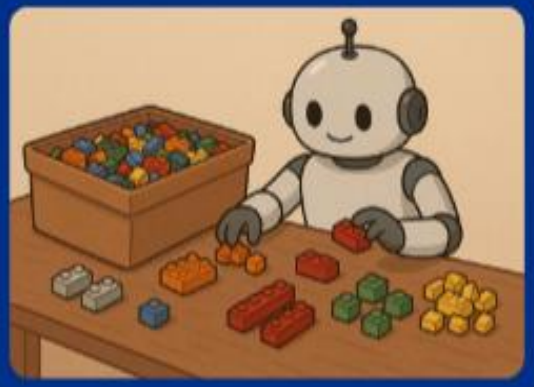

Machine Learning explores information where no answers are given. It tries to find interesting patterns or groups within the information by itself.

Learning Through Trial and Error (Reinforcement Learning)

Machine Learning learns by trying different actions. It gets 'rewards' for useful actions and 'penalties' for unhelpful ones, discovering the best approach over time.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

PEOPLE AND AI WORKING TOGETHER

01 Defining the Problem

Humans decide what goal machine learning should achieve.

02 Gathering & Preparing Data

Humans collect, clean, and prepare the data for the machine learning model. This is often the hardest part!

03 Choosing the Approach

Humans select the right type of learning (Supervised, Unsupervised, etc.) and tools.

04 Checking the Results

Humans evaluate if the output makes sense, is accurate, and importantly, is fair.

05 Making Decisions

Humans use the insights from machine learning models to make real-world decisions.

06 Monitoring

Humans need to keep an eye on the system to ensure it continues working well over time.

09

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

REAL-WORLD APPLICATIONS OF MACHINE LEARNING (ML)

Machine Learning is revolutionising various industries, impacting our lives in countless ways. These technologies offer solutions to complex problems, enhancing efficiency, productivity, and decision-making across different sectors.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

TRANSPORT & NAVIGATION

Smarter routes, safer journeys.

GOOGLE MAPS

Predicts traffic patterns and adjusts routes in real time to get you there faster.

UBER & RIDE-SHARING APPS

Matches drivers and passengers efficiently using real-time location and demand prediction.

SELF-DRIVING CARS

Machine Learning models detect pedestrians, other vehicles, and road signs to make safe driving decisions.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

FINANCE & FRAUD DETECTION

Protecting your money with real-time insight.

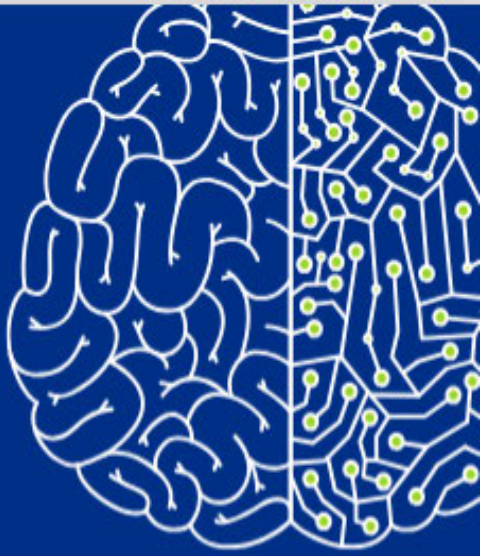

FRAUD DETECTION

Flags unusual transactions based on spending patterns learned across millions of users.

CREDIT SCORING

Machine Learning analyses many variables (not just credit history) to better predict repayment risks.

BANKING CHATBOTS

Learns from past queries to answer customer questions 24/7 with increasing accuracy.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

LANGUAGE & COMMUNICATION

Breaking barriers with smart language tools.

GOOGLE TRANSLATE

Learns how languages work by studying patterns in millions of real-world text examples, leading to better, more accurate translations.

LIVE TRANSCRIPTION/SUBTITLES

Converts speech into text in real time using voice recognition models.

PREDICTIVE TYPING

Suggests the next word or sentence based on context and past writing.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

BUSINESS & AUTOMATION

Behind the scenes in modern businesses.

DEMAND FORECASTING

Retailers use machine learning to predict what products will sell and when, based on past trends.

CUSTOMER SERVICE BOTS

Understand common queries and resolve them without human intervention.

CV SCREENING TOOLS

Identify ideal job candidates by learning patterns from successful hires.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CREATIVE & GENERATIVE AI

AI that writes, draws, and designs.

CONTENT WRITING (CHATGPT)

Learns from large text datasets to generate human-like responses and articles.

IMAGE GENERATION (DALL·E)

Turns plain-text prompts into visual artwork using deep neural networks.

GAME DEVELOPMENT TOOLS

Automatically generate scenery, levels, or animations to save time and enhance gameplay.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 **It Learns from Experience**

Remember, machine learning is not about writing instructions for every possibility. It is about learning patterns directly from data we provide.

02 **It Finds Patterns & Makes Predictions**

At its heart, machine learning is used to automatically discover hidden connections within data and then use that knowledge to make predictions or forecasts about new, unseen situations.

03 **You Use It Every Day, Often Without Realising**

This is not just futuristic technology. From your spam filter and streaming service recommendations to navigation apps and online searches, machine learning is already part of your daily digital life.

04 **Good Quality Data is Key**

Machine Learning systems are completely reliant on the data they learn from. If the data is poor, incomplete, or biased, the results will be too. Poor quality data leads to poor results.

05 **It is a Powerful Tool, Guided by People**

While incredibly powerful, machine learning is still a tool created and directed by humans. It does not understand things like we do, and it is crucial that people guide its development and use responsibly.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 7

Machine Learning – A Clinical Overview

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS MACHINE LEARNING?

Machine Learning is a type of Artificial Intelligence (AI), which aims to make computers perform tasks that typically require human intelligence. Think of machine learning as the 'learning' part of AI. While powerful at finding patterns, machine learning does not 'think' or 'understand' clinical context the way humans do.

01

Imagine teaching a junior doctor to spot signs of sepsis. You show them many patient records, eventually they start recognising patterns.

02

Machine learning does something similar; it automatically learns the key patterns and features associated with the known outcome (sepsis), without needing explicit rules for every single possibility.

03

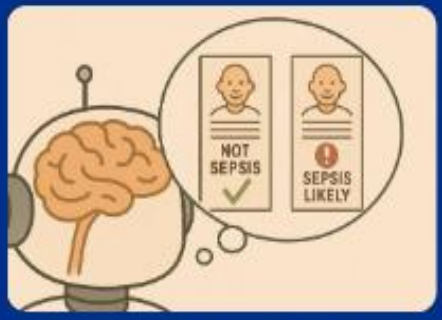

Similar to how the junior doctor learns to spot potential 'sepsis patterns' independently, the machine learning model automatically identifies these predictive patterns from examples. It does this without needing step-by-step instructions for every unique patient presentation.

04

Once machine learning has learned the patterns, it can assess new patient data. It can then make an accurate prediction or generate an alert, such as estimating the likelihood of sepsis. This could potentially flag high-risk patients earlier for prompt clinical review and action.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

HOW AI, MACHINE LEARNING & DEEP LEARNING FIT TOGETHER

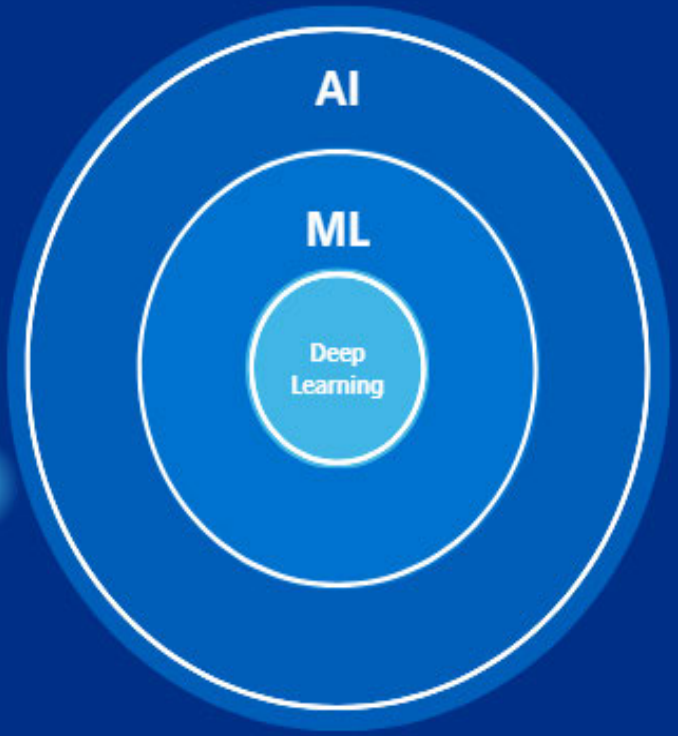

Artificial Intelligence (AI): The broad, overall idea of making machines smart or capable of human-like tasks, such as interpreting medical images or analysing clinical notes.

Machine Learning (ML): A specific way to achieve AI. Instead of programming rules, machines learn patterns directly from clinical data, like patient records or scans.

Deep Learning (DL): A popular and powerful type of machine learning, loosely inspired by how connections in the brain work. Deep Learning is particularly powerful for complex tasks like analysing medical images (X-rays, CT scans, pathology slides) or understanding clinical notes.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY USE MACHINE LEARNING IN HEALTHCARE?

01

Improving Accuracy

Machine learning can help doctors and other healthcare professionals by identifying diseases or risks more accurately, sometimes even at early stages. For example, AI can flag potential lung nodules on CT scans, assisting radiologists.

02

Increasing Efficiency

By automating routine tasks, AI can help streamline workflows, reduce waiting times, and free up healthcare staff to focus on more complex tasks. For instance, AI could help with tasks such as prioritising urgent cases.

03

Personalised Care

Machine learning can analyse individual patient data to predict how they might respond to different treatments, leading to more personalised and effective care plans.

04

Predicting and Preventing Illness

AI can help identify patients at higher risk of certain conditions (like diabetes or mental health crises) so that preventative measures can be put in place.

05

Better Resource Management

Machine learning can forecast patient flow and length of stay in hospitals, helping to manage resources like beds more effectively.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

DIFFERENT WAYS MACHINES LEARN

Learning with a Teacher (Supervised Learning)

Machine Learning learns by using examples where the correct answers are provided. It is like studying with a teacher who tells you if you are right or wrong.

Finding Patterns on Its Own (Unsupervised Learning)

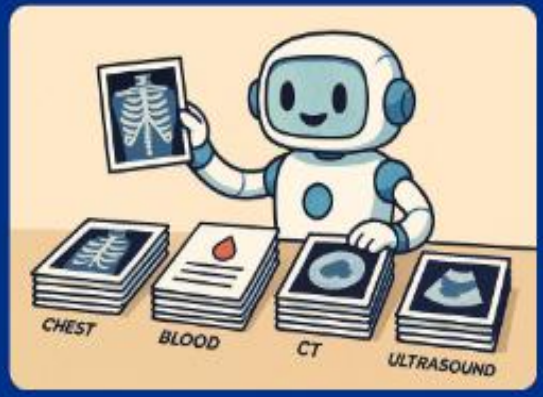

Machine learning explores information where no answers are given. It tries to find consistent patterns or groups within the information by itself.

Learning Through Trial and Error (Reinforcement Learning)

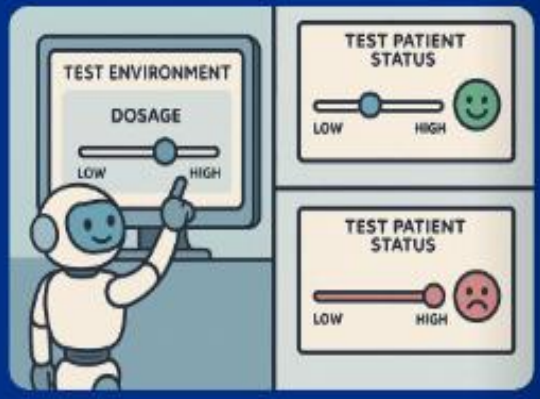

Machine learning learns by trying different actions. It gets 'rewards' for useful actions and 'penalties' for unhelpful ones, discovering the best approach over time.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

