Tutorial

DIGITAL HEALTH

DIGITAL HEALTH
Volume 12: 1–12

DOI: 10.1177/20552076261437226
journals.sagepub.com/home/dhj
Sage

A clinician's quick-start guide to implementing digital health innovations in the NHS – with lessons from a UK-deployed AI stroke imaging decision-support software

Anurup Mukherjee[1,2], Sukhi Shergill[3,4], and Chee Siang Ang[5,6]

Abstract

The successful implementation of digital health and artificial intelligence (AI) innovations in the National Health Service (NHS) requires more than technical development. Navigating regulation, generating decision-grade evidence, and meeting clinical safety, information-governance, and interoperability standards are critical steps that frequently delay or prevent adoption. This article presents a practical, implementation-focused roadmap designed to help clinicians, innovators, and healthcare leaders translate policy requirements into real-world NHS deployment. Drawing on guidance from the Medicines and Healthcare products Regulatory Agency (MHRA), the National Institute for Health and Care Excellence (NICE), and NHS Digital, we outline an eight-step pathway covering medical-device classification, value-proposition development, intended-purpose definition, regulatory approval, evidence generation, algorithmic fairness and generalisability, interoperability and information governance, and post-market surveillance. Unlike high-level digital health frameworks, the roadmap specifies minimum artefacts, typical ownership and sign-off responsibilities, and decision points aligned with NHS procurement and clinical governance processes. The roadmap is illustrated through a detailed case study of a UK-deployed AI stroke imaging decision-support software. Its progression from academic development to multi-site NHS deployment demonstrates how early regulatory engagement, robust real-world evaluation, and sustained clinical collaboration can support safe scaling and measurable service improvements, including increased access to reperfusion therapies and reduced inter-hospital transfer times. By distilling complex regulatory and evidence requirements into executable steps, this guide offers a clear route from idea to adoption. It emphasises that aligning regulation, evidence generation, bias mitigation, and interoperability from the outset is essential to sustainable digital health integration within the NHS.

Keywords

digital health, artificial intelligence, MHRA, NHS innovation, interoperability and DTAC, health technology assessment, NICE evidence standards framework

Received: 27 August 2025; Revised: 10 March 2026; Accepted: 13 March 2026

[1]Clinical Academic Training Office, Kent and Medway Medical School, Canterbury, UK
[2]Department of Medicine, Maidstone and Tunbridge Wells NHS Trust, Maidstone, UK
[3]Department of Research, Kent and Medway Medical School, Canterbury, UK
[4]Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK
[5]School of Computing, Kent and Medway Medical School, Canterbury, UK
[6]National Heart and Lung Institute, Imperial College London, London, UK

Corresponding author:
Anurup Mukherjee, Clinical Academic Training Office, Kent and Medway Medical School, Parkwood Road, Canterbury CT2 7FS, UK.
Email: kmmsam2878@kent.ac.uk

1. Why read this guide?

Digital health and AI tools promise faster diagnosis, more personalised care and lighter clinical workloads, but the NHS rightly insists on evidence of safety, effectiveness and sound data governance before innovation reaches patients.[1–4] This concise guide keeps every critical regulation in mind but tells the story in plain language so you can move from idea to deployment without wading through pages of legal jargon.

2. Evidence first

Well-designed digital programmes consistently cut cost and improve care quality.[5] National reviews still spotlight two stubborn gaps: many staff lack digital confidence, and too few products land with robust real-world evidence.[6–8] If you plan training and evidence collection from the outset, the rest of the journey becomes easier.

3. An eight-step roadmap

These steps weave together MHRA, NICE and NHS Digital guidance with lessons from real NHS implementations of digital therapeutics, community-based diagnostics, AI-enabled rehabilitation support, and imaging decision-support software.[1]

The eight-step process is illustrated schematically in Figure 1. Unlike high-level digital health frameworks, this roadmap is designed to be executable: it specifies minimum artefacts, typical ownership/sign-off responsibilities, and decision points aligned to NHS procurement, deployment, and post-market monitoring.

Examples are illustrative and described by function; where outcomes are referenced, they reflect externally reported evaluations of access, diagnostic performance, pathway efficiency, or cost impact.

3.1. Step 1: Decide whether you have a medical device

If your software influences a clinical decision, UK Medical Device Regulations (UK MDR 2002) apply.[9] Classify it honestly (Class I–III); under-classifying only postpones pain.[10]

The regulation groups devices by risk, from Class I (low risk) to Class III (highest risk). Correct classification is critical: under-classifying to "keep things simple" may speed early development, but regulators will force a costly re-route later. Annex IX of Directive 93/42/EEC lays out the decision rules, while MHRA "borderline" guidance helps where categories overlap with medicines, cosmetics or wellness products.[10,11]

A NICE-assessed digital CBT-I platform for insomnia illustrates how correct Class I classification can streamline the regulatory pathway.[12] If borderline, document the rationale and seek MHRA borderline advice early to avoid reclassification late in procurement.