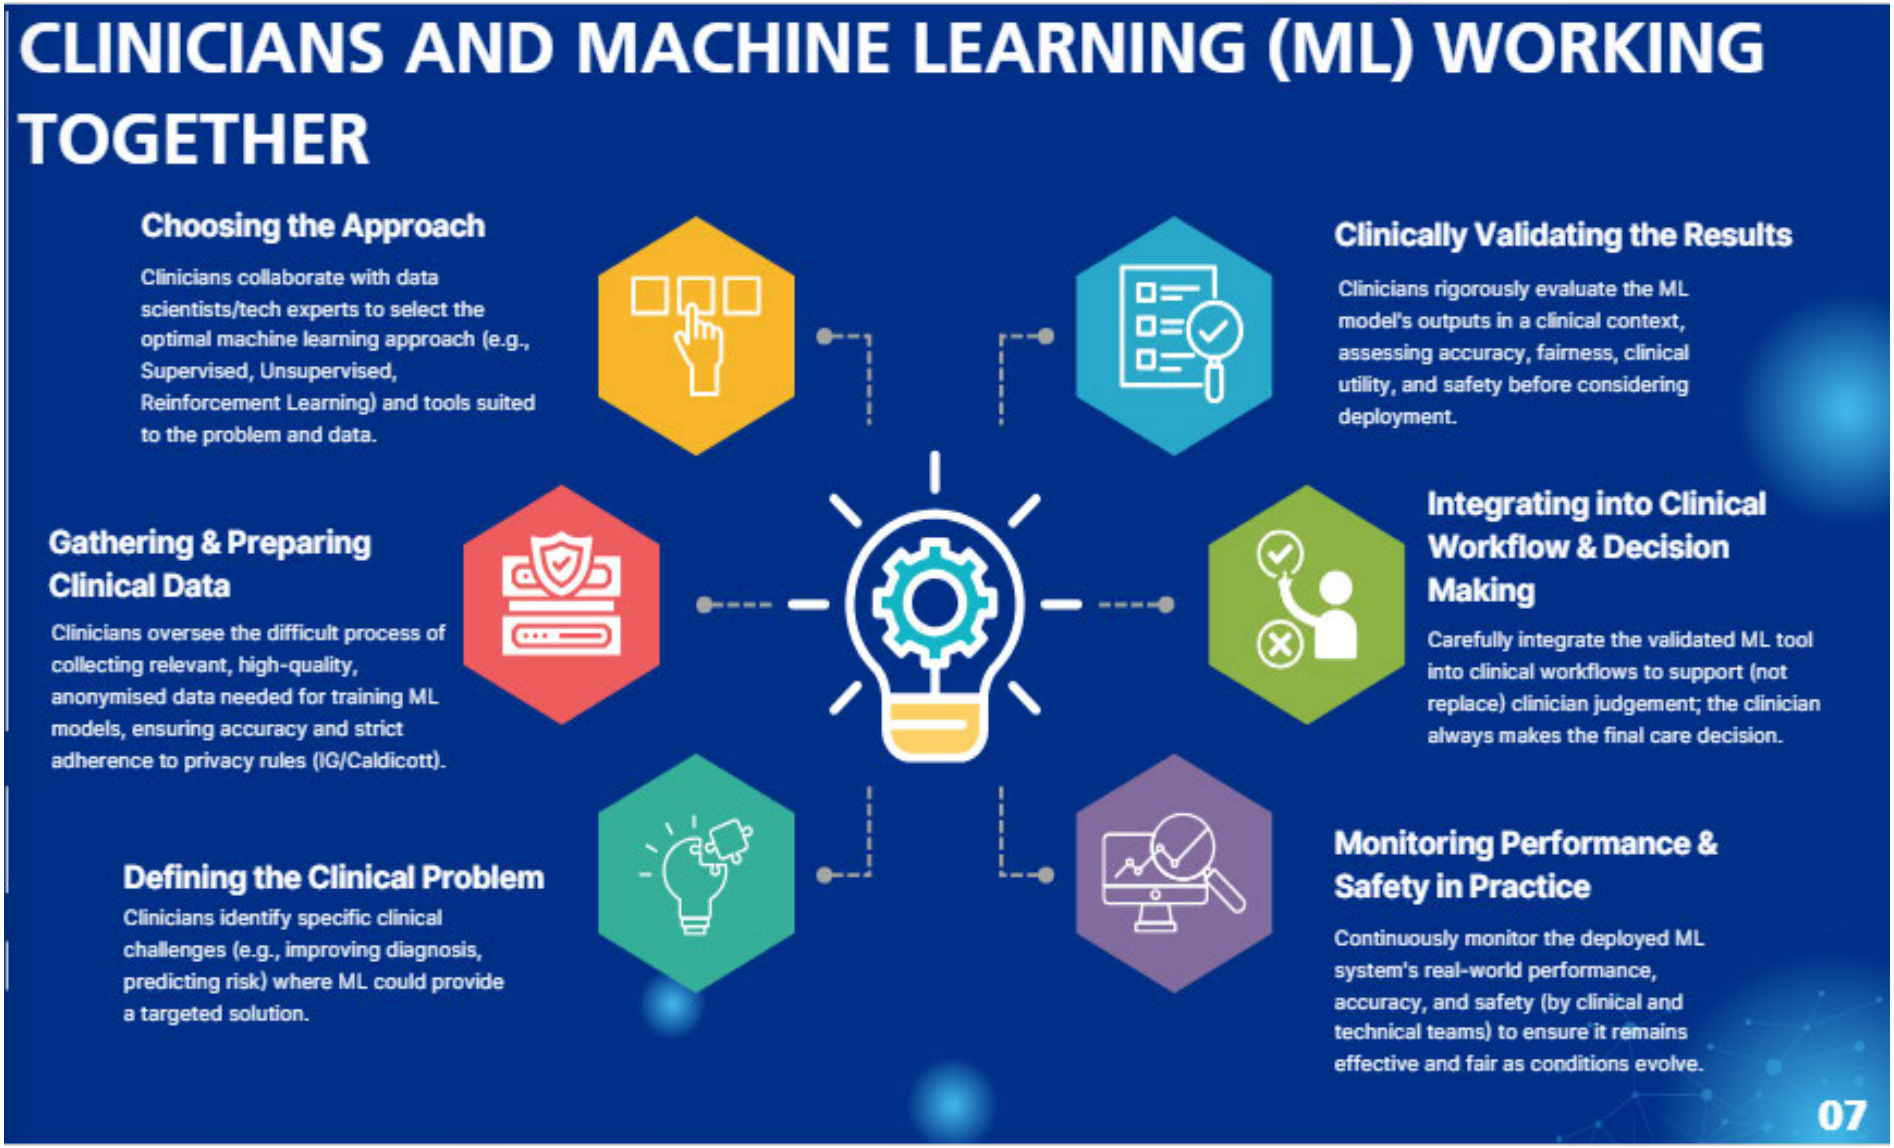

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CLINICIANS AND MACHINE LEARNING (ML) WORKING TOGETHER

01

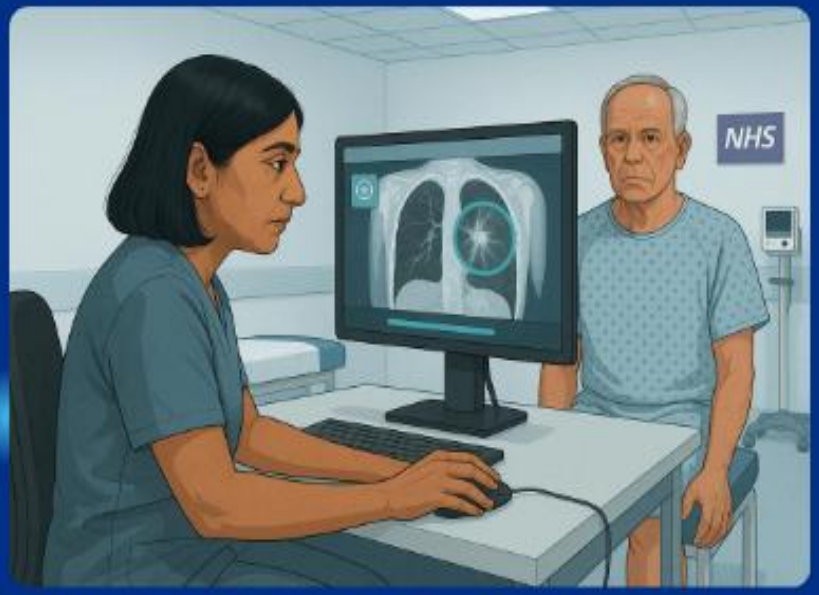

Working together with specialised software based on machine learning, the clinician reviews the patient's scan. The machine learning component is particularly helpful for identifying possible abnormalities, in this case highlighting a faint area that warrants further investigation.

02

The clinician applies their medical expertise to evaluate the area pinpointed by the machine learning software, considering it within the full picture of the patient's health. This leads to the decision that regular monitoring is the appropriate next step.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

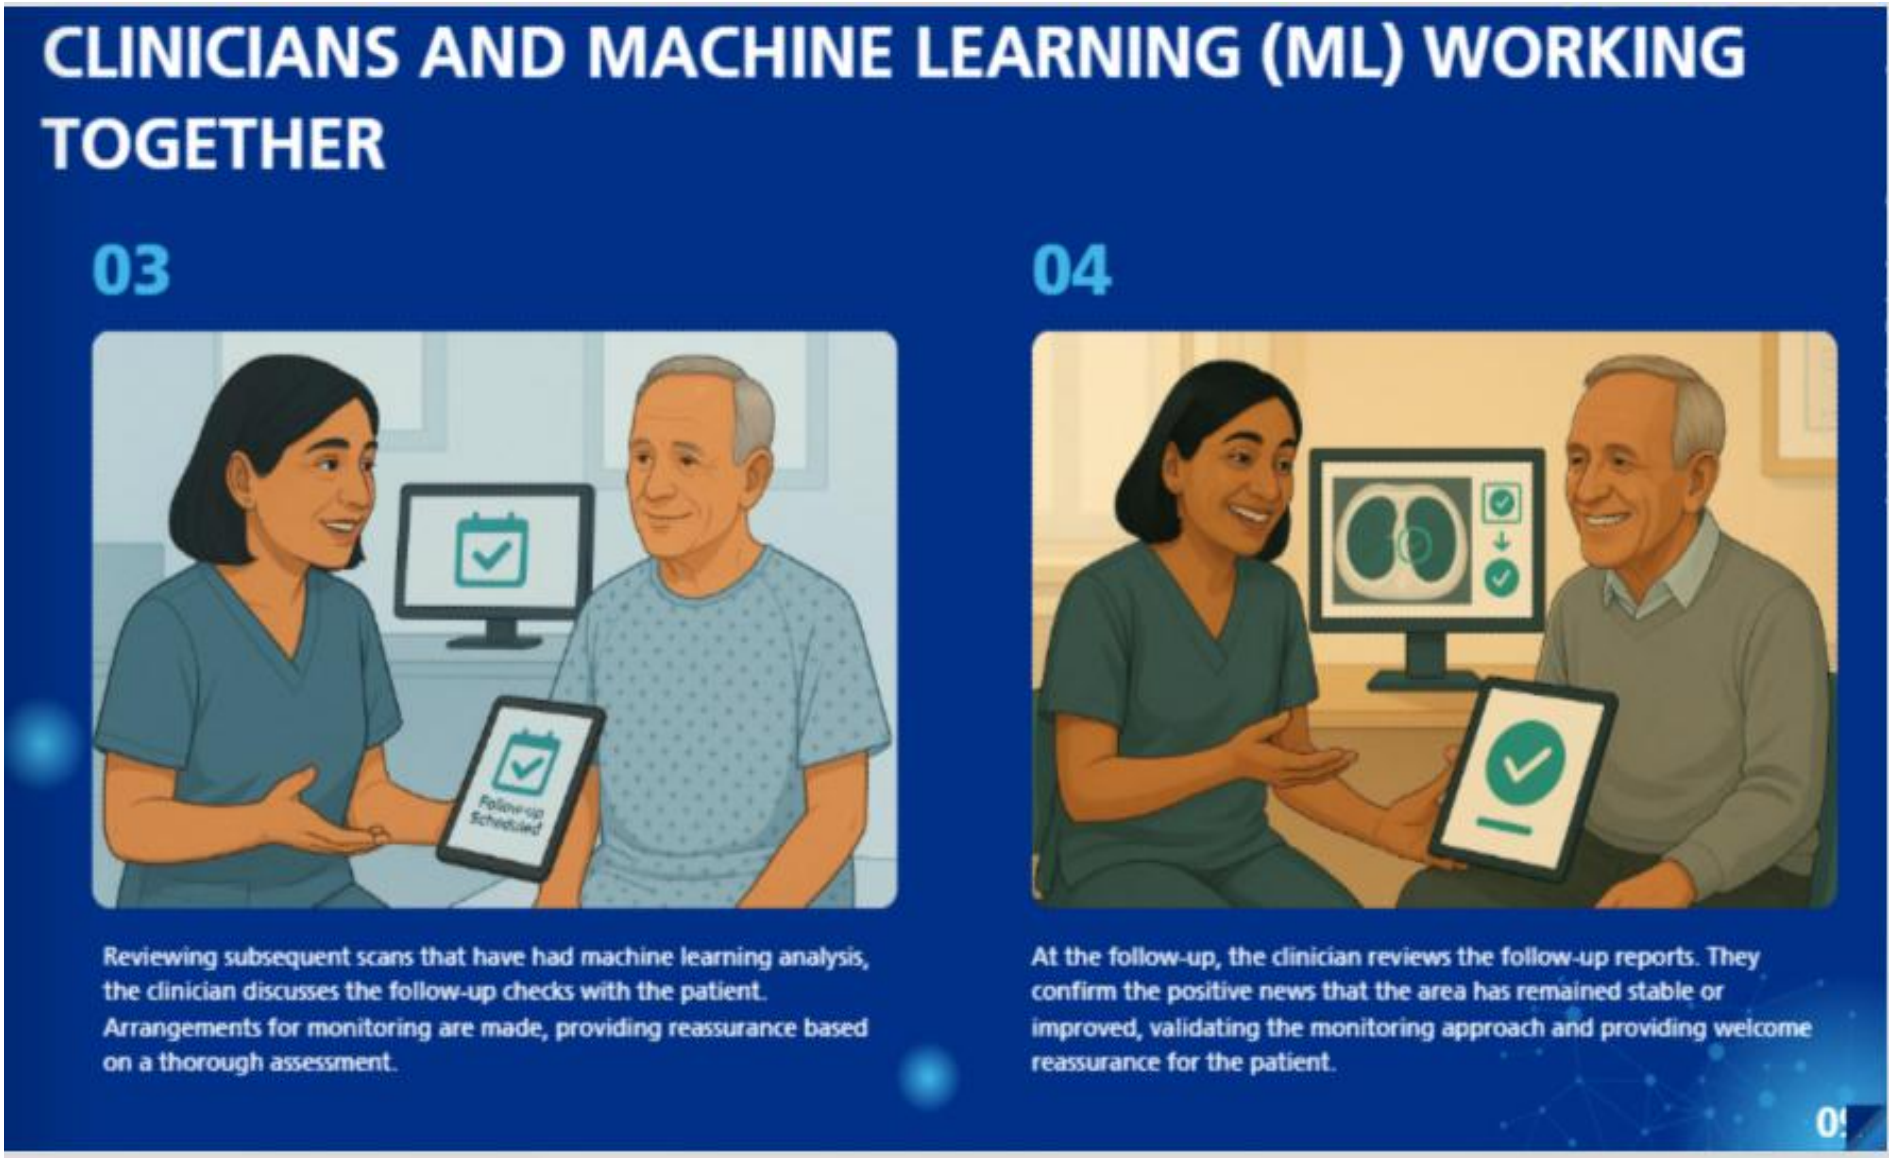

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

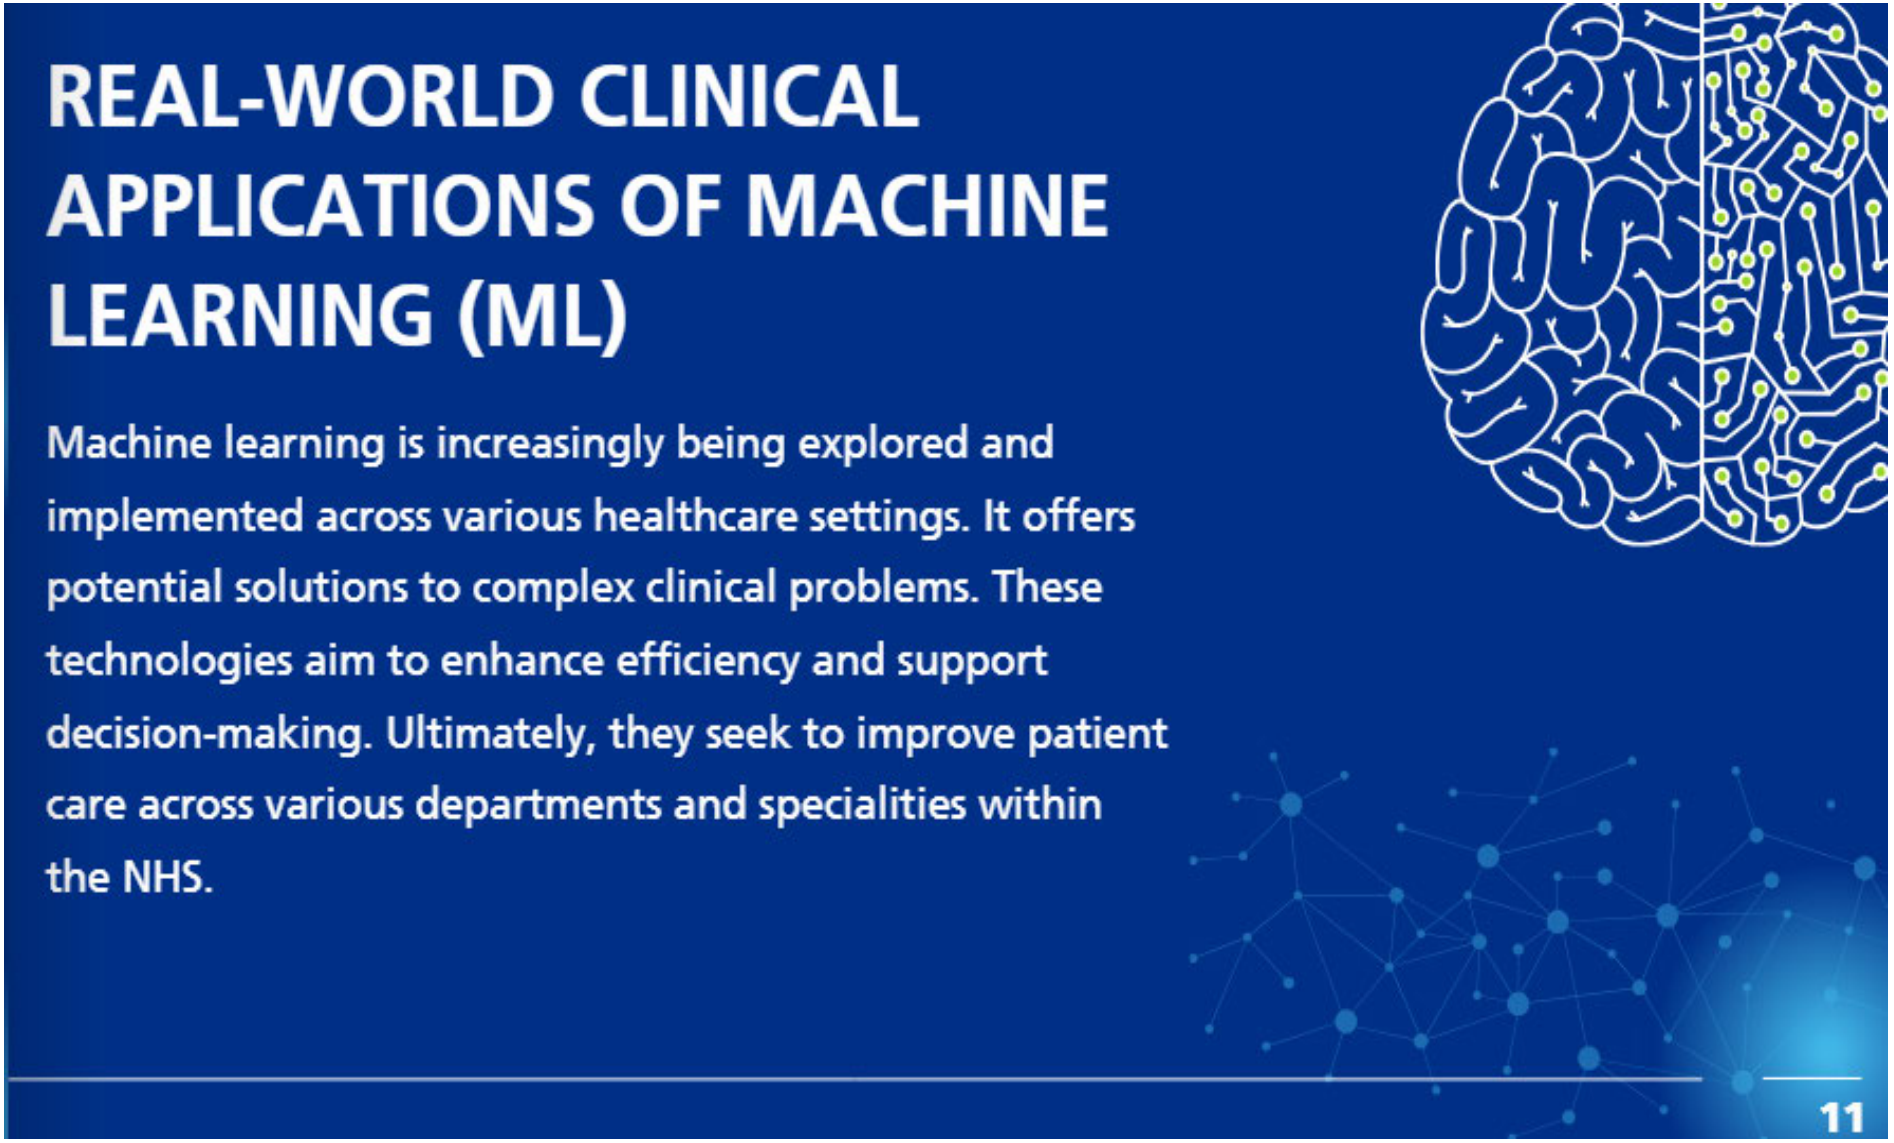

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

MEDICAL IMAGING & DIAGNOSTICS

Supporting Faster and More Accurate Interpretation

EARLY NODULE DETECTION

Assisting radiologists by automatically highlighting small lung nodules on CT scans that might need further management.

PATHOLOGY SLIDE ANALYSIS

Helping pathologists by automatically counting mitotic figures or identifying specific cell populations on digital pathology slides, speeding up analysis.

RETINOPATHY SCREENING PRIORITISATION

Screening retinal fundus images for signs of diabetic retinopathy, automatically grading severity to prioritise patients needing urgent ophthalmologist review.

STROKE IDENTIFICATION SUPPORT

Rapidly analysing non-contrast CT head scans to help detect subtle early signs of ischaemic stroke, aiding faster treatment decisions in ED.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