3.2. Step 2: Shape a value proposition everyone understands

Clinicians care about clinical benefit, commissioners about cost or time savings, and patients about outcomes. Frame your claims against the NICE Evidence Standards Framework (ESF) to satisfy all three audiences.[4] A community-based cardiovascular risk screening platform illustrates how clearly articulated clinical and service value can support adoption, including earlier identification of cardiovascular risk and potential reduction in reliance on traditional GP-led testing pathways.[13] The *NICE Real-World Evidence Framework* sets out how to do this.[14]

3.3. Step 3: Write an intended-purpose statement

Regulators expect a concise statement that specifies *what* the product does, *for whom* and *where*. Include only claims you can later substantiate; if a sentence cannot be evidenced, it is advisable to rewrite it. The signed statement belongs in your technical file and if an Approved Body is involved, it forms part of the submission pack. The MHRA now encourages publication of the text because transparency speeds conversations with Trust boards, commissioners and patient groups.[9] Treat intended purpose as a controlled document; significant changes to claims or workflow should trigger review of risk classification, clinical evaluation, and safety artefacts.

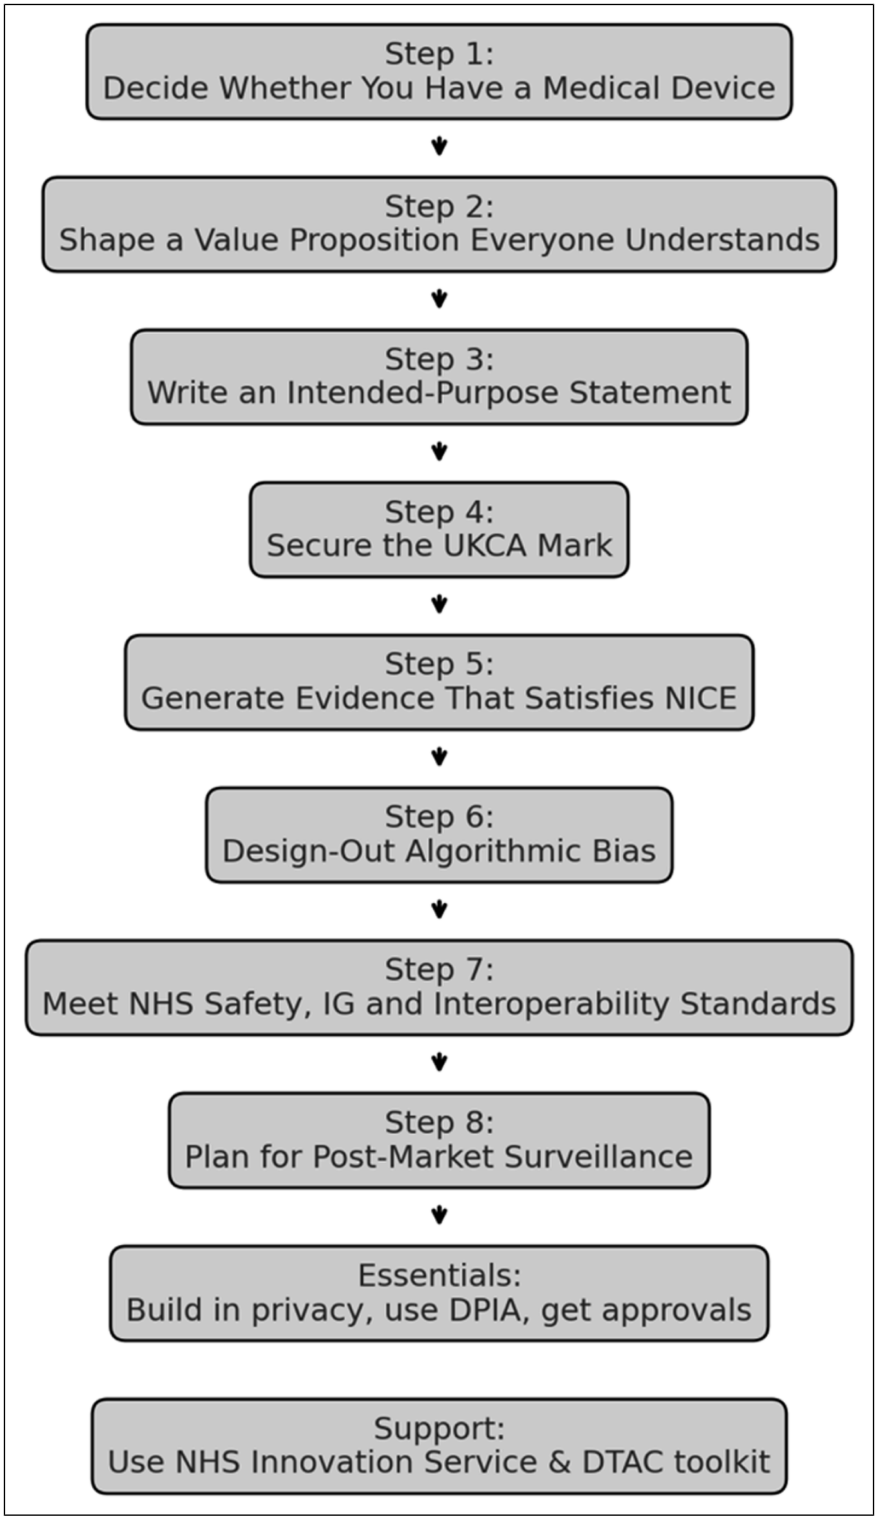

Figure 1. Flowchart to illustrate steps required for approval of digital health innovation.

3.4. Step 4: Secure the UKCA mark

Once you have confirmed the risk class, the hard regulatory work begins. Every device that reaches NHS patients must comply with **UK MDR 2002**.[9] In practical terms this means three deliverables:

1. **A complete technical file**—your evidence pack—covering risk analysis, usability testing and clinical validation.
2. **An ISO 13485 quality-management system** that shows you can keep the product safe and effective over its life-cycle.
3. **UKCA marking.** Class I devices can self-certify, whereas Classes IIa through III require an external audit by an Approved Body

3.4.1. One last checkpoint: Will you need a clinical investigation?

If your device cannot rely on published data or evidence from an equivalent UK-approved comparator, you must generate fresh clinical data through a prospective **clinical investigation**. That study needs twin approvals: *"no objection"* from the **MHRA** (normally within 60 days) and full **HRA/REC** ethical clearance (about 40 days). Prepare by notifying MHRA early, then submitting a single application via **IRAS**, and budgeting for the MHRA review fee. Where strong equivalence can be demonstrated—clinical, technical and biological—you may avoid a trial and save precious months; but for many AI-driven systems, differing training datasets make equivalence hard to prove, so an investigation is the safer route. Skipping this step can halt UKCA certification and ultimately block NHS deployment [MHRA Clinical-Investigation Guide, 2024].