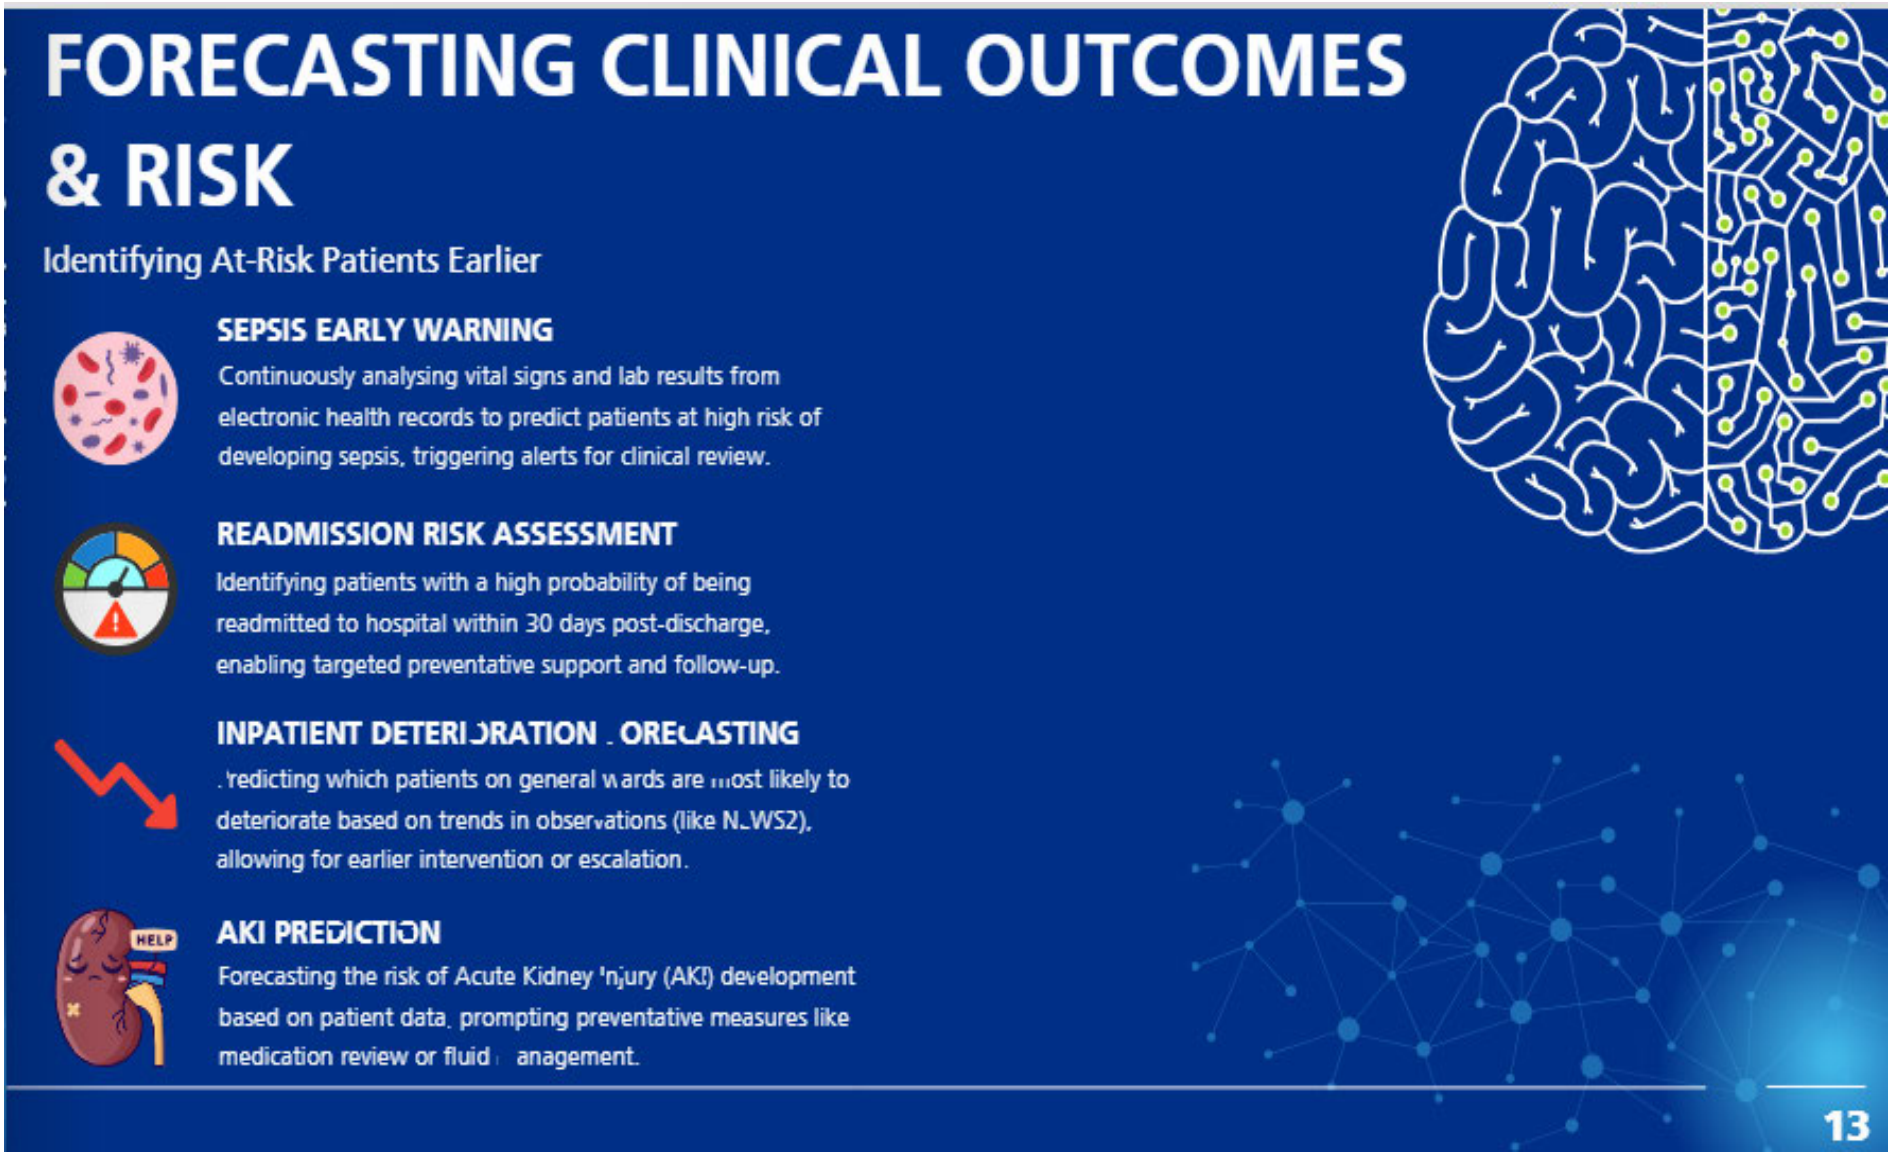

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

PERSONALISED MEDICINE & TREATMENT

Tailoring Care to the Individual

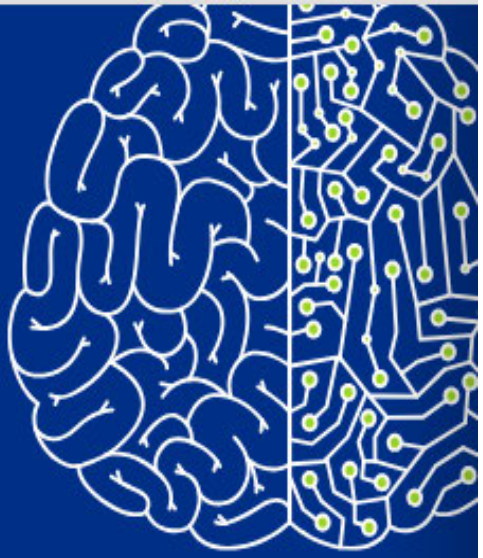

CANCER TREATMENT PATHWAY SUGGESTION

Analysing genomic data alongside clinical history and imaging to suggest potentially optimal cancer treatment regimens or identify targeted therapies for individual patients.

MEDICATION RESPONSE PREDICTION

Predicting how effectively an individual patient might respond to different classes of medication (e.g., specific antidepressants), aiding clinicians in selecting the best initial therapy.

DISEASE SUBTYPE IDENTIFICATION

Grouping patients with complex conditions like heart failure into distinct subtypes based on patterns in their data, leading to more tailored management plans for each group.

DRUG DOSAGE OPTIMISATION

Using patient characteristics e.g., genetics, kidney function, to predict individual metabolism and help optimise dosing for drugs with narrow therapeutic windows (e.g., warfarin).

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

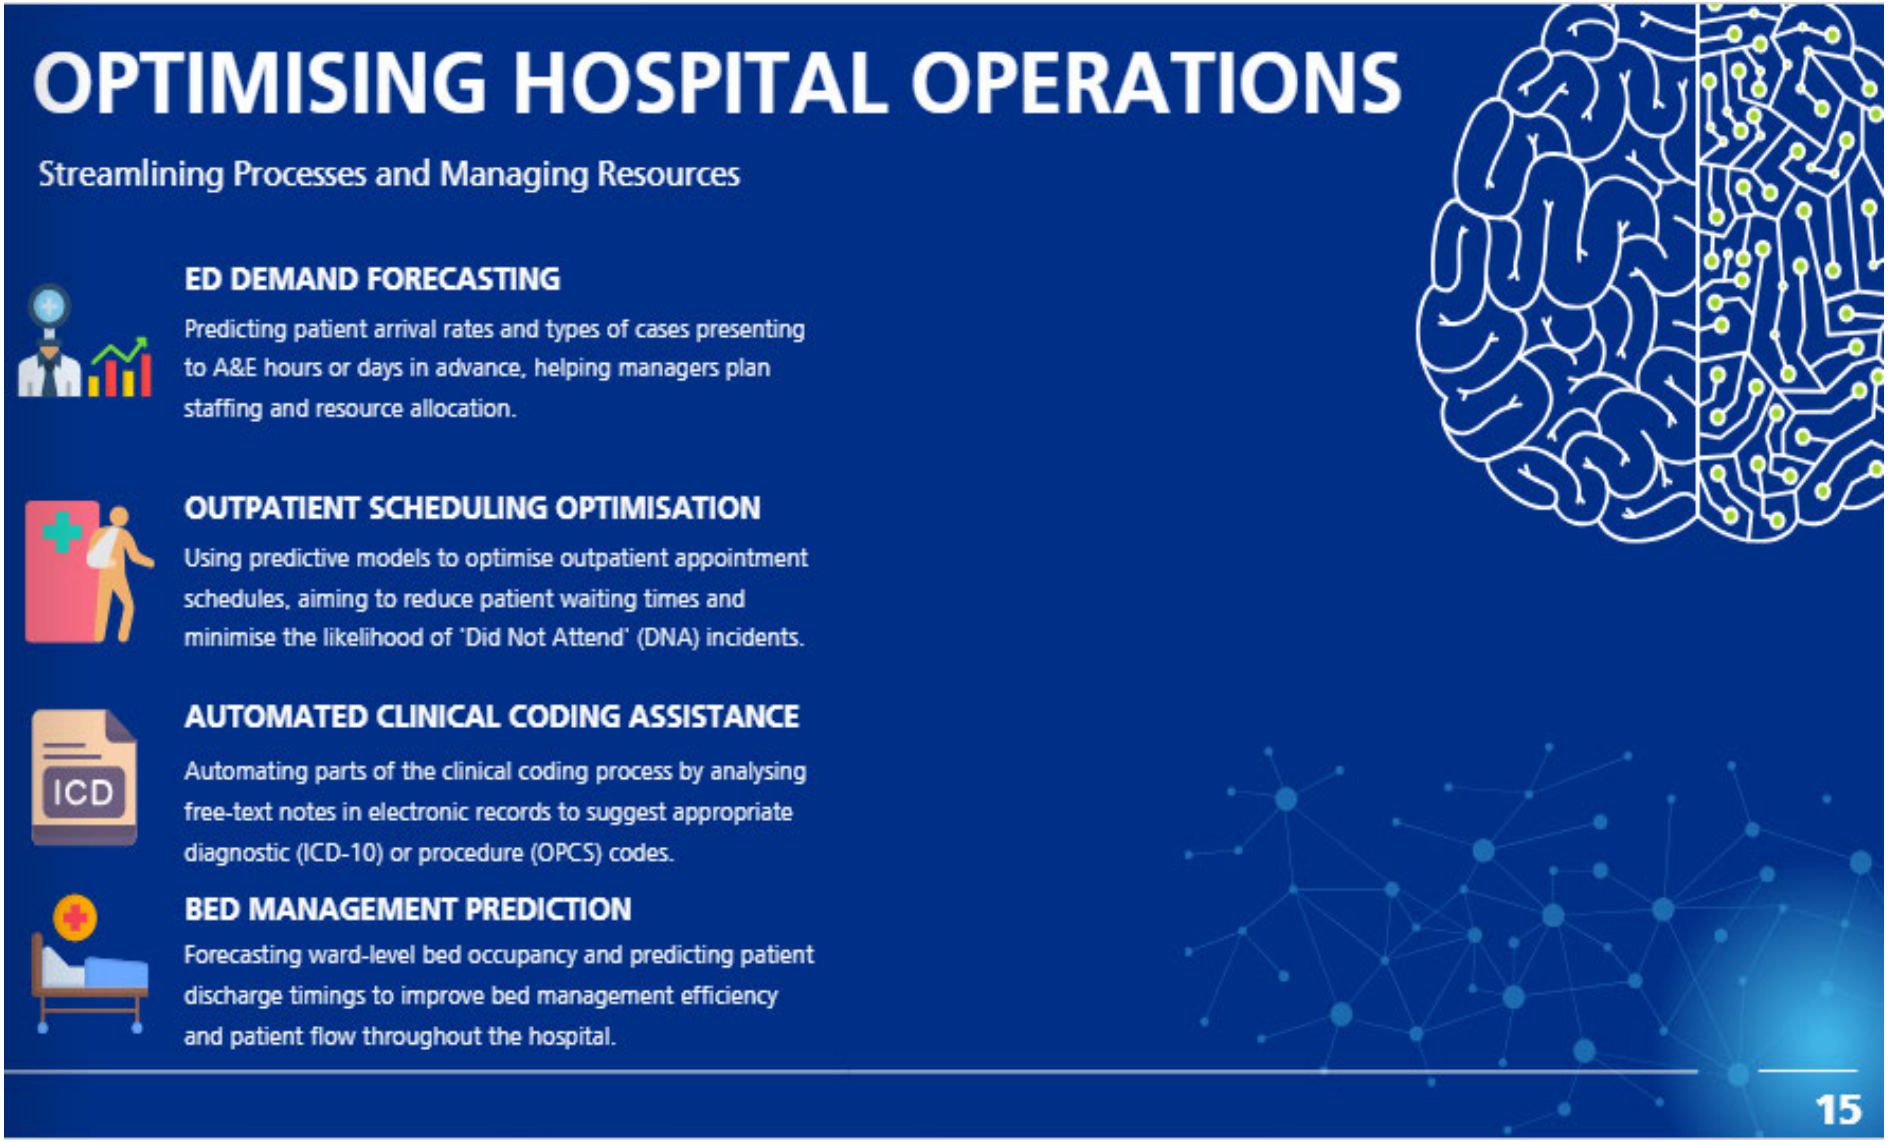

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

THE FUTURE OF MACHINE LEARNING IN HEALTHCARE

Expect machine learning to become increasingly integral to clinical practice. It will often work seamlessly within systems to enhance diagnostic capabilities, personalise treatments, and optimise healthcare operations. We are witnessing rapid advancements, particularly in analysing complex medical images with greater precision and interpreting clinical language in patient records. There is also emerging potential in areas like drug discovery. Crucially, alongside this innovation, there is a fundamental commitment to developing and deploying this technology responsibly. This involves focusing intensely on several key areas: robust validation for patient safety; proactive mitigation of bias to ensure fairness; and transparent governance to build and maintain trust among clinicians and patients.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 **It Learns from Clinical Data**

Unlike traditional programming, machine learning algorithms learn directly from vast amounts of relevant clinical data. This allows them to identify complex patterns and insights without following explicit instructions for every situation.

02 **It Finds Clinical Patterns & Predicts Outcomes**

At its core, machine learning excels at automatically discovering subtle, often hidden, relationships within patient data. It then uses this learned knowledge to make clinically relevant predictions about risk, disease progression, or potential outcomes.

03 **High-Quality Clinical Data is Foundational**

The safety, reliability, and fairness of any clinical machine learning system is entirely dependent on the quality, completeness, and representativeness of the patient data used for its training. Poor data inevitably leads to poor and potentially unsafe performance.

04 **ML Tools are Increasingly Supporting Care**

While development continues rapidly, validated machine learning applications are becoming integrated into clinical workflows, designed specifically to support clinicians in areas like diagnostics, risk identification, and operational efficiency.