Register the project on the **NHS Innovation Service** portal; it is now the fastest way to flag a new technology for NICE topic-selection and to be sign-posted to the right regulatory support.[15]

3.5. Step 5: Generate evidence that satisfies NICE

NICE advice hinges on two questions: **Does the product work?** and **Is it worth the money?** Randomised controlled trials remain the gold standard for demonstrating clinical efficacy, but they are not the only route. If your innovation affects workflows rather than direct patient outcomes, a well-designed service evaluation can be enough.[4]

Once your technical file is in hand, turn your attention to the **National Institute for Health and Care Excellence (NICE)**. A positive NICE recommendation not only reassures clinicians and commissioners but also accelerates national uptake. Building your evidence plan around NICE methods from day one avoids costly redesigns later.

3.5.1. How topics reach NICE

Digital, AI and other medical technologies flow into NICE through the HealthTech Programme. A prioritisation board short-lists topics that align with national health priorities; eligible products usually have a UKCA (or CE) mark, sit in Tier C of the Evidence Standards Framework (ESF), and already carry preliminary data to support their value proposition.

Depending on a product's maturity, the programme can run four main assessment routes as summarized in Table 1.

Registration on the NHS Innovation Service portal is the quickest way to check eligibility and flag a new topic for consideration.

3.5.2. Design studies that answer NICE's questions

Remember the distinction between regulatory and HTA evidence:

> **UKCA** focuses on *safety* and *efficacy*: does the product work under ideal conditions?
> **NICE** focuses on *relative effectiveness* and *value for money*: does it outperform current NHS practice and at what cost?

Table 1. Routes for NICE assessment of medical technologies.

NICE route	When it is used	Evidence depth
Early-Use Assessment	Promising, evidence-light tech	Initial clinical + economic signals
Routine-Use Guidance	Not yet widespread but with substantial data	Full relative-effectiveness + cost evidence
Late-Stage Assessment	Already in wide NHS use	Real-world effectiveness + budget impact
Interventional Procedures Guidance	Novel procedures or device–procedure combos	Safety, procedural success, outcomes

Practical implications:

Recruit populations that look like real NHS users.
Pick outcomes that matter to patients (symptom relief, function, survival), not just process metrics.
Follow-up long enough to capture downstream benefits or harms.
Pre-register protocols and use accepted risk-of-bias tools.

For digital tools with modest clinical risk but large workflow impact, well-designed service evaluations can suffice. NICE says so explicitly in its **Real-World Evidence Framework**.[14]

Economic evidence must match the clinical story:

Same or better outcomes at similar or lower cost? Provide a cost-consequence analysis.
Better outcomes at higher cost? Provide a full cost-utility model (£/QALY) with sensitivity tests, using the ESF economic tier as your checklist.