05 **ML Augments; Clinicians Interpret & Decide**

Machine learning provides powerful analytical support, but it is a tool, not a replacement for expertise. Critical thinking, contextual understanding, and final patient care decisions are always the responsibility of the healthcare professional.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

AI Information Guide 8

Natural Language Processing – A Clinical Overview

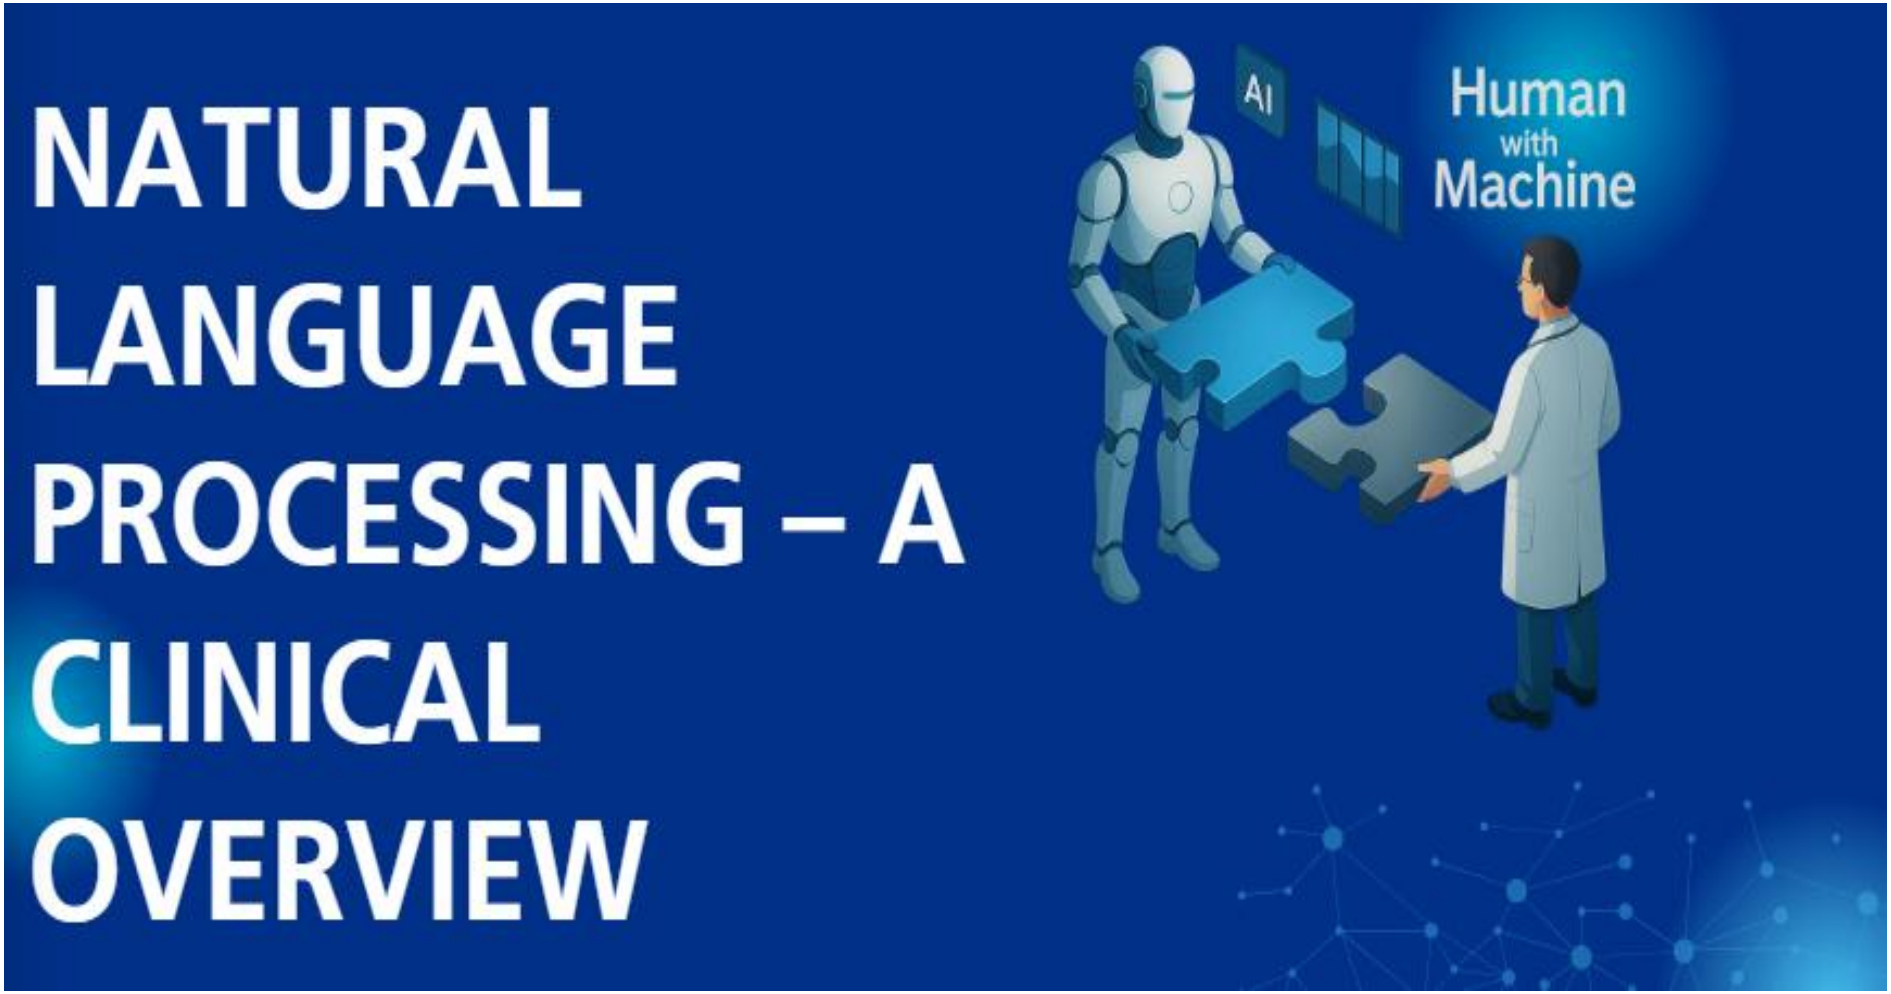

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHY ARE WE ALL TALKING ABOUT NATURAL LANGUAGE PROCESSING IN HEALTHCARE?

Natural language processing is a significant area within artificial intelligence that is becoming increasingly important for medicine. It provides computers with a remarkable ability to understand the complexities of medical language and to process information in ways that can support clinicians and patients. This involves more than simply recognising individual words; it means these systems are learning to grasp clinical context and the subtle differences in meaning found in medical notes and research.

- **An Essential Technology for Modern Medicine:** It is a key tool for making sense of the vast amounts of text-based information in healthcare.
- **Powerful Capabilities for Healthcare AI:** This technology represents a fundamental way for artificial intelligence to work with, interpret, and respond to the language of health.
- **Shaping the Future of Healthcare:** Natural language processing is already starting to influence how medical information is managed and how care is delivered.
- **Enabling Advanced Clinical Tools:** It is a key reason that computer systems can now engage with more complex medical information, from understanding patient records more deeply to helping summarise the latest research.

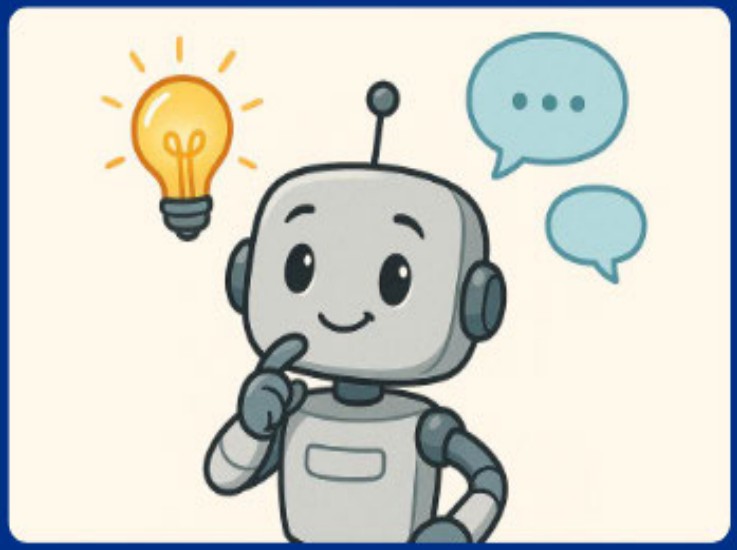

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT IS NATURAL LANGUAGE PROCESSING IN HEALTHCARE?

Natural language processing is an important area within artificial intelligence. In healthcare, its main focus is to enable computers to understand, interpret, and use the complex language of medicine, both spoken and written.

You can think of it as the process of teaching computers to become skilled with clinical words, sentences, and the way medical ideas are expressed. By learning from large amounts of text, such as clinical notes and research papers, these systems identify and understand patterns in how medical information is communicated.

This learning allows computers to make better sense of medical language and, as a result, to provide more useful support, for example by generating clear summaries or helping to answer clinical questions.

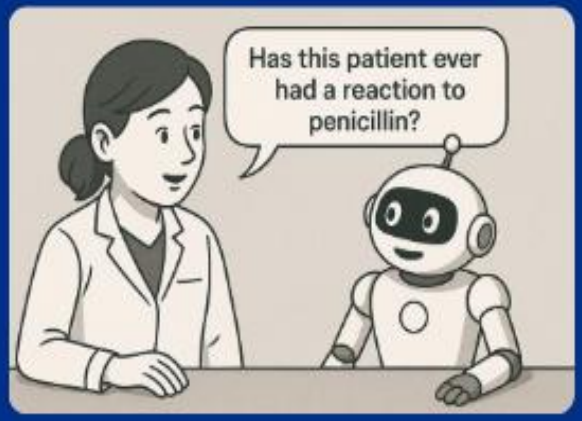

A clinical task often begins with a question asked in everyday language. For a system using natural language processing to help, the first step is to receive this question, either as spoken or typed words.

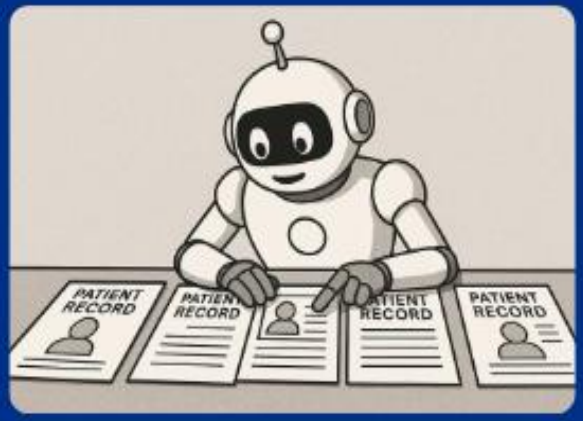

The system then analyses this question. It identifies key medical words, understands the grammar, and interprets the clinical meaning. This analytical process transforms the clinician's question into a structured search for specific information.

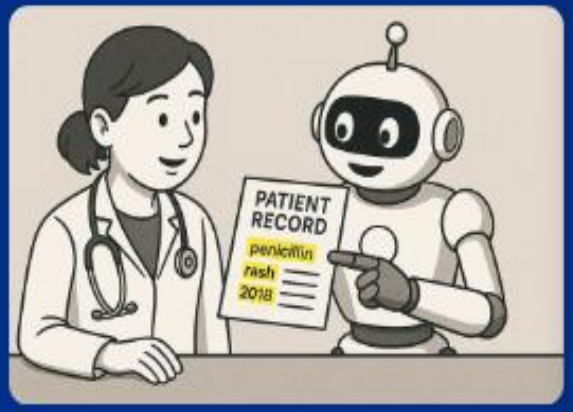

Once the system has a clear understanding of the original question, it can generate an appropriate and helpful response. In this instance, understanding the question leads to providing the specific, critical safety information that the clinician asked for.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

NATURAL LANGUAGE PROCESSING IN THE AI FAMILY

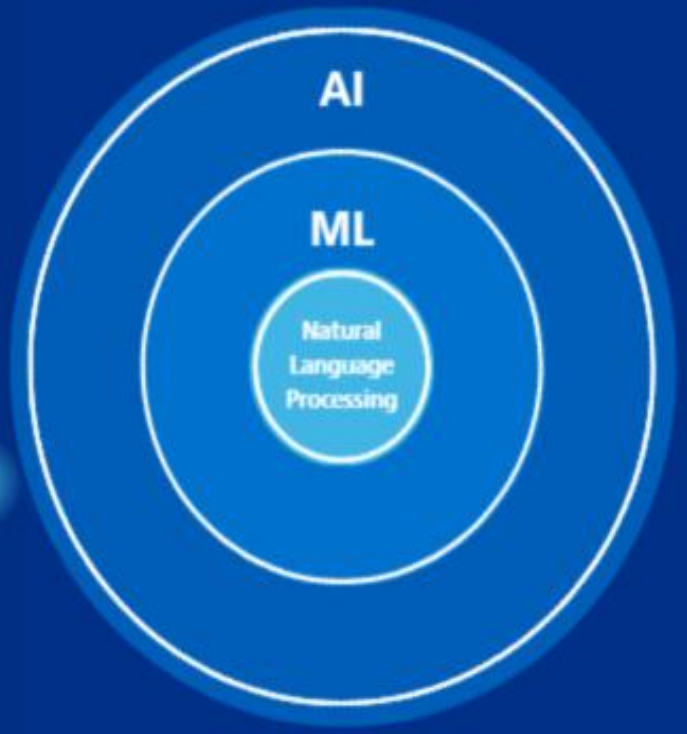

Artificial Intelligence (AI): AI is a wide ranging field focused on creating computer systems that can perform tasks normally requiring human intelligence, such as learning, reasoning, problem-solving, and decision-making. It is the overall goal of making machines intelligent.