3.5.3. Use NICE advice to de-risk the journey

If you are unsure whether your draft protocol will satisfy the committee, book an appointment with **NICE Advice**. The service will review study plans, highlight patient-centred outcomes, and map a faster route to guidance. Teams developing similar tools credit early Advice sessions with tightening their trial design and leading to a favourable Medtech Innovation Briefing.

3.5.4. Pack the perfect submission

By submission day aim to hand NICE a dossier that:

1. Demonstrates **clinical effectiveness** versus current practice.
2. Confirms **patient safety**, including the quality of any AI-training data.
3. Quantifies **cost impact** under realistic NHS resource flows.

Digital CBT provides a model. The digital cognitive behavioural therapy (CBT-I) platform mentioned above, secured a positive Medtech Innovation Briefing by publishing peer-reviewed data showing equivalence to face-to-face therapy and an acceptable cost-effectiveness ratio.[12]

3.6. Step 6: Design-out algorithmic bias and plan for drift

Adopters will judge an AI tool's trustworthiness primarily by the quality, representativeness and generalisability of the data used for training, testing and validation. Data should reflect: (i) the target population, (ii) the intended clinical setting and workflow, and (iii) technical quality (e.g., imaging acquisition quality and device/software heterogeneity). Inadequate representation (including "marginal cases") risks unstable models and uncertain performance estimates in under-represented groups.

3.6.1. Minimum artefacts (attach to the technical file/evidence pack)

1. **Dataset documentation ("health dataset factsheet")** written in plain language that states the dataset's origin and purpose, sampling strategy, aggregation decisions (including any changes to demographic coding), missingness and measurement validity, and known/expected **data shifts over time** (population, practice, devices, software).
2. **Pre-specified contextualised groups of interest** relevant to the use case (not limited to one protected characteristic; allow intersectionality and local context).
3. **Subgroup performance report**: demonstrate and transparently report performance *for each group*, not only overall averages; treat any clinically meaningful disparity as a patient-safety risk that requires mitigation actions and governance sign-off.

3.6.2. Minimum metrics to report (tailor to task)

Report discrimination and error within each contextualised group (e.g., AUROC/AUPRC for classification; sensitivity/specificity; PPV/NPV; calibration slope/intercept; and clinically relevant error trade-offs). Where disparities exist, document the likely mechanism (sampling, measurement, aggregation, label bias), the mitigation applied, and residual risk accepted by governance.

4. **External validation plan** using datasets/systems not used during development (e.g., other Trusts/regions/vendors) to evidence generalisability to routine NHS practice.
5. **Drift surveillance plan**: specify what will be monitored over time (input quality, case-mix, outcome rates, performance metrics), thresholds for investigation, and actions (recalibration, retraining, rollback). Dataset documentation should explicitly describe anticipated data shifts over time.

Typical minimum artefacts and sign-off responsibilities are summarised in Table 2.

Recommended reporting standards: Consult the 'STANDING Together consensus recommendations' and use CONSORT-AI/DECIDE-AI/TRIPOD+AI where applicable to make bias and generalisability claims auditable by reviewers and adopters.[16–19]

3.7. Step 7: Meet NHS safety, IG and interoperability standards

Long before an NHS Trust plugs your app into its network, its clinical-safety lead will open two documents: **DCB0129** and **DCB0160** and ask whether you have done your homework.[3] These companion standards, issued by NHS Digital, are the bedrock of clinical-risk governance for health IT:

DCB0129 applies to **developers**.
DCB0160 applies to **adopters** (the Trust, ICS or practice that will deploy your product).

3.7.1. What DCB0129 actually demands

As the developer you must put in place a formal **clinical-risk-management system**. In practice that means:

appointing named clinical-safety officers and documenting their governance arrangements;
running structured hazard analyses throughout the build, not just at the end;
recording evidence that team members have the right safety training and competence.

Risk management does not stop at release. The standard expects you to plan for **maintenance, monitoring and—eventually—safe decommissioning**. Think of the safety case as a living file that grows alongside every software update.

NHS Digital provides templates that map directly to the standard; use them. If you cannot show a complete, traceable safety case, an adopter simply cannot switch you on—DCB0160 forbids it.

Table 2. Minimun artefacts and typical ownership/sign-off.

Roadmap step	Minimum artefact	Typical owner	Typical sign-off/assurance
Classification	Classification rationale + intended use	Supplier/manufacturer	MHRA (if borderline)
Intended purpose	Intended-purpose statement	Manufacturer	Approved Body (if applicable)
UKCA	Technical file + Clinical Evaluation Report + ISO 13485 QMS	Manufacturer	Approved Body
NICE evidence	Protocol + analysis plan + economic model approach	Clinical lead/sponsor	NICE (advice/guidance route)
Bias + Drift	Dataset documentation ("factsheet") + predefined groups of interest + subgroup performance report + external validation summary + drift/PMS monitoring plan	Manufacturer (ML lead + clinical safety officer) + adopter clinical safety officer for local monitoring	Trust clinical governance/CCIO + (where applicable) safety case linkage (DCB0129)
Safety	DCB0129 safety case (supplier) + DCB0160 (adopter)	CSO (supplier/trust)	Trust governance/CCIO
IG	DPIA + lawful basis documentation + data processing agreements	Trust (controller)	DPO/IG lead
Interop/DTAC	DTAC evidence pack (cyber, FHIR, accessibility)	Supplier + Trust IT	Procurement/digital assurance
PMS	Post-market surveillance plan + KPI dashboard	Manufacturer + adopter	MHRA/Trust clinical governance

3.7.2. The DTAC: Baseline passport to procurement

Clinical-risk governance is only one slice of the approval pie. During procurement most NHS organisations now apply the **Digital Technology Assessment Criteria (DTAC)**—a national "passport" that rolls clinical safety together with four other pillars: data protection, technical assurance, interoperability, and usability & accessibility.

Meeting DTAC is not optional; it is the new baseline. The form will ask for:

proof of **UKCA marking** (or CE if still valid) and your underpinning risk-management system;
a data-protection-impact assessment
evidence that you exchange data via open standards such as **FHIR**
accessibility testing against WCAG and use of the NHS Design System for consistent look-and-feel.

Because DTAC echoes legal obligations you already face—UK GDPR, medical-device law, cyber-security essentials—building its requirements into your design phase saves painful retro-fits later.

3.7.3. Medical devices add ISO 14971 to the mix

If the product is a medical device, the safety bar rises again. You must also comply with **ISO 14971:2019**, the international risk-management standard for medical devices, and show how that dovetails with DCB0129. In other words: one harmonised risk process, two sets of evidence outputs.

3.7.4. Best-practice accelerators

NHS Digital Service Manual. Re-use the NHS Design System, Personal Demographics Service and NHS Login components; they satisfy parts of the usability, interoperability and identity checks "for free."
Early adopter engagement. Involve at least one Trust-side clinical-safety officer while you draft the hazard log; you will surface workflow hazards far sooner.
Version discipline. Every release—no matter how small—needs a mini risk-assessment and an update to the safety file. Automate as much of that paperwork flow as you can.

3.8. Step 8: Plan for post-market surveillance

Deployment is the start of a new phase, not a finish line. The NHS expects ongoing surveillance: track adverse events, monitor key performance indicators and collect user feedback. Push software updates promptly when safety signals emerge and remember that high-risk services may need Care Quality Commission registration. A transparent improvement cycle sustains patient safety and clinician trust.[20]

4. Data & information-governance essentials

Apply "data protection by design" at every phase. Prototype with anonymous or synthetic data where possible; where personal data is essential, complete a DPIA and document Article 6 and 9 lawful bases.[21–23] Research studies or un-marked device investigations need MHRA "no-objection" and HRA/REC clearance via IRAS.

5. Getting help faster

The NHS Innovation Service portal offers one-door sign-posting to regulatory support and NICE topic selection and even links to MHRA advice on tricky borderline-device questions.[11,15] Re-using components from the NHS Digital Service Manual can instantly satisfy parts of DTAC.

6. In a nutshell

Start with classification, value proposition and intended purpose; build your technical file alongside your code; collect evidence that satisfies both UKCA and NICE; design for fairness, interoperability and solid IG up-front; and treat deployment as day one of monitoring and improvement. Follow this path and your innovation will reach clinicians and patients more quickly and stay there.

7. Case study – UK-deployed AI stroke imaging decision-support software

The software is an AI decision-support platform that automatically analyses CT/MRI brain scans to identify stroke patients who could benefit from reperfusion therapies such as thrombectomy. Deployed across multiple NHS trusts, it illustrates the roadmap in action.