Machine Learning (ML): ML is a core method used to achieve AI. Instead of being explicitly programmed for every task, ML systems learn patterns and make predictions or decisions by analysing large amounts of data. It is how many AI systems acquire their intelligence.

Natural Language Processing (NLP): This is a specialised field of artificial intelligence focused on enabling computers to understand and use human language. Natural language processing often uses machine learning to help computers make sense of our words. Advanced tools, such as large language models, are a key part of modern natural language processing.

05

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

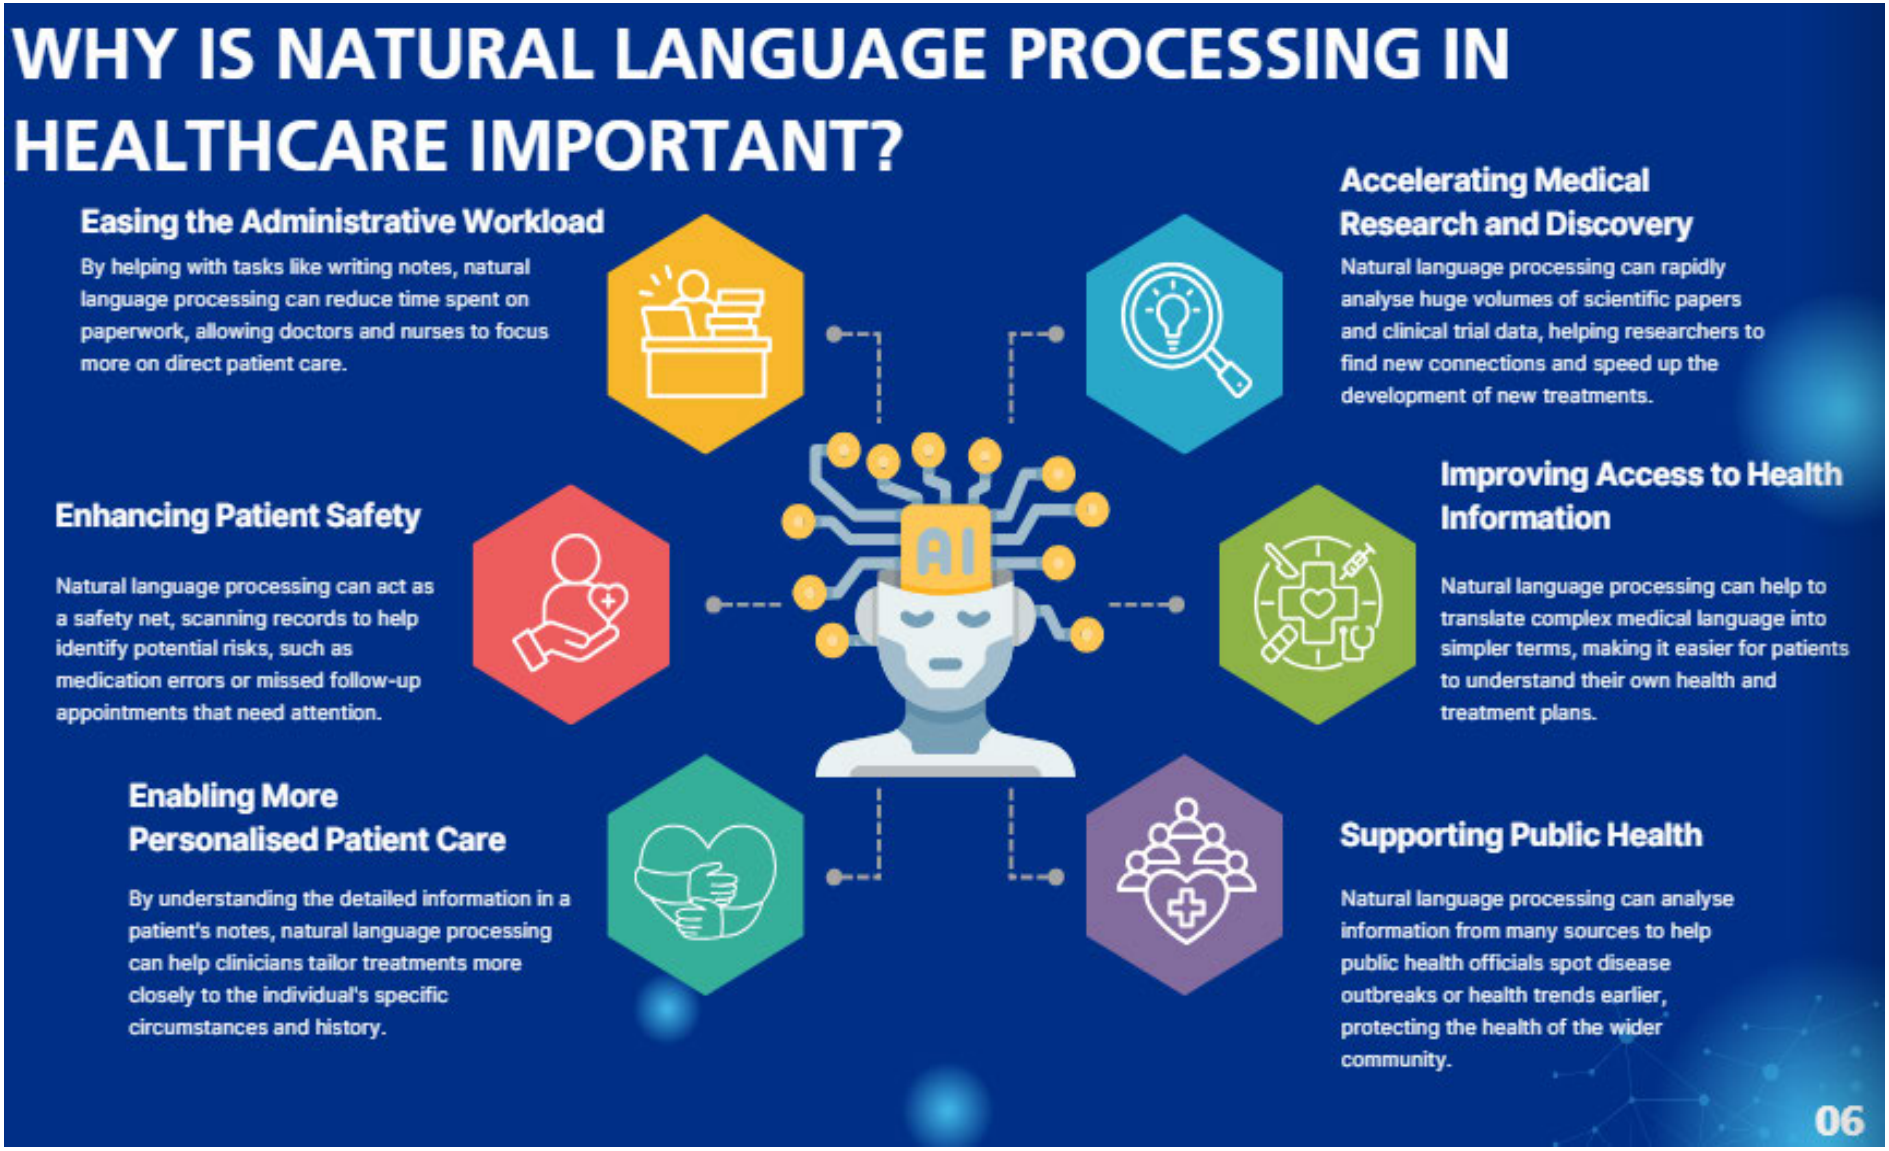

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

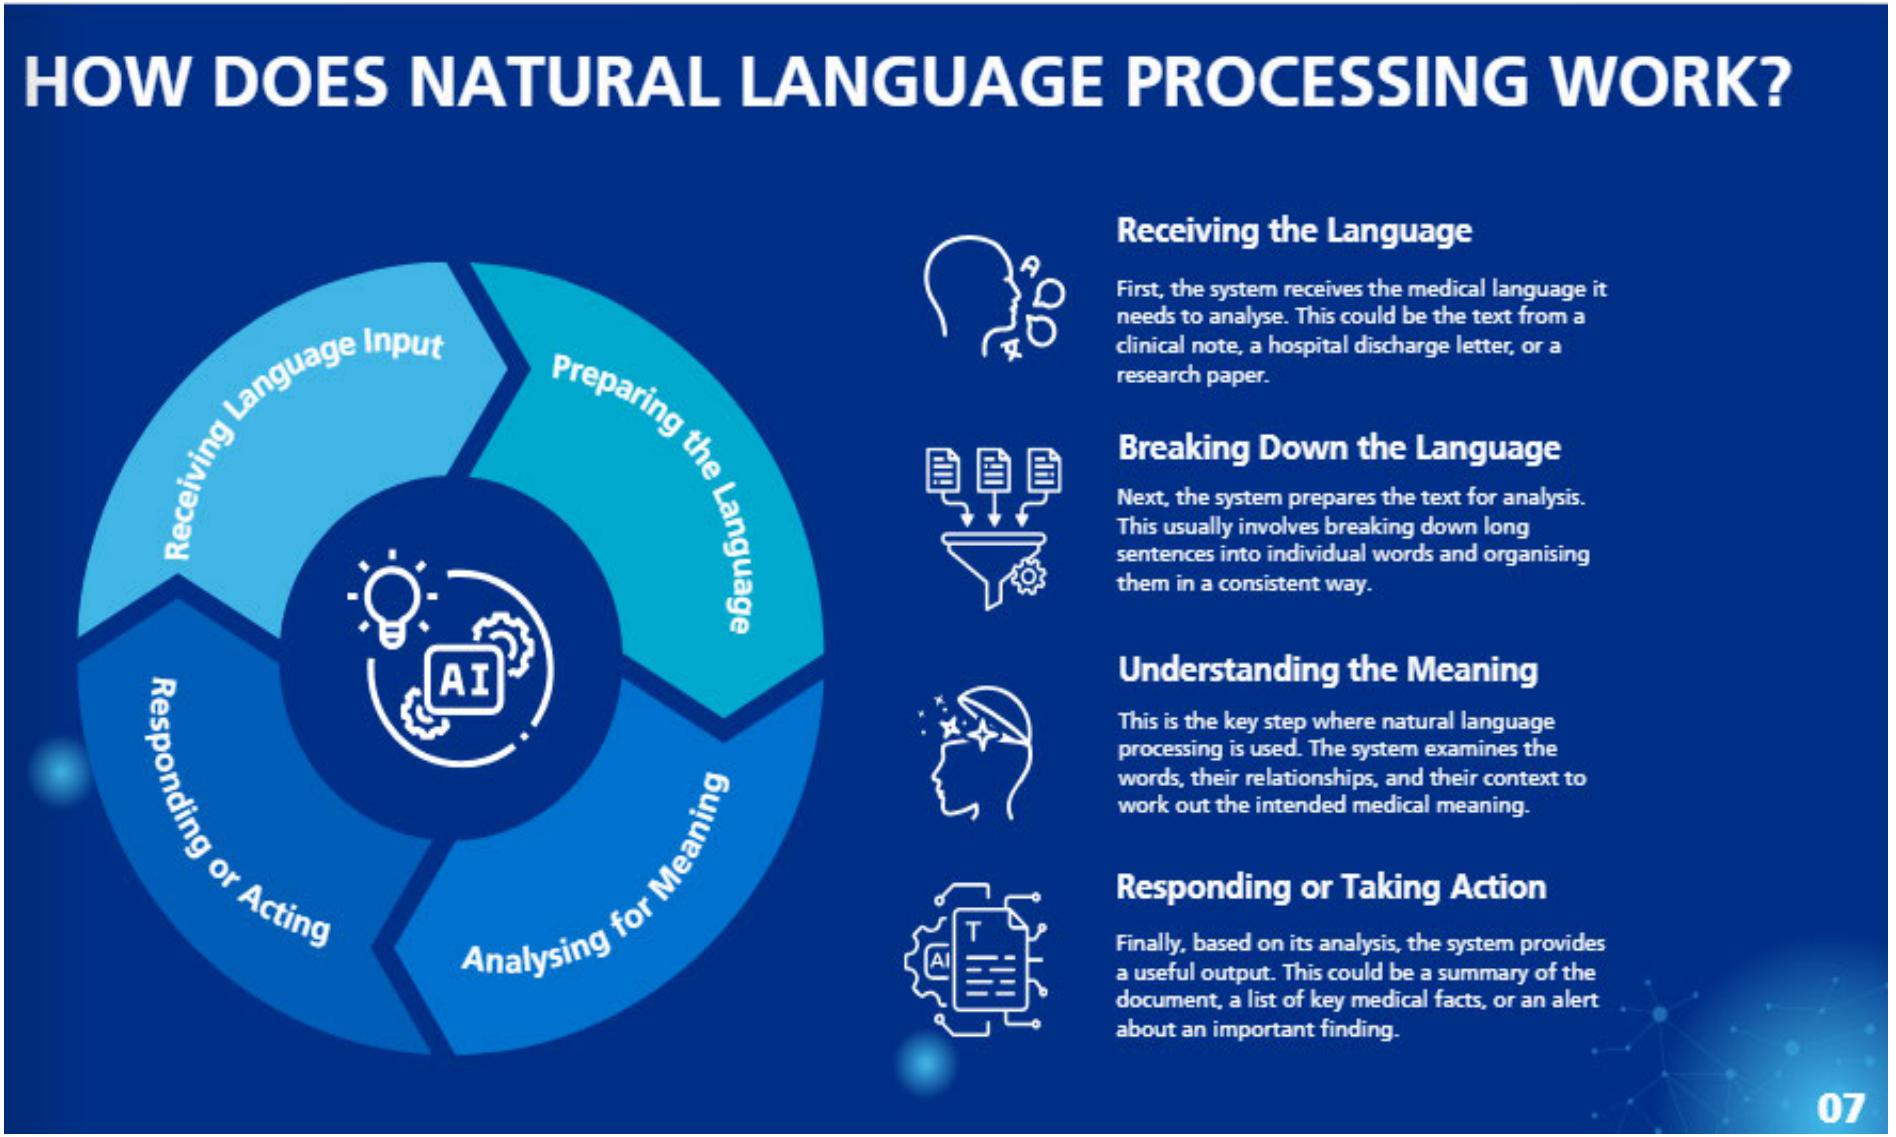

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES NATURAL LANGUAGE PROCESSING IN HEALTHCARE WORK?

Tokenisation - Breaking Language into Pieces

For natural language processing to begin understanding human language, it first needs the text to be broken down into manageable units. Tokenisation is like carefully taking that long stream of letters and neatly separating it into individual words and punctuation marks. This essential first step transforms a continuous flow of text into distinct pieces, or 'tokens'. Once the language is divided into these clear, individual tokens, natural language processing can then begin to analyse each part more effectively, which is fundamental for any further understanding.

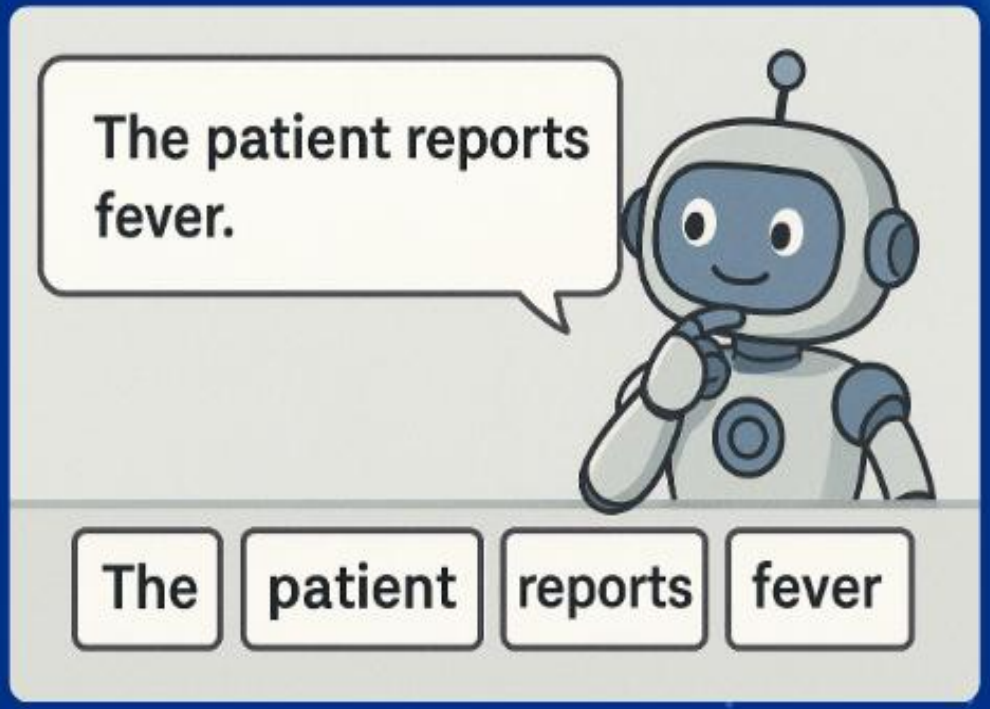

First, tokenisation breaks the sentence into individual words or 'tokens'. Our example, 'Patient reports fever.', becomes separate pieces.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES NATURAL LANGUAGE PROCESSING IN HEALTHCARE WORK?

Part-of-Speech Tagging - Identifying Word Types

After breaking text into 'tokens', natural language processing must understand each word's function. Part-of-speech tagging achieves this by assigning a grammatical category to every token, such as identifying nouns (things), verbs (actions), or adjectives (describers). This is like knowing the different roles of players in a team; understanding these word types helps the system to grasp sentence structure and begin to work out the meaning. This step is vital for deeper language comprehension.

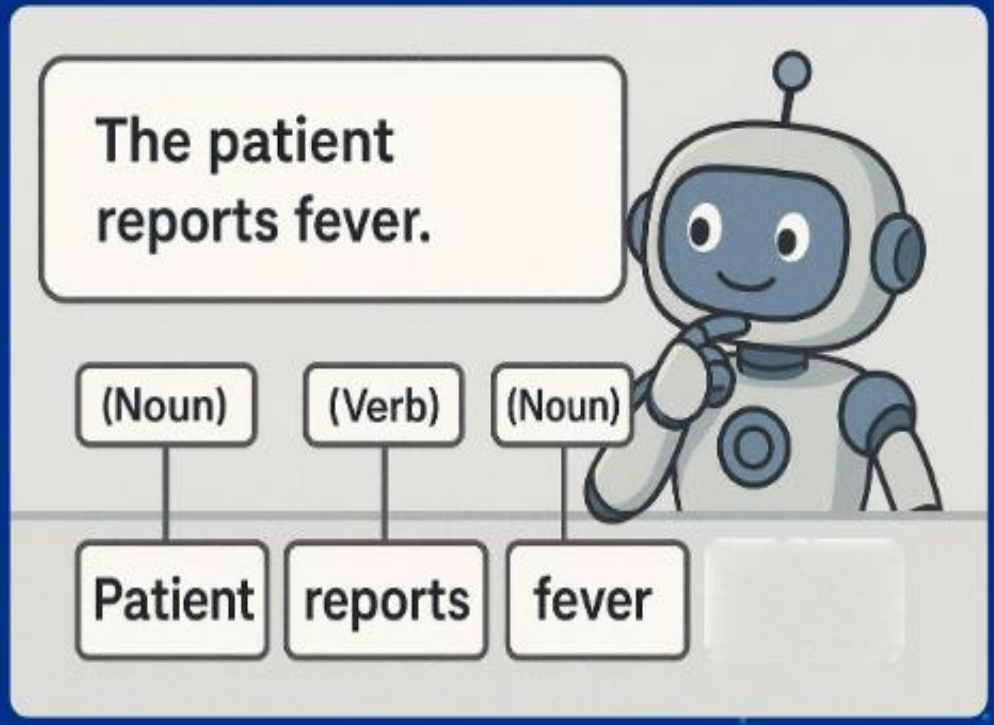

Then, part-of-speech tagging identifies the grammatical type of each token. Our example, 'Patient reports fever.', now has each word labelled by its role.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT MAKES NATURAL LANGUAGE PROCESSING IN HEALTHCARE WORK?

Understanding Meaning & Relationships

With words broken into 'tokens' and their grammatical types identified, natural language processing then works to understand how these pieces fit together to create meaning. This stage involves looking at the relationships between words in the sentence. For instance, it determines which adjective describes which noun, or which noun is performing the action of the verb. It is similar to understanding how individual musical notes (the words and their types) combine to form a melody (the meaning of the sentence). By analysing these connections, the system starts to build a picture of what the sentence is actually trying to say, moving beyond individual words to a more complete interpretation.