Step 1 Classification – Recognised early as a **Class IIa medical device** (Tier C digital tool), triggering a full CE/UKCA route and MHRA engagement.
Step 2 Value proposition – *The software translated its value proposition into commissioner-relevant outcomes. An independent Oxford AHSN service evaluation, using a real-world observational implementation design across 25 NHS stroke sites, demonstrated significant reductions in pathway times, increased rates of mechanical thrombectomy (relative increase ∼33–55%), and favourable cost–benefit ratios, with health benefits expressed in QALYs and overall benefits exceeding costs under base-case and optimistic scenarios.*[24]
Step 3 Intended purpose – 'AI software that flags large-vessel occlusion on acute stroke imaging to support, not replace, clinician judgement.' Clear scope guided design, risk files and training.
Step 4 Regulatory approval - The software achieved a CE mark in 2017, transitioned to UKCA marking in 2024 and simultaneously obtained FDA clearance, illustrating how an early focus on global regulatory alignment can shorten commercial timelines.[25] DCB0129 safety case and ISO 13485 QMS underpin product updates.
Step 5 Evidence & NICE – In 2024, NICE issued an "evidence generation" recommendation for the software after Oxford AHSN evaluation (≥70 000 scans) reported a relative increase of ∼55% relative thrombectomy-rate increase and 50-min faster transfers.[24] NICE (2024) endorsed it under its 'evidence generation' route.[25]
Step 6 Bias mitigation – Multisite dataset (>70 k scans) and continuous performance dashboards detect drift; clinicians supply feedback loops to refine algorithms.
Step 7 Interoperability & DTAC – The software worked with NHS Digital to host its image-processing engine in a Health-and-Social-Care Network (HSCN)-connected cloud, achieving DTAC compliance and allowing real-time image sharing between referring hospitals and neuroscience centres.[24] DICOM-to-cloud PACS links and encrypted mobile alerts met DTAC cyber, IG and FHIR requirements; go-lives typically <4 weeks per site.
Step 8 Post-market monitoring – Real-world outcome and safety data flow back via Oxford AHSN; The software uses them for PMS and ongoing algorithm validation.
Outcome: Reported outcomes from participating hospitals included higher MT rates (e.g., from ∼3.6% to ∼7%) and shorter door-in–door-out times (approximately 1 hour in some sites), as described in external evaluations; these findings should be interpreted in the context of service evaluation designs and local pathway variation.[26]

7.1. Behind-the-scenes insights

The software spun out of the University of Oxford in 2012, but in its first two years the founders discovered that capital and talent were scarcer than algorithms. A 2014 £1.35 million seed round from the Oxford Innovation Fund topped up by Innovate UK grants, finally let the team grow beyond a handful of academics.[27]

Turning research code into a clinical product brought two early pitfalls. First, robust evidence takes longer than code: the company had to generate multiple peer-reviewed validation studies before stroke physicians were comfortable relying on AI decisions.[26] Second, regulations kept moving; The software secured a Class IIa CE-mark in 2017 but still had to re-tool documentation for the tougher EU MDR and prepare for UKCA marking, all while pursuing FDA clearances for other markets.

Rolling the software out to NHS trusts surfaced very human hurdles. Integration teams found that "plug-and-play" often meant weeks of firewall tweaks and PACS scripting, so the software now insists on an up-front responsibilities matrix with each local IT department. Clinician buy-in followed a familiar curve: initial scepticism gave way to advocacy once real cases showed faster, safer transfers for thrombectomy. Identifying a local clinical champion helped, but programmes learned not to rely on one enthusiast; shared ownership and regular cross-site user forums sustain momentum.[28]

With clinical, economic and regulatory boxes ticked, the focus has shifted to scaling: an NHS-led business-case template and the AI in Health & Care Award funding now help new trusts procure the software under standard tariffs.[29]

7.2. Key lessons

Secure milestone-linked funding early; AI validation is a multi-year marathon.
Publish evidence continuously; credibility trumps hype.

Treat regulatory engagement as iterative, not a one-off hurdle.
Agree an IT responsibilities matrix before installation; it halves deployment time.
Cultivate multiple clinical champions and schedule refresher training to embed use.

8. Limitations and future work

This roadmap is intentionally UK-centred, reflecting current MHRA, NICE, and NHS Digital requirements. While the principles generalise, artefacts and approval routes vary internationally. The guide focuses on software-led digital health (including AI decision support) rather than hardware devices or consumer wellness applications, and it is not a substitute for formal regulatory, legal, or information-governance advice for higher-risk deployments. Future work should produce openly accessible templates for key artefacts (e.g., intended-purpose statements, DCB0129 safety cases, DTAC evidence packs, and post-market surveillance plans), specialty-specific extensions, and prospective evaluation of roadmap-guided implementations compared with ad hoc approaches.

9. Conclusion

Digital health adoption in the NHS rarely fails because of weak algorithms; it fails when regulation, evidence generation, information governance, and workflow integration are treated as late-stage hurdles rather than design requirements. This roadmap translates MHRA, NICE, and NHS Digital expectations into executable steps, helping teams define intended use, build the correct regulatory artefacts, generate decision-grade evidence, and plan monitoring from day one. The overarching principle is practical: align regulation, evidence, and clinical workflow continuously through the product lifecycle to achieve safe scaling rather than perpetual piloting.

Acknowledgements

The authors thank colleagues in digital health regulation, NHS clinical governance, and translational research environments for informal discussions that informed the development of this roadmap.

ORCID iD

Anurup Mukherjee https://orcid.org/0009-0007-8742-463X

Ethical considerations

This manuscript does not involve primary data collection, human participants, or patient-level data analysis. Therefore, ethical approval and participant consent were not required.

Author contributions

AM (Anurup Mukherjee) conceptualised the manuscript, conducted the literature synthesis, developed the implementation roadmap, and drafted the manuscript.
SS (Sukhi Shergill) provided senior academic oversight, contributed to critical review of the manuscript, and advised on research positioning and regulatory framing.
CA (Chee Ang) contributed expertise in digital systems and human–computer interaction and critically reviewed the manuscript with respect to implementation and technical governance considerations.
All authors reviewed, revised, and approved the final manuscript.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Data Availability Statement

No datasets were generated or analysed during the current study.