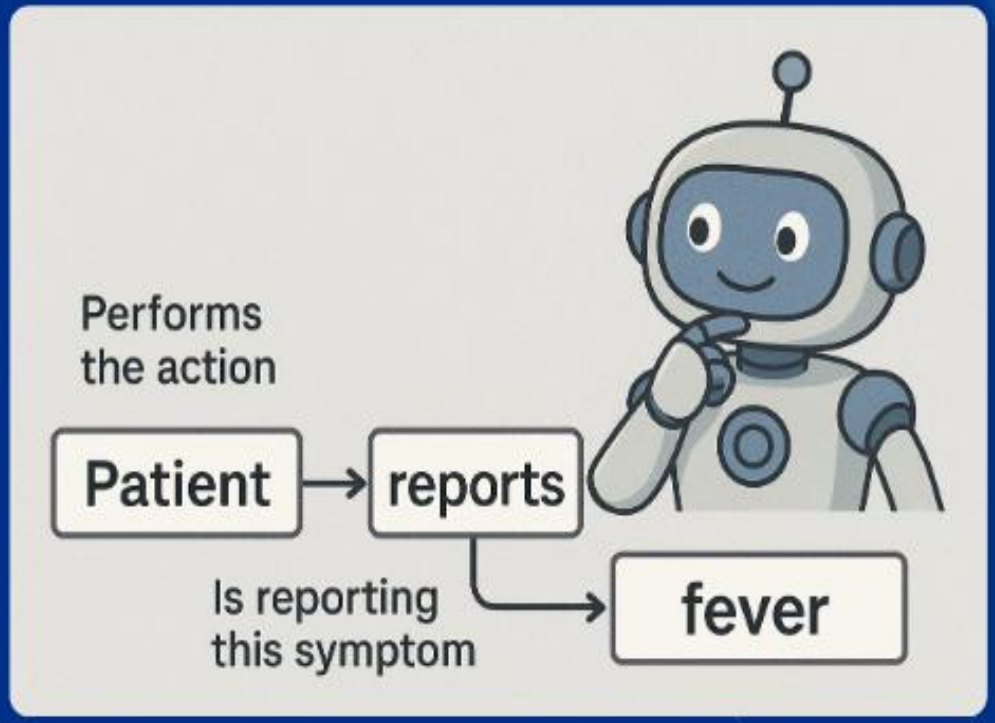

Finally, the system analyses how the labelled words connect. It understands that the 'Patient' performs the action of 'reports', and that 'fever' is the symptom being reported.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

WHAT CAN NLP ACTUALLY DO IN HEALTHCARE?

01 Medical Coding

Natural language processing can read clinical notes and help to assign the correct administrative codes for diagnoses and procedures. This is a vital task for hospital records and billing. Automating this process helps to improve both the accuracy and efficiency of this important work.

02 Clinical Trial Matching

Finding the right patients for medical research studies can be a slow process. NLP can quickly review thousands of anonymised patient records to identify individuals who may be suitable for a particular clinical trial, helping to speed up the pace of important research.

03 Monitoring Medicine Safety

After a new medicine is approved, it is important to continue monitoring its safety. NLP can analyse clinical notes to help detect potential side effects or adverse reactions that might not have been formally reported, providing a fuller picture of a medicine's real-world impact.

04 Automatic Summaries

A patient's medical record can be hundreds of pages long. NLP can read through a long and complex record and create a short, easy-to-read summary of the most important information, which can be very helpful for clinicians who need to understand a patient's history quickly.

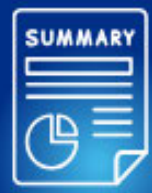

05 Medical Chatbots

NLP is the technology that powers helpful computer programs, or chatbots. In healthcare, these can be used to answer common patient questions about hospital services or to help with tasks like booking appointments, improving access to information for patients.

06 Analysing Medical Reports

Reports from tests like X-rays and biopsies are often written as text documents. NLP can read these reports and automatically pull out the key clinical findings, helping to ensure that important results are highlighted for the doctor's review.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CLINICIANS AND NLP WORKING TOGETHER

An Assistant, Not a Replacement

Natural language processing should be seen as an advanced assistant for clinicians. It can handle routine or time-consuming language tasks, but it does not replace the knowledge, skill, and judgement of a healthcare professional.

Clinician Oversight is Essential *(Human-in-the-Loop)*

It is vital that a clinician always reviews and checks any information or suggestion provided by an NLP system. The healthcare professional makes the final decisions about patient care.

Guiding the System for Best Results

Clinicians play an active role by giving the system clear instructions. Learning how to ask the right questions helps ensure the technology provides the most useful and relevant support.