AI declaration

Generative AI tools were used to assist with language refinement and structural editing. All content was critically reviewed and approved by the authors.

Guarantor

Anurup Mukherjee accepts full responsibility for the integrity of the manuscript and the accuracy of the content.

References

1. AI and Digital Regulations Service. *All Developers Guidance for Health and Social Care*. NHS England, 2023. Available at: https://www.digitalregulations.innovation.nhs.uk/regulations-and-guidance-for-developers/all-developers-guidance/(accessed 25 January 2026).
2. Health Education England. *Preparing the Healthcare Workforce to Deliver the Digital Future: The Topol Review – Interim Report*. Health Education England, 2018. Available at: https://www.hee.nhs.uk/our-work/topol-review/ (accessed 25 January 2026).
3. NHS Digital. *Review of digital clinical safety standards: DCB0129 and DCB0160*. NHS Digital, last edited 22 July 2025. Available at: https://digital.nhs.uk/data-and-information/information-standards/governance/latest-activity/standards-and-collections/review-of-digital-clinical-safety-standards-dcb0129-and-dcb0160 (accessed 25 January 2026).
4. National Institute for Health and Care Excellence (NICE). *Evidence standards framework for digital health technologies (ECD7). Corporate document*. NICE, first published 10 December 2018, last updated 9 August 2022. Available at: https://www.nice.org.uk/corporate/ecd7/ (accessed 25 January 2026).
5. Kraus S, Schiavone F, Pluzhnikova A, et al. Digital transformation in healthcare: Analysing the current state-of-research. *J Bus Res* 2021; 123: 557–567. https://doi.org/10.1016/j.jbusres.2020.10.030
6. Al-Haimi B, Ali F and Hujainah F. Digital transformation in healthcare: impact on organisations' strategies, future landscape, and required skills. In: El-Chaarani H, El Dandachi I, El Nemar S, et al. (eds) *Navigating the Intersection of Business, Sustainability and Technology*. Springer, 2023, pp. 61–74. https://doi.org/10.1007/978-981-99-8572-2_3.
7. Abou Hashish EA and Alnajjar H. Digital proficiency: assessing knowledge, attitudes, and skills in digital transformation, health literacy, and artificial intelligence among university nursing students. *BMC Med Educ* 2024; 24: 508. https://doi.org/10.1186/s12909-024-05482-3
8. Gomis-Pastor M, Berdún J, Borrás-Santos A, et al. Clinical validation of digital healthcare solutions: state of the art, challenges, opportunities. *Healthcare (Basel)* 2024; 12(11): 1057. https://doi.org/10.3390/healthcare12111057
9. Medicines and Healthcare products Regulatory Agency (MHRA). *Software and Artificial Intelligence (AI) as a Medical Device: Regulation*. UK Government, last updated 20 April 2025. Available at: https://www.gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device/software-and-artificial-intelligence-ai-as-a-medical-device (accessed 25 January 2026).
10. European Union. *Council Directive 93/42/EEC of 14 June 1993 concerning medical devices*. Official Journal of the European Communities L 169/1, 1993. Available at: https://data.europa.eu/eli/dir/1993/42/oj (accessed 25 January 2026).
11. Medicines and Healthcare products Regulatory Agency (MHRA). *Borderlines with medical devices and other products in Great Britain*. UK Government, last updated 22 October 2024. Available at: https://www.gov.uk/government/publications/borderlines-with-medical-devices/borderlines-with-medical-devices-and-other-products-in-great-britain (accessed 25 January 2026).
12. National Institute for Health and Care Excellence (NICE). *Sleepio to treat insomnia and insomnia symptoms (HealthTech guidance HTG624)*. NICE, published 20 May 2022, migrated December 2025. Available at. https://www.nice.org.uk/guidance/htg624 (accessed 25 January 2026).
13. Roest KN, Lee M, Rees J, et al. Smartphone-based digital point-of-care panel assay with enzymatic catalytic reaction. *Biosensors and Bioelectronics: X* 2024; 19: 100504. https://doi.org/10.1016/j.biosx.2024.100504
14. National Institute for Health and Care Excellence (NICE). *NICE real-world evidence framework for health technologies*. NICE, published 23 June 2022. Available at: https://www.nice.org.uk/corporate/ecd9 (accessed 25 January 2026).
15. NHS Innovation Service. *Developer portal*. NHS England, 2024. Available at: https://innovation.nhs.uk/ (accessed 25 January 2026).
16. Alderman JE, Palmer J, Laws E, et al. Tackling algorithmic bias and promoting transparency in health datasets: the STANDING Together consensus recommendations. *NEJM AI* 2025; 2(1). https://doi.org/10.1056/AIp2401088
17. Liu X, Cruz Rivera S, Moher D, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat Med* 2020; 26: 1364–1374. https://doi.org/10.1038/s41591-020-1034-x
18. Vasey B, Nagendran M, Campbell B, et al. Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *BMJ* 2022; 377: e070904. https://doi.org/10.1136/bmj-2022-070904
19. Collins GS, Moons KGM, Dhiman P, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ* 2024; 385: e078378. https://doi.org/10.1136/bmj-2023-078378

20. Saenz AD, Harned Z, Banerjee O, et al. Autonomous AI systems in the face of liability, regulations and costs. *npj Digital Medicine* 2023; 6(1): 185. https://doi.org/10.1038/s41746-023-00929-1
21. Regulation (EU) 2016/679 (General Data Protection Regulation), Article 6. *gdpr-info.eu*, 2016. Available at: https://gdpr-info.eu/art-6-gdpr/ (accessed 25 January 2026).
22. Regulation (EU) 2016/679 (General Data Protection Regulation), Article 9. *gdpr-info.eu*, 2016. Available at: https://gdpr-info.eu/art-9-gdpr/ (accessed 25 January 2026)
23. Health Research Authority (HRA). *Guidance for applicants to the Confidentiality Advisory Group (CAG)*. Health Research Authority, published 2023, Available at: https://www.hra.nhs.uk/about-us/committees-and-services/confidentiality-advisory-group/guidance-cag-applicants/ (accessed 25 January 2026)
24. Health Innovation Oxford and Thames Valley. *AI in Health and Care Award Phase 4: Final real-world evaluation report of AI stroke imaging decision-support software (e-Stroke)*. Health Innovation Oxford and Thames Valley, 2024. Available at: https://www.healthinnovationoxford.org/wp-content/uploads/2024/04/e-stroke-evaluation-final-report.pdf (accessed 25 January 2026).
25. National Institute for Health and Care Excellence (NICE). *Artificial intelligence (AI)-derived software to help clinical decision making in stroke (HealthTech guidance HTG708)*. NICE, published 23 January 2024. last updated 2 May 2024, migrated December 2025. Available at: https://www.nice.org.uk/guidance/htg708 (accessed 25 January 2026).
26. Mallon D, Fallon M, Blana E, et al. Real-world evaluation of Brainomix e-Stroke software. *Stroke Vasc Neurol* 2024; 9: e002859. https://doi.org/10.1136/svn-2023-002859
27. BioIndustry Association. *Pioneering AI in healthcare: Brainomix member spotlight*. BioIndustry Association, 2025. Available at: https://www.bioindustry.org/resource/pioneering-ai-in-healthcare-brainomix-member-spotlight.html (accessed 25 January 2026).
28. Health Innovation Oxford and Thames Valley. *Case study: Harnessing technology for quicker stroke care*. Health Innovation Oxford & Thames Valley, published18 July 2024. Available at: https://www.healthinnovationoxford.org/covid-19/covid-19-case-studies/ai-technology-speeds-up-stroke-care-and-reduces-costs/case-study-quicker-stroke-care-and-reduced-costs/ (accessed 25 January 2026).
29. Brainomix. *Brainomix awarded prestigious AI grant to demonstrate impact in NHS*. Brainomix, published 8 September 2020. Available at: https://www.brainomix.com/news/brainomix-awarded-prestigious-ai-grant-to-demonstrate-impact-in-nhs/ (accessed 25 January 2026).

Appendix

Acronyms and abbreviations

AI	Artificial Intelligence
AHSN	Academic Health Science Network
CBT	Cognitive Behavioural Therapy
CE	Conformité Européenne (European Conformity)
CQC	Care Quality Commission
CT	Computed Tomography
DCB0129	Data Coordination Board Standard 0129 (Clinical Risk Management Standard)
DG	Diagnostic Guidance
DICOM	Digital Imaging and Communications in Medicine
DPIA	Data Protection Impact Assessment
DTAC	Digital Technology Assessment Criteria
ECG	Electrocardiogram
ESF	Evidence Standards Framework
EU MDR	European Union Medical Device Regulation
FDA	Food and Drug Administration (U.S.)
FHIR	Fast Healthcare Interoperability Resources
GDPR	General data protection regulation
HRA	Health research authority
HSCN	Health and social care network
IG	Information governance
IRAS	Integrated research application system
ISO	International organization for standardization
ISDN	Integrated stroke delivery network
IT	Information Technology

MRI	Magnetic resonance imaging
MDR	Medical device regulation
MHRA	Medicines and healthcare products regulatory agency
MT	Mechanical thrombectomy
NICE	National institute for health and care excellence
NHS	National health service
PACS	Picture archiving and communication system
PMS	Post-market surveillance
QALM	Quality-adjusted life month
QMS	Quality management system
REC	Research ethics committee
UKCA	United Kingdom conformity assessed
UK GDPR	United Kingdom general data protection regulation
UK MDR	United Kingdom medical device regulations.