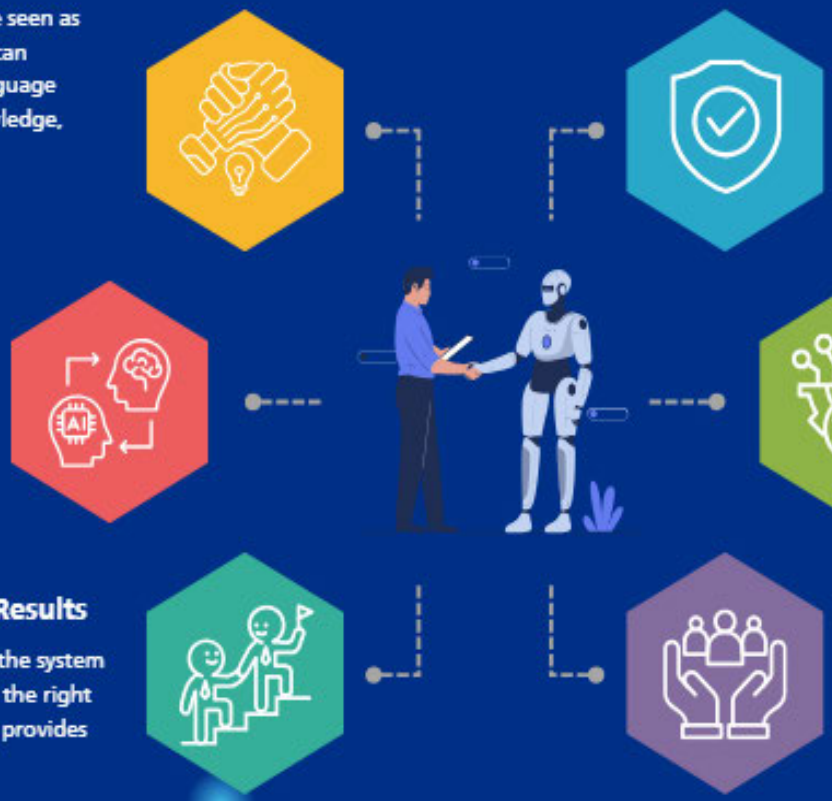

Training for Safe and Effective Use

For this partnership to be successful, healthcare staff need training. This helps them understand how the tools work, what their benefits are, and, just as importantly, what their limitations are.

Refining the Technology with Clinical Feedback

The partnership works both ways. When clinicians provide feedback on the performance of these tools, it helps developers to make them more accurate, safe, and useful in real clinical situations.

Focusing on Patient-Centred Care

This partnership aims to reduce the time clinicians spend on administrative work. This frees them up to focus on their most important responsibilities: direct interaction with patients and making complex care decisions.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT CONSIDERATIONS FOR USING NATURAL LANGUAGE PROCESSING IN HEALTHCARE

01 Patient Safety and Accuracy

In healthcare, accuracy is vital. Natural language processing tools must be reliable, and their outputs must always be carefully checked by a clinician to ensure patient safety.

02 Data Privacy & Cyber Security

Protecting sensitive patient health information is essential. When natural language processing uses patient data, strict rules like the UK GDPR and Caldicott Principles must be followed to ensure data is secure and used appropriately. Robust cybersecurity defences, including encryption and secure networks, are vital to prevent breaches or attacks that could compromise data.

03 Fairness and Avoiding Bias

Natural language processing learns from data. If this data contains biases, e.g., related to different patient groups, the language model may produce unfair or unequal outputs. Efforts are needed to ensure these tools benefit all patients fairly.

04 Transparency and Explainability

It is important to understand how these systems arrive at their conclusions. This "explainability" helps to build trust and allows clinicians to assess the reasoning behind a suggestion.

05 Accountability and Responsibility

Clear lines of responsibility must be established. While the technology can assist, healthcare professionals remain accountable for all clinical decisions and patient outcomes.

06 The Complexity of Medical Language

The specialised language of medicine, including jargon and abbreviations, is a constant challenge. Natural language processing systems must be able to handle this complexity correctly.

07 Regulatory Oversight

Natural language processing models used in clinical settings, especially those influencing diagnosis or treatment, must meet the stringent safety, efficacy, and quality standards set by official medical regulatory bodies like the MHRA to ensure they are fit for purpose.

08 Integration with Clinical Workflows

For natural language processing tools to be helpful, they must fit smoothly into the day-to-day work of clinicians and be easy to use with existing hospital computer systems.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

REAL-WORLD APPLICATIONS OF NATURAL LANGUAGE PROCESSING IN HEALTHCARE

Natural language processing is already being used in a variety of practical ways within healthcare, often working behind the scenes to improve how clinical information is managed and used. This field of artificial intelligence helps computers to understand and respond to medical language, leading to safer clinical systems, better access to health knowledge, and innovative solutions for patient care and medical research.

The real-world applications of this technology in medicine are continuously growing. They are helping healthcare organisations and professionals to manage information more effectively, support clinical decisions, and shape a future where technology can better understand and assist with the complexities of health and care.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

NLP FOR CLINICAL DOCUMENTATION

Easing the administrative burden and giving clinicians more time for patients.

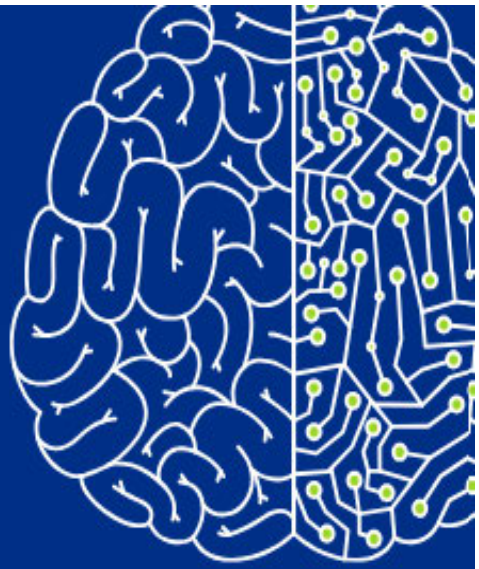

AUTOMATED CLINICAL NOTE DRAFTING

NLP can help to reduce the time spent on paperwork. For example, systems can listen to a doctor-patient conversation (with consent) and automatically draft a clinical note for the clinician to review and approve.

MEDICAL SPEECH-TO-TEXT

Advanced speech recognition tools, powered by NLP, allow clinicians to dictate their notes accurately in real time. This is often much faster and more convenient than typing, helping to streamline the documentation process.

GENERATING LETTERS AND SUMMARIES

Natural language processing can also assist by creating initial drafts of documents like referral letters or hospital discharge summaries, based on the information already in the patient's record.

ASSISTING WITH MEDICAL CODING

Natural language processing can analyse the text in a clinical note to suggest the correct administrative codes for diagnoses and treatments. This helps to improve the accuracy of essential record-keeping and speeds up the overall process.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

NLP FOR CLINICAL DECISION SUPPORT

Providing information to enhance patient safety and support care decisions.

IDENTIFYING PATIENT SAFETY RISKS

Natural language processing can scan patient records to find important information that might be hidden in the text, such as a medicine allergy mentioned in an old note, and can then alert the clinician to the potential risk.

SUPPORTING DIAGNOSTIC THINKING

By analysing a patient's symptoms and history as described in clinical notes, natural language processing can suggest a list of possible diagnoses for the clinician to consider, which is especially helpful in complex cases.

FLAGGING POTENTIAL DRUG INTERACTIONS

By analysing a patient's complete list of medicines, natural language processing can help to flag potentially harmful interactions between different drugs, which supports safer prescribing practices for clinicians.

INTERPRETING TEST RESULTS IN CONTEXT

Natural language processing can help to interpret laboratory results by considering them alongside the full context of the patient's record, including their symptoms and history, providing a more complete picture for the clinician.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

NLP FOR PATIENT COMMUNICATION AND ENGAGEMENT

Making health information more accessible and amplifying the patient voice.

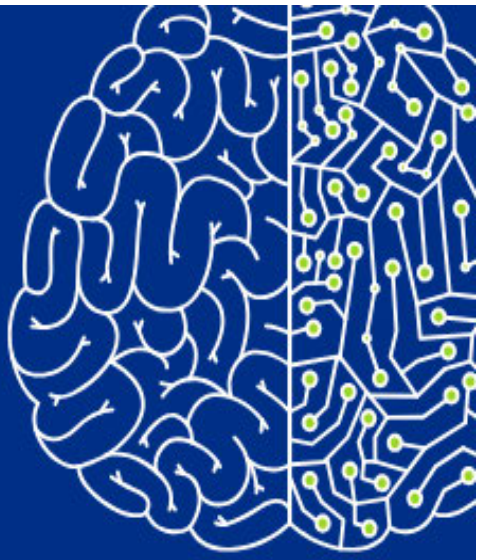

POWERING HEALTHCARE CHATBOTS

Natural language processing enables helpful chatbots that can answer common patient questions about appointments or hospital services, providing 24/7 access to routine information.

ANALYSING PATIENT FEEDBACK

The technology can automatically analyse thousands of patient survey comments or online reviews to identify common themes and concerns, helping hospitals to understand patient experiences and improve their services.

SUPPORTING MEDICATION REMINDERS

Natural language processing can power systems that send patients personalised reminders to take their medicines. The system can also be designed to understand simple patient replies, helping to track adherence or flag when a patient reports a problem.

CREATING CLEAR EXPLANATIONS FOR PATIENTS

Natural language processing can help to translate complex medical information from clinical reports into simpler, patient-friendly language, helping people to better understand their own health conditions and treatment plans.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

NLP FOR MEDICAL RESEARCH AND DISCOVERY

Speeding up scientific progress and the development of new treatments.

ANALYSING AND SUMMARISING RESEARCH

Natural language processing can quickly process huge volumes of medical research papers, helping scientists to stay updated with the latest findings and identify key information much more efficiently.

ASSISTING WITH DRUG DISCOVERY

By identifying patterns and connections in existing medical literature, natural language processsing can help researchers to form new ideas for investigation, potentially accelerating the discovery of new medicines.

SUPPORTING PUBLIC HEALTH SURVEILLANCE

By analysing large volumes of clinical notes or public health reports, natural language processing can help to identify emerging disease outbreaks or health trends, giving authorities an earlier warning to help protect community health.

IDENTIFYING PATIENTS FOR CLINICAL TRIALS

Natural language processing can help researchers by quickly screening anonymised patient records to find individuals who may be suitable for participation in clinical trials, helping to make recruitment faster and more effective.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

CHALLENGES AND IMPORTANT QUESTIONS

01 **Patient Privacy**

How do we ensure sensitive patient health information is always kept completely private and secure when using these powerful tools?

02 **Accuracy and Reliability**

National language processing systems are not perfect and can make mistakes. How do we ensure all outputs are carefully checked by a clinician to guarantee patient safety?

03 **Fairness and Bias**

If data used for training contains biases, the system can learn and repeat them. How do we prevent this from creating unfair disadvantages or inequalities in healthcare?

04 **The Complexity of Medical Language**

Medical language is full of jargon and ambiguity. A major challenge is building systems that can understand this complexity with a high degree of reliability.

05 **Accountability When Things Go Wrong**

If a tool contributes to a mistake, who is responsible? Defining clear lines of accountability for developers, hospitals, and clinicians is a critical issue.

06 **Understanding How Decisions Are Made**

It can be difficult to know exactly how an advanced natural language processing system arrived at a suggestion. How can we make their reasoning more transparent to build trust with clinicians?

07 **Regulatory Standards**

How should these advanced clinical tools be regulated to ensure they meet the high standards for safety and effectiveness required in medicine?

08 **Integration with Clinical Workflows**

How can we ensure that new natural language processing tools are designed to fit smoothly into the day-to-day work of busy clinicians and are easy to use?

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

THE FUTURE OF NATURAL LANGUAGE PROCESSING IN HEALTHCARE

Looking ahead, we can expect natural language processing to become even more seamlessly integrated into our clinical tools and digital health services. It will continue to refine how we manage medical information and communicate, making technology feel more like a helpful partner in delivering care. Ongoing advancements are consistently enabling these systems to understand more complex clinical requests and generate richer text. The long-term journey of teaching computers to truly understand the depth and nuance of medical language is very much still in progress. As these technologies become more powerful, there is a strong and growing focus on guiding their development responsibly, ensuring that progress aligns with human values and the core principles of ethical medical practice.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth

IMPORTANT THINGS TO REMEMBER

01 **Unlocks Written Medical Information**

It is the key tool for making sense of the vast amount of knowledge "locked away" in clinical notes and research papers.

02 **A Tool to Support, Not Replace, Clinicians**

This technology is best seen as a helpful assistant that can handle time-consuming tasks, giving clinicians more time to focus on direct patient care.

03 **The Clinician is Always in Charge**

A qualified healthcare professional must always use their expert judgement to review any information from a language processing system and make the final clinical decision.

04 **Responsible Use is Essential**

Patient safety, data privacy, and fairness must be the highest priorities when developing and using these tools in a healthcare setting.

05 **Partnership Between People and Technology**

For these tools to be successful and safe, they require ongoing collaboration between technology developers and clinical staff.

06 **The Goal is Always to Improve Patient Outcomes**

Ultimately, this technology is used to enhance patient safety, accelerate medical research, and make healthcare more effective and personalised for everyone.

Credit – University Hospitals Plymouth NHS Trust, Derriford Hospital Plymouth